

STRATEGIC BINARY CHOICE MODELS WITH PARTIAL OBSERVABILITY

Mark David Nieman

Iowa State University

Abstract: Strategic interactions among rational, self-interested actors are commonly theorized in the behavioral, economic, and social sciences. The theorized strategic processes have traditionally been modeled with multi-stage structural estimators, which improve parameter estimates at one stage by using the information from other stages. Multi-stage approaches, however, impose rather strict demands on data availability: data must be available for the actions of each strategic actor at every stage of the interaction. Observational data are not always structured in a manner that is conducive to these approaches. Moreover, the theorized strategic process implies that these data are missing not at random. In this paper, I derive a strategic logistic regression model with partial observability that probabilistically estimates unobserved actor choices related to earlier stages of strategic interactions. I compare the estimator to traditional logit and split-population logit estimators using Monte Carlo simulations and a substantive example of the strategic firm–regulator interaction associated with pollution and environmental sanctions.

Key words and phrases: Data missing not at random, partial observability, strategic choice models.

1. Introduction

Strategic interactions among rational, self-interested actors are of interest in the behavioral, economic, and social sciences. Commonly one assumes that each actor anticipates the expected action of other actors in order to choose its own best response. An important empirical implication of strategic interactions is that observable data are missing not at random (MNAR) (Signorino (1999, 2003); Signorino and Yilmaz (2003); Nieman (2015)). Consider, for example, the modeling of a defender’s ability to deter an attack on a protégé (Schelling (1960); Zagare and Kilgour (2000)). Successful deterrence is one possible outcome of a strategic interaction between the defender and the attacker. The likelihood that an attacker invades the protégé is determined, in part, by whether it expects the defender to intervene. Strategic models are a common approach for statistically modeling these type of interactions.

Strategic models provide a statistical theory of normal and extensive form non-cooperative games (McKelvey and Palfrey (1995, 1996, 1998)). The model assumes that actors choose strategies based on relative expected utility, and that other actors do as well. Strategic models build on random utility assumptions (e.g., McFadden (1974, 1976)), where an error term represents that (1) actor i follows a “bounded rationality” logic, in that, while generally correct, it sometimes errs when implementing its own actions or misperceives the other actor’s utility, or (2) actor i has private information (Signorino (1999, 2003)). Each approach is consistent with perfect Bayesian equilibrium (PBE), a Nash-based equilibria concept common in rational choice models (McKelvey and Palfrey (1995, 1998)).

Strategic models identify equilibria probabilities for actor actions at each information set in a game. McKelvey and Palfrey (1995, 1996, 1998) originally developed the approach to explain variation in subject responses in experimental settings. Signorino (1999) extended this to estimate coefficient parameters rather than the variance. In either case, covariates are specified at each information set, and the predicted probabilities from later stages are used to inform and condition estimates at earlier stages to account for the theorized strategic process. This process is similar to some imputation based approaches (e.g., Boehmke (2003); Liu et al. (2013)) which use the estimates from one equation to construct latent measures for censored cases. Strategic models differ from both traditional nonignorable selection models (e.g., Heckman (1979); Sartori (2003)) and auxiliary/imputation based approaches in that strategic models are designed for situations in which two or more actors are in a non-cooperative setting. Strategic models treat an actor’s choice in an earlier stage of the interaction as a function of both its own expected behavior *and the expected behavior of the other actor* in a later stage.

A drawback to this approach is that strategic models require data be available for each actor at each information set of the game (Signorino and Yilmaz (2003, pp.556–557); Nieman (2015)), but observational data often fail to meet this condition. Instead, data are often available only for the outcome of an interaction, with little or no data on the individual actor actions that lead to the observed outcome. This means that observational data are only partially observed, since they are the result of unobserved *joint* decisions of multiple actors, rather than those of a single decision-maker (Poirier (1980)). Moreover, the theorized strategic process implies that data reflecting the unobserved joint decisions are characterized by non-random missingness.

Figure 1 presents an example of the theoretical interaction between a firm and

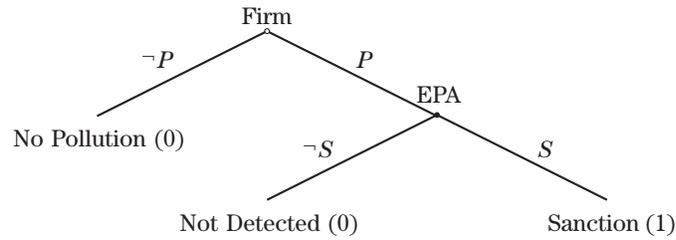


Figure 1. Interaction between firm and the environmental protection agency.

the environmental protection agency (EPA). Assume that the firm is interested in maximizing its profits for producing widgets, but in the course of doing so generates waste. The firm is presented with two choices: it can either stay within its regulated guidelines of waste production ($\neg P$), or it can exceed regulations and pollute illegally (P). If the firm chooses $\neg P$, then the game ends with the outcome *No Pollution*. If the firm chooses P , then the EPA will either detect the violation and issue a sanction (S), resulting in the outcome *Sanction*, or fail to detect the violation and not sanction ($\neg S$), leading to the outcome *Not Detected*. Observational data, however, contains only information on whether the interaction resulted in a sanction, but not the individual actions. The number in parentheses indicates how observational data would code each outcome.

Observational data combines two of the outcomes—*No Pollution* and *Not Detected*—and treats them as non-events (0s), while treating the other outcome—*Sanction*—as the observed event (1s). It is unlikely that unsanctioned, noncompliant firms are randomly distributed; rather, firms may intentionally exceed pollution standards when they believe that the EPA is less likely to sanction them, meaning data for a firm's actions are MNAR. Moreover, additional data collection would not resolve the nonrandom missingness problem: firms that pollute and evade EPA sanctions are unlikely to volunteer accurate information regarding their noncompliance with EPA standards. Ignoring the underlying strategic process and simply treating the data as binary when estimating predictors of noncompliance may lead to incorrect inferences regarding whether decreases in the number of sanctions indicates that noncompliance is decreasing or whether firms are increasingly evading detection.

In remainder of the paper, I derive a statistical solution to the theoretical puzzles driven by the limitations in many datasets by extending the strategic model to cases of partial observation. The estimator outperforms both traditional and binary choice mixture models in a set of Monte Carlo simulations. Finally, I apply the estimator to the firm-level data regarding compliance with

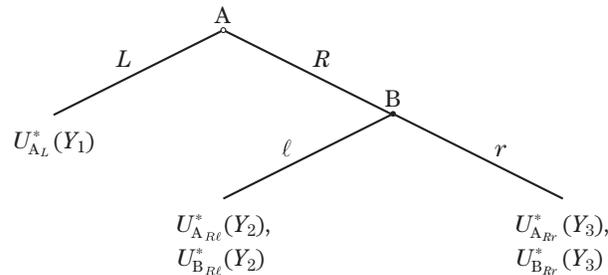


Figure 2. Strategic interaction with agent error.

environmental regulations from Konisky and Teodoro (2015).

2. A Strategic Model with Agent Error

A strategic discrete choice model where errors enter at information sets—agent error—is depicted in Figure 2. The game has two actors, [A, B], four actions [L , R , ℓ , r], and three possible outcomes, [Y_1 , Y_2 , or Y_3]. The actors move sequentially, with A moving first. If actor A chooses L , the game ends with Y_1 . If actor A chooses R , then actor B selects between ℓ and r . If actor B chooses ℓ , the game ends with Y_2 . If actor B chooses r , the game ends with Y_3 .

The utilities displayed under the terminal node reflect the actions of each actor, so that $U_{ij}^* = U_{ij} + \alpha_{ij}$, where U_{ij}^* is the true utility, i is the actor, j is the action, U is the observable utility (that is known by both actors and the analyst), α represents agent error associated with the action, and the value in parentheses is the outcome. Actor i knows the value of α , but actor $\neg i$ (and the analyst) only know its distribution. Each actor chooses the action where $U_{ij}^* > U_{i\neg j}^*$. Since the model is sequential, actor A must take into account the expected action of actor B in order to maximize its own utility. Thus, the game is solved by working backwards up the game tree.

If actor A chooses R , then actor B's utilities for selecting ℓ and r are

$$U_{B_{R\ell}}^* = U_{B\ell} + \alpha_{B\ell}, \quad (2.1)$$

$$U_{B_{Rr}}^* = U_{B_r} + \alpha_{B_r}. \quad (2.2)$$

Since the analyst and actor $\neg i$ know only the distribution of α_{ij} , they only have probabilistic estimates of i 's choice. If we assume that α_{ij} are i.i.d. type I extreme value, the resulting choice probabilities for actor B have the logistic distributions

$$p_\ell = \frac{e^{U_{B\ell}}}{e^{U_{B\ell}} + e^{U_{B_r}}}, \quad (2.3)$$

$$p_r = \frac{e^{U_{B_r}}}{e^{U_{B_\ell}} + e^{U_{B_r}}}. \quad (2.4)$$

Actor A must account for actor B's choices in order to calculate actor A's own utility. To do this, actor A conditions its own utilities by whether it expects actor B to choose ℓ or r . That is, actor A must calculate its expected utility and does so based on the observable portion of actor B's utility and the known distributions of α_{B_j} , which are the choice probabilities from (2.3) and (2.4). Actor A's utilities for selecting L and R are

$$U_{A_L}^* = U_{A_L} + \alpha_{A_L}, \quad (2.5)$$

$$U_{A_R}^* = U_{A_R} + \alpha_{A_R}. \quad (2.6)$$

Inserting the probability of each of actor B's choices from (2.3) and (2.4) in order to calculate the expected value for U_{A_R} in (2.6) yields:

$$EU_{A_R}^* = p_\ell U_{A_{R\ell}} + p_r U_{A_{Rr}} + \alpha_{A_R}. \quad (2.7)$$

The resulting choice probabilities for actor A, from the analyst's vantage, are

$$p_L = \frac{e^{U_{A_L}}}{e^{U_{A_L}} + e^{EU_{A_R}^*}}, \quad (2.8)$$

$$p_R = \frac{e^{EU_{A_R}^*}}{e^{U_{A_L}} + e^{EU_{A_R}^*}}. \quad (2.9)$$

Equations (2.8) and (2.9) differ from a traditional logit model as they account for the expected utility calculations of actor A, the endogeneity of strategic decision-making from (bounded) rational actors.

The probabilities for each outcome Y_1 , Y_2 , and Y_3 are simply the product of the choice probabilities following the sequence of the game in Figure 2

$$p_{Y_1} = p_L, \quad (2.10)$$

$$p_{Y_2} = p_R p_\ell, \quad (2.11)$$

$$p_{Y_3} = p_R p_r. \quad (2.12)$$

These outcome probabilities are the equilibria outcomes of the game.

The observable utilities U_{ij} can be specified as a set of regressors, such that $U_{ij} = X_{ij}\beta_{ij}$. In order for the model to be identified, the same variable cannot be specified for all j (Lewis and Schultz (2003)). The uncertainty introduced by each actor action allows us to directly estimate the theoretical model using maximum likelihood. Assuming there are data for each actor decision and regressors for the utilities, parameters are recovered by maximizing the likelihood function

$$L = \prod_{i=1}^n P(Y_{1,i} = 1)^{y_{1,i}} P(Y_{2,i} = 1)^{y_{2,i}} P(Y_{3,i} = 1)^{y_{3,i}}. \quad (2.13)$$

3. A Strategic Model with Partial Observability

Data on actor decisions are frequently not available. In the absence of such data, one often employs conventional binary choice models, such as a logistic regression, using a first order $\beta_{ij}X_{ij}$ specification. This specification ignores the endogenous and conditional choices made by strategic actors. Even the use of a mixture model, such as an endogenous switching or split-sample (e.g., zero-inflated) logit (Miranda and Rabe-Hesketh (2005); Greene (2010)), while helping to account for the two distinct processes leading to outcomes coded as 0, fails to model the strategic interaction depicted in Figure 2.

To address these data concerns, the likelihood in (2.13) can be re-written as

$$L = \prod_{i=1}^n P(Y_i = 1)^{y_i} P(Y_i = 0)^{1-y_i}, \quad (3.1)$$

where

$$P(Y_i = 1) = p_R p_r, \quad (3.2)$$

$$P(Y_i = 0) = 1 - p_R + p_R(1 - p_r) = 1 - p_R p_r. \quad (3.3)$$

The likelihood captures the theorized strategic interaction between actors A and B. The random variable $Y = 1$ when both actions R and r occur. Equation (3.2) is the probability of observing $Y = 1$ in the data. The random variable $Y = 0$ when either action L or actions R and ℓ occur. These events are pooled in (3.3) and are separated probabilistically using the observed portion of the utility.

4. Monte Carlo Analysis

I analyze the ability of the strategic logistic model with partial observability (SLPO) to recover parameter estimates using Monte Carlo simulations based on a data generating process (DGP) that assumes a strategic interaction between two actors, as shown in Figure 2. I compare the estimates to those from two other models that are commonly used when estimating theorized strategic interactions between (boundedly) rational actors: the split-sample logit (SSL) and the traditional logit. As the data are simulated and we know the actual parameter values, I also estimate a full information strategic logit to assess the loss of efficiency between the full information and partially observed models when data on actor

actions are missing.

Assume the DGP

$$Y_{\text{true}} = \begin{cases} Y_1 & \text{if } U_{A_L}^* \geq U_{A_R}^*, \\ Y_2 & \text{if } U_{A_R}^* > U_{A_L}^* \text{ and } U_{B_{R\ell}}^* \geq U_{B_{Rr}}^*, \\ Y_3 & \text{if } U_{A_R}^* > U_{A_L}^* \text{ and } U_{B_{Rr}}^* > U_{B_{R\ell}}^* \end{cases}$$

where, from the perspective of actor A and the analyst,

$$U_{B_{R\ell}}^* = X_{B_{\ell}}\beta_{B_{\ell}} + \alpha_{B_{\ell}}, \tag{4.1}$$

$$U_{B_{Rr}}^* = X_{B_r}\beta_{B_r} + X_c\beta_{B_{rc}} + \alpha_{B_r}, \tag{4.2}$$

and from the perspective of the analyst,

$$U_{A_L}^* = X_{A_L}\beta_{A_L} + \alpha_{A_L}, \tag{4.3}$$

$$U_{A_R}^* = p_{\ell}X_{A_{R\ell}}\beta_{A_{R\ell}} + p_r(X_{A_{Rr}}\beta_{A_{Rr}} + X_c\beta_{A_{rc}}) + \alpha_{A_R}, \tag{4.4}$$

with $p_{\ell} = \exp(U_{B_{R\ell}}^*) / \{1 + \exp(U_{B_{R\ell}}^*)\}$ and $p_r = \exp(U_{B_{Rr}}^*) / \{1 + \exp(U_{B_{Rr}}^*)\}$. Consistent with random utility assumptions, I normalize $\beta_{A_{R\ell}} = \beta_{B_{\ell}} = 0$. I set $\beta_{A_L} = \beta_{A_{Rr}}\beta_{A_{rc}} = \beta_{B_r} = \beta_{B_{rc}} = 1$. Fixed variables $X_{A_L}, X_{A_{Rr}}, X_{B_r}$, and X_c are uniformly distributed $[-2, 2]$. The inclusion of X_c , a regressor that is common to both actors, makes the simulations more realistic to applied research. The α_{ij} are i.i.d. type I extreme value, and the resulting choice probabilities following logistic distributions. I ran 1,000 simulations with 2,000 observations each. All estimates were performed using *Stata 14* statistical software.

I recoded the random variable Y_{true} into a binary variable, to reflect the situation where data on individual actor actions are missing, and the data only report whether actors A and B choose a specific joint outcome, but other outcomes are not distinguishable, so

$$Y_{\text{miss}} = \begin{cases} 1 & \text{if } Y_{\text{true}} = Y_3, \\ 0 & \text{otherwise.} \end{cases}$$

Y_{miss} is treated as the random variable for each of the SLPO, SSL, and traditional logit models. SLPO is specified consistent with the DGP and adheres to the likelihood in (3.1).

SSL is a mixture model where the random variable is a function of two processes, such that

$$P(Y_{\text{miss}} = 1) = SL, \tag{4.5}$$

$$P(Y_{\text{miss}} = 0) = (1 - S) + S(1 - L), \tag{4.6}$$

where S and L are logistic cumulative density functions. S is specified as $S =$

$\exp(S^*)/\{1 + \exp(S^*)\}$ where

$$S^* = X'_{A_L}\beta_{A_L} + X_c\beta_{A_{Rrc}} + \epsilon_1, \quad (4.7)$$

and L is specified as $L = \exp(L^*)/\{1 + \exp(L^*)\}$ where

$$L^* = X_{A_{Rr}}\beta_{A_{Rr}} + X_{B_r}\beta_{B_r} + X_c\beta_{B_{rc}} + \epsilon_2, \quad (4.8)$$

and the ϵ_i follow logistic distributions. Equation (4.7) represents the “selection” equation and (4.8) represents the traditional logit equation.

The traditional logit is specified as

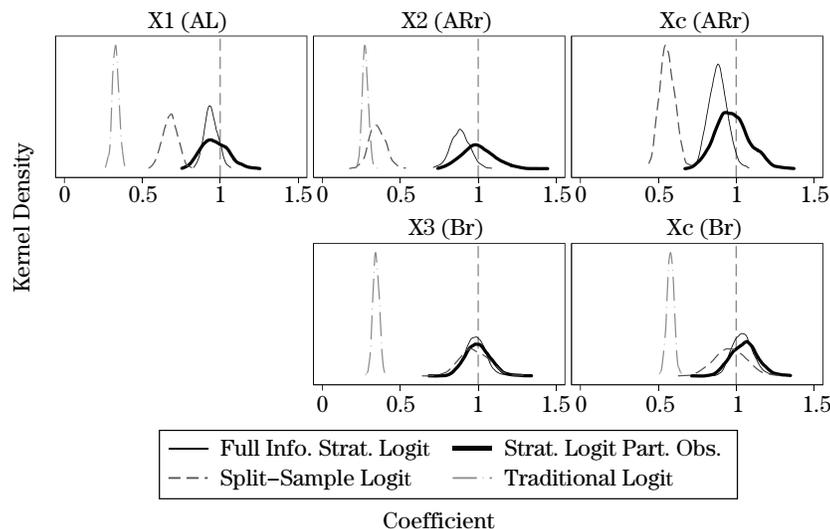
$$Y_{\text{miss}}^* = X'_{A_L}\beta_{A_L} + X_{A_{Rr}}\beta_{A_{Rr}} + X_{B_r}\beta_{B_r} + X_c\beta_{B_{rc}} + \epsilon, \quad (4.9)$$

where ϵ follows a logistic distribution, and $Y_{\text{miss}} = 1$ if $Y_{\text{miss}}^* > 0$ and 0 otherwise.

The full information strategic logit (FISL) used all of the data from the DGP, treating Y_{true} as the random variable. The FISL was estimated using the likelihood in (2.13).

Figure 3 provides a visual display of the parameter estimates from the FISL (thin solid line), SLPO (thick solid line), SSL (dashed line), and traditional logit (dash dot line). Both of the strategic models are able to capture the true parameter estimates, while the split-sample model is able to capture parameters associated with actor B’s utility, but not those associated with actor A, and the traditional logit produces only biased parameter estimates. The difference between the strategic models and SSL reflects the conditional nature of actor A’s actions, which SSL does not account for. As one would expect, when more information is available, FISL produces more efficient estimates than SLPO, as indicated by the taller densities.

Table 1 reports the estimated coefficients, standard errors, and root mean squared error (RMSE) from the Monte Carlos. The strategic models are able to approximate the true values for the parameters. SLPO recovers slightly less biased estimates of the first stage coefficients than FISL, but is less efficient for both first and second stage estimates in terms of RMSE. The slight conservative bias associated with FISL stems from the greater precision of the second stage estimates, which carries through the predicted probabilities that condition the estimated coefficients in the first stage (Leeman (2014)). Comparing SLPO to SSL, the former outperforms the latter, especially when estimating coefficients associated with actor A, whose actions are conditioned by the expected actions of actor B. SSL is able to recover unbiased estimates of coefficients associated with actor B’s actions, but is less efficient than SLPO. The traditional logit recovers biased estimates of all coefficients.



Note: Dashed vertical line represents the equation's true coefficient. Results of 1,000 simulations with 2,000 observations each. B_{ij} represents the estimated coefficient for regressor X_{ij} , where i is the actor and j is the outcome. X_c is a common regressor that appears in each actor's utility. Because the traditional and split-sample logit models estimate less parameters, they do not appear in each graph.

Figure 3. Kernel densities of estimated coefficients.

I compare model fit statistics for the three binary choice models (SLPO, SSL, logit) in Table 2. Since the models are non-nested in terms of their functional form, I used Clarke's distribution-free test and the Vuong test. Clarke's distribution-free test is informed by the median logged ratio of the likelihood for the individual observations of two empirical models (Clarke (2007)). If the first model is closer to the true specification, then the median logged ratio of the two likelihoods is positive. If the second model is closer to the true specification, then the ratio is negative. More formally,

$$H_0 : Pr_0 \left\{ \ln \frac{f(Y_i|X_i; \beta_*)}{g(Y_i|Z_i; \gamma_*)} > 0 \right\} = 0.5, \quad (4.10)$$

where the numerator is the estimated model f that predicts Y_i from a set of covariates, X_i , and estimated parameters, β_* ; the denominator is the estimated model g , that predicts Y_i from a set of covariates, Z_i , and estimated parameters, γ_* . The null hypothesis is that the median logged ratio of the likelihoods between the two models is equal to 0. If d_i is set equal to $\ln f(Y_i|X_i; \beta_*) - \ln g(Y_i|Z_i; \gamma_*)$, the test statistic is

$$B = \sum_{i=1}^n I_{(0,+\infty)}(d_i), \quad (4.11)$$

Table 1. Comparison of estimated coefficient, standard error, and root mean squared error.

Estimates	FISL	SLPO	SSL	Logit
β_{AL}	0.942	0.966	0.679	0.329
SE	0.043	0.085	0.055	0.026
RMSE	0.218	0.305	0.399	0.691
β_{ARr}	0.885	1.002	0.353	0.275
SE	0.061	0.099	0.054	0.026
RMSE	0.279	0.330	0.688	0.742
β_{ARrc}	0.882	0.976	0.559	—
SE	0.062	0.105	0.059	—
RMSE	0.281	0.341	0.505	—
β_{Br}	0.984	0.999	0.976	0.344
SE	0.064	0.079	0.094	0.027
RMSE	0.261	0.292	0.321	0.676
β_{Brc}	1.041	1.054	0.963	0.578
SE	0.645	0.083	0.102	0.030
RMSE	0.265	0.304	0.337	0.456

Note: The $RMSE = \sqrt{Bias^2 + Variance}$. Because the traditional probit is a single equation model, it estimates only one parameter for X_c , which is displayed with β_{Brc} .

where I is a dichotomous indicator equal to 1 if $n_i > 0$ in (4.10), and 0 if $n_i \leq 0$. Equation (4.11) is the sum of positive differences and is a Binomial distribution with n trials and a mean equal to 0.5.

The Vuong test compares the mean log-likelihood ratios of two models. If the first model is closer to the true specification, then the mean log-likelihood ratio is positive and statistically significant, similarly, the second model. As is common practice, I use Schwarz's correction to the Vuong test

$$LR_n(\tilde{\theta}_n, \tilde{\gamma}_n) - \left\{ \left(\frac{p}{2} \right) \ln n - \left(\frac{q}{2} \right) \ln n \right\}, \quad (4.12)$$

where LR is the log-likelihood ratio, $\tilde{\theta}$ and $\tilde{\gamma}$ are the model estimates, and p and q are the number of estimated parameters for model f and g , the two models being compared (Vuong (1989)). I also assessed in-sample goodness of fit by comparing correctly predicted and false positive rates across models.

Table 2 demonstrates that the various model fit statistics identify SLPO as the best fit to the strategic DGP. Both the Clarke and the Vuong tests indicate that SLPO reflects the DGP better than SSL and traditional logit; we can reject the null hypothesis that either SSL or logit models are equal to SLPO, as the p -value for each comparison is less than 0.001 using either test. SLPO is

Table 2. Comparison of average model fit with a strategic data generating process.

	SLPO	SSL	Logit
Clarke Test			
$\sum_i^n l_{\text{SLPO},i} - l_{\text{alternative},i} > 0$	—	1,191.254	1,622.903
Positive, 1-side test (p -value)	—	< 0.001	< 0.001
Negative, 1-side test (p -value)	—	> 0.999	> 0.999
Equal, 2-side test (p -value)	—	< 0.001	< 0.001
Vuong Test			
Vuong	—	94.595	374.007
SE	—	1.385	1.120
t -statistic	—	73.460	336.481
p -value	—	< 0.001	< 0.001
In-sample Predictions			
% predicted Y_1 (True = 44.7%)	45.5	52.1	—
% predicted Y_2 (True = 24.5%)	23.4	16.1	—
% predicted Y_3 (True = 30.8%)	31.0	31.8	50.5
% of Obs. Correctly Classified Y_{true}	78.5	70.8	—
% of Obs. Correct Classified Y_{miss}	87.3	83.8	67.4
% of Obs. Correct if $\hat{Y}_{\text{miss}} = 1 Y_{\text{miss}} = 1$	79.9	75.5	79.2
% False Positive ($\hat{Y}_{\text{miss}} = 1 Y_{\text{miss}} = 0$)	9.3	12.4	37.8
% Correctly Predicted Unobserved ($\hat{Y}_2 Y_{\text{true}} = Y_2$)	67.9	67.9	—

able to identify the true outcomes, an important consideration when calculating predicted values. Both SLPO and SSL are able to correctly classify over 80% of cases with binary outcomes, while logit correctly classifies less than 70% of such cases. SLPO has a higher rate of correctly identifying cases where $Y_{\text{miss}} = 1$ than SSL, and a lower rate of false positives. Logit is able to correctly identify cases where $Y_{\text{miss}} = 1$ at approximately the same rate as SLPO, but it does so with a higher rate of false positives. Both SLPO and SSL are able to correctly identify separate the two types of 0 cases in Y_{miss} data.

These differences between SLPO, SSL, and traditional logit have ramifications when testing rational choice model using aggregated data that is missing information about individual actor choices. Estimates using SSL are downward biased on actions associated with actor A. The effect of this is that while inferences related to the direction of an effect may be correct, any subsequent substantive effects estimated from the model are not. Turning to the traditional logit, it is unable to isolate the effects of X_c for each actor. On balance, SLPO is the model best able to capture rational choice theories.

5. Application

I apply SLPO to the example discussed in the introduction of firm compliance with EPA regulations. In particular, I use data from Konisky and Teodoro (2015) to explore the determinants of compliance with the Clean Air Act (CAA) for 54,206 US firms during the period 2000–2011. Firms are coded as noncompliant when they (1) violate the CAA and (2) are detected by the EPA (or state regulators). Firms that either follow the standards set in the CAA, or that the EPA fail to detect, are both coded the same in that they have not been found to be noncompliant with environmental regulations.

Firms are expected to follow a “calculated motivation” (e.g., Winter and May (2001); Konisky and Teodoro (2015)) in that they act as rational decision-makers who comply with regulations if the net cost of doing so is less than the net cost of noncompliance. The interaction of firms and regulators is strategic in that firms are more likely to adhere to CAA standards when they expect regulators to identify firms that violate the law, and are less likely to meet CAA standards otherwise. Regulators, meanwhile, have limited budgets, and detecting violations and issuing sanctions is costly. Regulators must therefore prioritize inspections for some locations as the expense of others. The result of the firm–regulator interaction is the same as that depicted in Figure 1, where firms are treated as the first actor and the EPA (and state regulators) as the second actor.

The random variable, *noncompliance*, is a binary variable equal to 1 if a firm is officially sanctioned by the EPA or state regulators, and 0 otherwise. Approximately 9% of observations are coded as noncompliant in the sample. I specify a firm’s utility for the action to stay within CAA regulations (U_{F-P}^*) with the binary variable *major air source*, which captures firm size and is coded 1 if a stationary source emits pollutants above a certain threshold (approximately 100 tons/year of air pollutants, 10 tons/year of a single hazardous air pollutant, or 25 tons/year of combined hazardous air pollutants). Given their more complex regulations, the parameter for major air source is expected to be negative. I include several county-level demographic and economic characteristics to represent the utility a firm receives from acting to run afoul of CAA regulations (U_{F-P}^*). The percent of the population that is *African American* or *Hispanic* are expected to be associated with a reduction in a firm’s utility to comply with the CAA, higher *median household incomes* should increase compliance, and higher *poverty* and *unemployment rates* should decrease compliance. Finally, I control for whether a firm was identified as a noncomplier in the previous year. The expected utility for U_{F-P}^* , of course, accounts for the expected actions of the EPA.

Table 3. Summary statistics.

Variable	Mean	Std. Dev.	Min.	Max.	Obs.
Noncompliance	0.090	0.286	0	1	650,472
Major Air Source	0.305	0.460	0	1	650,472
Percent African American	11.905	14.062	0	86.489	650,472
Percent Hispanic	9.601	12.136	0.094	97.539	650,472
Poverty Rate	14.850	5.509	2.117	50.888	650,472
Unemployment Rate	6.117	2.612	1.1	29.700	650,472
Median Household Income	45,055.472	11,782.858	16,271	119,075	650,472
Public	0.050	0.217	0	1	650,472

Finally, I include several variables to capture the utility for regulators to detect and sanction noncompliant firms (U_{Es}^*). Public-sector firms can use political back channels to pressure regulators and evade sanctions. *Public* is a binary variable coded as 1 if a firm is publicly owned. I also include country-level demographic and economic characteristics. Higher rates of minorities, poverty, and unemployment are expected to be associated with lower probabilities of sanctioning noncompliant firms, while higher median household incomes may lead to higher probabilities of sanctioning noncompliant firms. Descriptive statistics of each variable are reported in Table 3.

Table 4 presents the results from a SLPO and logit. I compare estimates from the SLPO to a traditional logit model, as either logit or probit are typically used in applied work with binary dependent variables. The models differ in that logit treats the random variable as an additive function that combines the theorized causal mechanisms attributed to a variable for each actor, while SLPO models it as a strategic interaction and can isolate casual mechanisms attributed to different actors by placing the same variable in multiple equations associated with different actors.

The results between the two models differ in a number of ways. The logit model identifies *percent African American*, *percent Hispanic*, and *poverty rate* as statistically significant predictors of noncompliance, with the two former having positive coefficients and the latter a negative coefficient. SLPO provides a more nuanced interpretation: the increase in noncompliance associated with increases in *percent African American* from the logit model is attributed to a lack of deterrence, as the coefficient is negatively associated with the EPA's utility for sanctioning. *Percent Hispanic* has a positive coefficient in the equation for a firm's utility for polluting and a negative coefficient in the equation for the EPA's utility for sanctioning, i.e. increases in the percent of the Hispanic population are associated with both an increase in the incentive of a firm to pollute and a

Table 4. Comparison of logit and strategic logit with partial observability using data on firm noncompliance with clean air act.

Traditional Logit		SLPO	
		<u>Firm (\negPollute):</u>	
Major Air Source	0.807* (0.012)	Major Air Source	-1.184* (0.015)
Percent African American	0.001 (0.001)	<u>Firm (Pollute):</u>	
Percent Hispanic	0.004* (0.001)	Percent African American	0.001 (0.001)
Poverty Rate	-0.020* (0.001)	Percent Hispanic	0.009* (0.001)
Unemployment Rate	-0.003 (0.003)	Poverty Rate	-0.040* (0.005)
Median Household Income	-0.001 (0.001)	Unemployment Rate	-0.004 (0.005)
Public	0.011 (0.027)	Median Household Income	-0.001 (0.001)
Lagged Noncompliance	4.371* (0.013)	Lagged Noncompliance	8.617* (0.244)
Constant	-3.423* (0.059)	Constant	-4.152* (0.116)
		<u>EPA (Sanction):</u>	
		Percent African American	-0.004* (0.001)
		Percent Hispanic	-0.007* (0.001)
		Poverty Rate	-0.022* (0.004)
		Unemployment Rate	0.040* (0.004)
		Median Household Income	-0.001 (0.001)
		Public	0.004 (0.035)
		Constant	1.435* (0.100)
Observations	650,472		650,472
Log-Likelihood	-111,859.60		-110,710.25

Note: * $p < 0.001$, two-tailed. Standard error in parentheses.

decrease in the deterring effect from EPA sanctions. The effect of *poverty rate* is nonlinear, as the signs on the coefficients work in opposite directions. Increases in the poverty rate decrease a firms utility for polluting, but they also decrease the EPA's utility to sanction. In addition to these variables, the logit model identifies *unemployment rate* as statistically insignificant, while SLPO indicates

Table 5. Model fit of logit and strategic logit with partial observability on firm noncompliance estimates.

Clarke Test		
$\sum_i^n ll_{SLPO,i} - ll_{Logit,i} > 0$	603,207	
Positive, 1-side test (<i>p</i> -value)	< 0.001	
Negative, 1-side test (<i>p</i> -value)	> 0.999	
Equal, 2-side test (<i>p</i> -value)	< 0.001	
Vuong Test		
Vuong	1,189.531	
SE	0.277	
t-statistic	4,287.058	
<i>p</i> -value	< 0.001	
In-sample Predictions		
	<u>Logit</u>	<u>SLPO</u>
% of Obs. Correct	95.0	95.0
% of Obs. Correct if $Y = 1$	67.5	67.5
% False Positive	2.3	2.3
Pred. % Y_2 (<i>Pollute</i> \neg <i>Sanction</i>)	—	2.2

Note: Estimates from Table 4. SLPO is model 1 and logit model 2 in Clarke and Vuong tests.

that it increases the EPA's utility from sanctioning, exerting a deterring effect.

Table 5 reports model fit comparisons of the SLPO and logit models. One sees that SLPO allows us to estimate the percent of observations in which firms pollute but are not identified by regulators as noncompliant—predicted unobserved Y_2 outcomes—something that a traditional logit model cannot do. In this case, the model indicates that 2.2% of observations fit into this category. Since the probability of each outcome can be calculated for every observation, this feature of SLPO has potential practical benefits to regulators and watchdog organizations. In this specific case, the model indicates that in 1,253 of the observations, a firm has a probability greater than 30% of being non-compliant yet avoiding sanctions.

The bottom of the table compares the fit of the two models to one another. The Clarke and Vuong tests each offer strong support for SLPO over the logit model. SLPO has a greater log-likelihood in over 90% of the individual observations, while the mean log-likelihood is also greater. On balance, the fit statistics suggest that the strategic model is a better fit to the observed data than the logit model.

6. Conclusion

Strategic interactions among actors are commonly theorized across the be-

havioral, economic, and social sciences. Common approaches to empirically modeling these interactions are made difficult in the presence of nonignorable missing data on individual actor choices. I provide a solution to this problem by using a strategic logit with partial observability that is able to capture the theorized underlying strategic process in the presence of data on actor actions missing not at random. I use Monte Carlo simulations to demonstrate that the strategic logit with partial observability outperforms other binary choice estimators, such as traditional logit and split-sample logit. The Monte Carlo simulations show that strategic logit with partial observability recovers the same parameter estimates, and is only slight less efficient, as a full information strategic logit does with complete information on individual actor choices.

I apply the estimator to the interaction of firms and regulators concerning noncompliance with the Clean Air Act. Model fit statistics demonstrate the strategic model better explains the observed data. The multiple equation nature of the strategic model, moreover, is able to test specific causal mechanisms as they pertain to each actor, something that traditional binary choice models cannot. The estimator has numerous applications, such as criminal behavior, lending practices by financial institutions, international and domestic conflict onset and escalation, among others.

Supplementary Materials

Replication data and code are available at www.marknieman.net.

Acknowledgment

I would like to thank special issue guest editor Jae-Kwang Kim and co-editor Zhiliang Ying. I would also like to thank Olga Chyzh, Walter Mebane, Fred Boehmke, and Curt Signorino.

References

- Boehmke, F. J. (2003). Using auxiliary data to estimate selection bias models, with an application to interest group use of the direct initiative process. *Political Anal.* **11**, 234–254.
- Clarke, K. A. (2007). A simple distribution-free test for nonnested model selection. *Political Anal.* **15**, 347–363.
- Greene, W. (2010). A stochastic frontier model with correction for sample selection. *J. Productivity Analysis* **34**, 15–24.
- Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica* **47**, 153–162.
- Konisky D. M. and Teodoro, M. P. (2015). When governments regulate governments. *Am. J.*

- Political Sci.* Forthcoming. DOI: 10.1111/ajps.12221.
- Leeman, L. (2014). Strategy and sample selection: A strategic selection estimator. *Political Anal.* **22**, 374–397.
- Lewis, J. B. and Schultz, K. A. (2003). Revealing preferences: Empirical estimation of a crisis bargaining game with incomplete information. *Political Anal.* **11**, 345–367.
- Liu, J., Meng, X., Chen, C. and Alegria, M. (2013). Statistics can lie but can also correct for lies: Reducing response bias in NLAAS via Bayesian imputation. *Stat. Interface* **6**, 387–398.
- McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In *Frontiers in Econometrics*, (Edited by P.). Academic Press, New York.
- McFadden, D. (1976). Quantal Choice Analysis: A Survey. *Annals of Econ. Social Measurement.* **5**, 363–390.
- McKelvey, R. D. and Palfrey, T. R. (1995). A statistical theory of equilibrium in games. *Games Econom. Behav.* **10**, 6–38.
- McKelvey, R. D. and Palfrey, T. R. (1996). A statistical theory of equilibrium in games. *Jpn. Econ. Rev.* **47**, 186–209.
- McKelvey, R. D. and Palfrey, T. R. (1998). Quantal response equilibria for extensive form games. *Experimental Econ.* **1**, 9–41.
- Miranda, A. and Rabe-Hesketh S. (2005). Maximum likelihood estimation of endogenous switching and sample selection models for binary, count, and ordinal variables. *Technical report, Centre for Economic Research.* Keele University.
- Nieman, M. D. (2015). Statistical analysis of strategic interaction with unobserved player actions: Introducing a strategic probit with partial observability. *Political Anal.* **23**, 429–448.
- Poirier, D. J. (1980). Partial observability in bivariate probit models. *J. Econometrics* **12**, 209–217.
- Sartori, A. E. (2003). An estimator for some binary-outcome selection models without exclusion restrictions. *Political Anal.* **11**, 111–138.
- Schelling, T. C. (1960). *The Strategy of Conflict.* Harvard University Press, Cambridge.
- Signorino, C. S. (1999). Strategic interaction and the statistical analysis of international conflict. *Am. Political Sci. Rev.* **93**, 279–297.
- Signorino, C. S. (2003). Structure and uncertainty in discrete choice models. *Political Anal.* **11**, 316–344.
- Signorino, C. S. and Yilmaz, Y. (2003). Strategic misspecification in regression models. *Am. Political Sci. Rev.* **47**, 551–566.
- Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica.* **57**, 307–333.
- Winter, S. and May, P. (2001). Motivation for compliance with environmental regulations. *J. Policy Anal. Manag.* **20**, 675–698.
- Zagare, F. C. and Kilgour, D. M. (2000). *Perfect Deterrence.* New York: Cambridge University Press.

Department of Political Science, Iowa State University, 537 Ross Hall, Ames, IA 50011, USA.
E-mail: mdnieman@iastate.edu

(Received June 2016; accepted August 2017)