International Interactions, 28:93–115, 2002
Copyright © 2002 Taylor and Francis
0305-0629/02 $12.00 + .00

STRATEGY AND SELECTION IN INTERNATIONAL RELATIONS

CURTIS S. SIGNORINO

Department of Political Science, University of Rochester, Rochester, New York, USA

(Received for publication September 21, 2001)

It is widely recognized that many of the samples we use for statistical analysis in international politics are the result of some selection process. Not surprisingly, selection models are becoming increasingly popular. At the same time, the role of strategic interaction has begun to play a more important role in statistical analyses. However, it has not been clear how statistical strategic models and selection models relate to each other, or what the effects are of employing one when the other is the more appropriate model. In this article, I 1) clarify why international relations scholars cannot shield themselves from selection bias simply by assuming their results are limited to a given sample; 2) show how recent statistical strategic models relate to traditional selection models and generalize the two sets of models by deriving a correlated strategic model; and 3) examine the effects of misspecifying either correlated errors or strategic interaction. My results indicate that failure to model the strategic interaction produces worse specification error than failure to account for correlated disturbances. In fact, traditional bivariate probit models appear to be superior only when states are almost completely uncertain about each others' preferences.

KEY WORDS: Competitive, conflict, econometric, game theory, international relations, selection, statistical, strategic.

The author would like to thank Tasos Kalondrakis, Jack Levy, Jon Pevehouse, Ahmer Tarar, Kuzey Yilmaz, and especially John Duggan for their comments and very helpful discussions. Support from the National Science Foundation (Grant # SES-9817947) and from the Peter D. Watson Center for Conflict and Cooperation are gratefully acknowledged.

Address correspondence to Curtis S. Signorino, Department of Political Science, University of Rochester, Rochester, NY 14627, USA. E-mail: sign@troi.cc.rochester.edu

INTRODUCTION

Selection models are becoming increasingly popular in international relations. It is widely recognized that many of the samples we have used for our statistical analysis are the result of some selection process: e.g., states selecting to escalate disputes or enter into war. Indeed, if we accept the premise that politics is the result of individuals making choices, then the entire history of political interaction is one of repeated selection. Put another way, history is one grand selection model.

As the econometrics literature has demonstrated for some time, if the selection process is ignored, our inferences may be biased. Hence, international relations researchers have begun employing typical econometric selection models. At the same time, the role of strategic interaction has begun to play a more important role in statistical analyses. However, it has not been clear how statistical strategic models and selection models relate to each other, or what the effects are of employing one when the other is the more appropriate model.

As we will see later in this article, strategic models are inherently selection models, but not of the form traditionally implemented in statistical analyses. They are selection models because the actors select themselves and others into "subsamples" based on their choices. However, the strategic models implemented so far by Signorino (1999, 2000) differ from traditional bivariate probit selection models in that they assume independence of error terms across all decisions. This is by no means a necessary condition of strategic models. Rather, the assumption is made to simplify estimation.

Although strategic models are inherently selection models, in practice the converse is not true. Typical selection models are almost never implemented as strategic, even when the analyzed behavior is assumed to be. Traditional selection models assume the "selection equation" takes a linear $X\beta$ form. However, our theories generally make explicit the functional form of the selection process, and strategic decision making implies that this relationship will be a nonlinear expected utility calculation.

In this article, I more fully examine the relationship of traditional selection models and more recent strategic models. To do so, I present a general selection model for which traditional bivariate probit selection models and more recent statistical strategic models are special cases. I then synthesize the "best" aspects of these models and derive a "correlated strategic model"—one that directly models the strategic interaction but also accounts for correlated private information (or disturbances) as in typical selection models.

Following that, Monte Carlo analysis is used to analyze the misspecification in typical bivariate probit selection models and the (independent) strategic probit models, when our data is generated by the behavior underlying the correlated strategic model. My results indicate that failure to model the expected utility calculations produces worse bias than failure to account for correlated disturbances. In fact, the typical bivariate probit selection model appears to be superior only when 1) the underlying behavior is not strategic or 2) decision makers are so uncertain about each other's preferences that the uncertainty dwarfs all other systematic aspects of the decision making process.

DIFFERENT FORMS OF SAMPLE SELECTION

Before proceeding to the general selection model, it is important to differentiate the problem addressed here from other forms of sample selection. The term "selection" is often used to refer to a number of very different issues concerning political science data. All of these pertain to how a particular sample is generated, or selected from a population. In my experience, however, scholars sometimes confuse the forms of selection and their corresponding effects on our inferences. In part, this seems to be purely because they fall under the "selection" rubric. In part, it also seems to be because the issues are raised in different methodology courses, but sometimes never ultimately related to each other. Depending on the level of the course and whether the emphasis is on research design or data analysis, the selection issues tend to be presented as one of the following: 1) systematically selecting a sample based on an explanatory variable; 2) threats to external validity of inferences; and 3) systematically selecting a sample based on a dependent variable. I will briefly address the first two. The focus of this article is on the last, and especially on the role of strategic interaction in that selection process.

Sample Selection Related to an Explanatory Variable

The advice often given to students studying censoring and truncation in more advanced statistics courses is that there is no problem if the selection mechanism is systematically related only to (or correlated only with) an explanatory variable (King et al., 1994, pp. 137–149; King, 1998, p. 208). Consider a researcher in the natural sciences who is designing an experiment. For a continuous and unbounded explanatory variable, it is not possible to take measurements for every value of that variable. The researcher inevitably must determine some range of values over which to conduct the experiment. Similar constraints are often placed on social scientists conducting field research. Even though they may not be conducting an experiment, time and resources may constrain them to collect data for only a limited range of a key explanatory variable.

It is well known that, under certain (commonly assumed) conditions, neither of these two researchers need fear threats to their inferences due to the selection. Assume that Y is linearly related to X and that the disturbance has conditional mean 0:

$$Y = X\beta + \varepsilon,$$

$$E[\varepsilon|X] = 0.$$

Regardless of whether X is considered fixed or stochastic, the expected value of the dependent variable Y conditional on X is the same as for the true population

$$\begin{aligned} E[Y|X] &= E[X\beta|X] + E[\varepsilon|X] \\ &= X\beta. \end{aligned}$$

Hence, the selection has no effect on our inferences concerning the relationship of X

and Y, even for those values of X that lie outside the range of X in the given sample. In other words, on average, the linear regression we estimate for the truncated sample should produce the same coefficients—and therefore the same inferences—as those of the "true" model that generated the data.[1]

Threats to External Validity, or Generalizing Beyond a Given Sample

In introductory statistics courses, we generally study different research designs, the samples they generate, and the effect on the validity of our inferences. Selection bias is often raised in this context as a threat to the "external validity" of our inferences—i.e., the ability to generalize beyond our sample (Campbell and Stanley, 1963, p. 19). The logic is that if our sample is not representative of the larger population, then even if our inferences are valid for the sample they may not be for the larger population. An example of this might be a study of drug efficacy conducted solely on men. Even if our inferences concerning its effect on men are correct, inferring that the drug will have the same effect on women may not be.

This selection issue actually points to an often overlooked limitation of the aforementioned advice on selecting based on an explanatory variable: it assumes that the functional form of X and Y's relationship is correctly specified and that it will therefore hold outside the range of the given sample. For example, suppose $Y = |X| + \varepsilon$, and consider two samples truncated on X: one sample consisting only of observations where $X > 0$ and another consisting of observations where $X \leq 0$. If we linearly regressed Y on X for each sample, we would arrive at exactly opposite conclusions about the relationship of Y and X. The relationship found for the region $X > 0$ would be correct for that region, but would not hold for "the rest of the population"—i.e., those with values of $X < 0$. Returning to the drug efficacy example, if the drug had the opposite effect on women as it did on men, then there would exist an implicit interaction of drug dosage with another explanatory variable, here gender. Truncating the sample based on the gender explanatory variable (e.g., to males), would yield correct inferences for males, but incorrect inferences when generalized to the rest of the population, here women.

The discussion highlights that selecting on an explanatory variable is in practice not much of a problem. At best, when the functional form is correct, it is no problem at all and we can generalize beyond the values of X in our sample. At worst, the researcher need only state that the conclusions are "conditional on the given sample" and they are shielded from criticisms concerning external validity—of course, this is precisely because they are relinquishing any claims to external validity. Yet, even in this worst case, the typical international relations researcher sees no problem at all, since many of our data sets (i.e., samples) consist of *all* observed cases of some phenomenon over some lengthy period of time—e.g., the outcomes of crises or disputes from 1816 to 1990. With the exception of forecasting the future, there is often no remaining sample to which one might generalize. Even if we limit the sample to some region of the world, valid inferences for that sample alone are not only good enough in the minds of many scholars, but the primary goal itself. For example, finding a valid general explanation "limited" to European crises from the Concert of Europe to the end of the Cold War would be an incredible feat, and well worth the limitation.

Sample Selection Related to the Dependent Variable(s)

With the growing trend for political scientists to take at least introductory statistics courses, most researchers are now familiar with the selection issues concerning generalizing beyond a given sample and with the magic words that shield them from criticism. However, there appears to be some confusion concerning when these words are appropriate. It is not uncommon to hear or read research where the author, who is obviously cognizant that "selection" may be an issue in their data, state that their analysis is "only valid within the context of their given sample." The statement often appears to be made 1) to display an awareness of selection issues, 2) to suggest that there are "limitations" to the author's research, but also 3) to imply that those limitations have no real effect on the inferences, especially since the data almost always consist of all realizations. Unfortunately, what the author is often not aware of is that the selection issue present in their data is of a more nefarious form that is not so easily dispelled with the words "conditional on this sample."

Since Heckman (1976, 1979) it has been well known, at least within economics, that if the sample selection is systematically related to or correlated with the dependent variable, and if this is not incorporated into the statistical model, then resulting inferences will be biased. Consider the same situation as in the above section "Sample Selection Related to an Explanatory Variable":

$$Y = X\beta + \varepsilon$$
$$E[\varepsilon|X] = 0,$$

where Y is linearly related to X, and ε has conditional mean 0. Assume now that we select only those observations of Y such that $Y > 0$. The effect on the conditional expected value of Y is

$$\begin{aligned} E[Y|X, Y > 0] &= E[X\beta|X, Y > 0] + E[\varepsilon|X, Y > 0] \\ &= X\beta + E[\varepsilon|Y > 0] \\ &= X\beta + E[\varepsilon|\varepsilon > -X\beta]. \end{aligned}$$

Notice that the conditional expectation of the error term is no longer 0 (except when $X\beta = \infty$). Therefore, if we run a linear regression, which assumes the error term has conditional mean 0, then our estimate of β will be biased—in fact, it will be attenuated towards 0 here. The greater the truncation, the greater the attenuation, and the more it will appear that X has no effect on Y.[2]

The preceding is just one example of how selection related to the dependent variable can bias inferences. More generally, for any type of dependent variable, whether continuous or discrete, if the selection mechanism is systematically related to or correlated with the dependent variable of the observed sample, failure to account for that selection mechanism will lead to biased inferences. This is where the "selection" issue differs so dramatically from the cases in the previous two sections. There, one could sustain internal validity and avoid the external validity critique simply by claiming that one's inferences were limited to the given sample. However, *when the selection mechanism is correlated with the dependent variable and the researcher does not correctly model it, then the inferences are biased, even for the given sample.*

STRATEGY AND CORRELATION IN SAMPLE SELECTION

The previous section was not meant to suggest that all international relations scholars (and political scientists more generally) are unaware of selection models where sample selection is correlated with the dependent variable. Rather, it was intended to highlight that there are a number of forms of sample selection, and to clarify that they are not all addressed in the same way. Indeed, selection models of the last form are increasingly appearing in international relations research (for example, Smith, 1996; Reed, 2000). In addition, scholars have become increasingly aware of the role that strategic interaction plays in our statistical analyses.

These parallel (statistical) modeling efforts raise the question of how traditional selection models relate to the more recent statistical strategic models. Traditional selection models are often employed in analyses of strategic behavior by states—e.g., of disputes and crises. Do these selection models capture the same thing as the strategic models? Or are they different models entirely? When would one approach versus the other be more appropriate? What is the effect of using one when the other is the correct model? These are the questions I address here.

Although much of the analysis will be conducted with reference to a general, abstract "selection" model, it may be helpful to motivate it with an example. Figure

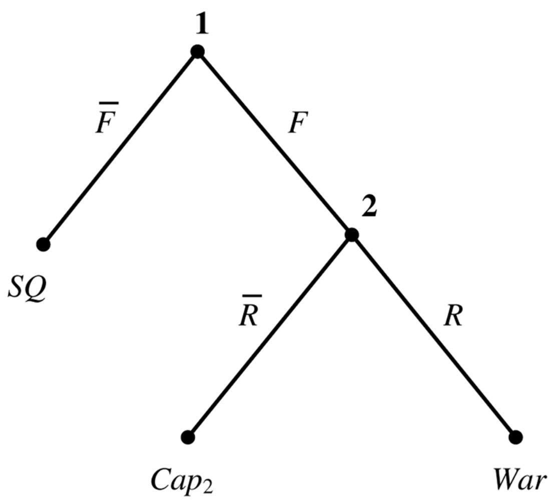

Figure 1. Deterrence model as a selection model. Strategic models are inherently selection models. Each player chooses—or "selects"—which action to take and, therefore, selects the rest of the players into a particular subsample. In the above deterrence model, state 1 must choose whether to attack or not. If 1 attacks, she selects state 2 into the sample of "having been attacked." State 2 must then choose between resisting or not.

1 depicts a simple deterrence situation, where state 1 must choose between attacking (F) or not attacking ($\bar{F}$) state 2. If state 1 does not attack, the status quo is maintained. If state 1 attacks state 2, then state 2 must decide whether to resist (R) or not resist ($\bar{R}$), leading to war or capitulation by 2, respectively.

Although it has been stated in different ways before (Fearon, 1994; Smith, 1999), it bears repeating more precisely in this context: *strategic models are inherently selection models.* The choices made by individuals or states force the world (or at least some set of states) down another path in the "game." In other words, each state's choice self-selects that state into a particular subsample. The "path play" in strategic situations results in the observed sample. Everything else off the path of play can be thought of as the counterfactuals that comprise the rest of the strategic model.

Consider again the deterrence situation in Figure 1. In the terminology of the selection literature, state 1's decision to attack "selects" state 2 into the subsample where state 2 must choose whether to resist or not. Suppose the only data available for our analysis consisted of state 2's decision of whether to resist or not once it was attacked. When we discussed selection effects in the section "Sample Selection Related to the Dependent Variables", the sample was truncated based on values of the dependent variable. Here, the sample of state 2's actions is selected based on state 1's decision. Even though these reflect two different "dependent variables" (i.e., state 1's and state 2's decisions), the logic (and math) of the selection effect is the same. If the dependent variable in our sample (state 2's decision) is correlated with the selection mechanism (state 1's decision), then failure to account for the selection process will result in biased inferences. It is actually immaterial whether our data consists only of state 2's decision or whether we have data on both state 1's and state 2's decisions. If state 1's decision is correlated with 2's decision and we fail to model that, then our inferences will be biased.

Of course, the traditional "selection" critique can be extended backwards infinitely many times. Even if we correctly model the correlation in decisions by states 1 and 2 in the deterrence game, there was some decision prior to state 1's that led to the deterrence game. If that decision was correlated with either 1's or 2's choices in the deterrence game, then a selection problem again arises. We could go back a step before that, and a step before that, and a step before that, infinitely many times—or at least until the dawn of social interaction. To suggest that all of history is one grand selection model is not to say that statistical analysis is impossible. The bias induced by ignoring the correlated errors is a matter of degree, not of kind—the extent of the bias depends on the degree of the correlation among the decisions. As we will later see, that may not be as troublesome as we might expect.

Finally, there is the issue of the strategic interaction between states 1 and 2 in the deterrence model. For the moment, let us set aside the issue of correlated decisions. What is the strategic aspect to the deterrence model? It would seem natural to think of state 2's decision as being based on the difference in its utility for war versus capitulation. In the model, state 2 does not have to condition its behavior on state 1's, so state 2's decision is actually not strategic in any way. Therefore, if we only had data on state 2's choices, and if those choices were uncorrelated with state 1's, then neither a strategic nor a selection model would be required. However, state 1's decision *is* strategic. It would also seem natural to assume state 1's decision is based on the difference in its utility for attacking versus not attacking. Although its utility for not attacking is just its utility for the status quo (see Figure 1), its utility for attacking depends on what state 2 will do in response to being attacked. State 1 must therefore condition its behavior on what it expects state 2 will do if attacked. As Signorino (1999) and Signorino and Yilmaz (2000) have shown, how we characterize that conditioning has important implications for the statistical model. The expected utility calculation implies a particular functional form for state 1's decision as it relates to regressors in the model. Failure to account for that will lead to biased inferences concerning state 1's behavior. Note that this is completely separate from the issue of whether state 1's and state 2's decisions are correlated. We could certainly include both in the model. Moreover, as we will see shortly, the correlation does not have to be separate from strategic considerations.

A General Selection Model

To examine the issues of strategy versus correlation in selection models, it will be useful to set the specific models developed in this paper in the context of a more general model. In that regard, each of the three specific models I analyze in this paper can be viewed as a special case of the general selection model depicted in Figure 2. Selection models of this form have been used to analyze granting and defaulting on bank loans (cf. Greene, 2000), market entry by firms, alliance reliability (Smith, 1996), crisis escalation (Reed 2000), extended deterrence (Signorino and Tarar, 2001), and voter turnout (Dubin and Rivers, 1989; Timpone, 1998).

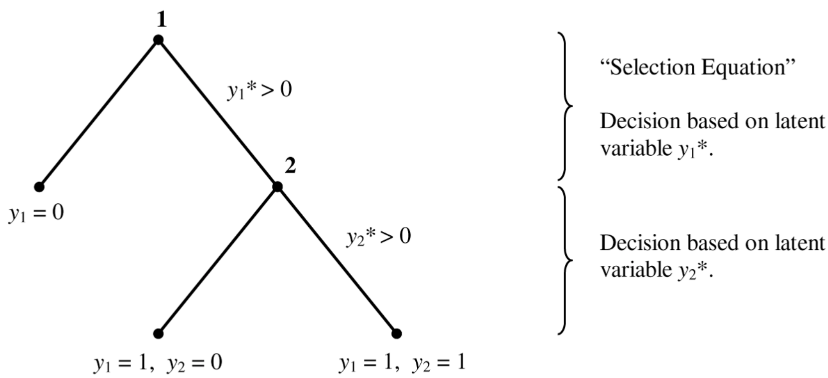

Figure 2 General selection model. All three of the selection models presented in this paper are special cases of the above general model, which consists of a sequence of two choices. The decisions at stages 1 and 2 are based on the latent (or unobservable) variables y_1^* and y_2^*. We observe only y_1 and y_2.

The general model in Figure 2 consists of a sequence of two choices. In the first stage, the decision maker chooses one option versus another depending on the value of some underlying latent variable y_1^*. The latent variable is usually assumed to be a function of regressors and is often referred to as the "selection equation," since it determines whether the observations in the second stage are observed. If our sample is truncated—i.e., consists only of those second stage observations—then y_1^* determines how that sample was selected. The left option, $y_1 = 0$, represents the case where $y_1^* \leq 0$. The right option, $y_1 = 1$, represents the case where $y_1^* > 0$. In the context of our prior deterrence model, $y_1 = 0$ would represent state 1 not attacking and $y_1 = 1$ would represent attacking. We tend to associate y_1^* with state 1's propensity to choose $y_1 = 1$. From a rational choice perspective, we would more specifically think of y_1^* as being the difference between state 1's utilities for the $y_1 = 1$ option and the $y_1 = 0$ option. Because we cannot perfectly know or observe state 1's utility or propensity y_1^*, it is considered unobservable, or latent. All we observe is whether state 1 chooses $y_1 = 1$ versus $y_1 = 0$.

If the decision maker chooses $y_1 = 1$, then there is a second choice over two options, which depends on the latent variable y_2^*. Similar to the prior stage, the left option, $y_2 = 0$, represents the case where $y_2^* \leq 0$. The right option, $y_2 = 1$, represents the case where $y_2^* > 0$. Again, in the context of the deterrence model, $y_2 = 0$ would represent state 2 not resisting and $y_2 = 1$ would represent resisting. Again, we observe y_2, but not the true propensity or utility y_2^*. Note that in traditional selection models the decision maker in this second stage may be the same as in the first, or may be a different decision maker altogether. In our deterrence model, we had a different attacker and defender. However, selection models have been constructed of the same individual making a sequence of decisions—of self-selecting into a sample. Similarly, selection models have been constructed where "dyads" select into crises and then decide whether to go to war (Reed, 2000).

The full probability model for Figure 2 is determined by the assumptions we make 1) about whether y_1^* and y_2^* are correlated and 2) about the functional form of y_1^* and y_2^*. Data may be available on all outcomes of the model, or only for state 2's actions, in which case the sample is considered "truncated." Regardless of the data at our disposal, failing to account for correlation or strategy will bias our inferences. For the purposes of this paper, it does not matter which type of data is assumed given, so long as the probabilities of the outcomes in the data are consistent with the assumptions of the underlying model.[3]

For simplicity, suppose data is available on all actions, then we might denote the outcomes of the model as shown in Figure 2: ($y_1 = 0$), ($y_1 = 1$; $y_2 = 0$), and ($y_1 = 1$; $y_2 = 1$). In that case, statistical analysis could be undertaken via maximum likelihood using the likelihood function

$$L = \prod_{i}^{N} \Pr(y_1 = 0)^{1-y_1} \Pr(y_1 = 1, y_2 = 0)^{y_1(1-y_2)} \Pr(y_1 = 1, y_2 = 1)^{y_1 y_2},$$

where the i index is dropped for notational convenience. This is the general likelihood for all the "selection" models in this paper. Where they differ is in the specification of the probabilities, which in turn depend on the assumptions concerning whether y_1^* and y_2^* are correlated, and on the functional form of y_1^* and y_2^*. I now turn to each of the three models.

Correlated Errors and Selection

The most commonly implemented selection model for binary data dates back to the early 1980s (Van de Ven and Praag, 1981; Meng and Schmidt, 1985; Dubin and Rivers, 1989). Its popularity among political scientists is due in part to the availability of such data and also to the fact that this method is now a standard feature in most statistical computer programs.

Assume the general choice structure as in Figure 2. Further, assume that the latent variables are linear:

$$y_1^* = X_1\beta_1 + \varepsilon_1, \tag{1}$$

$$y_2^* = X_2\beta_2 + \varepsilon_2, \tag{2}$$

where X_1 and X_2 are sets of regressors that affect the propensity for choosing $y_1 = 1$ and $y_2 = 1$, respectively, and $(\varepsilon_1, \varepsilon_2)$ are unobserved "errors," distributed bivariate normal with mean 0, variance 1, and correlation ρ. It is easy to show that these assumptions lead to the following bivariate probit probabilities:

$$\Pr(y_1 = 0) = \Phi_n(-X_1\beta_1), \tag{3}$$

$$\Pr(y_1 = 1, y_2 = 0) = \Phi_{bn}(X_1\beta_1, -X_2\beta_2, -\rho), \tag{4}$$

$$\Pr(y_1 = 1, y_2 = 1) = \Phi_{bn}(X_1\beta_1, X_2\beta_2, \rho), \tag{5}$$

where $\Phi_n(\cdot)$ is the cumulative standardized univariate normal distribution and $\Phi_{bn}(\cdot)$ is the cumulative standardized bivariate normal distribution.

I will later derive another selection model that uses the bivariate probit distribution. However, for purely historical reasons, I will refer to the model just presented as the "bivariate probit model"—and the reader should interpret this as the bivariate probit model with linear latent variables. In other words, the typical selection model for binary data.

Strategic Model with Independent Errors

More recently, Signorino (1999) and Signorino and Yilmaz (2000) have shown that when analyzing data representing strategic behavior, failure to statistically model that strategic interaction results in a misspecified model, leading to biased estimates. Signorino (2000) provides a number of methods for developing statistical strategic models. Signorino and Tarar (2001) use these methods to analyze extended immediate deterrence and find that the statistical strategic model does very well at explaining the outcomes.

The model in Figure 3 is identical to the deterrence model in Figure 1, but will be used henceforth as the referent model to maintain notational consistency across the different selection models. Here, we assume state 1 has a choice between actions a_1 and a_2. If state 1 chooses action a_2, then state 2 must choose between actions a_3 and a_4. Denote the outcomes of this interaction as Y_1, Y_3, and Y_4. Further assume that each

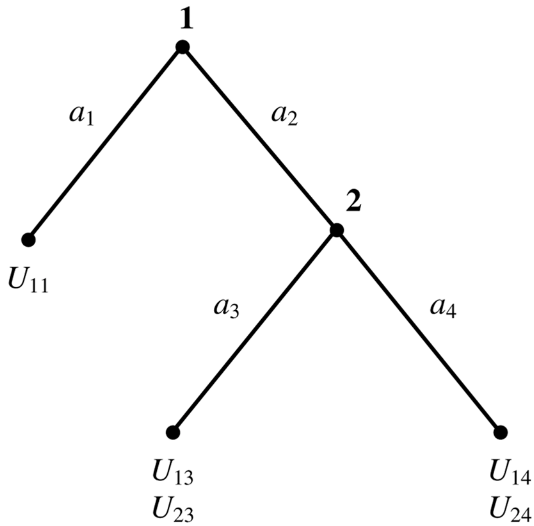

Figure 3. Referent strategic selection model. State 1 must choose between actions a_1 and a_2. If state 1 chooses a_2, then state 2 must choose between a_3 and a_4. U_{jk} is state j's observable utility for outcome Y_k. The strategic model can be reformulated equivalently as a selection model.

player's true utility for actions is composed of an observable component and a private component. U_{jk} corresponds to state j's observable utility for outcome Y_k. As in Signorino (2000), I will assume the private components are independent of each other.

Denoting $p_3 = \Pr(y_2^* \leq 0)$ and $p_4 = \Pr(y_2^* > 0)$, the strategic model in Figure 3 can be written equivalently as a selection model, where

$$y_1^* = p_3\, U_{13} + p_4\, U_{14} - U_{11} + \varepsilon_1, \tag{6}$$

$$y_2^* = U_{24} - U_{23} + \varepsilon_2. \tag{7}$$

Notice that y_1^* reflects the expected utility calculation one would expect in a strategic model. We will observe state 1 choosing $y_1 = 1$ when $y_1^* > 0$—i.e., when its expected utility for $y_1 = 1$ is greater than its utility for $y_1 = 0$ (which is U_{11} here). State 1's uncertainty about state 2's decision is represented by the probabilities $\Pr(y_2^* \leq 0)$ and $\Pr(y_2^* > 0)$ in the expected utility calculation. Similarly, state 2's decision is based on a comparison of its utility for $y_2 = 1$ versus $y_2 = 0$.

Assuming ε_1 and ε_2 are each independent and identically distributed N(0,1), then the above assumptions result in the strategic probit probabilities (Signorino, 2000)

$$\Pr(y_1 = 0) = \Phi_n[U_{11} - (p_3\, U_{13} + p_4\, U_{14})], \tag{8}$$

$$\Pr(y_1 = 1, y_2 = 0) = \Phi_n[p_3\, U_{13} + p_4\, U_{14} - U_{11}]\, \Phi_n[U_{23} - U_{24}], \quad (9)$$

$$\Pr(y_1 = 1, y_2 = 1) = \Phi_n[p_3\, U_{13} + p_4\, U_{14} - U_{11}]\, \Phi_n[U_{24} - U_{23}], \quad (10)$$

where

$$p_3 = \Phi_n[U_{23} - U_{24}], \quad (11)$$

$$p_4 = \Phi_n[U_{24} - U_{23}]. \quad (12)$$

Notice that because of the independence assumption, the joint probabilities for the outcomes reduce to the product of the choice probabilities along the path. For example, the first term in Equation 10 is the probability that state 1 chooses $y_1 = 1$, and the second term is the probability that state 2 chooses $y_2 = 1$. Note too, that, in contrast to the traditional bivariate probit selection model, the strategic model specifically incorporates the expected utility calculation by state 1.

The one area in which this strategic selection model is deficient relative to the traditional bivariate probit selection is in the assumption that errors or private information are independent. As implemented here, the strategic model does not capture correlation in the disturbances associated with each state's decision. Substantively, this implies that states learn nothing about each other's incentives when viewing their own private information. Of course, it is not yet clear which misspecification—functional form versus correlated errors—is worse. That will be the subject of Section 4. The tradeoff is not necessary, however. We could very well specify a strategic selection model that allows for correlated errors—and that is exactly what we will do next. To differentiate the present model from one with correlated errors, I will henceforth refer to this model as the "independent strategic model," and will refer to the model of the next section as the "correlated strategic model."

Strategic Model with Correlated Errors

To derive a strategic selection model that allows states to condition on their own private information—i.e., allows for correlated errors in the decision equations—we make the exact same assumptions as for the independent strategic model, with the exception that $(\varepsilon_1; \varepsilon_2)$ are now assumed to be jointly distributed bivariate normal with mean 0, variance 1, and correlation ρ. In short, we are synthesizing the strategic probit model and the traditional bivariate probit model.

Although the derivation would appear to be relatively straightforward, there is a new twist. In the strategic models of Signorino (1999, 2000), because the errors are assumed to be independent, the assumption is also made that the analyst and players all share the same information concerning the errors. For example, refer to Figure 3. If the private information for state 1 and 2 are independent, then state 1's draw of an "error" provides it no information concerning state 2's disturbance. Because the analyst is not omniscient, we assume she shares the same information (or assumption concerning the distribution) for state 2's disturbance. The analyst and state 1 therefore estimate (or have the same belief) about state 2's probability of choosing one option versus another.

When we now make the assumption that state 1 and 2 have correlated disturbances—and both know this—then they have information concerning the other's probabilities that the analyst does not. In viewing her draw, state 1 knows something about the distribution of state 2's draw. When state 1 forms a belief about state 2's probability of choosing $y_2 = 1$ versus $y_2 = 0$, state 1 conditions on its own draw. However, because the analyst does not see state 1's draw, the analyst cannot condition similarly. The fact that the analyst and players no longer share the same beliefs (or estimates) of the probabilities must now be factored into the statistical model.

Allowing for state 1 to condition on its draw, the correlated strategic model can also be written as a special case of the general selection model, where

$$y_1^* = \Pr(y_2^* \leq 0|\varepsilon_1)\, U_{13} + \Pr(y_2^* > 0|\varepsilon_1)\, U_{14} - U_{11} + \varepsilon_1 \tag{13}$$

$$y_2^* = U_{24} - U_{23} + \varepsilon_2 \tag{14}$$

Because $(\varepsilon_1,\varepsilon_2)$ are distributed bivariate normal, the resulting outcome probabilities will take the same bivariate probit form as for the typical selection model in the section "Correlated Errors and Selection," but with the expected utility calculation reflected in the arguments. The correlated strategic selection model is then

$$\Pr(y_1 = 0) = \Phi_n[-\varepsilon_1^\circ], \tag{15}$$

$$\Pr(y_1 = 1, y_2 = 0) = \Phi_{bn}[\varepsilon_1^\circ, U_{23} - U_{24}, -\rho], \tag{16}$$

$$\Pr(y_1 = 1, y_2 = 1) = \Phi_{bn}[\varepsilon_1^\circ, U_{24} - U_{23}, \rho], \tag{17}$$

where ε_1° solves

$$\Pr(y_2^* \leq 0|\varepsilon_1)\, U_{13} + \Pr(y_2^* > 0|\varepsilon_1)\, U_{14} + \varepsilon_1 = U_{11} \tag{18}$$

for ε_1.[4] Notice that when $\rho = 0$, the conditional probabilities reduce to marginal probabilities and the bivariate normal distributions reduce to the product of the univariate distributions. The independent strategic model is therefore a special case of the correlated strategic model. However, the typical bivariate selection model is not a special case of the correlated strategic model, because it does not capture the strategic behavior a la the expected utility calculations.

As specified, the correlated strategic model incorporates the "best" features of the typical selection model and the independent strategic model: it directly models the strategic interaction of the states and it allows for correlated error terms (or private information). If one believed the interaction in one's data was represented by this form of strategic selection model, then Equations (15–17) would be used as the basis for maximum likelihood estimation. One would specify the utilities in terms of regressors—e.g., $U_{jk} = X_{jk}\beta_{jk}$—and then estimate the coefficient parameters β and the correlation ρ.

Given the proliferation of the bivariate probit selection model and given the introduction of Signorino's (1999, 2000) strategic models, we might ask how bad the misspecification would be if the data were generated by the behavior underlying the correlated strategic model but then analyzed with the bivariate probit and independent strategic models. This would give us some sense as to how concerned we should be about previous studies using these techniques. It might also be the case that if the misspecification is not severe, then researchers could very well use these previously existing techniques without jeopardizing their results. From a practical perspective, this would be nice, since solving Equation (18) for ε_1 is not trivial and involves numerically finding the roots of the equation. When this is implemented in maximum likelihood estimation, it becomes time consuming, since it has to be done for every observation in a data set, and then repeatedly until the maximum likelihood estimate is found. If the average user had a simpler and faster method, one that did not affect their inferences too much, that might be a reasonable and attractive option.

MONTE CARLO ANALYSIS OF STRATEGY VERSUS CORRELATION

The purpose of this section is to analyze how well the typical bivariate probit selection model and the independent strategic model will perform when used to analyze data resulting from the strategic behavior underlying the correlated strategic model. My rationale for employing the correlated strategic model is that it combines the functional form of the strategic model with the correlated errors of the traditional selection model. It therefore allows us to assess how much functional form matters for traditional selection models and how much correlated errors matter for independent strategic models, without privileging one versus the other.

Before analyzing the misspecification of the two models, however, we should verify that the correlated strategic model based on Equations (15–17) correctly recovers the true parameters that generate the data. To do this, I conducted a Monte Carlo analysis, based on the strategic model in Figure 3, and with the utilities specified as

$$U_{11} = 0, \tag{19}$$

$$U_{13} = X_{13}\beta_{13}, \tag{20}$$

$$U_{14} = X_{14}\beta_{14}, \tag{21}$$

$$U_{23} = 0, \tag{22}$$

$$U_{24} = X_{24}\beta_{24}, \tag{23}$$

where the X_{jk} consist of a single regressor, and where $\beta_{13} = \beta_{14} = \beta_{24} = 1$. This specification keeps the model as simple as possible, but allows the states to have different preference orderings over all outcomes. The X_{jk} were randomly generated using a uniform distribution over [-1,1]. The $(\varepsilon_1,\varepsilon_2)$ were randomly generated bivariate normal with mean 0, variance 1, and correlation $\rho = .8$. For each iteration, a total of 2000 observations were generated using the underlying behavioral (i.e., choice-based) model of the correlated strategic model. The regressors and outcomes of the behavior (i.e., y_1 and y_2) were then saved to a data set and maximum likelihood estimation was conducted using the correlated strategic statistical model based on Equations (15–17). The resulting parameter estimates for the βs and for ρ were then also saved. This process was repeated for a total of 350 Monte Carlo iterations, producing 350 estimates of β_{13}, β_{14}, β_{24}, and ρ. The averages of $\hat{\beta}_{13}$, $\hat{\beta}_{14}$, and $\hat{\beta}_{24}$ were .9966, 1.0016, and 1.0066, respectively. The average value of $\hat{\rho}$ was .8028. The correlated strategic model therefore correctly recovers the true parameter estimates in expectation.

To analyze the misspecification of these models, I conducted a number of Monte Carlo analyses as above, but varying the model parameters. In each Monte Carlo, the β_{jk} were all set to one, a sample size of N = 2000 was generated each iteration, and 500 iterations were conducted to provide the densities of the parameter estimates. For a given iteration, both the bivariate probit and independent strategic models were estimated on the same data, and the estimates were saved. Two elements were allowed to vary in the research design. First, Monte Carlos were run for values of $\rho = \{0, .5, .9\}$, signifying no correlation, moderate correlation, and high correlation, respectively. Second, I wanted to assess the impact of the model misspecification when either 1) the utilities "mattered" more or 2) the error term mattered more. To examine this, I assumed in each Monte Carlo that the variance of the ε remained at one (generated by a standardized bivariate normal distribution). However, I set the range of the X_{jk} such that the ratio of the utility variance to the error variance would be $V(U)/V(\varepsilon) = \{1/10, 1/3, 1, 3, 10\}$. In other words, $V(U)/V(\varepsilon) = 1/10$ represents a case where the variance of the error term is 10 times as large as that of each observable utility U_{jk}. In other words, this means that the true utilities of the states are dominated by private information (the disturbances), rather than by the observable components—i.e., the states are highly uncertain about each other's true utilities. In contrast, $V(U)/V(\varepsilon) = 10$ represents the case where there is great certainty on the part of the states—the variance of the observable utility is 10 times larger than that of the error term.

Monte Carlo analyses were conducted for each pairing of the correlation levels $\rho = \{0, .5, .9\}$ and the variance ratios $V(U)/V(\varepsilon) = \{1/10, 1/3, 1, 3, 10\}$. It would seem reasonable to expect that 1) the larger the correlation level, the worse the independent strategic model will perform and 2) the less the error term matters, the worse the bivariate probit model will perform and the better the independent strategic model will perform.

Figures 4a and b display the expected estimates of β_{13}, β_{24} and ρ, when using a typical bivariate probit selection model to analyze strategic behavior. Figure 4a shows that the mean estimate of β_{24} is unaffected by the level of correlation or the ratio of the utility and error variances. The expected estimate of β_{13} is also unaffected by the correlation level. However, $\hat{\beta}_{13}$ and $\hat{\beta}_{14}$ are attenuated as the utility variance grows relative to the error variance. Figure 4b shows that when $\rho > 0$, the larger the utility variance to the error variance, the greater the attenuation in the expected estimate of ρ. The intuition behind this is that when the true utilities are dominated by the error term rather than the observed utilities (e.g., when $V(U)/V(\varepsilon)$ is small), it is the bivariate distribution and the correlation that drive the probabilities rather than the functional form of the expected utilities. Hence, the estimate of ρ tends to be on target. In contrast, when the error term matters much less in the true utilities (e.g., when $V(U)/V(\varepsilon)$ is large), then the functional form of the expected utility has a much larger

(a)

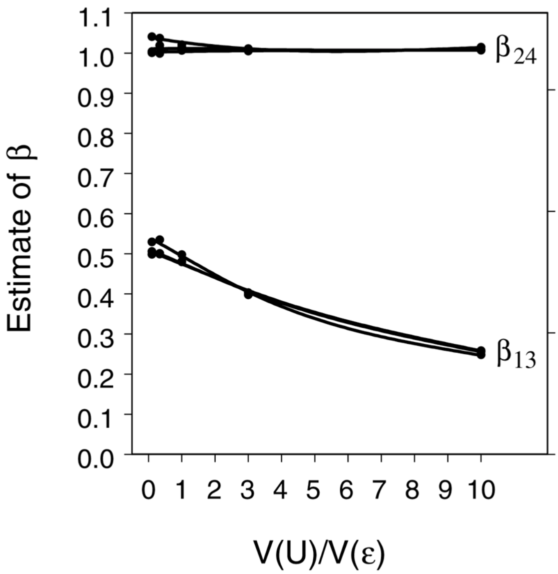

(b)

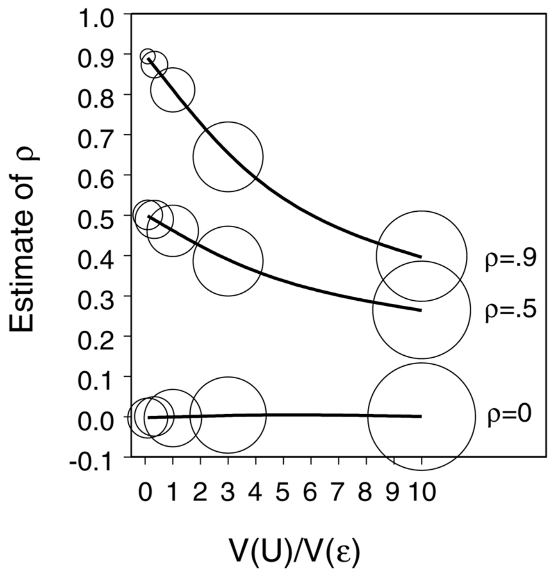

Figure 4. When bivariate probit is used to estimate correlated strategic behavior. Figure (a) shows that the average bivariate probit estimate of β_{24} is unaffected by the correlation level or variance ratio. However, the estimates of β_{13} and ρ are attenuated as the variance of the error term decreases. Moreover, as Figure (b) shows, when the error term matters less, the estimates of ρ are very inefficient.

(a)

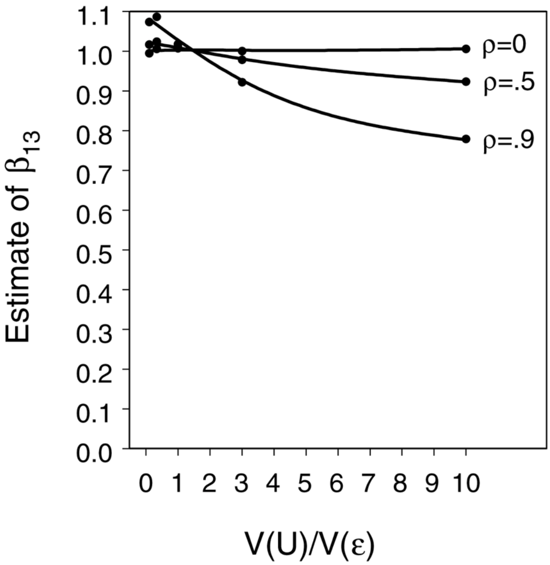

(b)

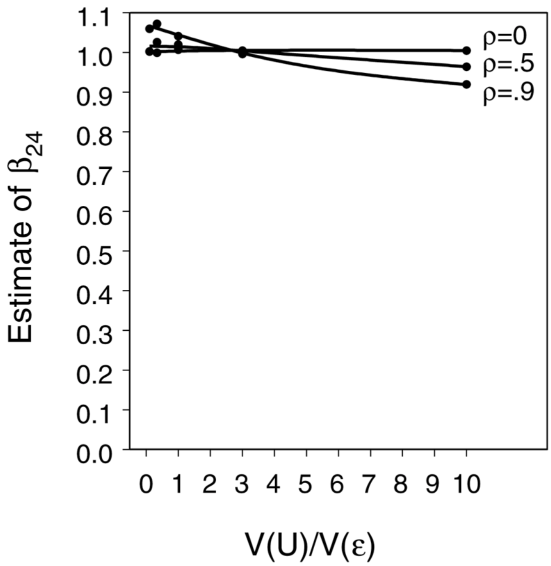

Figure 5. When the independent strategic model is used to estimate correlated strategic behavior. (a) Estimates of β_{13} and β_{24} are generally attenuated as the variance ratio increases. (b) When $\rho = 0$, there is no effect, because the correlated strategic model is equivalent to the independent strategic model.

impact. Since the typical bivariate probit selection model does not implement the expected utility calculation, the estimate of ρ is biased. Finally, the circles around the means depict the relative sizes of the variance of $\hat{\rho}$ at those points. The radius of each circle is exactly one standard deviation, so the vertical diameter about the mean represents a two standard deviation interval. Notice that as $V(U)/V(\varepsilon)$ increases, the variance about the mean increases quite a bit. This suggests that as $V(U)/V(\varepsilon)$ increases, not only is $\hat{\rho}$ biased, it is extremely inefficient.

We turn now to the case where an independent strategic model is used to analyze data generated by the correlated strategic model. Figures 5a and b indicate that both the correlation level and the variance ratio affect the estimates of β_{13} and β_{24}. As expected, in both cases, the larger the correlation and the larger the error variance, the greater the attenuation in the estimates. That said, the extent of the bias is surprisingly small, even when the correlation is extremely high.

Analyzing the bias or efficiency of the parameter estimates only tells us so much about how our inferences would be affected by the model. Of more importance is how the misspecification affects the quantities of interest we would generally use to interpret the model. In discrete choice models such as this, the most common quantity of interest is simply the predicted probability of a particular action or outcome. The problem here is that we would like some concise way of summarizing how close the two models are to the "truth" in general. For example, does the bivariate probit's probability of War tend to match the true probability of War? Or does it generally diverge greatly?

To assess this, I implement a measure of average distance between the true and predicted probabilities. Consider the true correlated strategic model. For a given set of utilities (defined in terms of the true parameters and values for the regressors), we can calculate a "true" predicted probability for all three of the outcomes. Similarly, given a set of average estimates for the bivariate probit model, we can calculate the average bivariate probit model's predicted probability for the same values of the regressors. The extent to which the bivariate probit prediction differs from the true probabilities tells us something about how well the former approximates the latter. To measure this difference, I use the Euclidean distance between the two sets of probabilities. Consider the two sets of probabilities as two points in a three dimensional probability simplex. Then the Euclidean distance is simply the distance between these two points. The farthest two points can be is $\sqrt{2} \approx 1.414$—e.g., if one model predicted probabilities (1,0,0) and the other predicted (0,1,0).

Those are the steps for only one possible set of values for the regressors. Of course, I could pick or choose those to give any impression I wanted. Rather than doing so, we might prefer to know how well one model approximates another on average. In that regard, I constructed a three-dimensional grid of 125,000 points, evenly distributed throughout the entire regressor space (i.e., 50 points along each dimension). I then calculated the distance between the two models' predicted probabilities for each of the 125,000 points in the grid. The average of these 125,000 distances can then be used as an indicator of how well the two models approximate each other on average throughout the regressor space. In general, the smaller the average distance, the better one model approximates another. The larger the average distance, the worse one model approximates another. This procedure was carried out for both the bivariate probit and

(a)

(b)

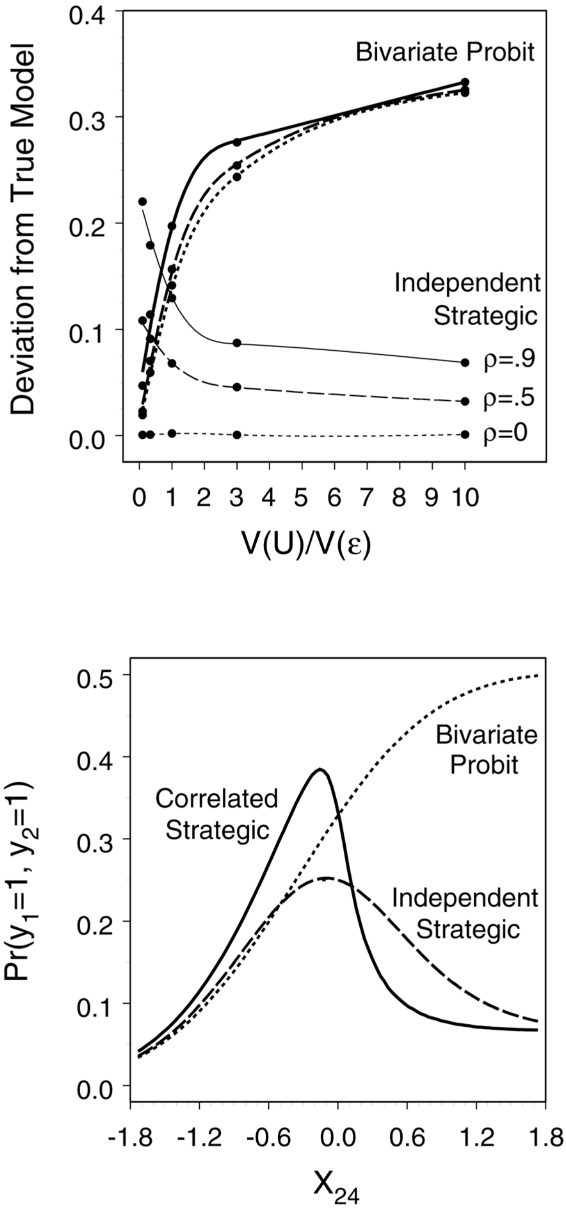

Figure 6. Comparison of bivariate probit and independent strategic models. Figure (a) shows that the independent strategic model is a better approximation when the errors matter less or are uncorrelated. The bivariate probit model is a better approximation when the true utilities are almost entirely dominated by the error term—i.e., provide the states with no information concerning each other's preferences. Figure (b) demonstrates that even when the independent strategic model diverges from the (true) correlated strategic model, it can still capture the nonmonotonic expected utility calculations, whereas the bivariate probit does not.

independent strategic models, assuming the correlated strategic model as the true model.

Figure 6a displays the extent to which the bivariate probit and independent strategic models diverge on average from the correlated strategic model. When $\rho = 0$, the correlated strategic model is equivalent to the independent strategic model, so the average deviation is 0 across the variance ratio. The greater the correlation, the more the independent strategic model deviates from the correlated strategic model. Moreover, the independent model deviates more when the error component dominates the true utility. As the error term decreases in size relative to the observable utility, the correlation matters less and the independent model deviates less. The exact opposite is true for the bivariate probit model. There, the correlation level does not appear to affect the extent to which the bivariate model deviates from the correlated strategic model. Of primary importance is the variance ratio. When the error term dominates the true utility, the more the correlated bivariate distribution matters relative to the expected utility structure and the less the bivariate probit model deviates on average. Notice, however, that even when $\rho = 0$ and the variance ratio favors the error term, the bivariate probit still is not equivalent to the correlated strategic model, since it does not incorporate the expected utility calculation by state 1. As the variance ratio increases in favor of the observable utility, the expected utility structure matters more relative to the error terms and their correlation, leading to greater deviation from the correlated strategic model.

Can we say whether the bivariate probit model or independent strategic model is generally the lesser evil when the data is correlated strategic? Clearly there are situations where one does better than the other on average in terms of the predicted probabilities. Neither completely dominates the other in all situations. However, we can make a few general observations. First, there appears to be greater potential for damage by the bivariate probit model—i.e., its worst is worse than the independent strategic model's worst.

Second, it is only in the case where the error variance is 10 times that of the observable utilities (i.e., the points farthest to the left) that the bivariate probit does significantly better than the independent strategic model, but that difference is less than the reverse situation (when $V(U)/V(\varepsilon) = 10$). Even when $V(U)/V(\varepsilon) = 1/3$, the independent strategic model does not fare much worse than the bivariate probit model. Moreover, for all other variance ratios—across all correlation levels—the independent strategic model is a better approximation than the bivariate probit. Consider the implications of this for the international relations scholar using typical selection models. It suggests that traditional selection models would be the better option only 1) when there is no strategic interaction in the first place or 2) when the strategic aspect to the model is swamped by the error term—i.e., when the correlation is very high and when there is almost no information in the utilities. Most international relations researchers employing selection models are usually interested in finding general "processes" by which states escalate crises or enter into war. Indeed, the rationale for using the selection models is to better model the process itself. Ironically, the results here suggest that the typical selection model is most appropriate when the "process" essentially does not matter.

Third, the independent strategic model has some hope of being the correct model—e.g., when ρ truly is zero—whereas the bivariate probit model does not. The typical bivariate probit implementation does not capture nonmonotonic relationships produced by expected utility calculations. This is due to the specification of y_1^*, which assumes a linear $X\beta$ functional form. As an example, consider Figure 6b. The figure shows the probability of ($y_1 = 1, y_2 = 1$) as a function of X_{24}, for the case where $X_{13} = 1.5$, $X_{14} = -1.5$, $\rho = .5$ and $V(U)/V(\varepsilon) = 1$—i.e., when there is moderate correlation and the variance of the observable utilities is equal to the error variance. This would not seem to be an ideal case for the independent strategic model. Yet, as Figure 6b shows, even though there is some divergence between the correlated and independent strategic models, the independent model at least captures the nonmonotonic equilibrium dynamic, while the bivariate probit does not—indeed cannot.

CONCLUDING REMARKS

I have attempted to accomplish three tasks in this article: 1) to clarify that international relations scholars cannot shield themselves from selection bias simply by assuming their results are limited to a given sample; 2) to show how recent statistical strategic models relate to traditional selection models and to generalize the two sets of models by deriving a correlated strategic model; and 3) to examine the effects of misspecifying either the correlated errors or the strategic functional form in this class of selection models.

Clearly, misspecifying either correlated errors or the strategic functional form will result in biased inferences. But what should international relations scholars do with that information? The correlated strategic model derived here is intuitive and can be implemented in many standard statistics packages (e.g., Gauss or Stata). However, it is still not trivial and is certainly more computationally intensive than current "canned" methods.[5] The analysis conducted here suggests that if one has to "get by" with one canned method versus the other, the better option would be to use a method that models the strategic interaction but forgoes the correlated errors.

Of course, one might argue that the independent strategic model fares better here simply because I have stacked the deck in its favor, making it a special case of the correlated model and then using that to generate the data. However, it remains that much of international relations theory is predicated on strategic interaction, especially deterrence models such as the one in Figure 1. The correlated and independent strategic models are equilibrium models, and there is a theory of strategic choice underlying them. Although traditional selection models are based on an underlying theory of choice, as implemented they are not consistent with *strategic* theories of social interaction.

Strategy and selection also point to a number of other issues concerning our statistical inferences that were not addressed in the body of this article. First, it is not uncommon in analyses of international conflict to use a selection model where the choice made at each step is not by an individual state but by the *dyad* of states involved in the crisis. If the two states are behaving strategically, we would assume there is some "game" occurring between them. Yet, that raises the question: why would we ever expect aggregating that game to be consistent with modeling the states as jointly making a decision—and one that is specified as a linear $X\beta$? No proofs of this have ever been offered. Given the deleterious effects of misspecifying the strategic relationship, it is not clear at all what we should make of the inferences in these studies.

To be fair, the model that is implemented is often a matter of what data is available. Researchers do the best with what they have. Given the importance of strategic behavior in international relations theory, and given the importance of theory in our statistical analyses, it would certainly be helpful if future data collection was conducted with an eye towards these issues. An alternative would be to continue using the same data, but to model the underlying strategic behavior and then derive the probability model based on it for the data on hand. The probability model for the data would be consistent with the underlying theory and could then be used in maximum likelihood estimation. A similar approach, although not involving strategic interaction, is taken by Alt, King, and Signorino (2001).

Finally, I began this paper by suggesting that the history of political interaction can be thought of as having been generated by one huge selection model. That, of course, is a problem for statistical analysis if it implies that all of history is a single data point! Rather than suggesting that the cup is 99.99 percent empty, I would argue that there is an obvious analogy to time series. Individuals cannot peer infinitely into the future—or likely even more than a few moves ahead. That would imply that the disturbances at a particular decision are correlated with those in only a small number of past and present decisions. An interesting avenue of future research would be to examine what the implications are of this and how we account for it in our statistical analysis.

NOTES

1. This, of course, ignores the inefficiency of having a smaller sample, and assumes the researchers have correctly specified the model in the first place.
2. For other examples, see Heckman, 1976,1979; King, 1989.
3. For the truncated data, we would simply derive the conditional probabilities for state 2's decisions and use those in statistical estimation.
4. The intuition for ε_1° is as follows. For $\varepsilon_1 > \varepsilon_1^\circ$ state 1 chooses $y_1 = 1$. For values $\varepsilon_1 < \varepsilon_1^\circ$, state 1 chooses $y_1 = 0$. Therefore, ε_1° defines the value of ε_1 that makes state 1 indifferent between moving right or left. Since we do not know ε_1, but only that it must be greater than or less than this value, the cumulative probability is taken with respect to ε_1°. I am indebted to John Duggan for helpful discussions concerning this.
5. The bivariate probit model is available in almost all standard statistics programs. The independent strategic model employed here is now available in STRAT, a program for estimating statistical strategic models, which is available at http://www.rochester.edu/College/PSC/signorino.

REFERENCES

Alt, James E., Gary King, and Curtis S. Signorino (2001). "Aggregation Among Binary, Count, and Duration Models: Estimating the Same Quantities from Different Levels of Data." *Political Analysis,* Vol. 9, No. 1, pp. 21–44.

Dubin, Jeffrey A. and Douglas Rivers (1989). "Selection Bias in Linear Regression, Logit, and Probit Models." *Sociological Methods and Research,* Vol. 18, Nos. 2,3, pp. 360–390.

Fearon, James D. (1994). "Signaling Versus the Balance of Power and Interests." *Journal of*

Conflict Resolution, Vol. 38, No. 2, pp. 236–269.
Greene, William H. (2000). *Econometric Analysis,* 4th ed., Prentice Hall, Upper Saddle River, NJ.
Heckman, James J. (1976). "The Common Structure of Statistical Models of Truncation, Sample Selection, and Limited Dependent Variables and a Simple Estimator for Such Models." *Annals of Economic and Social Measurement,* Vol. 5/4, pp. 475–492.
Heckman, James J. (1979). "Sample Selection Bias as a Specification Error." *Econometrica,* Vol. 47, No. 1, pp. 153–162.
King, Gary (1998). *Unifying Political Methodology,* Michigan University Press, Ann Arbor, MI.
King, Gary, Robert O. Keohane, and Sidney Verba (1994). *Designing Social Inquiry,* Princeton University Press, Princeton, NJ.
Meng, Chun-Lo, and Peter Schmidt (1985). "On the Costs of Partial Observability in the Bivariate Probit Model." *International Economic Review,* Vol. 26, pp. 71–86.
Reed, William (2000). "A Unified Statistical Model of Conflict Onset and Escalation." *American Journal of Political Science,* Vol. 44, No. 1, pp. 84–93.
Signorino, Curtis S. (1999). "Strategic Interaction and the Statistical Analysis of International Conflict." *American Political Science Review,* June.
Signorino, Curtis S. (2000). "Theoretical Sources of Uncertainty in Statistical Strategic Models." Working Paper, University of Rochester
Signorino, Curtis S. and Ahmer Tarar (2001). "A Unified Theory and Test of Extended Immediate Deterrence." Working Paper, University of Rochester.
Signorino, Curtis S. and Kuzey Yilmaz (2000). "Strategic Misspecification in Discrete Choice Models." Working Paper, University of Rochester.
Smith, Alastair (1996). "To Intervene or Not to Intervene: A Biased Decision." *Journal of Conflict Resolution,* Vol. 40, No. 1.
Smith, Alastair (1999). "Testing Theories of Strategic Choice: The Example of Crisis Escalation." *American Journal of Political Science,* Vol. 43, No. 4.
Timpone, Richard J. (1998) "Structure, Behavior, and Voter Turnout in the United States." *American Political Science Review,* Vol. 92, Vol. 1.
Van de Ven, Wynand P.M.M. and Bernard M.S. Van Praag (1981). "The Demand for Deductibles in Private Health Insurance." *Journal of Econometrics,* Vol. 17, pp. 229–252.

CONTRIBUTOR

Curtis S. Signorino is an Assistant Professor of Political Science at the University of Rochester. He has published articles in the *American Political Science Review*, *International Studies Quarterly, Journal of Conflict Resolution*, and *Political Analysis*.