Paper 448-2013

Estimating Censored Price Elasticities Using SAS/ETS®: Frequentist and Bayesian Approaches

Christian Macaro, Jan Chvosta, Kenneth Sanford, James Lemieux
SAS Institute Inc.

ABSTRACT

The number of rooms rented by a hotel, spending by "loyalty card" customers, automobile purchases by households—these are just a few examples of variables that can best be described as "limited" variables. When limited (censored or truncated) variables are chosen as dependent variables, certain necessary assumptions of linear regression are violated. This paper discusses the use of SAS/ETS® tools to analyze data in which the dependent variable is limited. It presents several examples that use the classical approach and the Bayesian approach that was recently added to the QLIM procedure, emphasizing the advantages and disadvantages that each approach provides.

INTRODUCTION

Economic theory suggests that as the price of a product or service rises, the demand for it falls, if all other factors are equal. This theory is often applied in the field of revenue management. As part of a company's optimization strategy, revenue management analysts are often asked to build models that allow prices to influence the number of units sold and therefore affect the company's revenue. A necessary component of this effort is the development of price elasticities. Price elasticity measures how quantities respond to changes in price. The analyst usually uses a data set that contains past prices and quantities, along with additional explanatory variables, and applies regression techniques to estimate this historical relationship. Price elasticities are used in many revenue management applications. Sophisticated revenue management systems for clothing retailers, hotels, and airline and sports ticket agencies incorporate price elasticities into their optimization routines. Analysts must give special consideration to the statistical or, more precisely, the econometric methods that they use to estimate these elasticities. Neglecting certain aspects of the data and the model can severely bias the elasticity estimate and lead to nonoptimal policy recommendations. Among the important considerations is awareness of the possible values that the data might realize.

This paper focuses on hotel data. The maximum number of hotel rooms that can be occupied on a given day is limited by the capacity of the hotel. This capacity is fixed over any reasonable period of time. Therefore, the maximum number of rooms that can be booked for any given day must be less than or equal to the capacity of the hotel. This type of data is commonly referred to as censored data (Maddala 1983). In the case of censored data, analysts must use special considerations to estimate price elasticities.

This paper presents a classic case of censoring and the econometric concerns that arise from ignoring censoring. It then presents a real-life example of censoring in the hotel industry. Finally, the paper shows how to estimate the model by using the newest Bayesian estimation techniques that are available in SAS/ETS 12.1.

ELASTICITY ESTIMATION

Suppose a management analyst is charged with estimating the price elasticity ϵ^p for a product or service. More formally, the price elasticity of demand is

$$\epsilon^p = \frac{\Delta\text{Quantity(Price)}/\text{Quantity(Price)}}{\Delta\text{Price}/\text{Price}}$$

where Quantity(Price) is the dependent variable in the regression of quantity on price of the product and Price is the price of the product. Based on this regression model, ordinary least squares (OLS) or some other method of estimation can provide consistent estimates of the price elasticity if the error in the regression is not correlated with included regressors and if the dependent variable is not censored. The problem for a hotelier (or for any analyst in a capacity-constrained industry) is far more complex. Specifically, the hotelier is limited by the number of rooms available. Therefore, the upper bound on the number of rooms violates the assumption of normality of the data and the assumptions of the classic normal regression model. This problem of a censored distribution is relatively minor when the data rarely reach the threshold. However, for goods whose capacity is highly constrained, such as tickets and hotel rooms, sellouts are fairly common. In these cases, the upper-censoring problem creates meaningful challenges for estimating the elasticity of demand.

This paper examines two upper-censoring cases. The first example simulates a typical censored regression problem with a fixed threshold, such that the bias in the elasticity is negative and causes the value of the elasticity to be smaller in absolute value than is accurate. This leads the analyst to conclude that consumers are less price sensitive than they have actually demonstrated themselves to be, which, in turn, leads to a higher-than-optimal price being set. In the second example, a real data set is analyzed. These data are characterized by a non-fixed upper censoring that causes a positive bias in the elasticity estimation. This leads the analyst to conclude that consumers are more price sensitive than they have actually demonstrated themselves to be. This result might suggest setting a lower-than-optimal price.

BIAS DUE TO CENSORING: A SIMULATED EXAMPLE

Econometric problems that are associated with the classic case of right censoring are prevalent in the literature (Maddala 1983). To illustrate the bias that is associated with ignoring upper censoring, this example generates data that have a coefficient on the price variable of –0.5 and have errors distributed normally. It estimates price elasticities under two scenarios: OLS on the uncensored data and OLS on the censored data. Figure 1 shows the price elasticities of the two estimates. A large bias is associated with ignoring the censoring. The magnitude of the bias in this case is approximately 30% in the downward direction. This is the typical effect on the magnitude of the estimates that are affected by right censoring. Figure 1 graphically presents the magnitude of the bias.

Figure 1 Effect of Censoring on Linear Regression

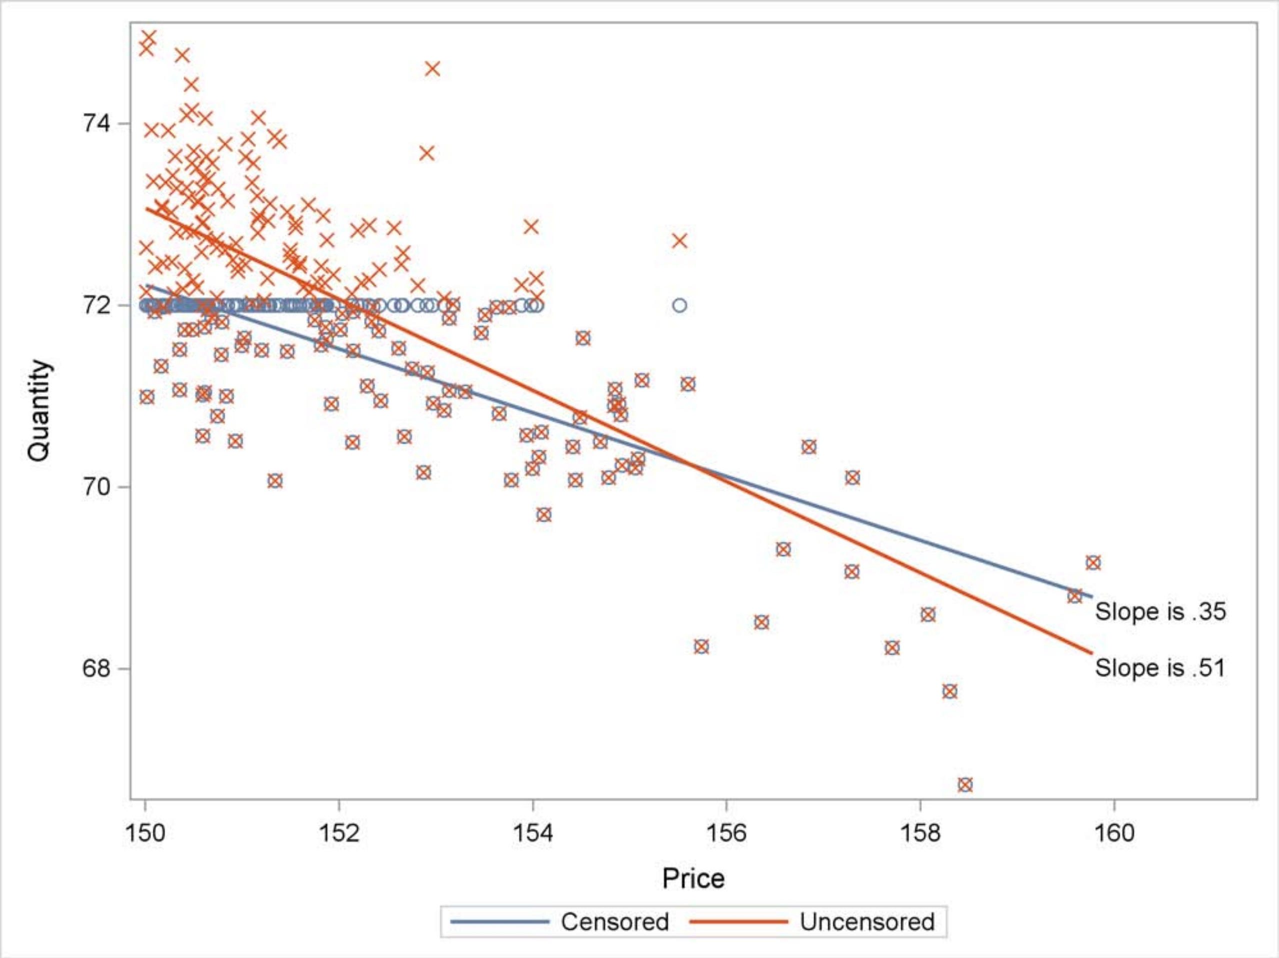

You can see the magnitude of these differences in Table 1, which evaluates the results of a $10 and $20 increase in the price of a room. Under the biased estimates in this simple example, a $10 increase in the price of a room leads to a $4,155 increase in the hotel's revenue for one night. This scenario suggests that the hotel guests are not very price sensitive and that an increase in price will actually cause revenue to increase. In this case guests still respond by booking fewer rooms, but the decrease in booking is offset by the higher prices. The revenue policy would be to increase the price that is charged in order to maximize revenue. The unbiased coefficients tell a different story. Based on the unbiased coefficient of –0.51, a $10 increase in price actually causes overall hotel revenue to fall by $91, and the drop is even greater for a $20 increase in price. Table 1 illustrates the importance of the coefficient in the statistical portion of a revenue management problem.

Table 1 Effect of Censoring on Revenues

	Biased (β=–0.35)	Unbiased (β=–0.51)
$10 price increase	$4,155	–$91
$20 price increase	$4,483	–$304
Means: Q=72; Price=$152		

In the classic case of right censoring (the type displayed in Figure 1), the coefficient estimates that are produced by OLS are asymptotically biased estimates of the true price elasticity of demand. The important factor for the revenue manager is that these estimates will lead to incorrect conclusions about the price responsiveness of customers.

The correction to this econometric problem must acknowledge that without this artificial barrier on the upper bound of rooms, more rooms would have been booked. Formally, this means that actual demand for rooms

(Quantity) is observed when room occupancy is below capacity, but as the demand exceeds the upper bound (C^*, a constant), only C^* is observed instead of the true demand for the rooms. Econometrically, this amounts to maximizing a modified log-likelihood function that takes censoring into consideration.

The next section shows the implications of ignoring censoring by using a more realistic example in the hotel industry. It demonstrates how censoring might actually occur in practice and how accurately estimating price responsiveness leads to more accurate pricing recommendations.

BIAS DUE TO CENSORING: A REAL-DATA EXAMPLE

The following data were collected for four different hotel properties (A, B, C, D) over a period of one year:

- price of booking
- occupancy (the number of booked rooms)
- occupancy rate (the overall occupancy percentage of the hotel)
- date when occupancy occurred

The data were collected daily for one-night stays only, for the same type of room, and with all bookings taking place relatively close to the guest's arrival. It is important to recognize that the occupancy and the occupancy rate are inversely related (see Figure 2). The censored observations (in red) are determined by the 90% threshold of "OccupancyRatio."

Figure 2 Occupancy versus Occupancy Rate and Occupancy versus Price Scatter Plots for Property A

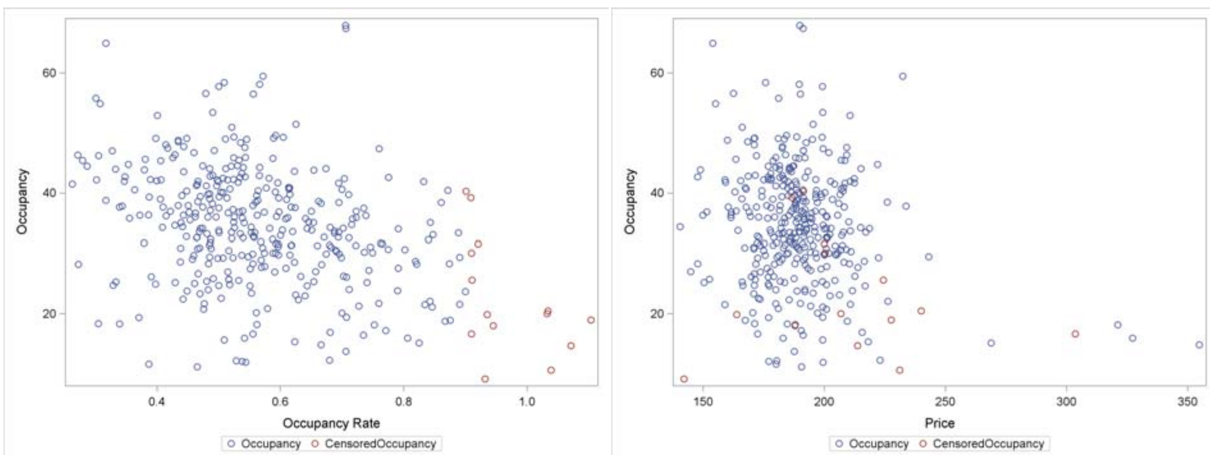

Figure 2 shows that higher occupancy rate is associated with fewer booked rooms. Although this behavior might look unusual, it is consistent with the type of room under analysis because one-night stays, booked relatively close to the arrival time, are the most likely to be affected by the overall availability of the hotel. The limited availability of one-night stays when the hotel is nearly full results in a form of right censoring, with a higher chance of the observation being censored when the observed booking rate is low. You might think that this phenomenon is the result of a large number of "group sales" for this night.

Seasonal Adjustment

The data in the analysis need to be seasonally adjusted. The bookings for one-night stays represent a time series of data that are collected over a one-year period, which is temporal in nature. The analyst needs to realize that the series contains temporal components such as seasons, trends, and cycles that must be removed from the series before proceeding with the analysis of demand for one-night hotel stays. For example, a hotel near the ocean is likely to be fully booked during the peak summer travel season. If the seasonal component that is related to the peak summer travel season is left in the series and the analyst proceeds with the analysis, the results are likely to be biased. Seasonal adjustment filters out particular temporal events that inflate or deflate the series of interest and bias the analysis.

The decomposition of the series can be accomplished in several ways. In many cases, a simple deterministic adjustment is sufficient. For example, a dummy variable for the peak summer season can be sufficient to explain why the demand for hotel rooms is higher in the peak summer months than in the rest of the year. When the seasonal cycle is more complex, a more flexible approach is needed in which the fluctuations and shifts in seasonal patterns are modeled dynamically. For example, the top left panel in Figure 3 shows the temporal predictability of price over time. Clearly, the price is higher in May and October than in the rest of the year. A closer look also reveals that the price shows a day-of-the-week effect, which seems to evolve dynamically over time.

Figure 3 Smoothing of Price

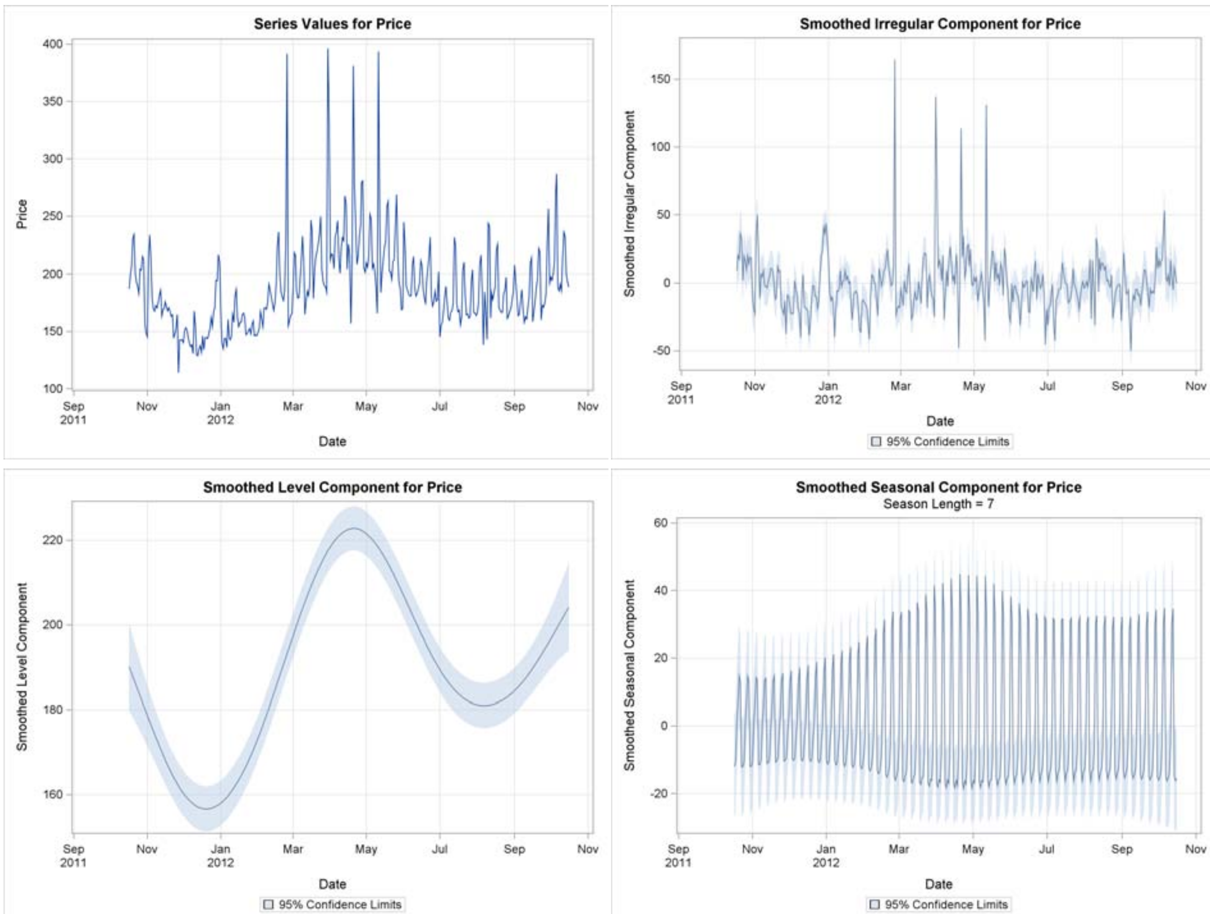

The UCM procedure in SAS/ETS allows for both deterministic and stochastic smoothing of a time series and produces an adjusted series that can be used in the demand analysis (Harvey 1989).

This example focuses on a decomposition in which the observed price is decomposed into three components:

$$\text{observed}_t = \text{level}_t + \text{seasonal}_t + \text{irregular}_t$$

The LEVEL statement in the UCM procedure captures the long-term fluctuations (following Hodrick and Prescott 1997, for simplicity), whereas the SEASON statement captures the day-of-the-week effect through a trigonometric functional approach:

```
proc ucm data=PropertyA;
     id Date  interval=day;
     model price;
     irregular plot=(smooth);
     level var=0 noest plot=smooth;
     slope var=0.001 noest;
     season length=7 type=trig plot=(smooth);
```

```
run;
```

Figure 3 shows the estimates of the three unobserved components. Notice that the irregular component is centered around 0, whereas the level component retains the original magnitude of the series. For interpretational purposes, you can shift the irregular component back to the original magnitude of the series by adding a term that represents the mean of the series.

Ordinary Least Squares and Censored Linear Regression Analysis

You can use the standard linear regression model to analyze the relationship between occupancy and price. Standard economic theory suggests that the two quantities are inversely related: high prices reduce the quantity demanded, and vice versa. Figure 4 shows the summary of the linear regression of occupancy versus price.

Figure 4 Regression of Occupancy versus Price

Parameter Estimates

Parameter	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	52.515435	4.399096	11.94	<.0001
Price	1	-0.095683	0.022981	-4.16	<.0001
_Sigma	1	9.816725	0.362840	27.06	<.0001

The coefficient on price is negative as expected and statistically significant at the 5% level, indicating that a $1 increase in price will reduce occupancy by 0.10 rooms. Moreover, if the price of the room were $180 and the hotel booked 35 rooms, the corresponding elasticity would be

$$\epsilon^{p} = -0.095683 \times \frac{180}{35} = -0.492084$$

This means that if the price increases from $180 to $181.80 (a 1% increase), then the corresponding demand decreases by almost 0.5%. Nevertheless, this estimate might not be accurate because of censoring: when the hotel is completely booked, a decrease in price will not increase the occupancy. You can use the QLIM procedure in SAS/ETS to estimate censored linear regression models.

Varying Hotel Censoring Point

As mentioned earlier, censoring occurs in hotel data because a hotel's capacity is restricted. In theory, censoring would occur at the maximum number of rooms in the hotel and would remain fixed as long as capacity remains the same. However, in practice this capacity does not remain fixed. The number of rooms available for sale each night depends on several factors, including the number of multinight stays (people already in the hotel), the number of rooms booked as part of a group reservation, and the number of rooms that are not in service.

You can formalize the theoretical number of rooms available for sale as

$$\text{Number of One-Night Stays} = \text{Total Hotel Capacity} - \text{Number of Rooms Not Booked as One-Night Stays}$$

The analysis here considers that all three sources of capacity restrictions influence the number of available rooms for one-night stays, which is the dependent variable in the model. These influences are not directly observable by the researcher. The analysis uses the overall capacity rate of the hotel across all booking types on any given night as a proxy for whether the number of rooms is upper censored. Overall hotel capacity is defined as

$$\text{Overall Capacity} = \frac{\text{Number of Rooms Sold across All Types}}{\text{Total Physical Capacity of Hotel}}$$

This value is bounded between 0 and 1. You can use the following statement to assign the upper bound to the capacity for each night:

```
proc qlim data=PropertyA;
     model Occupancy =  Price / censored(lb=0 ub=up);
     run;
```

The CENSORED option enables you to introduce an upper-censoring series (“up”) to detect which data were observed at a full occupancy rate. The CENSORED option also enables you to set a physical lower bound at 0: demand cannot be negative.

Figure 5 shows the summary of the censored linear regression of occupancy versus price when the upper censoring is determined by the 100% threshold of occupancy rate.

Figure 5 Regression of Occupancy versus Price with Censoring at 100% of Occupancy Rate

Parameter Estimates

Parameter	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	49.896157	4.402882	11.33	<.0001
Price	1	-0.080730	0.023046	-3.50	0.0005
_Sigma	1	9.712454	0.360665	26.93	<.0001

The coefficient on price is once again negative and significant at the 5% level. The addition of censoring at the hypothetical price of $180 and the occupancy equal to 35 rooms reduce the elasticity by approximately 15% and can be calculated as follows:

$$\epsilon^{p} = -0.0.080730 \times \frac{180}{35} = -0.4151829$$

You can further investigate the concept of upper censoring by considering a lower threshold. In Figure 5 an observation is considered censored if the occupancy rate is 100%. In practical terms, this means that no more rooms will be booked if the hotel is full. In many situations, this threshold can be lower than 100%. Usually, not all rooms at a hotel are evenly exchangeable. For example, it is unlikely that a luxury suite will be deeply discounted and booked as an economy room. Figure 6 shows the summary of the censored linear regression of occupancy on price when the upper censoring is determined by the 90% threshold of occupancy rate instead of the 100% threshold. This indicates that lowering the threshold has very little impact.

Figure 6 Regression of Occupancy versus Price with Censoring at 90% of Occupancy Rate

Parameter Estimates

Parameter	DF	Estimate	Standard Error	t Value	Approx Pr > |t|
Intercept	1	50.230885	4.497878	11.17	<.0001
Price	1	-0.080776	0.023562	-3.43	0.0006
_Sigma	1	9.599270	0.360472	26.63	<.0001

Bayesian Analysis of the Censored Linear Regression

The previous example analyzes the demand for hotel rooms by using censored linear regression, in which you obtain point estimates by maximizing the likelihood function (frequentist approach). For practical purposes, it might not be sufficient to obtain a point estimate that reveals how occupancy and price are related. In many situations, assessing the uncertainty can be just as important. From a hotel revenue management perspective, questions such as the following are very important:

- What is the probability that the hotel will be completely booked at a given price per room?
- What is the price that will predict a sold-out hotel with 70% probability?
- What is the probability that the occupancy will remain the same if the price is increased by 5%?

You can easily answer these questions through Bayesian analysis (Berger 1985). For example, Figure 7 shows the posterior predictive distribution of the censored regression of occupancy versus price when the focus is on the upper tail of price. Points outside the shaded region occur with low probability. The plot clearly shows that the probability of selling out the hotel is very low if the price of a room is greater than $200 and 50 rooms are available.

Figure 7 Posterior Predictive Distribution of the Censored Regression Model, Focusing on Price Greater Than $200

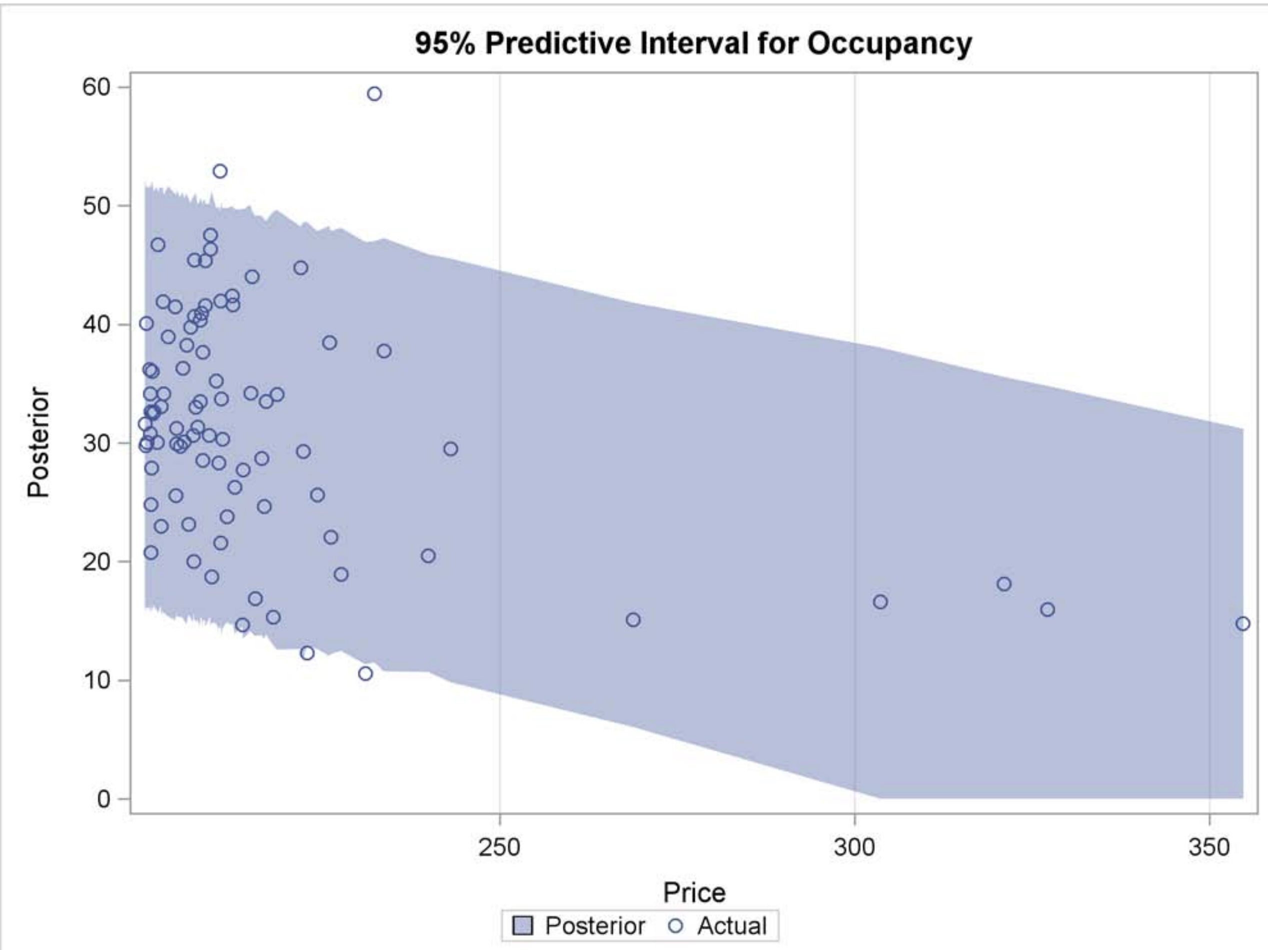

Bayesian analysis is theoretically quite different from its frequentist counterpart. Bayesian analysis focuses on analyzing the posterior distribution, which you can obtain by combining the information that is provided by the data (likelihood function) and the information that is not included in the data but is obtained from other sources (prior distribution). This framework is very general and flexible. You can use prior distribution to enhance theoretical and statistical properties of the model by introducing constraints that are often desirable in economic models—results from related studies (meta-analysis) and personal insight. Therefore, you should not judge the quality of a prior distribution based on how similar the conclusions of Bayesian and frequentist analyses are, even though some priors can determine similar conclusions. The next section shows how Bayesian analysis and the use of different priors can help improve the demand analysis for one-night hotel stays in which you observe a limited history for one property.

Analysis of Property B with Several Priors

Up to this point, all the analysis has focused on the same hotel, Property A. This section analyzes the data for a different hotel, Property B. As in the previous section, the analyst faces the problem of studying the demand curve for a hotel (Property B). Likewise, censoring should be an important aspect to take into account. The analyst would expect a price response for Property B to be qualitatively similar to that for Property A—that is, a negative price response, as economic theory suggests. However, a preliminary analysis shows contradictory results: price seems to have a positive effect on the number of booked rooms. This behavior is unusual and contradicts what the analyst learned from Property A. The data are inspected in Figure 8.

Figure 8 Scatter Plot of Occupancy versus Price for Property B

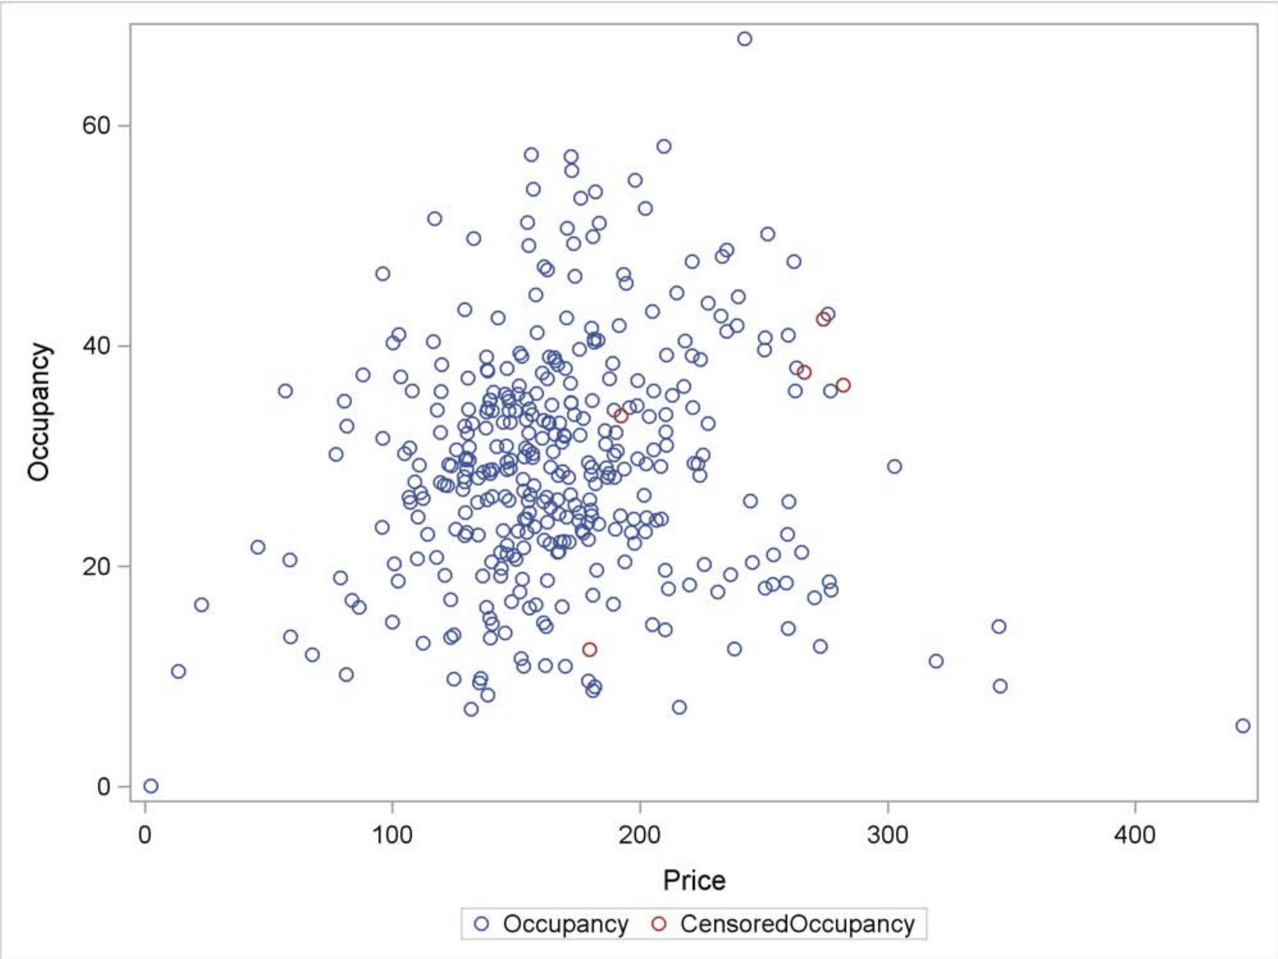

Although some of the observations look unusual (there are days when rooms are booked almost for free), the overall pattern of Property B is very similar to that of Property A (see Figure 2). A preliminary Bayesian analysis with diffuse priors confirms the anomalies (see Figure 9):

```
proc qlim data=PropertyB;
     model Occupancy =  Price / censored(lb=0 ub=UpperBound);
     bayes nmc=30000 nbi=10000;
     run;
```

The BAYES statement in PROC QLIM performs a Bayesian analysis of the model. If you do not specify a PRIOR statement, the QLIM procedure selects a set of diffused priors by default. The NMC=30000 option determines how many samples should be drawn to approximate the posterior distribution, and the

NBI=10000 option specifies how long the burn-in period should be (how many initial samples to discard so that the posterior distribution can be approximated well).

Figure 9 Censored Regression of Occupancy versus Price

Posterior Summaries

Parameter	N	Mean	Standard Deviation	Percentiles 25%	50%	75%
Intercept	30000	25.3699	1.9387	24.0749	25.3371	26.7223
Price	30000	0.0241	0.0111	0.0165	0.0241	0.0315
_Sigma	30000	10.7970	0.3968	10.5235	10.7833	11.0549

Economic theory clearly suggests that the relationship between occupancy and price must be negative, and yet Figure 9 shows a positive relationship (notice that similarly contradictory results are obtained with a frequentist approach). One way to correct the unexpected behavior of the censored linear regression for Property B is to include additional information through a prior distribution. The remainder of the paper discusses four different prior implementations:

- pseudo-hierarchical prior
- constrained diffuse prior
- constrained pseudo-hierarchical prior
- strong informative prior

Pseudo-hierarchical Prior

The original data set contains four properties. In the first part of the paper, the analysis was carried out for Property A. In this section, the goal is to improve the analysis of Property B, because the preliminary results with a diffuse prior provided contradictory results. This contradiction might be caused by several outliers in the data for property B or by the demand being inelastic at low prices.

One way to improve this analysis is to include what the analysis of Properties A, C, and D has revealed. (The analyses of properties C and D have been omitted from this paper.) You can use the following PRIOR statement in the QLIM procedure to specify a normal prior with mean and variance that reflect the information obtained from the other properties (Gelman et al. 1995):

```
Prior Price ~ Normal(mean= -0.037, var=0.00018);
```

Figure 10 shows the posterior summary of the censored linear regression that uses this pseudo-hierarchical prior.

Figure 10 Censored Regression of Occupancy versus Price: Pseudo-hierarchical Bayes

Posterior Summaries

Parameter	N	Mean	Standard Deviation	Percentiles 25%	50%	75%
Intercept	30000	29.5714	1.5587	28.5171	29.5671	30.6258
Price	30000	-0.00091	0.00866	-0.00673	-0.00086	0.00506
_Sigma	30000	10.8504	0.4146	10.5636	10.8311	11.1184

The posterior mean that is associated with the price coefficient is negative (–0.00076), as the economic theory suggests. This improvement is the result of introducing information from the analysis of Properties A, C, and D. The idea of introducing extra-experimental information through a prior distribution is essential for Bayesian analysis and, in many cases, can correct for anomalies in the data.

Constrained Diffuse Prior

Another way to ensure a negative posterior mean of the regression coefficient that is associated with price is to assume a uniform prior with an upper bound set at 0. This approach is justified by standard economic theory: demand is inversely related to price. You can use the following PRIOR statement in PROC QLIM to set such a prior distribution:

```
Prior Price ~ Uniform(max=0);
```

Introducing a sign restriction on a parameter through a prior distribution is common in Bayesian analysis and has a natural interpretation. Figure 11 shows the posterior predictive distribution of occupancy. This plot shows that occupancy and price are inversely related. The plot also shows that the posterior predictive distribution cannot reproduce some of the observations. Finally, you can use the posterior predictive distribution to get a general idea of how things might behave in an out-of-sample analysis.

Figure 11 Posterior Predictive Distribution of the Censored Regression Model When a Constrained and Diffuse Prior Is Used

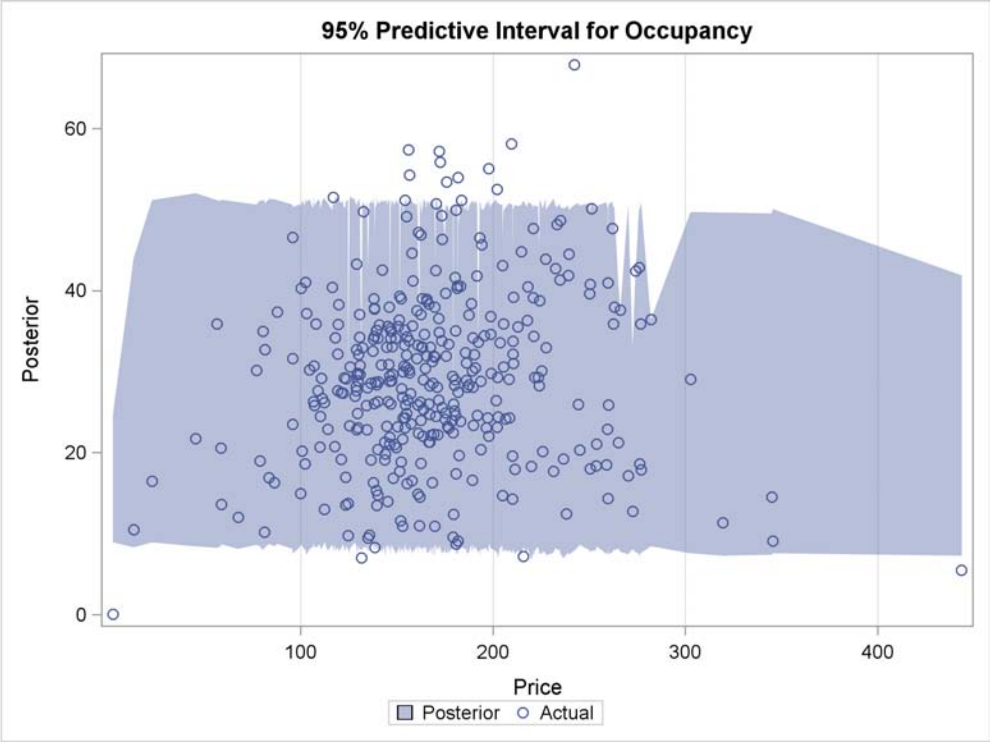

Constrained Pseudo-hierarchical Prior

In the example of the pseudo-hierarchical approach, a normal prior distribution with mean = –0.037 and variance = 0.00018 was chosen. This choice is reasonable if you are interested only in the mean of the posterior distribution, but it can be highly misleading if you are interested in the upper tail of the posterior distribution. The normal prior distribution has an unbounded support; the possibility that price can have a positive effect on occupancy is not completely excluded. This can lead to positive percentiles in the posterior distribution of the price regressor. One way to ensure both a negative sign for the price elasticity and the ability to borrow information from the analyses of Properties A, C, and D is to use a beta distribution that is defined on (–1,0) with mode located around –0.037. This approach is very similar to the approach that uses a normal prior with mean = –0.037 and variance = 0.00018, except that the functional is restricted on the support (–1, 0) while the mode is still centered around –0.037. You can accomplish this by specifying the following SAS code:

```
Prior Price ~ Beta(shape1=10, shape2=1.35, min=-1, max=0);
```

Figure 12 shows the posterior and prior predictive distributions of the regression of occupancy versus price when a constrained and informative prior distribution is considered (beta distribution with shape1 = 10, shape2 = 1.35, min = –1, and max = 0).

Figure 12 Diagnostic Plot for Price and Posterior Predictive Distribution of the Censored Regression Model

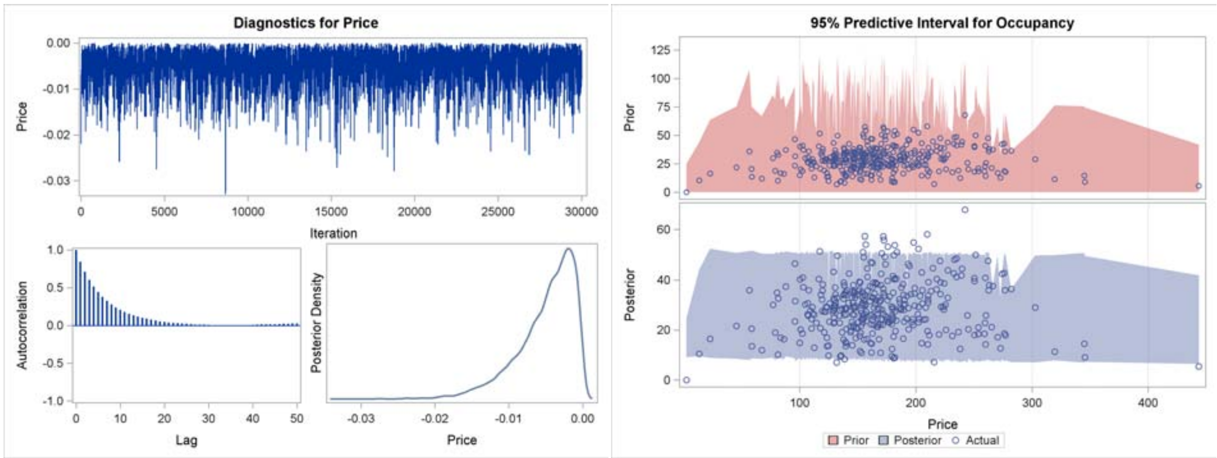

These plots are useful in analyzing the information that is provided by the data and by the prior. That most of the data are covered by the posterior predictive distribution (bottom right panel) and by the prior predictive distribution (top right panel) means that the data and the prior distribution provide coherent information. Figure 12 also displays a diagnostic plot that helps show whether the posterior distribution is being well approximated by the samples of the MCMC algorithm. A good convergence is usually characterized by a stationary time series (top left panel) and by an autocorrelation (bottom left panel) that goes to 0 quickly.

Strong Informative Prior

The previous three examples show the implementation of three different priors that return similar results. An analysis is robust to the prior specification if the posterior analysis does not change dramatically when different priors are used. The following example shows how the results are affected when a highly informative prior contradicts the information in the data:

```
prior Intercept ~ Normal(mean=50, var=0.0001);
prior Price ~ Normal(mean=-0.01, var=0.0001);
prior _Sigma ~ Gamma(shape=1, scale=0.1);*mean around 0.01;
```

The choice of the priors does not seem to be incoherent with the data. This is confirmed by the first panel of Figure 13, which compares the posterior and prior distributions to the information that is provided by the likelihood.

Figure 13 Posterior Predictive Distribution of the Censored Regression Model That Uses a Very Informative Prior

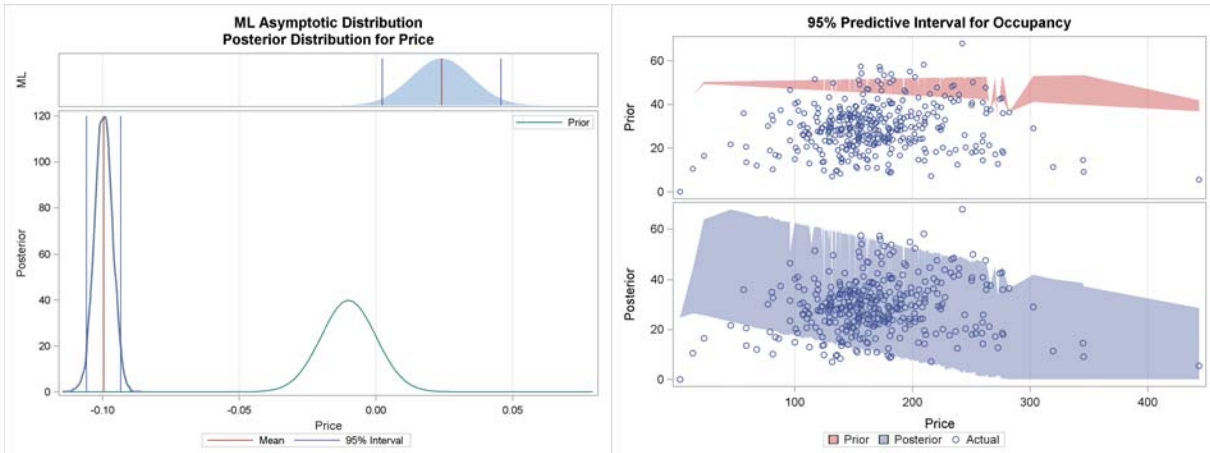

The maximum likelihood estimate (MLE) for the price coefficient is close to the mode of the prior distribution, but the posterior mode is far away from both the MLE and the prior mode. Figure 13 shows that the data in the sample are not well predicted by either the prior or posterior distribution, because many observations (represented by circles) fall outside the central 95% of the posterior (blue) and prior (red) predictive regions. This result does not necessarily mean that the prior is misspecified, but it indicates that the informational content of the prior distribution is very different from that of the data. If the analyst has good reasons to believe that this is a correctly specified informative prior, it can have a beneficial impact on the analysis. On the other hand, if the predictive behavior of the prior and posterior distributions looks suspicious, the analyst might have to rethink this approach and revisit the prior distribution choice.

CONCLUSION

This paper discusses the practical complexities of accurately estimating price elasticity when the dependent variable is censored because of capacity restrictions. It shows that ignoring censoring can bias coefficients and lead to suboptimal revenue maximization. Accounting for the censoring by using a censored regression specification corrects the bias of these estimates. In practice the problem is likely to be even more complex because of the nature of the data collection process and the occasional lack of good data.

New Bayesian estimation options in the QLIM procedure in SAS/ETS 12.1 allow potentially informative priors to assist in estimating price effects. These methods enable data to be used in ways that were previously not possible. This paper presents four examples of using priors in different ways to make price elasticity estimation possible when typical frequentist methods do not. These Bayesian estimation options are especially useful in situations where the sample size is small and the estimation can incorporate external information from other properties.

To learn more about these methods, see the *SAS/ETS 12.1 User's Guide*.

REFERENCES

Berger, J. O. (1985), *Statistical Decision Theory and Bayesian Analysis*, 2nd Edition, New York: Springer-Verlag.

Gelman, A., Carlin, J., Stern, H., and Rubin, D. (1995), *Bayesian Data Analysis*, London: Chapman & Hall.

Harvey, A. C. (1989), *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge: Cambridge University Press.

Hodrick, R. J. and Prescott, E. C. (1997), "Postwar U.S. Business Cycles: An Empirical Investigation," *Journal of Money, Credit, and Banking*, 29, 1–16.

Maddala, G. S. (1983), *Limited-Dependent and Qualitative Variables in Econometrics*, New York: Cambridge University Press.

ACKNOWLEDGMENTS

The authors are grateful to Anne Baxter and Ed Huddleston at SAS Institute Inc. for their valuable assistance in the preparation of this paper.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author:

Christian Macaro
SAS Institute Inc.
SAS Campus Drive
Cary, NC 27513
919-531-4547
christian.macaro@sas.com