

## Chapter 329

# Zero-Inflated Poisson Regression

## Introduction

The zero-inflated Poisson (ZIP) regression is used for count data that exhibit overdispersion and excess zeros. The data distribution combines the Poisson distribution and the logit distribution. The possible values of  $Y$  are the nonnegative integers: 0, 1, 2, 3, and so on.

The results presented here are documented in the books by Cameron and Trivedi (2013) and Hilbe (2014) and in Garay, Hashimoto, Ortega, and Lachos (2011).

This program computes ZIP regression on both numeric and categorical variables. It reports on the regression equation as well as the confidence limits and likelihood. It performs a comprehensive residual analysis including diagnostic residual reports and plots.

## The Zero-Inflated Poisson Regression Model

Suppose that for each observation, there are two possible cases. Suppose that if case 1 occurs, the count is zero. However, if case 2 occurs, counts (including zeros) are generated according to a Poisson model. Suppose that case 1 occurs with probability  $\pi$  and case 2 occurs with probability  $1 - \pi$ . Therefore, the probability distribution of the ZIP random variable  $y_i$  can be written

$$\Pr(y_i = j) = \begin{cases} \pi_i + (1 - \pi_i)\exp(-\mu_i) & \text{if } j = 0 \\ (1 - \pi_i)\frac{\mu_i^{y_i}\exp(-\mu_i)}{y_i!} & \text{if } j > 0 \end{cases}$$

where  $\pi_i$  is the logistic link function defined below.

The Poisson component can include an exposure time  $t$  and a set of  $k$  regressor variables (the  $x$ 's). The expression relating these quantities is

$$\mu_i = \exp(\ln(t_i) + \beta_1 x_{1i} + \beta_2 x_{2i} + \cdots + \beta_k x_{ki})$$

Often,  $x_1 \equiv 1$ , in which case  $\beta_1$  is called the *intercept*. The regression coefficients  $\beta_1, \beta_2, \dots, \beta_k$  are unknown parameters that are estimated from a set of data. Their estimates are symbolized as  $b_1, b_2, \dots, b_k$ .

This logistic link function  $\pi_i$  is given by

$$\pi_i = \frac{\lambda_i}{1 + \lambda_i}$$

where

$$\lambda_i = \exp(\ln(t_i) + \gamma_1 z_{1i} + \gamma_2 z_{2i} + \cdots + \gamma_m z_{mi})$$

## Zero-Inflated Poisson Regression

The logistic component includes an exposure time  $t$  and a set of  $m$  regressor variables (the  $z$ 's). Note that the  $z$ 's and the  $x$ 's may or may not include terms in common.

## Solution by Maximum Likelihood Estimation

The regression coefficients are estimated using the method of maximum likelihood. The logarithm of the likelihood function is

$$\mathcal{L} = L1 + L2 - L3$$

where

$$L1 = \sum_{\{i:y_i=0\}} \ln[\lambda_i + \exp(-\mu_i)]$$

$$L2 = \sum_{\{i:y_i>0\}} \{y_i \ln(\mu_i) - \mu_i - \ln(y_i!)\}$$

$$L3 = \sum_{i=1}^n \ln(1 + \lambda_i)$$

The gradient of  $\mathcal{L}$  is

$$\frac{\partial \mathcal{L}}{\partial \beta_r} = \sum_{\{i:y_i=0\}} \left[ \frac{-x_{ir}\mu_i}{\lambda_i \exp(\mu_i) + 1} \right] + \sum_{\{i:y_i>0\}} [y_i - \mu_i] x_{ir}, \quad r = 1, 2, \dots, k$$

$$\frac{\partial \mathcal{L}}{\partial \gamma_r} = \sum_{\{i:y_i=0\}} \left[ \frac{z_{ir}\lambda_i \exp(\mu_i)}{\lambda_i \exp(\mu_i) + 1} \right] - \sum_{i=1}^n \left[ \frac{\lambda_i}{1 + \lambda_i} \right] z_{ir}, \quad r = 1, 2, \dots, m$$

The second derivatives are

$$\frac{\partial^2 \mathcal{L}}{\partial \beta_r \partial \beta_s} = \sum_{\{i:y_i=0\}} \frac{x_{ir}x_{is} \mu_i [(\mu_i - 1)\lambda_i \exp(\mu_i) - 1]}{[\lambda_i \exp(\mu_i) + 1]^2} - \sum_{\{i:y_i>0\}} \mu_i x_{ir}x_{is}, \quad r, s = 1, 2, \dots, k$$

$$\frac{\partial^2 \mathcal{L}}{\partial \gamma_r \partial \gamma_s} = \sum_{\{i:y_i=0\}} \frac{z_{ir}z_{is} \lambda_i \exp(\mu_i)}{[\lambda_i \exp(\mu_i) + 1]^2} - \sum_{i=1}^n \frac{z_{ir}z_{is} \lambda_i}{(\lambda_i + 1)^2}, \quad r, s = 1, 2, \dots, m$$

$$\frac{\partial^2 \mathcal{L}}{\partial \beta_r \partial \gamma_s} = \sum_{\{i:y_i=0\}} \frac{x_{ir}z_{is} \lambda_i \mu_i \exp(\mu_i)}{[\lambda_i \exp(\mu_i) + 1]^2}, \quad r = 1, 2, \dots, k; s = 1, 2, \dots, m$$

---

## Distribution of the MLE's

The asymptotic distribution of the maximum likelihood estimates is multivariate normal as follows

$$\begin{bmatrix} \hat{\beta} \\ \hat{\gamma} \end{bmatrix} \sim N \begin{bmatrix} \beta \\ \gamma \end{bmatrix} \left( \begin{array}{cc} \frac{\partial^2 \mathcal{L}}{\partial \beta_r \partial \beta_s} & \frac{\partial^2 \mathcal{L}}{\partial \beta_r \partial \gamma_s} \\ \frac{\partial^2 \mathcal{L}}{\partial \beta_r \partial \gamma_s} & \frac{\partial^2 \mathcal{L}}{\partial \gamma_r \partial \gamma_s} \end{array} \right)^{-1}$$

---

## Akaike Information Criterion (AIC)

Hilbe (2014) mentions the Akaike Information Criterion (AIC) as one of the most commonly used fit statistics. It is calculated as follows

$$AIC = -2[\mathcal{L} - k]$$

Note that  $k$  is the number of predictors including the intercept.

---

## Residuals

As in any regression analysis, a complete residual analysis should be employed. This involves plotting the residuals against various other quantities such as the regressor variables (to check for outliers and curvature) and the response variable.

### Raw Residual

The raw residual is the difference between the actual response and its expected value estimated by the model. Because we expect the variances of the residuals to be unequal, there are difficulties in the interpretation of the raw residuals. However, they are still popular. The formula for the raw residual is

$$r_i = y_i - \hat{\mu}_i(1 - \hat{\pi}_i)$$

### Pearson Residual

The Pearson residual corrects for the unequal variance in the residuals by dividing by the standard deviation of  $y$ . The formula for the Pearson residual is

$$p_i = \frac{y_i - \hat{\mu}_i(1 - \hat{\pi}_i)}{\sqrt{\hat{\mu}_i(1 - \hat{\pi}_i)[1 + \hat{\mu}_i \hat{\pi}_i]}}$$

---

## Variable Selection

Because of the complexity of the model, this routine does not have a direct variable selection capability. A reasonable stepwise strategy is as follows: remove the model term (other than the intercepts) with largest  $p$ -value over 0.200 and rerun. Repeat until all  $p$ -values are less than a threshold such as 0.20.

---

## Data Structure

At a minimum, datasets to be analyzed by ZIP regression must contain a dependent variable and one or more independent variables. Long (1990) presents a dataset of 915 rows that he uses as an example in his regression book: Long (1997). This dataset contains five independent variables (Female, MentorArts, Prestige, Married, Children) and one dependent variable (Articles).

### Long 1990 Dataset

Female	MentorArts	Prestige	Married	Children	Articles
0	8	1.38	1	2	3
0	7	4.29	0	0	0
0	47	3.85	0	0	4
0	19	3.59	1	1	1
0	0	1.81	1	0	1
0	6	3.59	1	1	1
0	10	2.12	1	1	0
0	2	4.29	1	0	0
0	2	2.58	1	2	3
0	4	1.8	1	1	3

---

## Missing Values

If missing values are found in any of the independent variables being used, the row is omitted. If only the value of the dependent variable is missing, that row will not be used during the estimation process, but its predicted value will be generated and reported on.

## Example 1 – Zero-Inflated Poisson Regression using the Long 1990 Dataset

Long (1997) discusses a dataset used as an example of Zero-Inflated Poisson regression. This dataset contains five independent variables (Female, MentorArts, Prestige, Married, Children) and one dependent variable (Articles). These variables are defined as follows

- Articles**      Number of articles published during the last 3 years of Ph.D.
- Female**        1 if female scientist; 0 if male scientist.
- MentorArts**    Number of articles published by the scientist mentor during the last 3 years.
- Prestige**        Prestige of the scientist’s Ph.D. department.
- Married**        1 if married; 0 otherwise.
- Children**       Number of children 5 or younger.

The dataset can also be used to validate the program since the results of this model are given in Long (1997), page 246.

In this example, we will fit a Zero-Inflated Poisson regression model to these data.

### Setup

To run this example, complete the following steps:

- 1 Open the Long 1990 example dataset**
  - From the File menu of the NCSS Data window, select **Open Example Data**.
  - Select **Long 1990** and click **OK**.
- 2 Specify the Zero-Inflated Poisson Regression procedure options**
  - Find and open the **Zero-Inflated Poisson Regression** procedure using the menus or the Procedure Navigator.
  - The settings for this example are listed below and are stored in the **Example 1** settings file. To load these settings to the procedure window, click **Open Example Settings File** in the Help Center or File menu.

Variables Tab

---

Dependent Y.....**Articles**  
 Numeric X's (Poisson Model Variables).....**Female, Married, Children, Prestige, MentorArts**  
 Numeric X's (Logistic Model Variables) .....**Female, Married, Children, Prestige, MentorArts**

Model Tab

---

Terms (Poisson Regression Model).....**1-Way**  
 Terms (Logistic Regression Model) .....**1-Way**

Zero-Inflated Poisson Regression

Reports Tab

---

All Available Plots ..... **Checked** (click the *Check All* button)  
 Incidence ..... **Checked**  
 Incidence Counts ..... **0 1 2 3 4**  
 Exposure Value ..... **1**

Plots Tab

---

All Available Plots ..... **Checked** (click the *Check All* button)

**3 Run the procedure**

- Click the **Run** button to perform the calculations and generate the output.

**Run Summary**

Run Summary			
Item	Value	Rows	Value
Dependent Variable	Articles	Rows Processed	915
Exposure Variable	None	Rows Used in Estimation	915
Frequency Variable	None	Observations with Y = 0	275 (30.1%)
Number of Parameters in Model	12		
Log-Likelihood	-1604.7739		
AIC(1)	3233.5478		
Number of Likelihood Iterations	8 of 100		
Convergence Setting	1E-09		
Relative Log-Likelihood Change	1.310594E-13		

This report provides several details about the data and the MLE algorithm.

**Dependent, Exposure, and Frequency Variables**

These variables are listed to provide a record of the variables that were analyzed.

**Number of Parameters in Model**

This is the total number of parameters in the model. It includes those in the Poisson portion and in the logistic portion. Note that some variables may be in both portions, but they will of course have different parameters.

**Log-Likelihood**

This is the value of the log-likelihood that was achieved for this run.

**AIC(1)**

This is Akaike’s information criterion discussed above. It has been shown that using AIC to compare competing models with different numbers of parameters amounts to selecting the model with the minimum estimate of the mean squared error of prediction.

## Zero-Inflated Poisson Regression

**Number of Likelihood Iterations**

This is number of iterations used by the estimation algorithm.

**Convergence Setting**

When the relative change in the log-likelihood is less than this amount, the maximum likelihood algorithm stops. The algorithm also stops when the maximum number of iterations is reached.

**Relative Log-Likelihood Change**

This is the relative change of the log-likelihoods from the last two iterations.

**Rows Processed**

This is the number of rows read from the database. Rows with missing values and filtered rows are not included in the analysis.

**Rows Used in Estimation**

This is the number of rows used by the estimation algorithm. Rows with missing values and filtered rows are not included. Always check this value to make sure that you are analyzing all of the data you intended to.

**Sum of Frequencies**

This is the number of observations used by the estimation algorithm if you specified a Frequency Variable.

**Observations with Y = 0**

The gives the number and percentage of the observations in which Y is zero. Since this procedure is for the case in which there are too many zeros in the dataset, this value is important to consider.

---

**Means**

<b>Means</b>			
<b>Variable</b>	<b>Mean</b>	<b>Minimum</b>	<b>Maximum</b>
Articles	1.692896	0	19
Ps_Female	0.4601093	0	1
Ps_Married	0.6622951	0	1
Ps_Children	0.495082	0	3
Ps_Prestige	3.103372	0.755	4.62
Ps_MentorArts	8.767213	0	77
Lg_Female	0.4601093	0	1
Lg_Married	0.6622951	0	1
Lg_Children	0.495082	0	3
Lg_Prestige	3.103372	0.755	4.62
Lg_MentorArts	8.767213	0	77

This report gives the mean, minimum, and maximum of each variable. These values let you quickly determine if any of the data values are outside a reasonable range.

## Regression Coefficients

### Regression Coefficients

Parameter	Regression Coefficient b(i)	Standard Error Sb(i)	Z-Test of H0: $\beta(i) = 0$		95% Confidence Interval Limits for $\beta(i)$	
			Z-Statistic	Two-Sided P-Value	Lower	Upper
<b>Poisson Regression Model</b>						
Ps_Intercept	0.64031	0.12131	5.28	0.0000	0.40254	0.87808
Ps_Female	-0.20914	0.06340	-3.30	0.0010	-0.33341	-0.08487
Ps_Married	0.10379	0.07111	1.46	0.1444	-0.03559	0.24316
Ps_Children	-0.14331	0.04743	-3.02	0.0025	-0.23627	-0.05034
Ps_Prestige	-0.00600	0.03101	-0.19	0.8466	-0.06678	0.05478
Ps_MentorArts	0.01809	0.00229	7.88	0.0000	0.01360	0.02259
<b>Logistic Regression Model</b>						
Lg_Intercept	-0.57792	0.50935	-1.13	0.2565	-1.57624	0.42039
Lg_Female	0.10974	0.28009	0.39	0.6952	-0.43922	0.65870
Lg_Married	-0.35398	0.31762	-1.11	0.2651	-0.97651	0.26855
Lg_Children	0.21716	0.19648	1.11	0.2691	-0.16794	0.60226
Lg_Prestige	0.00158	0.14526	0.01	0.9913	-0.28313	0.28628
Lg_MentorArts	-0.13414	0.04526	-2.96	0.0030	-0.22285	-0.04543

### Estimated Models

#### Poisson Regression Model

Exp( 0.640306740121054 -0.209138863236085\*Female + 0.103785018739315\*Married -0.143305425134039\*Children -0.00599920296959671\*Prestige + 0.0180945143077501\*MentorArts )

#### Logistic Regression Model

Exp( -0.577923020403549 + 0.109741976031919\*Female -0.353982199937225\*Married + 0.217158482157067\*Children + 0.00157507191560552\*Prestige -0.134139805328132\*MentorArts )

#### Transformation Note:

Regular transformations must be less the 255 characters. If this expression is longer the 255 characters, copy this expression and paste it into a text file, then use the transformation FILE(filename.txt) to access the text file.

This report provides the estimated coefficients of the ZIP regression and associated statistics. It provides the main results of the analysis.

### Variable Selection

This report can be used to reduce the number of terms in the model. One variable selection strategy is to remove the model term (other than the intercepts) with largest p-value over 0.200 and rerunning. This can be repeated until all p-values are less than a threshold such as 0.20. In our example, we would definitely remove Prestige since its p-value is very high.



## Zero-Inflated Poisson Regression

**Parameter**

This item provides the name of the parameter shown on this line of the report. Parameters that begin with "Ps" are in the Poisson portion of the model. Parameters that begin with "Lg" are in the logistic portion of the model. The *Intercept* refers to the optional constant term. The *Alpha* value is the estimated value of the dispersion coefficient.

Note that whether a line is skipped after the name of the independent variable is displayed is controlled by the *Stagger label and output if label length is ≥* option in the Format tab.

**Regression Coefficient b(i)**

These are the maximum-likelihood estimates of the regression coefficients. Their direct interpretation is difficult because the formula for the predicted value involves the exponential function.

**Standard Error Sb(i)**

These are the asymptotic standard errors of the regression coefficients. They are an estimate the precision of the regression coefficient. The standard errors are the square roots of the diagonal elements of this covariance matrix.

**Z-Statistic**

This is the z-test statistic for testing the null hypothesis that  $\beta_i = 0$  against the two-sided alternative that  $\beta_i \neq 0$ . This is a Wald-type statistic. This test has been found to follow the normal distribution only in large samples.

The test statistic is calculated using

$$Z = \frac{b_i}{s'_{b_i}}$$

**Two-Sided P-Value**

The probability of obtaining a z value greater in absolute value than the above. This is the significance level of the test. If this value is less than some predefined alpha level, say 0.05, the variable is said to be statistically significant.

**95% Confidence Interval Limits for  $\beta(i)$  (Lower and Upper)**

These provide a large-sample confidence interval for the values of the coefficients. The width of the confidence interval provides you with a sense of how precise the regression coefficients are. Also, if the confidence interval includes zero, the variable is not *statistically significant*. The formula for the calculation of the confidence interval is

$$b_i \pm z_{1-\alpha/2} s'_{b_i}$$

where  $1 - \alpha$  is the confidence coefficient of the confidence interval and z is the appropriate value from the standard normal distribution.

**Estimated Regression Models**

These give the Poisson and logistic models in standard, full-precision format.

## Rate Ratios

Parameter	Regression Coefficient b(i)	Rate Ratio Exp(b(i))	95% Confidence Interval Limits for the Rate Ratio	
			Lower	Upper
<b>Poisson Regression Model</b>				
Ps_Female	-0.20914	0.811	0.716	0.919
Ps_Married	0.10379	1.109	0.965	1.275
Ps_Children	-0.14331	0.866	0.790	0.951
Ps_Prestige	-0.00600	0.994	0.935	1.056
Ps_MentorArts	0.01809	1.018	1.014	1.023
<b>Logistic Regression Model</b>				
Lg_Female	0.10974	1.116	0.645	1.932
Lg_Married	-0.35398	0.702	0.377	1.308
Lg_Children	0.21716	1.243	0.845	1.826
Lg_Prestige	0.00158	1.002	0.753	1.331
Lg_MentorArts	-0.13414	0.874	0.800	0.956

This report is mainly for binary (0-1) variables.

This report provides the rate ratio for each independent variable.

### Parameter

This item provides the name of the parameter shown on this line of the report. Parameters that begin with "Ps" are in the Poisson portion of the model. Parameters that begin with "Lg" are in the logistic portion of the model.

### Regression Coefficient b(i)

These are the maximum-likelihood estimates of the regression coefficients,  $b_1, b_2, \dots, b_k$ . Their direct interpretation is difficult because the formula for the predicted value involves the exponential function.

### Rate Ratio Exp(b(i))

These are the exponentiated values of the regression coefficients. The formula used to calculate these is

$$RR_i = e^{b_i}$$

The rate ratio is mainly useful for interpretation of the regression coefficients of indicator variables. In this case, they estimate the incidence in the given category relative to the category whose indicator variable was omitted (usually called the *control* group).

Zero-Inflated Poisson Regression

**95% Confidence Interval Limits for the Rate Ratio (Lower and Upper)**

These provide a large-sample confidence interval for the rate ratios. The formula for the calculation of the confidence interval is

$$\exp(b_i \pm z_{1-\alpha/2} s'_{b_i})$$

where  $1 - \alpha$  is the confidence coefficient of the confidence interval and  $z$  is the appropriate value from the standard normal distribution.

---

**Residuals**

Row	Articles (Y)	Conditional Mean of Y E(Y X,Z)	Residual		(T)
			Raw Y - E(Y X,Z)	Pearson Raw/ $\sigma$	
1	3	1.4989	1.5011	1.0704	1
2	0	1.7189	-1.7189	-1.1162	1
3	4	4.3346	-0.3346	-0.1603	1
4	1	2.4237	-1.4237	-0.8746	1
5	1	1.4924	-0.4924	-0.3197	1
6	1	1.6306	-0.6306	-0.4236	1
7	0	1.9122	-1.9122	-1.2391	1
8	0	1.6319	-1.6319	-1.0449	1
9	3	1.0998	1.9002	1.4729	1
10	3	1.5070	1.4930	1.0162	1
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.

This report provides the conditional mean (predicted value), the raw residual, and the Pearson residual. Large residuals indicate data points that were not fit well by the model.

## Predicted Means

Predicted Means							
Row	Articles (Y)	Poisson Mean $\mu$	Logit CDF Pr(Y = 0) $\pi$	Conditional Mean of Y E(Y X,Z)	Standard Error of E(Y X,Z) $\sigma$	95% Confidence Interval Limits for E(Y X,Z)	
						Lower	Upper
1	3	1.8111	0.1724	1.4989	1.4025	-1.2499	4.2477
2	0	2.0985	0.1809	1.7189	1.5399	-1.2994	4.7371
3	4	4.3390	0.0010	4.3346	2.0866	0.2449	8.4242
4	1	2.5169	0.0370	2.4237	1.6278	-0.7667	5.6141
5	1	2.0818	0.2831	1.4924	1.5401	-1.5262	4.5110
6	1	1.9894	0.1804	1.6306	1.4885	-1.2868	4.5480
7	0	2.1576	0.1138	1.9122	1.5432	-1.1125	4.9368
8	0	2.1266	0.2326	1.6319	1.5618	-1.4292	4.6930
9	3	1.6132	0.3183	1.0998	1.2901	-1.4288	3.6283
10	3	1.9394	0.2230	1.5070	1.4692	-1.3726	4.3866
.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.

This report provides the predicted values along with their standard errors and confidence limits. It also provides the mean of the Poisson portion of the model ( $\mu$ ) and the probability that  $Y = 0$  from the logistic portion of the model.

If you want to generate predicted values and confidence limits for  $X$  values not on your database, you should add them to the bottom of the database, leaving  $Y$  blank (if you are using an exposure variable, set the value of  $T$  to a desired value). These rows will not be included in the estimation algorithm, but they will appear on this report with estimated  $Y$ 's.

## Incidence when Exposure = 1

Incidence when Exposure = 1						
Row	Average Incidence Rate	Probability that Count is				
		0	1	2	3	4
1	1.4989	0.4668	0.1897	0.1222	0.0787	0.0507
2	1.7189	0.4453	0.1790	0.1213	0.0821	0.0556
3	4.3346	0.1881	0.1521	0.1236	0.1004	0.0816
4	2.4237	0.3109	0.1960	0.1402	0.1004	0.0718
5	1.4924	0.5157	0.1571	0.1061	0.0717	0.0484
6	1.6306	0.4545	0.1825	0.1214	0.0808	0.0538
7	1.9122	0.3944	0.1918	0.1310	0.0895	0.0612
8	1.6319	0.4781	0.1669	0.1135	0.0772	0.0525
9	1.0998	0.5791	0.1611	0.0994	0.0614	0.0379
10	1.5070	0.4873	0.1744	0.1151	0.0759	0.0501
.	.	.	.	.	.	.
.	.	.	.	.	.	.
.	.	.	.	.	.	.

This report gives the average incidence rate and estimated probabilities of various counts.

Zero-Inflated Poisson Regression

**Row**

The row number of the item. If you have excluded some rows by using a filter or if some of the rows had missing values, the row number identifies the original row on the database.

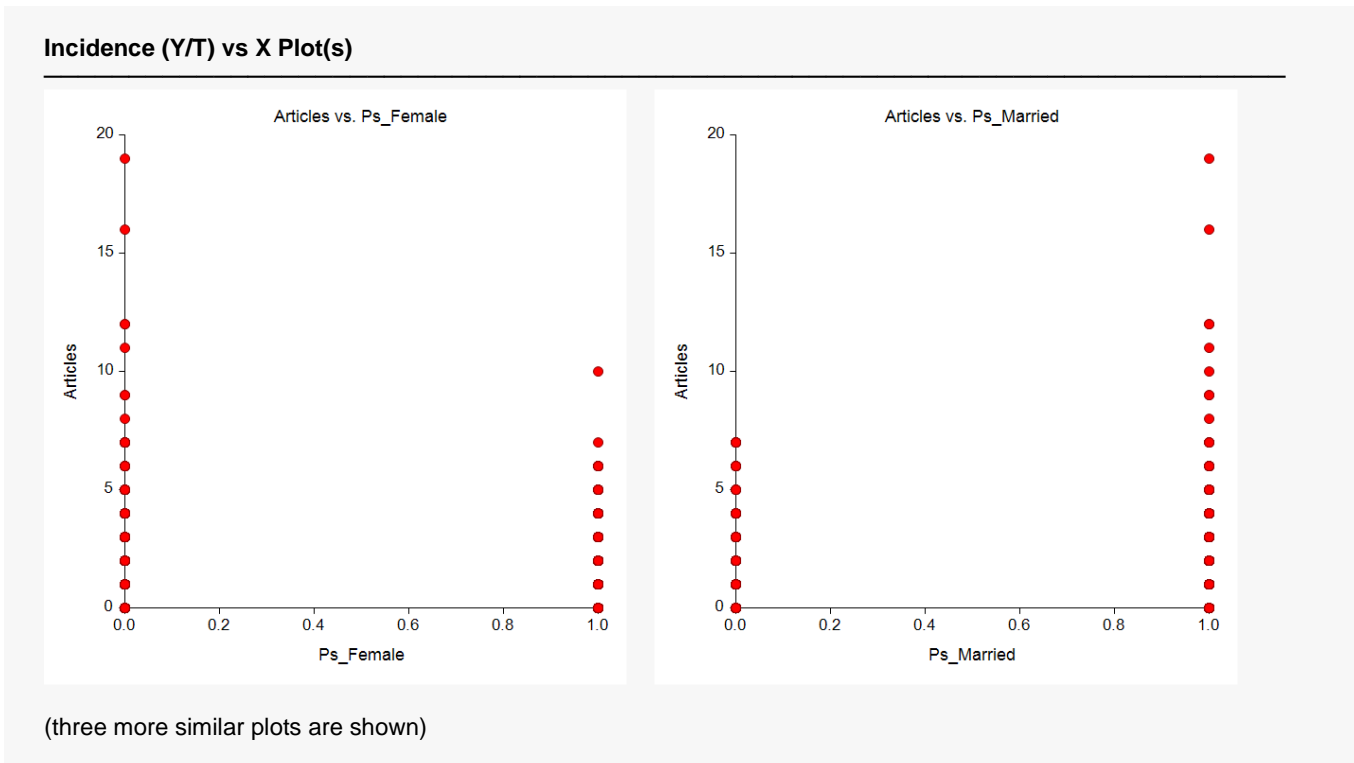
**Average Incidence Rate**

This is the predicted incidence rate. Note that the calculation is made for the specified exposure value, not the value of *T* on the database. This allows you to make valid comparisons of the incidence rates.

**Probability that Count is Y**

Using the ZIP model, the probability of obtaining exactly *Y* events during the exposure given in the Exposure Value box is calculated for the values of *Y* specified in the Incidence Counts box.

**Incidence (Y/T) vs X Plot(s)**

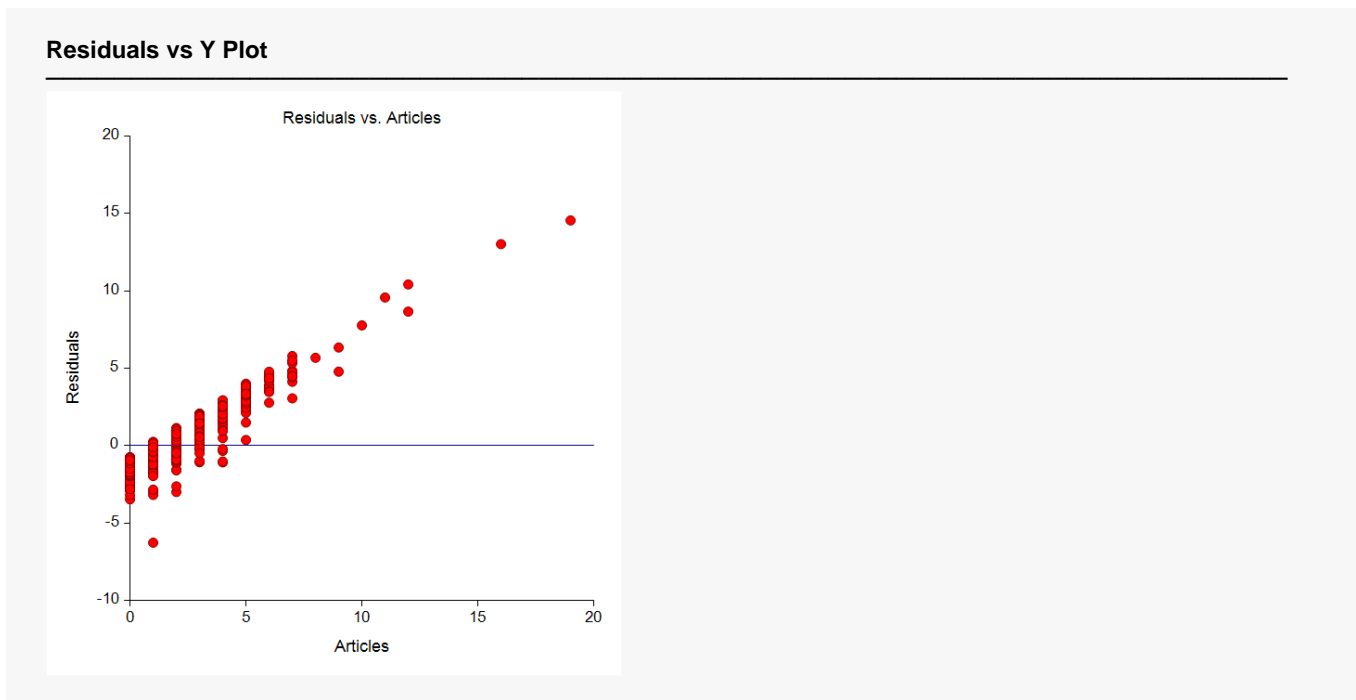


These plots show each of the independent variables plotted against the incidence as measured by *Y/T*. They should be scanned for outliers and curvilinear patterns.

**Incidence (Y/T) vs Z Plot(s)**

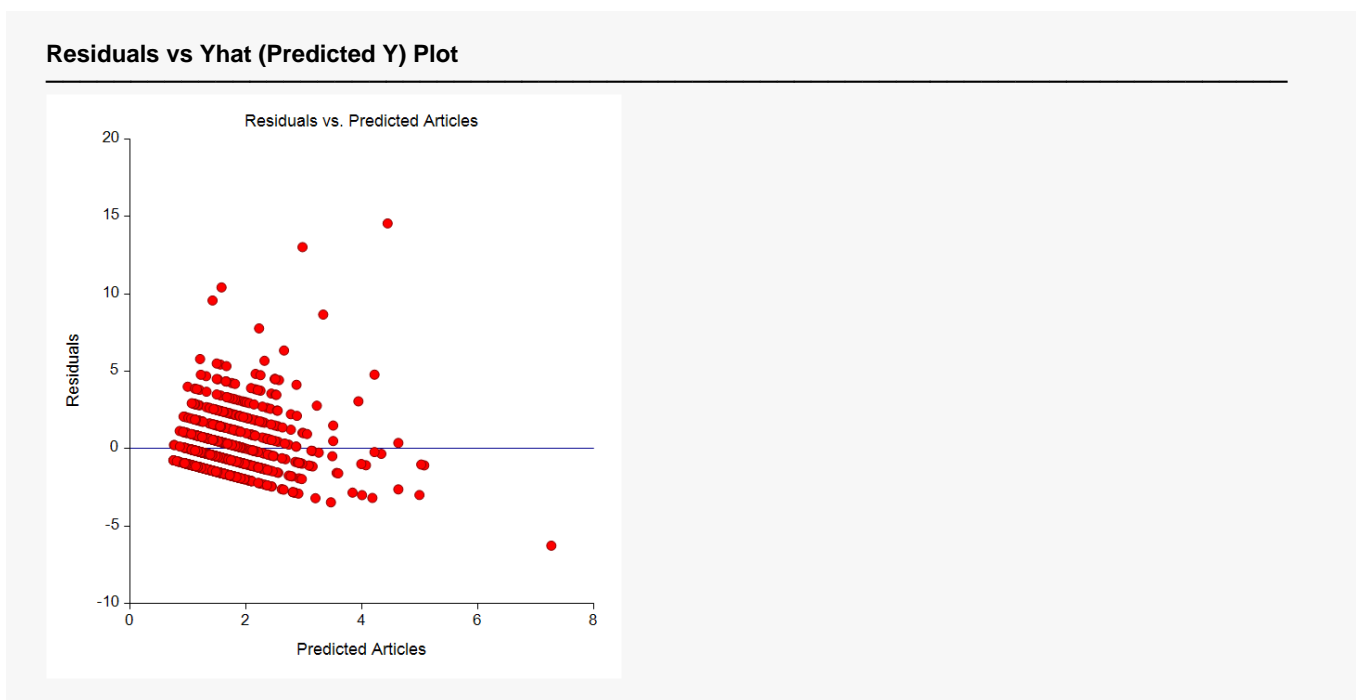
These plots are similar to the incidence versus *X* plots and are used for the same purpose, so we have not shown them here. They should be scanned for outliers and curvilinear patterns.

## Residuals vs Y Plot



This plot shows the residuals versus the dependent variable. It can be used to spot outliers.

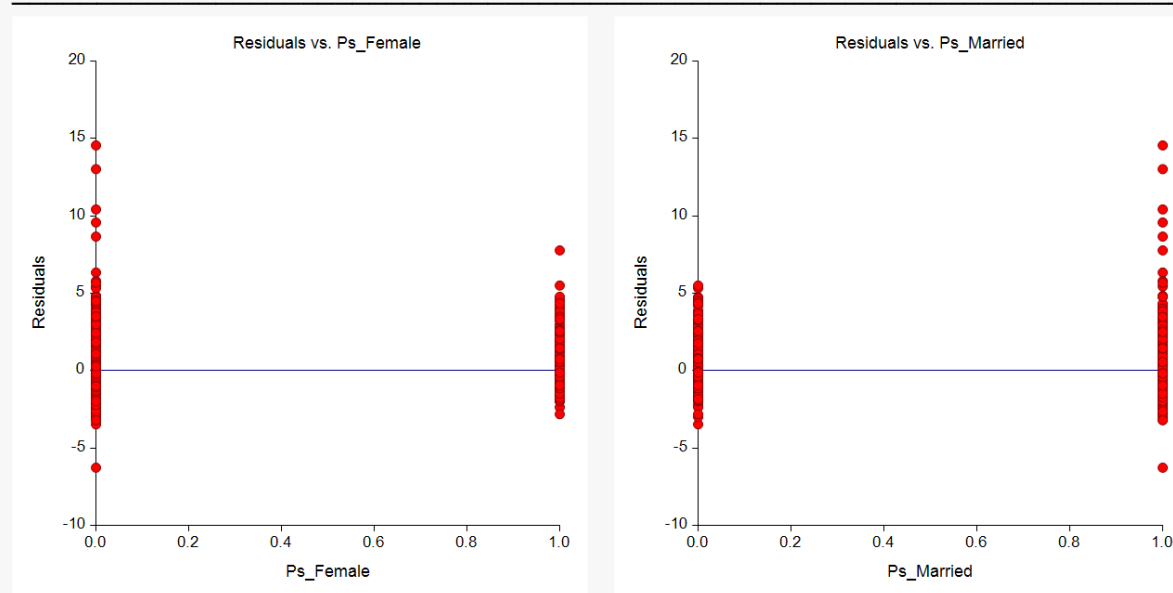
## Residuals vs Yhat (Predicted Y) Plot



This plot shows the residuals versus the predicted value (Yhat) of the dependent variable. It can show outliers.

## Residuals vs X Plots

### Residuals vs X Plot(s)




(three more similar plots are shown)

These plots show the residuals plotted against the independent variables. They are used to spot outliers. They are also used to find curvilinear patterns that are not represented in the regression model.

## Residuals vs Z Plot(s)

These plots are similar to the residual versus X plots and are used for the same purpose, so we have not shown them here. They should be scanned for outliers and curvilinear patterns. They are also used to find curvilinear patterns that are not represented in the regression model.

## Residuals vs Row Plot



This plot shows the residuals versus the row numbers. It is used to quickly spot rows that have large residuals.