County Food Stamp Program Participation to Poverty Ratio During Welfare Reform

Geof Gee and Robin Fisher[$]
US Census Bureau

Key Words: Poverty, Food Stamps, Small Area Estimates

1. Introduction

The Small Area Income and Poverty Estimates (SAIPE) program at the US Census Bureau estimates county poverty[1] as a function of administrative records[2] and the previous decennial census. Recent developments in SAIPE experimental models (Fisher, 2003) derive a measure of "true" poverty as a function of several independent measures of poverty, all of which are assumed to possess non-negligible variances; i.e., an "errors-in-variables" model. Examination of this work revealed a considerable amount of variability in the variance of participation in the Food and Nutrition Service (FNS) Food Stamp Program (FSP) across counties given poverty. In conjunction with other research, our research suggests that we could obtain more information on poverty from FSP participation by modeling its across-county variation.

The model gives us the opportunity to examine a fundamental government program to fight poverty and hunger in the US. In particular, we estimate the ratio of FSP recipients to number of poor in counties. Although FSP eligibility is more inclusive than the standard for being impoverished, the ratio should be a good indicator of the FSP efficacy in reaching those in the greatest need. We extend the analysis by estimating poverty and FSP participation in 1995 and 1999. These two years conveniently bridge the enactment of the Personal Responsibility and Work Opportunity Reconciliation Act of 1996 (PRWORA). Hence, we can measure the impact of welfare reform on FSP participation relative to the number of poor.

The provisions of PRWORA were not intended to affect FSP participation among most low-income families despite new restrictions for immigrants and able-bodied adults without dependents. Nonetheless, there is empirical evidence that welfare reform affected FSP participation and there are anecdotal and theoretical explanations for such an effect. Mathematica Policy Research, Inc. (MPR) estimates state FSP participation rates for the FNS (Castner and Schirm, 2003) using the Food Stamp Program Quality Control (FSPQC) and CPS data. In a summary of the estimates and data, Cunnyngham (2002) showed that FSP participation rates among those eligible have decreased significantly from 1994 to 2000; there are large differences in participation rates across race/ethnicity and age; and little difference in the participation rate across sex. Participation by immigrants and able-bodied adults without children also fell dramatically over the period.

A confounding factor in determining the influence of PRWORA on FSP participation is the economic growth that took place during the mid to late 90s. To summarize the literature, there is agreement that economic growth did reduce participation in the FSP and other entitlement programs; however, the magnitude of its effect is in question. Wallace and Blank (1999) estimate that changes in the economy can account for at most 44% of the drop in FSP caseloads. A report by Wilde, *et al.* (2000) attributed 35% of the decline in number of FSP participants to increased income of former participants.

McConnell and Ohis (2001) found that participation rates in rural and urban areas increased and decreased respectively during the period of welfare reform. However, the percentage decrease in the number of people eligible for the FSP was larger in rural areas. Haider, Schoeni, and Jacknowitz (2002)

[$] We have benefited from the comments of several people. In particular we thank, Paul Siegel, John Iceland, and Elizabeth Huang. This paper reports the results of research and analysis undertaken by the U.S. Census Bureau staff. It has undergone a Census Bureau review more limited in scope than that given to official Census Bureau publications. This report is released to inform interested parties of ongoing research and to encourage discussion of work in progress.
Contact information for the authors: e-mail Geoffrey.M.Gee@Census.gov ; Robin.C.Fisher@Census.gov; telephone (301) 763-3193.

[1] The Current Population Survey (CPS) Annual Social and Economic Supplement (ASEC) measure of poverty is the official poverty measure. Direct ASEC tabs of county poverty are used as the dependent variable of the model.

[2] The administrative records include tabulations from Internal Revenue Service 1040 tax returns and the Food and Nutrition Service's Food Stamp Program.

investigate why participation rates among the elderly are much lower than other adults. After controlling for a host of factors associated with the cost and benefits of FSP participation, they still found a large component of participation associated with age. Kabbani and Wilde (2002) found that shorter recertification periods lead to lower participation rates. A foundation of PRWORA was to encourage welfare recipients to work. As a consequence, states had an incentive to shorten recertification periods.[3] The initial application and recertification for the FSP are not costless activities: A FNS study determined that the initial application and recertification required an average of five and 2.5 hours respectively. Shorter recertification periods, therefore, represent a significant increase in the cost of participating in the FSP.

The paper proceeds as follows: First, we begin by describing the data, model, and the priors for the Bayesian estimates. Second, we present the model fit and results. Third, we discuss future research and present some concluding remarks.

2. The Model and Data

2.1 The Data

We restricted the sample to minimize the model's computational and memory demands. Consequently, the analysis is based on 1995 and 1999 county-level data for 9 states: Arizona, California, Colorado, Connecticut, Florida, Montana, New Jersey, Oregon, and Wyoming. There are 407 combined counties in the sample states. Some states were chosen as a consequence of past research on the behavior of the number of FSP recipients. Others were chosen to get a good geographic variety of states in the sample. Technically, we can only claim that our findings apply to the counties in these nine states.

The measures of poverty used are many of the same measures in the official SAIPE county poverty model. They are the following:

1. Direct survey estimates of county poverty from the Annual Social Economic (ASEC) Supplement of the Current Population Survey (CPS).
2. Estimates of county poverty from the previous or concurrent decennial census.
3. Tabulations of Internal Revenue Service (IRS) 1040 tax returns with income below the poverty threshold (henceforth referred to as "tax poor").
4. The county number of FSP recipients.

We also use demographic resident total population estimates from the US Census Bureau Population Division.[4] To model the variable relationship between FSP participation and poverty, we obtained data from several sources. The Bureau of Economic Analysis (BEA) provides county data on wage and salary earnings. Demographic county population estimates by age, race, sex and ethnicity were acquired from the US Census Bureau Population Division. We used the US Census Bureau 1990 Census to classify each county as rural or urban.

2.2 The Model

Our strategy is to use 1989 data, analogous to data described in Section 2.1, to develop a model and its priors before applying the model to 1995 and 1999 data. The resulting model of that preliminary process is described here. Some of the interesting findings from the preliminary analysis—for instance what was insignificant in the model—will be discussed in the results section of the paper.

The model assumes that there is an unobserved true log number of poor (LNP), m_i, for every county i. LNP, conditioned on parameters h and s_m, has a normal distribution given by

$$m_i \mid h, s_m \sim \mathrm{N}\left(h + TP_i, s_m^2\right)$$

where TP_i is the log total population of county i. The parameter h can then be interpreted as the national log poverty rate. The parameter h itself has a normal distribution given by $h \sim \mathrm{N}\left(m_h, s_h^2\right)$.

True poverty is unobservable. We only observe measures of LNP, X_{ij}, where i designates the county and $j = 1, \ldots, 4$, refers to the

[3] States face large penalties for failing to ensure that families obtain only the entitled amount of food stamps. Working families' income is more volatile compared to families that are unemployed or out of the labor force. As income changes, a family's entitled level of benefits will rise or fall. Hence, as more families work, state agencies have a greater incentive to shorten recertification periods.

[4] See the website http://eire.census.gov/popest/estimates.php for more information on demographic total resident population estimates.

aforementioned measures of LNP. Hence, we define X_{i1} as the ASEC direct estimate of LNP in county i. We define the other measures of LNP as X_{i2}, the LNP in the previous decennial census; X_{i3}, the log number of tax poor; and X_{i4}, the log number of FSP recipients. The relationship between our measures and LNP is described by the following equations

$$X_{ij} \mid m_i, a_j, b_j, s_{ij} \sim N\left(b_j m_i + a_j, s_{ij}^2\right)$$
$$j = 1,2,3$$
$$X_{i4} \mid m_i, a_{i4}, b_4, c_4, s_{i4} \sim N\left(c_4 m_i^2 + b_4 m_i + a_{i4}, s_4^2\right)$$
$$j = 4$$

The parameters a_j and b_j describe the linear relationship between m_i and X_{ij}; c_4 captures a quadratic relationship between m_4 and X_{i4}; s_{ij}^2 is the conditional variance of X_{ij}. The official measure is defined as the level of poverty measured by the ASEC. By definition, it is an unbiased estimate of true LNP so we restrict $a_1 = 0$ and $b_1 = 1$. The other parameters describe how our other measures of LNP are biased with respect to true poverty. We refer to these terms as the bias parameters.

With the exception of the FSP bias parameters, the bias parameters are constant across all counties. As mentioned earlier, previous SAIPE research in conjunction with outside research suggests that the relationship may vary across counties. Consequently, we allow the FSP additive bias term to vary according to the following function:

$$a_{i4} = Z_i g$$
$$X_{i4} \mid m_i, a_{i4}, b_4, c_4, s_{i4} \sim N\left(c_4 m_i^2 + b_4 m_i + a_{i4}, s_4^2\right)$$

Z_i is a vector of county characteristics described below including a fixed constant and g is a vector of parameters drawn from a single distribution across all counties. We control for county characteristics such as (1) the annual percentage change in wage and salary income, the percentage of total population that is (2) non-Hispanic white and (3) Hispanic, and (4) the rural/urban status of the county. In our preliminary model we tested and rejected any effect from measures about age and sex differences across counties. Much of the literature on FSP participation discusses differences in participation rates among those eligible. Therefore, it seemed natural to model the across-county heterogeneity through the additive term: if c_4 is small, then the additive term becomes a measure of the log FSP participation to poverty ratio.

From prior work, we know that the conditional variances for the measures of LNP often exhibit heteroskedasticity. We model the variances in the following manner.

$$s_{i1}^2 = u_1 \left(k_i\right)^{-1/2}$$
$$s_{i2}^2 = u_2^A * GVF_i + u_2^B$$
$$s_{i3}^2 = u_3 \left(TP_i\right)^{-1/2}$$
$$s_{i4}^2 = u_4$$

The u_j parameters are identically distributed across all counties. k_i is the sample size of the direct county ASEC estimate of poor. GVF_i is the decennial census generalized variance function, which captures the sampling variance for the decennial census estimate of poor.

The advantage of this model over other SAIPE production models is that the contribution of a measure of poverty to the derived estimate of true county LNP—defined as the expectation of ASEC county LNP—is based on the measure's precision.[5] In other words, the more confident we are of the measure and its relationship to LNP, the greater the measure's contribution to the estimate of LNP.

2.3 Priors

An advantage of the present SAIPE production model is its cross-sectional nature. In other words, there is no assumption that the relationship between poverty and its predictors are the same across time. We tried to preserve that advantage by purposely maintaining vague priors despite the exploratory model runs on 1989 data. The priors for the 1995 and 1999 models are identical.

[5] Although we have not calculated the exact distribution yet, we do know that the posterior distribution of LNP conditioned on the data and parameters is a N_3 density as discussed by Cobb, Koppstein, and Chen (1983).

than the observed value. We use two functions to evaluate the model for each measure of LNP.

Table 1: Priors for 1995 and 1999 Models[6]

Parm.	Prior	Parm.	Prior
h	N(-0.5,0.5)	g_2	N(0,1)
a_2	N(0,2)	g_3	N(0,1)
a_3	N(1,2)	g_4	N(0,1)
b_2	N(1,0.2)	$1/u_1$	Γ(0.03,0.1)
b_3	N(1,0.2)	u_2^A	N(1,0.2)
b_4	N(1.5,0.5)	$1/u_2^B$	Γ(5,0.2)
c_4	N(0,0.2)	$1/u_3$	Γ(0.1,0.1)
g_0	N(-0.5,1)	$1/u_4$	Γ(2.5,0.5)
g_1	N(0,1)		

3. Model Fit and Results

3.1 Model Fit

We use two methods to evaluate model fit. First, we examine scatterplots of standardized errors based on the mean of the posterior—designated as standardized Bayesian residuals[7]—instead of the linear prediction (as in a regression) against various model inputs and total population. Second, we examine posterior predictive p-values (PPP-values) in a fashion similar to the aforementioned "residuals."

In general, the standardized Bayesian residuals revolve around the line $y = 0$, show little sign of bias, and are of an acceptable magnitude. However, there are exceptions. The variance estimate of 1995 ASEC poor may be too small. The variance estimate of 1999 tax poor appears to shrink too slowly with total population. Lastly, there may be a slight upward bias of the 1999 ASEC poor posterior mean relative to the observed ASEC poor.

A PPP-value is defined as

$$p = \Pr\left(T\left(X_{obs}, q_{rep}\right) < T\left(X_{rep}, q_{rep}\right) \mid data\right)$$

where $T(\ )$ is some function chosen to evaluate an aspect of the model. More simply, a PPP-value compares characteristics of the replicated data—data drawn from the hypothetical posterior distribution of the model—to observed data. A simple example would be $T(x,q) = x$. One would calculate the probability that a replicated value of some variable in the model is greater than the observed value. We use two functions to evaluate the model for each measure of LNP.

$$T_1(x,q) = x$$
$$T_2(x,q) = \left(x - E\left[X \mid q\right]\right)^2$$

The straightforward interpretation of these functions allows one to examine T_1 and T_2 to make inferences about the first and second moments of the model, respectively. If a large majority of PPP-values based on the function T_1 is either above or below ½, say with respect to the estimated log number of poor, then we can infer that the model is biased upward or downward respectively. A similar observation with the function T_2 implies that estimates of the variable's variance are suspect.

Except as noted below, in scatterplots available upon request, the PPP-values exhibit many characteristics of a well-fitting model: the mean and median PPP-value is close to 50 percent and the distribution of points tends to be uniform over the space. There is evidence that the functional form of the 1995 and 1999 tax poor variance decreases too slowly with total population and that the model overestimates the conditional variance in general. In the T_1 based PPP-value scatterplots, there is a relatively high density of points—the density increases as population increases—near the line $y = ½$. The T_2 based scatterplots provide more evidence of some mis-specification of the tax poor variance model. There is an upward trend in the plot as total population increases. The mean of the 1995 and 1999 T_2 PPP-values is .72 and .81 respectively.

3.2 The Results

The first question to address is, "Does modeling across-county FSP participation conditioned on poverty improve the performance of the model?" A simple comparison between the homogenous model, where $a_{i4} = a_4$, and the heterogenous model, where $a_{i4} = Z_i g$, is to compare their Deviance Information Criterion (DIC) (Spiegelhalter et al., 1998) values. Quite like the Akaike Information Criterion (AIC), the DIC is a scale-free likelihood criterion with a penalty for the number of parameters a model contains. Using the 1995 data, the heterogenous model results in approximately a 90-point

[6] The parameterization for the gamma distribution used here is such that the mean is equal to a/b and the variance is equal to a/b^2.

[7] The definition of a standardized Bayesian residual is analogous to a standardized residual from regression. See Carlin and Louis (2000) for more details.

improvement[8] over the homogenous model. Furthermore, the posterior mean of the FSP participation variance conditioned on LNP, s_4^2, is 66 % larger in the homogenous model compared to the heterogenous model. The LNP point estimates barely change in the two models: 98% of the point estimates are within ± 1.5% of each other. In regards to the variance of LNP estimates, the variance estimates in the heterogenous model are smaller in small counties but bigger in large counties.[9]

The estimated parameters of X_{i4}, log FSP participation, are interesting as well. We controlled for the percentage of total population that is non-Hispanic white (g_1) and the percentage of total population that is Hispanic (g_2) using US Census Bureau's total resident population estimates: omitting the percentage of total population that is non-Hispanic black or "other." In addition, we included the annual percentage increase in wage and salary income (g_3) and a dummy variable that is one if the county is classified as rural (g_4). We also included an intercept (g_0). Table 2 contains the parameter estimates for the 1995 and 1999 poverty models.

Table 2: FSP participation parameter estimates

Parm.	1995 Model		1999 Model	
	Mean	SD	Mean	SD
c_4	-0.038	0.005	-0.039	0.005
b_4	1.677	0.081	1.697	0.088
g_0	-1.767	0.392	-1.981	0.419
g_1	-1.809	0.176	-1.775	0.188
g_2	-0.888	0.233	-0.937	0.249
g_3	-0.491	0.299	-2.425	0.403
g_4	0.169	0.066	0.261	0.076
$1/s_4$	6.779	0.623	4.948	0.379

There are several obvious conclusions drawn from Table 2. One, conditioned on the level of poverty, counties that had high percentages of non-Hispanic whites and/or Hispanics participated much less in the FSP relative to other counties. Two, other than the percentage growth of wage and salary income, the relative behavior of counties did not change much from 1995 to 1999; i.e., none of the first-moment parameters are significantly[10] different from each other. Three, the precision of FSP participation conditioned on the level of poverty in a county was more disperse—once again, true in a statistical sense—in 1999 than in 1995. Four, there is a large and statistically significant difference in the g_3 parameter—the effect of percentage growth in wage and salary income—from 1995 to 1999.

Finding differences across race is expected given the literature on FSP participation by those eligible. We also find that the differences across race and ethnicity remained consistent despite changes in welfare and the economy. Note that this consistency is conditioned upon the effect of percentage growth in wage and salary income, which became correlated with race and ethnicity after welfare reform. Tables 3 and 4 present the correlation matrix of the across-county heterogeneity variables. In 1995, the correlation coefficients between percentage growth in wage and salary income and the percentage Hispanic and non-Hispanic White variables are small and statistically insignificant. In 1999, they are larger and statistically significant. Tables 5 and 6 present the correlation statistics of the g parameters. One will immediately notice that the parameters g_0, g_1, and g_2 are highly correlated. Given the research on participation rates by non-Hispanic whites, non-Hispanic blacks, and Hispanics, and that the sum of the three demographic groups is very close to total population, such high correlation coefficients is not surprising.

Table 3: Pearson Correlation Coefficients for 1995 Across-County FSP Variables. Here are a Pearson correlation coefficients for the variables in the across-county FSP heterogenous participation model; i.e., the Z matrix of the model excluding the intercept. Correlation coefficients that are statistically **in**significant are *italicized*.

	% Hisp.	% White	% Growth	Rural
% Hisp.	1.000	-0.663	*0.016*	-0.206
% White	-0.663	1.000	*-0.029*	0.273
% Growth	*0.016*	*-0.029*	1.000	-0.165
Rural	-0.206	0.273	-0.165	1.000

[8] The "total" DIC criteria for the heterogenous and homogenous model were 412 and 506, respectively. The improvement in the FSP participation model, as measured by the DIC criterion, was 150 points. However, the improvement for ASEC poverty model was much more modest. There was only a four- to five-point improvement.

[9] Although, one would expect this to change as the tax poor variance model is refined.

[10] Here, two paramet ers q_1 and q_2 are significantly different if min(P($q_1 < q_2$ |data), P($q_2 < q_1$ |data)) < .1; q_1 is significantly greater than q_2 if P($q_1 > q_2$ |data) > 0.9.

Table 4: Pearson Correlation Coefficients for 1999 Across-County FSP Variables. Here are a Pearson correlation coefficients for the variables in the across-county FSP heterogenous participation model; i.e., the Z matrix of the model excluding the intercept. All correlation coefficients are statistically significant at the 1% level.

	% Hisp.	% White	% Growth	Rural
% Hisp.	1.000	-0.673	0.233	-0.214
% White	-0.673	1.000	-0.143	0.283
% Growth	0.233	-0.143	1.000	-0.338
Rural	-0.214	0.283	-0.338	1.000

Table 5: Correlation of g parameters in 1995 model. Correlations in *italics* cannot reject the hypotheses that $r = 0$. Correlation estimates with an exponent of "#" are only significant at the 10% level. The remaining correlation estimates are significant at the 1% level.

	g_0	g_1	g_2	g_3	g_4
g_0	1.000	-0.439	-0.318	0.165	-0.145
g_1	-0.439	1.000	0.605	*-0.022*	$0.063^{\#}$
g_2	-0.318	0.605	1.000	*0.017*	-0.135
g_3	0.165	*-0.022*	*0.017*	1.000	$0.055^{\#}$
g_4	-0.145	$0.063^{\#}$	-0.135	$0.055^{\#}$	1.000

Table 6: Correlation of g parameters in 1999 model. Correlations in *italics* cannot reject the hypotheses that $r = 0$. Correlation estimates with an exponent of "#" are only significant at the 10% level. The remaining correlation estimates are significant at the 1% level.

	g_0	g_1	g_2	g_3	g_4
g_0	1.000	-0.419	-0.313	$0.072^{\#}$	-0.166
g_1	-0.419	1.000	0.586	$-0.076^{\#}$	*-0.037*
g_2	-0.313	0.586	1.000	-0.141	-0.128
g_3	$0.072^{\#}$	$-0.076^{\#}$	-0.141	1.000	0.176
g_4	-0.166	*-0.037*	-0.128	0.176	1.000

The precise connection between percentage growth of wage and salary income and FSP participation conditioned on poverty is unclear. We have several hypotheses, some of which we summarize here. High growth counties are likely to possess higher rates of employment. The higher rates of employment would include those who are impoverished.[11] The working poor tend to have greater income—hence, they qualify for a lower amount of food stamps—and greater opportunity cost than the unemployed poor. This has the effect of lowering the benefit and increasing the cost of FSP participation leading to fewer participants among those still eligible. The causal relationship might instead work through state responses to greater income volatility of FSP participants. As discussed earlier, states have an incentive to keep FSP error rates low. Ceteris paribus, greater income variability logically leads to higher error rates, which would give states an incentive to shorten the recertification period. Another alternative is that impoverished residents of fast-growing counties see many opportunities, have a better outlook of the future, and, consequently, see less need to participate and suffer any negative effects associated with participation.

The increase in the conditional variance of FSP participation given poverty corresponds well with anecdotal evidence that the efficacy of state programs/offices in reaching impoverished people is more variable after welfare reform than prior to its enactment. During the period of welfare reform, states were given greater authority in how Temporary Aide for Needy Families (TANF), FSP, and other welfare programs were implemented. As state programs have become more different from each other, one might expect welfare participation rates—and consequently FSP participation rates—to vary more across states. Furthermore, as more FSP participants entered the workforce, states have reacted differently to the associated increase in error rates. These strategies include supervisory review, greater training, verifying eligibility and benefits with outside administrative data, and the aforementioned length of the recertification period. While these different strategies may all reduce a state's error rate, the strategies may also affect FSP participation among the poor differently; leading to greater variability of FSP participation conditioned on poverty.

Conclusion

We find that conditioned on the level of poverty, there are large differences in FSP participation across counties. Modeling these differences improves the estimates of FSP participation considerably; but modeling the across-county variation provides a modest enhancement to county estimates of poverty.

Conditioned on county poverty, we find that race, ethnicity, wage and salary income growth, and rural status all contribute to explaining across-county differences in FSP participation. Variables related to age and sex composition of the county were ineffective in explaining in cross-sectional FSP participation differences. These effects, with the exception of wage and salary income growth whose effect increased, are remarkably consistent in 1995 and

[11] In fact, the economics literature supports the notion that poor/unskilled workers are more affected by economic swings than affluent/skilled workers.

1999. The conditional variance of FSP participation increased in 1999 compared to 1995. Although we cannot prove a causal relationship between the 1995 to 1999 differences and welfare reform, there are several reasonable hypotheses on why such a relationship does exist.

3.3 Conclusion

In the immediate future, there are several possible improvements to the model. One is to find more variables that describe across county heterogeneity in FSP participation given poverty. A specific example in mind is to include the length of the recertification period as a predictor of FSP participation. A second improvement is to fine-tune the variance model for tax poor. A third improvement, although a fairly complicated one, is to investigate whether there is a time consistent effect that can be exploited; i.e., whether the errors in the model are consistent over time. We observe that a county's standardized Bayesian residuals from 1995 are highly correlated with those from 1999. This suggests that searching for more predictors can still be productive and that investigating models associated with longitudinal data might prove fruitful.

4. References

Blank, R. and Wallace, G. (1999), "What Goes Up Must Come Down? Explaining Recent Changes in Public Assistance Caseloads." In *Economic Conditions and Welfare Reform*. Edited by Sheldon Danziger. Kalamazoo, MI: Upjohn Institute.

Castner, L. and Shirm, A. (2003), "Empirical Bayes Shrinkage Estimates of State Food Stamp Participation Rates for 1998-2000," MPR Reference No. 8659-211, Washington, DC, Mathematica Policy Research, Inc..

Carlin, B. and Louis, T. (2000), *Bayes and Empirical Bayes Methods for Data Analysis*, 2nd ed., New York, Chapman and Hall/CRC.

Cobb, L., Koppstein, P., and Chen, N. (1983), "Estimation and Moment Recursion Relations for Multimodel Distributions of the Exponential Family," *Journal of the American Statistical Association*, Vol. 78, No. 381. (Mar., 1983): Alexandria, VA., pp 124-130.

Cunnyngham, K. (2002), "Trends in Food Stamp Program Participation Rates: 1994 to 2000," MPR Reference No. 8659-213, Washington, DC, Mathematica Policy Research, Inc..

Fisher, R. (2003), "Errors-in-Variable Model for County Level Poverty Estimation," SAIPE Working Paper, Washington, DC, U.S. Census Bureau.

Haider, S., Schoeni, R., and Jacknowitz, A. (2002), "Food Stamps and the Elderly: Why is Participation So Low?" *Focus*, Vol. 22, No. 2, Madison, WI, Institute for Research on Poverty, pp 66-67.

Kabbani, N. and Wilde, P. (2002), "Short Recertification Periods in the US Food Stamp Program: Causes and Consequences," *Focus*, Vol. 22, No. 2, Madison, WI, Institute for Research on Poverty, pp 64-65.

McConnell, S. and Ohis, J. (2001), "Food Stamp Participation Rate Down in Urban Areas But Not in Rural," *Food Review*, Vol. 24, No. 1, Washington, DC, USDA, pp 8-12.

Spiegelhalter, D., Best, N., and Carlin, B. (1998), "Bayesian Deviance, the Effective Number of Parameters, and the Comparison of Arbitrarily Complex Models," Research Report 98-009, Division of Biostatistics, University of Minnesota.

Wilde, P, Cook, P., Gundersen, C., Nord, M., and Tiehen, L. (2000), "The Decline in Food Stamp Program Participation in the 1990's," Food Assistance and Nutrition Research Report No.7, USDA, ERS.