

Implementing and Interpreting Sample Selection Models

By Kevin Sweeney
Political Research Lab

We will kick off the methods lunch today with my presentation on sample selection models. This is an appropriate topic because sample selection problems are pervasive in social science research, and confusion reigns about how and when to use the statistical tools to deal with those problems. I'm going to do my best to explain the intuition behind sample selection models and how to implement those models. I will cover a fair amount of material. Since I am the first person to do this, I'm not really sure how it will go. I will talk for about 40 minutes, and near the end of the presentation you will be estimating the models along with me on the machines. After that we can throw it open for a discussion, although it is not clear to me that I know any more than I am about to say. We'll see.

Outline of the Presentation

- The Intuition Behind Selection Models
 - Tobit
 - Heckman's Original Estimator
 - The Likelihood Function
 - An Empirical Example (Stata 7.0)
 - Censored Probit
 - An Empirical Example (Stata 7.0)
 - Some cool programs (Stata 7.0)
 - Related Models
 - Applications in Political Science
-

We're going to begin by getting a sense of the intuition behind sample selection models. Here I am going to detail the analysis in the original paper that brought selection questions into focus. Although tobit is not a sample selection model, it is a short leap from there to true selection models.

We will then shift focus to James Heckman's original sample selection estimator, which is an important twist on the tobit model (at least the nobel prize folks thought so). After I describe the model, we will hit the machines and estimate one in stata 7.

After that we will describe the censored probit model, which is the same as heckman's original estimator except that the dependent variable in the outcome equation is binary. After describing that model, we will estimate one on the machines.

After the fun on the computers I will talk very briefly about some related models, particularly about event count models.

And, finally I will list some applications in political science. I know lots of them in IR and have done some searching around for the other subfields as well. This will serve two purposes. First, you can go out and read in your own field, perhaps an application will make more sense to you than what I am about to say. Second, the list of references will serve to underscore the point that mastering this methodology is important because it is becoming increasingly popular

How Did This All Start?


"What do you mean, *less* than nothing?
Replied Wilbur. "I don't think there is any such thing as *less* than nothing. Nothing is absolutely the limit of nothingness. It's the lowest you can go. It's the end of the line. How can something be less than nothing? If there were something that was less than nothing then nothing would not be nothing, it would be something – even though it's just a very little bit of something. But if nothing is *nothing*, then nothing has nothing that is less than *it* is.

E.B. White, *Charlotte's Web*

Tobin began his seminal paper with this quote from Charlotte's Web. [read]

Although this is funny because it is confused, it highlights almost perfectly the substance of the problem encountered by Tobin.

Intuition of Selection Models: Censored Example



Tobin wanted to explain the relationship between household income and household luxury expenditures.

He figured that the more income a household gained the more they would spend on luxury goods, but noticed that there was a large concentration of households who spend exactly zero dollars on luxury goods. This presented somewhat of a conundrum.

Tobin posited that he needed to take account of the concentration of observations at zero when testing hypotheses about the relationship between household income and expenditure because an explanatory variable might have been expected to both influence the probability of whether a household spent 0 dollars on luxury items and how much they actually spent, given that they spent something.

If the probability of \$0 luxury expenditures were the only thing to explain, probit analysis would provide a suitable statistical model, but it is inefficient to throw away information on the value of the dependent variable when it is available. That is the case here because if a household spent something on luxury goods, we know how much they spent. If, on the other hand, if there were no concentrations at a lower limit, and we only cared to explain the amount of household luxury expenditure, multiple regression would be the appropriate statistical technique. But, since there is a concentration of values of the dependent variable at a limit (in this case \$0) OLS estimates are biased because the dependent variable is not continuous and unbounded. Tobin proposed a hybrid of these two methods as a solution to the problem, which now bears his name.

Before moving on to exactly how this model is calculated, we'll need to define some terminology. First, note that the dependent variable is Y^* , not Y . This is because the dependent variable is latent, it is not observed. In theory, household luxury expenditure extends along the length of the Y axis, in other words below \$0, but we do not observe those. If you are having trouble wrapping your mind around this you are not alone, it turns out that Tobit is not the right model to apply to this example, but we will stick with what Tobin did. Think of Y^* as the desire to spend on luxury items. Perhaps you have to reach a certain level of desire before you spend any money on luxury goods. Be that as it may, what we do observe is Y , which is how much the household spent, given

Tobit I

The Latent Model:

$$y_i^* = x_i' \mathbf{b} + u_i^*$$

But, we have censoring at $C = 0$:

$$y_i = y_i^* \text{ if } y_i^* > C;$$

$$y_i = C \text{ if } y_i^* \leq C$$

So, The Observed Model:

$$y_i = x_i' \mathbf{b} + u_i \text{ if } y_i > 0$$

$$y_i = 0 \text{ otherwise.}$$

The tobit model is generally represented in this way. First, we have a Latent model where the dependent variable is Y^* , has some independent variables and coefficients and a disturbance term that is normally distributed with a mean of zero. But, we have censoring at point C , which in our example is zero. Thus we have an observed Y that equals Y^* if the value of Y^* is greater than C , but equals C if the value of the unobserved Y^* is less than or equal to C .

The observed model, therefore, has a dependent variable Y , with some independent variables and coefficients, and an error term. Because of the censoring, however, the lower tail of the distribution of Y_i , and of u_i , is cut off and the probabilities are piled up at the cut-off point. The implication is that the mean of Y_i is different from that of Y^* and the mean of u_i is different from the mean of u^* (which is zero).

When we estimate that model we need to take account of the censoring. We note that the entire sample consists of two different sets of observations. The first set contains the observations for which the value of Y is zero. For these observations we know only the values of the X variables and the fact that Y^* is less than or equal to 0. The second set consists of all observations for which the values of both X and Y^* are known. The likelihood function of the Tobit is going to consist of each of these two parts.

Tobit II

In the Case of OLS: $E(y_i|x_i) = x_i'b$

$$L = \left(-\frac{n}{2} \right) \log(2\pi s^2) - \frac{1}{2} \sum_{i=1}^n \left[\frac{(y_i - bX_i)}{s} \right]^2$$

If we censor y at C=0:

$$= E(y_i|x_i) = \text{pr}(y_i > 0 | x_i) E(y_i | y_i > 0, x_i)$$

$$\prod_0 (1 - \Phi_i)$$

$$\prod_1 \Phi_i$$

$$\prod_{i=1}^n \frac{1}{s} \frac{f\left(\frac{y_i - x_i b / s}{s}\right)}{\Phi_i}$$

$$L = \sum_{y_i > 0} -\frac{1}{2} \left[\log(2\pi s^2) + \left[\frac{(y_i - bX_i)}{s} \right]^2 \right] + \sum_{y_i = 0} \log \left[1 - \Phi \left(\frac{bX_i}{s} \right) \right]$$

First, it is useful to consider the case of straight OLS. If we were using maximum likelihood to estimate the ols equation (say if we had actual luxury expenditure amounts for every household), we would be trying to estimate the unknown parameters of the model so that the probability of observing the Ys we observed is as high as possible. We find the maximum of the log likelihood function. And here it is.

This may look odd to you for one of two reasons. First, the lime green background is disconcerting, but I couldn't figure out how to get microsoft's equation editor to print lime green text, oh well. Second, you might not be familiar with maximum likelihood. If that is the case, let me reassure you this is correct. We know the Yi's are distributed normally with a mean of BX and a variance of sigma squared, and we know the density function of the a normally distributed variable. This equation is what results when you substitute that density function into the joint probability density function of the Y's and take the log of the equation. I present it like this, however, because it is clear what Tobit is doing when we look at the likelihood functions.

If, instead, we were censoring Y at 0 the expectation of Y given X would be a little different. It would be equal to the probability of Y exceeding zero, given the various covariates, multiplied by the expectation of Yi given that it exceeds zero and given the covariates. This turns out to be rather simple. Considering what we know about the sample... First, we know we assume that the disturbance has a normal distribution, that the errors for different observations are independent, and that the error terms is independent of the explanatory variables. Second, for all the sample households, we know whether or not they spent something on luxury goods. Third, for the noncensored observations, we know how much they spent. We use these three pieces of information to form the likelihood function.

First, for all observations, we know whether or not they were censored, so they contribute the likelihood, taken over all observations, of the probability they were censored.

Second, the uncensored observations contribute the product, taken over all uncensored observations, of the probability that they were uncensored.

Finally, for the uncensored observations, we know the amount of their expenditure, hence they contribute the density function for a truncated normal distribution. Putting these three terms together, doing a little math, and taking the log, we get the log likelihood of the tobit model. Notice that (CLICK) the part of the log likelihood summed over the uncensored observations is identical to the likelihood function for a simple OLS. The other term (CLICK) accounts for the likelihood the censored observations were censored.

Tobit Model: Caveat Emptor

Interpreting Coefficients

1. Expected value of the underlying latent variable (Y^*)

$$E(Y^* | x_i) = x_i' \mathbf{b}$$

2. Estimated probability of exceeding C

$$pr(y_i > C) = \Phi\left(\frac{x_i' \mathbf{b}}{\sigma}\right)$$

3. Expected, unconditional value of the realized variable (Y)

$$E(y_i | x_i) = \Phi\left(x_i' \mathbf{b} + \sigma \frac{f_i}{\Phi_i}\right) + (1 - \Phi_i)C$$

4. Expected Y_i conditional on exceeding C

$$E(y_i | y > C, x_i) = x_i' \mathbf{b} + \sigma \frac{f_i}{\Phi_i} + C$$

If we decide to estimate a tobit model, there are a couple of things we need to be aware of. First, is the tricky interpretation of the coefficients.

1. Most statistical packages estimate coefficients that are related to the latent variable Y^* . Thus, taken by themselves, each shows the effect of a change in a given x variable on the expected value of the latent variable, holding all other x variables constant. In other words, with respect to the latent variable, tobit betas can be interpreted in just the same way as the betas are from a regular OLS model. That said, this is not very useful because we do not observe the latent variable. If we did, we would not be estimating a tobit model. Other interpretations are more relevant.
2. Alternatively we could use the coefficient to calculate the probability that observations will exceed C . In this case the interpretation is the same as the interpretation in a probit model except that the coefficients need to be divided by σ . This is because, whereas σ is not estimable separately from the betas in probit, it is separately estimable in a tobit model.
3. The expected value of the observed y is equal to the relevant coefficient weighted by the probability that an observation will be uncensored. The greater this probability the bigger is the change in the expectation of y for a fixed change in a particular x .
4. Finally, we could calculate the expected value of y conditional on y exceeding the censoring threshold. All of these expectations look a little complicated, but we can easily get 1, 3, and 4 from postestimation commands in stata. I will tell you how to do next.

There are some other caveats that apply to the tobit model (having to do with the assumptions inherent in regression analysis). They apply equally to censored and sample selected models, so I am going to discuss them at the end of the

Estimating Tobit in Stata 7.0

Estimating the model

tobit y x1 x2..., ll and/or ul

Post Estimation Commands

Predict *newvar* → $E(y_i | x_i) = \Phi_i \left(x_i' b + s \frac{f_i}{\Phi_i} \right) + (1 - \Phi_i) C$

e *newvar* → $E(y_i | y > C, x_i) = x_i' b + s \frac{f_i}{\Phi_i} + C$

ystar *newvar* → $E(Y^* | x_i) = x_i' b$

The command to estimate a tobit model in stata is tobit. The model equation is laid out as the equations are laid out for other regression models. That is first the dependent variable then a list of the independent variables. You can, but do not have to, tell stata where your data is censored with the lower limit and upper limit commands. It is thus possible to estimate a model on data that is left censored, right censored, or both.

Three of the four interpretations of the coefficients from the last slide can be estimated with the postestimation commands.

First, the usual predict command gives you the third type of interpretation from the previous slide. That of the observed y conditional on x.

Second, e calculates the expected value of the observed y conditional on it being uncensored, the fourth interpretation from the previous slide.

Finally, the command ystar calculated the expected value of the latent dependent variable y* conditional on the xs.

Sample Selection Models

☐ Tobit Model Limitations

- Same set of variables, same coefficients determine both $pr(\text{censored})$ and the DV
- Lack of theory as to why obs. are censored


☐ Selection Models

- Different variables and coefficients in censoring (selection equation) and DV (outcome equation)
 - Allow theory of censoring, obs. are censored by some variable Z
 - Allow us to take account of the censoring process because selection and outcome are not independent.
-

The Tobit model has some notable limitations that can be remedied with the use of a sample selection model in its place. (CLICK) First, in the tobit model the same set of variables and coefficients determine both the probability that an observation will be censored and the value of the dependent variable. (CLICK) Second, this does not allow a full theoretical explanation of why the observations that are censored are censored. It is easy to see why this may be important, and I will demonstrate with Tobin's original example in a moment.

(CLICK) Sample selection models address these shortcomings by modifying the likelihood function. (CLICK) First, a different set of variables and coefficients determine the probability of censoring and the value of the dependent variable given that it is observed. These variables may overlap, to a point, or may be completely different. (CLICK) Second, sample selection models allow for, in my opinion, greater theoretical development because the observations are said to be censored by some other variable, which we call Z . (CLICK) This allows us to take account of the censoring process, as we will see, because selection and outcome are not independent.

Sample Selection Models I



Recall the example Tobin used to motivate the development of the tobit model. Here we were trying to explain Household Luxury Expenditure as a function of Household income. We only observed Expenditure for those households that spent more than \$0. There could be a very clear theoretical reason why some households do not purchase luxury items. (CLICK) My take on this is that perhaps the censoring occurs at the poverty line. You could easily have a theory that specifies this... households below the poverty line are primarily concerned with subsistence and have little or no money to spend on luxury items. In the framework of the sample selection model, you could specify one equation for whether or not a household is at or below the poverty line, and a different equation for how much that household spent on luxury items, given that it is above the poverty line. In fact, as Heckman demonstrated, if the processes are related, estimating a model of luxury expenditure with out first estimating an equation of whether or not the household was below the poverty line, would lead to biased results. To see this, lets consider the inner workings of the Heckman Model.

The Form of Sample Selection Models

$$z_i^* = w_i' \mathbf{a} + e_i$$

$$z_i = 0 \text{ if } z_i^* \leq 0;$$

$$z_i = 1 \text{ if } z_i^* > 0$$

Selection Equation

$$y_i^* = x_i' \mathbf{b} + u_i$$

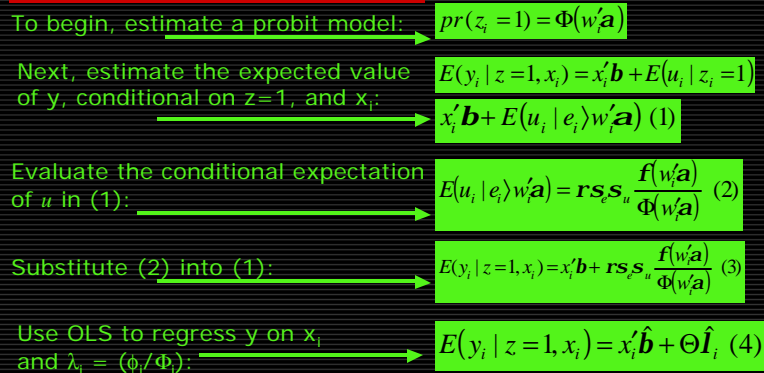
$$y_i = y_i^* \text{ if } z_i = 1$$

$$y_i \text{ not observed if } z_i = 0$$

Outcome Equation

The basic idea of a sample selection model is that the outcome variable, y , is only observed if some criterion, defined with respect to a variable z , is met. The common form of the model has two stages. In the first stage, a dichotomous variable z determines whether or not y is observed, y being observed only if $z=1$ (and you estimate a model with some matrix of independent variables w and get some coefficients α , the model is estimated, of course, with an error term, e); in the second state, we model the expected value of y , conditional on its being observed. So, we observe z , a dummy variable, which is a realization of an unobserved (or latent) continuous variable z^* , having a normally distributed, independent error, e , with a mean zero and a constant variance σ^2 . For values of $z=1$, we observe y , which is the observed realization of a second latent variable (and model that with some independent variables X and get a vector of coefficients β), y^* , which has a normally distributed, independent error, u , with a mean zero and a constant variance σ^2 . The two errors are assumed to have a correlation ρ . The joint distribution of u and e is bivariate normal.

Where does the Bias Come From?



In order to see where this bias comes from, let us consider heckman's selection model is slightly more detail.

(CLICK)To begin with, we estimate a probit model for the probability that $z=1$, in our example, for the probability that a household's income is above the poverty line. This model is estimated with all of our observations using a set of covariates called w and yielding a coefficient vector α .

(CLICK)The second step is to estimate the expected value of the outcome dependent variable, y , conditional on $z=1$ and the variables denoted x . This yields a coefficient vector β . Skipping ahead a few steps, we end up with equation (1).

(Click)To evaluate the conditional expectation of U in equation (1) we make use of the fact that the expected value of one of the variables in a bivariate distribution (in this case U) censored with respect to the value of the other variable (in this case e) is given by equation (2).

(CLICK)Inserting equation (2) into equation (1) we get equation (3), which gives the expected value of y given that $z=1$ – exactly what we are looking for in our outcome equation.

(CLICK)To estimate the OLS, we first that the probit results and, for the subsample for whom $z=1$, we compute the estimate of little ϕ over big Φ , the inverse mills ratio, symbolized by λ . Then, for this same subsample, we use OLS to regress y on X and on our estimate of λ . This will yield estimates of of the familiar vector of coefficients (β), and of θ , which is the covariance between u and e . Equation (4) shows that the resulting estimates of the vector β , in general, will be biased if the variable λ has been omitted. The problem of sample selection bias thus becomes equivalent to a misspecification problem arising through the omission of a regressor variable. There are only two cases where bias will not be a problem: First, if $\rho = 0$, second, if the correlation between the estimate of λ and any x variable is zero. We will come back to this last point when we estimate one of these models on the machines.

The Likelihood Function

$$L = \sum_0 \log(1 - \Phi_i) + \sum_1 \log \Phi \left[\frac{w'_i \mathbf{a} + r \left(\frac{y_i - \mathbf{x}'_i \mathbf{b}}{s_u} \right)}{(1 - r^2)^{\frac{1}{2}}} \right] + \sum_1 -\frac{1}{2} \left[\log(2ps_u^2) + \left[\frac{(Y_i - \mathbf{b}'X_i)}{s_u} \right]^2 \right]$$

If $\rho = 0$

Probit

OLS

$$L = \sum_0 \log(1 - \Phi_i) + \sum_1 \log \Phi(w'_i \mathbf{a})$$

$$L = \sum_1 -\frac{1}{2} \left[\log(2ps_u^2) + \left[\frac{(Y_i - \mathbf{b}'X_i)}{s_u} \right]^2 \right]$$

The likelihood function for the Sample selection model is quite complicated, Heckman did win a nobel prize for this, but showing it illustrates the same point in a different way. (CLICK) Note here that if $\rho = 0$ the likelihood function can be split into two parts: a probit for the probability of being selected and an OLS regression for the expected value of Y in the selected subsample.

Furthermore, because these two parts share no common parameters, they can be estimated separately. This shows that if there is no residual correlation between e and u , the simple OLS approach is all we need to explain Y. Herein lies the most important fact about sample selection models: it is not the fact that observations on Y are only available for a selected sample that causes the difficulty; rather, it is that the selection is not random with respect to Y. We will revisit some of these points when we interpret the results from the Heckman model we are about to estimate.

Onto the Machines!

1. Log in, if you are not already logged in.
 2. Open Stata
 3. Open a log file so you can save the stuff we do today.
 4. Type: `setmem 100m`
 5. Open: `"I:\general\Methods Lunch\heckman.dta"`
 6. Open Notepad
 7. Open: `"I:\general\Methods Lunch\heckcode"`
 8. Copy the first line of text in notepad into the Stata command line to estimate a Heckman Model.
-

Alright, now we officially know enough about sample selection models to be able to estimate one. So, (CLICK) log in, if you have not already done so. If you do not have a network account raise your hand and I will come log you in. (CLICK) Open Stata, there should be an icon on the desktop. If not do it from the start menu. (CLICK) The first thing you should do is open a log file so you can save all we do today. If you don't have a network place to save, save it to the desktop and we will get you a disk to transfer the file to. (CLICK) We need to increase the memory stata is using, so type `setmem 100m`. (CLICK) Now get the data, some real live dissertation data I would point out, from the I: drive at the location on the screen here. (CLICK) After the data comes up in stata, open a notepad. (CLICK) Once the notepad is open, find and open the text file called `heckcode`. It is in the same file as the data was on the I: drive. We'll be cutting and pasting some stuff from here to minimize the amount of typing you need to do. (CLICK) Copy that first line of text from the `hackcode` file into the command line in stata. Note that you will need to use `edit... paste` in stata, right mouse clicking does not work. After that is pasted, hit enter. Stata is now estimating a sample selection model on the data.

Interpreting the Coefficients

$$\frac{\partial E(y | z^* > 0, x)}{\partial x_k} = b_k - a_k \rho \sigma_u d(-w a)$$

For LN_CAPRATIO...

$$\beta_k = 139.95 \quad \alpha_k = -.48 \quad \rho = -.194 \quad \sigma_u = 48.89$$

$$\text{gen real_incap} = 139.95 - (-.48 * -.194 * 48.89 * \text{Dpr})$$

Variable	Obs	Mean	SD.	Min	Max
real_incap	31170	135.9	.162	135.6062	137.0012

Once the model is estimated we are probably interested in substantively interpreting the coefficients in the outcome equation. You may think that, because the outcome equation is analogous to an OLS, interpretation is easy. But, this is not the case for some variables in the model. If a variable appears ONLY in the outcome equation the coefficient on it can be interpreted as the marginal effect of a one unit change in that variable on Y. If, on the other hand, the variable appears in both the selection and outcome equations the coefficient in the outcome equation is affected by its presence in the selection equation as well. Sigelman and Zeng give us the marginal effect of the kth element of x on the conditional expectation of y as this formula. (CLICK) Beta is the coefficient in the outcome equation, alpha is the corresponding coefficient in the selection equation, rho (which stata spits out) is the correlation between the errors in the two equations and sigma is the error from the outcome equation (which stata also spits out) and d is a function of the inverse mills ratio (remember that from a few slides back), which we can compute.

An example is useful, because this rather daunting formula can be pretty easily computed in stata. Notice the variable LN_Capratio appears in both the outcome and selection equations in the model we just ran, so interpreting that beta as an OLS coefficient would be wrong. We have all of the components to this equation right on our output (CLICK), except for the function d. In order to get that function we first generate the probability that an observation will be selected. Go to the heckprob text file and copy the next line of code (predict selxbpr, xbs) and paste it into the command line. Hit return. Then we need to generate the inverse mills ratio, so go back to the text file and copy the next line of code (gen testpr = normden(selxbpr)/norm(selxbpr)) into the command line, and hit return. Finally, we need to get the proper function, which is the inverse mills ratio multiplied by the inverse mills ratio plus the probability of being selected, so go back to the text file and copy the next line of code (gen Dpr = testpr*(testpr+selxbpr)) and paste it into the command line. That is it, you have everything you need to adjust the coefficients in the outcome equation, so you can calculate the real ln capratio coefficient with this formula (CLICK) which you can also copy out of the text file and paste into the command line so you don't have to type it. This will give you an estimate of the adjusted beta for every observation in the data, so what sigelman and zeng say to do is use the average of these, and assess the sensitivity. So, sum the new variable (which you can do by copying the final line of code from the text file into the command line), and you should get something that looks like this (CLICK). Not too bad, the average beta is close to the estimated beta and the sensitivities are tight.

Interpreting Rho

Is Risky...

Generally if ρ is negative... any unobserved thing that makes selection more likely makes y less, and the converse. E.G. Resolve

The nature of ρ makes it very sensitive to model specification

Some people who use sample selection models do not even mention rho, others try to interpret it. I think interpretation is a risky strategy (CLICK).

Generally speaking if rho is negative, any component of the error that makes selection more likely makes y less. So in the equation we just estimated recall that dyads were selected if they had a militarized interstate dispute and Y was the severity of the dispute. I could think of a possible explanation for this. There are plenty of things we cannot observe (or not observe well) in international relations. Erik Gartzke wrote an article in the journal international organization about this called, war is in the error term. One such thing in this case could be resolve. So the interpretation would go something like this... when two states have a disagreement they are more likely to militarize that disagreement the more resolved they are to winning it, but once the dispute is militarized both sides realize the resolve of the other and, not wanting to have a long, costly conflict, work out a solution short of war (so Y is less).

But, I don't interpret rho in my dissertation because its nature makes it extremely sensitive to model specification. (Click) Remember rho is the correlation between the errors in the selection and outcome equations. Errors are, of course, necessarily tied up with model specification. So, alternative specifications change the errors, which in turn changes rho. Moreover, the correlation, I think, should be thought of as intrinsic to the model. In other words, we assume rho does not equal zero in the theoretical model that we posit for the population and not simply for the sample in which we may have omitted some variable common to x and z . Thus, whatever is the cause of the correlation between u and e should be inherently unmeasurable.

Censored Probit

...When the dependent variable in the outcome equation is dichotomous

1. Type: *clear*
2. Open: *"I:\general\Methods Lunch\censored probit.dta"*
3. Copy the next line of code from the heckcode notepad file, notice the command for censored probit is *heckprob*.
4. Your output should look exactly the same as the heckman model, note the location of the selection and outcome equations, the censoring, and the estimate of ρ .

It is possible to estimate a sample selection model when both dependent variables (selection and outcome) are dichotomous, but we would not run a regular heckman model (this would be akin to running an ols on a binary dependent variable), we would do a censored probit. I will dispense with the likelihood function stuff, you can trust me on this (or consult Greene or Dubin and Rivers), but the censored probit is like running two probits linked by correlated errors. If the errors are uncorrelated, you can unlink the likelihood and run separate models on the selection and outcome equations.

We can provide a stata example, but you will need to open a new dataset. (Click) First, type *clear* in the command line. (Click) Now, go back to the I general folder where you found the first dataset and open the dataset entitled censored probit. While it is opening let me tell you a little about the data. It is from a paper that Paul Fritz and I have that tests whether great power actually balance as often as balance of power theory says they do. We find they do not, and we are going to estimate an equation to explain the likelihood of bandwagoning behavior (which is the opposite of balancing). The only reason I bring this up is because I want to give you access to two very cool programs I have written to aid in the interpretation of this type of model.

(Click) Once the data is opened, go back to the heckcode notepad and copy the next line of code into the stata command line. Note that the command for a censored probit is *heckprob* (see if you win a nobel prize, the folks at stata will name not one, but two, commands after you.)

Censored Probit: How Well Does the Selection eq. predict?

A major assumption of selection modeling is that you have a robust prediction of whether an observations will be censored.

- Capture Matrices from your estimation and draw uncertain estimates from the normal distribution - DON' T DO THIS!!!

matrix params = e(b)

matrix P = e(V)

drawnorm b16-b33, means(params) cov(P)

2. In stata, click file and then do.

3. A dialog box will pop up, you should change the folder to:

"I:\general\Methods Lunch\PCPtest.txt"

Note: you will need to change the file type to all files

One of the most important things in selection modeling is having a robust selection equation. We can see from the censored probit estimates that you just got, that the selection equation in this case has lots of very statistically significant coefficients. And I can tell you that they are all in the correct direction. Turns out we know a lot about whether great powers will sign alliances. Yet, having significant coefficients in and of itself is not enough to say that the selection equation is robust. One thing we might think about doing it running the selection equation by itself (as a probit) and seeing what percentage of cases it correctly predicts.

Generally, what people do in this case is to calculate xb and say those predictions with likelihoods greater than .5 are 1's and those predictions with likelihoods less than .5 are zeros, then compare this to their dependent variable and see how many they got right. Michael Herron (Political Analysis 1999?) argued (in a very clarify sort of way) that this is less than correct because we are not incorporating the uncertainty that is associated with our model estimates (the standard errors) into the prediction. He argued, instead, we should use simulation to calculate the percent correctly predicted, and here is how. (CLICK)

We need to get estimates from our model that reflect some sort of uncertainty. We can do this by capturing the betas and the variance-covariance matrix post-estimation in stata, they we can draw uncertain estimates from the normal distribution where the mean for each beta is its estimate from the probit model we just ran and the variance for each beta is taken directly from the variance-covariance matrix we just capture. A couple of notes: DON'T DO THIS NOW. It would take about 20 minutes because the default in stata is to draw as many estimates as there are observations in your data (in this case 10980). I have already drawn them, and they are in the data set. Second, the b16-b33 are the random names I named these new variables, they are 16th through the 33rd variables in the censored probit if you started counting from the top. (Click)

I have written a do file to do the rest, so in stata, click file and do. When the dialog box pops up change the folder to I/General/Methods Lunch/ePCPtest.txt. Note you will have to change the file type to all files because this is a text file and not a stata .do file. Stata will now do a 10 loop simulations, that basically calculate linear prediction for each observation in the data set, factoring uncertainty into the estimates, according the the Herron method. I would encourage you, if you are interested, to compare this do file to the herron article. If the programming and code looks unfamiliar, I would encourage you to come back for my PRL brownbag on programming in stata.

The rest of what you need to do is in a do file in the same folder on the I drive that the rest of the stuff is in.

Results!

Variable	Obs	Percentile	Centile	[95% Conf. Interval]	
p_ally	10	2.5	.9133	.9133	.9133
		50	.9153	.9133	.9168
		97.5	.9181	.9171	.9181

Percent Correctly Predicted
91.53%
(91.33%, 91.81%)

Not too shabby!

In general, we would want to run more simulations than just ten, usually a thousand, but that would take a long time. What we see here, is that our selection equation really predicts very well. You could say something like... it predicts 91.53% of Great Power alliance onsets between 1816 and 1992 correctly, and we are 95% sure that it predicts at least 91.33% correctly. Cool, huh?

Uncertainty and First Differences

You have already got all of the stuff that you need to do
Clarify-like things to your results!

```
drop b16-33
```

```
matrix params = e(b)
```

```
matrix P = e(V)
```

```
drawnorm b1-b33, means(params) cov(P)
```

For instance, calculating first differences with confidence intervals:

1. In stata, click *file* and then *do*.

2. A dialog box will pop up, you should change the folder
to: "I:\general\Methods Lunch\clarify.txt"

Note: you will need to change the file type to all files

One of the most influential stats articles in recent years was King et al.'s "Making the Most of Statistical Analysis" in which they introduced their software clarify. One of the central arguments of the article (as we've sort of touched on a couple of times up until now) was that while it is laudable that political scientists are increasingly turning to substantive interpretation of their model results through predicted probabilities, first differences, and the like; they often do it wrong. Calculating a predicted probability requires, basically, multiplying Xb – which is what most people do. However, the betas are uncertain – remember they have standard errors attached to them! (Click)

Clarify is a great program that supports a number of statistical models. Unfortunately sample selection models are not among them, but we have everything we need to do clarify like stuff on our results already in the data. The first thing you would need to do, but don't do here, is parallel to what we did to calculate the percent correctly predicted. We need to reestimate the censored probit model, capture the vector of parameters and the variance-covariance matrix, and draw uncertain betas for all parameters in the model. Actually, if you were doing this on the exact dataset in front of you, you would have to drop the betas you generated for the percent correctly predicted estimate. So, don't do this (CLICK).

Now we can run a simulation program to generate first differences, with clarify-like uncertainty, for, say, the first variable in the model. (CLICK) In stata, click file... do..., change the file type to all files and surf over to our folder on the I: drive, and select the clarify text file. What this program is doing, and see Rich Timpone's 2002 Political Analysis article for details, is generating an uncertain base model prediction for the probability of bandwagoning (remember, that was our dependent variable in the outcome equation). Then, it is calculation an uncertain model prediction for when we add one standard deviation to that first variable. We will be able to take the results from this simple .do file and calculate one first difference.

Results!

Variable	Obs	Percentile	Centile	[95% Conf. Interval]
base_bwagon	10	2.5	.0046109	.0046109 .0054579
		50	.0100439	.005462 .0192198
		97.5	.025488	.0216371 .025488
intsim1_m2sd	10	2.5	.0000729	.0000729 .000089
		50	.0008054	.0001188 .0025754
		97.5	.0051542	.0031591 .0051542

The loop should take about 30 seconds to run. (CLICK) What I have included in the loop is a program to generate an uncertain XB, here the median of that prediction (the likelihood of a bandwagoning alliance) is about .01 with 95% confidence intervals of .0046 and .025. Then the program subtracts two standard deviations from the first variable in the model and recalculates XB, which you can again centile to get the median prediction (in this case around .00085) and the 95% confidence intervals around that (in this case $7.29e^{-5}$ and .005). You can then calculate the percentage differences from the base to the altered model. Just eyeballing this one, we can see that a two standard deviation decrease in this variable is going to have a large substantive effect on the likelihood of bandwagoning – which is good, because that is what the paper was about. In practice, I should note, we would do a larger number of loops (generally 1000) and the confidence intervals would be tighter.

Results II--for the entire model

Minus 2 SD	Minus 1 SD	Continuous Variables	Plus 1 SD	Plus 2 SD
-94.9 (-98.9, -85.5)	-71.1 (-87.1, -51.7)	Interest Similarity _{int}	220.3 (126.7, 259.7)	469.1 (217.6, 691.1)
41.5 (23.3, 49.3)	9.9 (-3.1, 10.0)	Interest Similarity _{instat}	-35.2 (-57.0, -16.1)	-51.0 (-72.5, -24.6)
-12.1 (-15.3, -4.7)	-6.4 (-8.3, -2.2)	Regime Dissimilarity _{int}	7.1 (-2.9, 10.4)	15.5 (-0.6, 23.7)
15.2 (13.1, 15.4)	7.4 (4.0, 8.1)	Regime Dissimilarity _{instat}	-6.2 (-7.0, -0.5)	-10.4 (-12.6, -0.1)
-9.1 (-52.0, 29.9)	-6.4 (-39.6, 14.4)	Common Threat _{int}	-0.2 (-4.1, 7.5)	4.0 (-12.4, 7.4)
-5.9 (-21.4, 6.6)	-4.0 (-14.0, 3.4)	Common Threat _{instat}	0.8 (0.2, 1.2)	4.2 (2.6, 7.6)
86.0 (48.0, 116.2)	39.1 (23.1, 56.2)	Military Capability	-24.9 (-27.6, -11.7)	-46.5 (-55.5, -25.3)
12.9 (-10.7, 30.0)	7.0 (-4.8, 14.1)	Threat	-1.0 (-1.9, -1.1)	-5.1 (-5.3, -2.3)
44.8 (-8.9, 66.7)	17.2 (16.9, 21.8)	Security	-24.6 (-46.7, 4.2)	-49.1 (-78.7, 19.1)
36.5 (24.4, 37.3)	25.0 (15.9, 25.9)	Free Great Powers	-16.2 (-22.2, -9.6)	-32.2 (-42.3, -19.8)
0		Dichotomous Variables	1	1
18.3 (9.0, 19.9)		Current War _{int}	-21.8 (-33.6, -12.6)	
-1.1 (-3.1, 0.0)		Current War _{instat}	2.2 (-0.1, 3.3)	
4.4 (-3.3, 10.0)		Colonial Contiguity _{int}	-3.5 (-9.3, 3.4)	
-0.3 (-0.6, 0.2)		Colonial Contiguity _{instat}	0.3 (-0.1, 0.4)	

If we wrote a much longer program, we could calculate these types of uncertain first differences for all of the variables in the model, for say plus or minus one or two standard deviations. If you did that, you would end up with a table that looked like this, except you would be able to read it.

Related Models: Event Counts

Event Count Models, in general...

The Hurdle Model

Reference: Mullahy. 1986. "Specification and Testing of Some Modified Count Data Models." *Journal of Econometrics* 33:341-65

The ZIP Model

Reference: King. 1989. "Event Count Models for International Relations: Generalizations and Applications." *International Studies Quarterly* 33:123-47

Or King. 1989. Unifying Political Methodology.

There are probably models related to the sample selection model in a number of different estimators of which I do not know. Sorry about that, you'll have to go out and find the stuff yourself. Consult Greene first.

(CLICK) One area in which there are clear links, however, is event count models. Event count models, in general, deal estimate the number of occurrences of an event, where the count of occurrences is non-negative and discrete. Typically they employ the poisson distribution to estimate the count. Such models have been used to account for a number of diverse events, such as the famous study that analyzed the number of soldiers kicked to death by horses in the Prussian Army (Bortkewitsch 1898). Sometimes there is no reason to believe any of the possible counts are qualitatively different, but sometimes there is reason to believe that the zero counts are different from the non-zero counts.

(CLICK)

Hurdle models account for this qualitative difference in the data generating process for zeros and nonzeros. In this model, a binary probit model determines whether a zero or a nonzero outcome occurs, then, in the latter case, a truncated poisson distribution describes the positive outcomes. The key here is that the nonzero outcomes in the second stage are exactly that, not zero. This presents conceptual problems when we begin to think of suitable applications, as will become clear in a minute. To be honest, I don't know of a statistical package that estimates hurdle models (perhaps `limdep`?), but this is the appropriate reference. (CLICK)

(CLICK) Far more popular are the zero inflated models, like the zip model (this stands for zero inflated poisson). In the ZIP model the outcome can arise from one of two processes. In the first the outcome is always zero. In the other the poisson process is at work and the outcome can be any nonnegative number, zero included. Basically this model is akin to running a logit or a probit, linked to the count model for the nonnegative numbers. Consider an example from the stata manual... we may count how many fish arch visitor to a park catches. A large number of visitors may catch zero, because they do not fish (as opposed to being unsuccessful). We may be able to model whether a person fishes depending on a number of covariates related to fishing activity (camper, child, male...) and we may model how many fish a person catches depending on a number of covariates having to do with fishing (lure, bait, time of day, temperature...). This type of model is estimable in stata with the command `zip`

Applications in Political Science

- | | |
|---|--|
| <ul style="list-style-type: none">❑ American Politics■ Rich Timpone <i>APSR</i> 1998, <i>PA</i> 2002■ Kevin Grier et al. <i>APSR</i> 1994■ McCarty and Rothenberg <i>AJPS</i> 1996■ Jay Goodliffe <i>AJPS</i> 2001■ Adam Berinsky <i>AJPS</i> 1999 | <ul style="list-style-type: none">❑ International Relations■ Paul Huth, <i>Standing Your Ground</i> 1996■ Bill Reed <i>AJPS</i> 2000■ Bill Reed and Doug Lemke <i>AJPS</i> 2001■ Special Issue Of <i>IJ</i> 2002■ Poe and Meernick <i>JoPR</i> 1995 |
|---|--|
-

There are lots of applications of these types of models in American Politics and International Relations. Here are a few. Generally, I would say that there are roughly twice as many American politics applications as there are IR applications, but this is really a hot methodology in IR. Also, notice the quality of the journals...

Censored, Sample Selected, and Truncated Variables

Sample	y variable	x variables
Censored	y is known only if some criterion df. in terms of y is met.	x variables are observed for the entire sample.
Sample Selected	y is observed only if a criteria df. in terms of another variable (Z) is met.	x and w are observed for the entire sample.
Truncated	y is known only if some criterion df. In terms of y is met.	x variables are observed only if y is observed.