ANALYSIS OF SIPP REINTERVIEW DATA: A POISSON MODEL OF RESPONSE ERROR

Daniel H. Hill[1]
Opinion Research Institute, University of Toledo, Toledo, Ohio 43606

Introduction

As part of its ongoing quality control program the Field Division of the Census Bureau conducts reinterviews monthly with small samples of the Survey of Income and Program Participation (SIPP) respondents.[1] The purpose of this reinterview program is to evaluate individual interviewer performance to determine if retraining or dismissal is necessary. In addition to ascertaining whether the interview was actually conducted with the correct unit and whether the proper procedures were employed, the reinterview contains a small set of questions of substantive content. While it was never the intent of the reinterview program designers, the existence of the reinterview data makes estimation and analysis of nonsampling error in the SIPP possible. Such analysis is potentially important because it is quite *apparent*[2] that data from the SIPP are far from perfect.

The purpose of the present research is to assess this potential by merging the reinterview data with public release data and analyzing the combined data. The paper is organized in three sections. In Section 1 the SIPP reinterview program is described in some detail. Section 2 presents a question-by-question description of response, procedural and overall interview/reinterview discrepancies. Finally, in Section 3, two classes of multivariate models are developed and estimated.

1. The SIPP Reinterview Program

The SIPP reinterview program is an ongoing systematic operation which is intended to monitor data quality by checking the interviewers' work. The sample to be reinterviewed each month is a multistage probability sample of current SIPP respondents. Figure 1 illustrates the question flow for the first five questions of the Reinterview Questionnaire. The questions actually asked of the respondent in both the interview and reinterview are printed in bold, while the Office Check Items which are transcribed to the Reinterview Questionnaire from the original appear in normal print. Unless otherwise indicated, questions are asked in sequence. In most cases, however, respondents are skipped around certain questions and these skips are indicated in the figure by lines and arrows. If, in response to question 1, for instance, the respondent said he had a job for at least part of the reference period ('yes' on item 1.), he is skipped around the questions about whether he spent any time looking for a job (2a.), or whether he wanted a job (3a.), and is asked about whether he had a job each week of the reference period instead (4.). In Figure 1, a skip such as this which results from a response to a question asked in the reinterview study is depicted with a dotted line. Skips from Office Check Items, being automatic from the reinterviewer's point of view, are depicted as solid lines.

It does not take a great deal of study of Figure 1 to see that the skip sequences employed in the SIPP can be quite complicated. Indeed, a major goal of the reinterview program is to see if individual interviewers are following these skip sequences properly.

2. Inconsistency Rates and Simple Response Variance Estimates

With the two independent observations provided by the interview and reinterview responses it is possible to estimate[3] the simple response variance for the various questions. To do so, we first confine our attention to that portion of the reinterview sample where a) the reinterview was successfully conducted and b) it was determined that the interviewer had visited the proper sample unit in conducting the original interview. We also eliminate from our sample those cases where the date of the original interview as recorded in the interview failed to match the date coded in the public release files, and those few cases where, even though the reinterview was conducted, no substantive questions were re-asked. These restrictions leave us with a sample of 1,559 cases of interview/reinterview data for waves 2 and 3 of the 1984 panel.

Figure 1.
SIPP Reinterview Question Flow

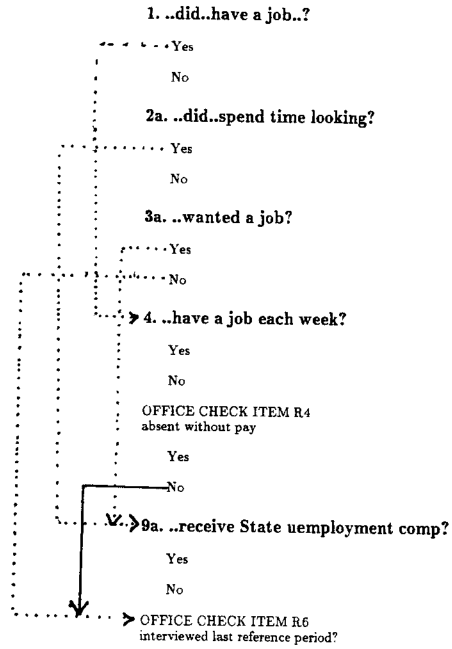

We can distinguish two types of inconsistencies when the interview and reinterview reports do not agree—response inconsistencies and procedural or 'skip' inconsistencies. We will reserve the term response variance or 'response inconsistencies' for estimates involving cases where the question was actually asked of the respondent in both the interview and reinterview and where a response was recorded. Given the complicated skip sequences employed, it should not be surprising that there are differences between the two reports not just in responses, but in whether or not the question was asked each time. Discrepancies between the interview and reinterview arising because a question was skipped in one and not the other will be referred to as 'procedural discrepancies'.

An example may be useful in clarifying these distinctions. Table 1 presents the recorded responses for the interview and reinterview for Item 4.—the question regarding receipt of state unemployment compensation. Actual responses in both interviews were recorded for only some thirteen percent (=100*207/1559) of the cases. Of these 2.9% (=100*(3+3)/207) of the reports were different. The simple response variance for this question is, therefore, .0145, or half the gross difference rate among those respondents who answered the question in both the interview and reinterview. We will define the procedural discrepancy rate as the simple gross difference rate for whether the question was skipped. For the unemployment compensation question results in Table 1, the procedural discrepancy rate is 6.54 percent (= 100*(7+59+7+29)/1559). The overall discrepancy rate is simply the fraction of the entire sample for which the interview and reinterview reports differ.

Table 1
Whether Received State Unemployment Compensation As Recorded in the Reinterview by How Recorded in Original Interview

	Reinterview			
Original Interview	Blank	1 'Yes'	2 'No'	Total
Blank	1,250	7	59	1,316
1 (Yes)	7	29	3	39
2 No	29	3	172	204
Total	1,286	39	234	1,559

Table 2 presents these discrepancy rates for each of twelve substantive questions asked in the SIPP reinterview. There is considerable variation in the overall discrepancy rates for these questions ranging from less than two percent for questions on employment during the reference period (1) and continued Medicaid coverage (26b) to about seven percent for the Health Insurance coverage (27a) and the employer's contribution to Health Insurance (27f) questions. This pattern is quite similar to that reported by the Census Bureau's Reinterview Evaluation Section (see e.g. Smith, 1987). While it does vary from question to question, the majority of the discrepancies in the data as a whole are procedural rather than response discrepancies. Given the skip patterns used in the study, it is not surprising that virtually all of the discrepancies on the Medicare coverage question were procedural in nature—i.e. the result of the question being skipped in one interview and not in the other. There are, after all, three distinct ways in which a respondent can be routed around question 23a and four ways in which he could be routed to it.

Procedural discrepancies also accounted for most of the overall discrepancies in all the remaining questions except for the initial employment and health insurance questions. That these are the initial questions in a sequence which all respondents are to be asked is significant and points to the fact that some of the procedural inconsistencies are the result of response inconsistencies in earlier portions of the interview.

Response inconsistencies also vary widely from a low of less that three-tenths of one-percent for the Foodstamp authorization question to more than seven and a half percent for the employer health insurance contribution question. The high response variances of health insurance coverage and employer contribution of .03 (=.5*6.03/100) and .038 (=.5*7.62/100), respectively, would suggest that there was something wrong with these questions. Independent analysis of the reinterview data by Bureau staff uncovered similar results and the questions have been modified. Similarly, the high response error variance for the discouraged worker question (3a) has lead the Bureau to drop it.

In summary, simple comparisons of interview and reinterview reports from the reinterview data are sufficient to highlight some questions and procedures that are particularly problematic in the current SIPP instrument. Considerable error is probably being introduced to the data, for instance, because the skip sequences are sometimes quite complex and may not always be successfully followed. Additional errors occur because not all the questions are as clearly worded as we would like, and the reinterview data reflect these glitches in the form of high response variance.

Table 2
Discrepancy Rates for the Substantive Reinterview Questions

Question	Discrepancy Rates (percent)		
	Overall†	Procedural	Response
1. Have job?	1.89	0.26	1.63
2a. Look for job?	2.20	1.28	2.55
3a. Want job?	2.81	1.54	4.04
4. Each week?	3.84	1.92	3.11
9a. U.I. Comp?	3.15	2.76	2.99
23a. Medicare?	5.13	5.07	*
24. Food Stamps?	1.93	1.67	.28
26a. Mcaid now?	3.08	2.44	0.71
26b. Mcaid B4?	1.68	1.48	*
27a. Health Ins?	6.78	0.77	6.03
27e. Via emplyr?	6.00	4.68	2.35
27f. Emplyr pay?	7.35	4.11	7.62

*Rate suppressed due to the small number of cases in the denominator.

†The dual response rate can be obtained by subtracting column 2 from column 1 and dividing by column 3.

3. Correlates of Inconsistency

If the procedural and response variability is the same for all respondents, then its existence is relatively benign. If, on the other hand, the extent of response or procedural variance differs systematically from one respondent to the next, all manner of problems can be expected to arise in bivariate or multivariate analysis. The purpose of this section is to explore the extent to which response and procedural variance differs systematically with characteristics of respondents and interviewers.

Traditionally, analysts have chosen some form of logit model (see e.g. O'Muircheartaigh and Wiggins, 1981) in investigating the association of respondent and interviewer characteristics with response discrepancies on a question-by-question basis. An alternative modeling approach is to analyze the reinterview data, not on a question-by-question basis, but as single experiment in which the outcome is the number of discrepancies occurring in the course of the reinterview. Each question asked in the reinterview can be thought of as a Bernoulli trial with a 'success' being defined as a report being given which differs from that provided in the original interview. If we assume that these trials are independent, then the reinterview process itself would be a series of Q(i) Bernoulli trials where Q(i) is the total number of questions put to the i^{th} respondent. One testable hypothesis is that discrepancies are independent and rare, in which case the Poisson distribution would be a natural choice for describing the data. Accordingly, the probability of exactly n inconsistencies occurring is:

$$p(n) = \exp(-\lambda)\, \lambda^n/n! \qquad 3.1)$$

where λ is the mean number of inconsistencies observed (i.e. λ = Qp). Both the mean and variance of the Poisson distribution are λ. Figure 2 presents the actual and theoretical Poisson distributions of response errors in the SIPP reinterview data. The theoretical distribution fits the data like a glove. The mean *and variance* of the observed data are .171, which is yet further confirmation of the extremely good fit of the Poisson to the response inconsistency data. Since respondents were asked, on average 6.3 questions per reinterview, this would imply an average response discrepancy rate of 2.7% (=(.171/6.3)*100) and an average response variance of .0135.

Conceptually, the nearly perfect fit of the response inconsistency data to the Poisson suggests that if respondents were asked a reinterview question repeatedly (and their memories of their previous responses were wiped clean) inconsistent reports would appear infrequently, randomly and independently in time. Indeed, the Poisson can be shown to be the maximum entropy or disorder process. One might think that given the skip sequences used in the SIPP that errors in one variable would lead to errors in subsequent ones, and the independence aspect would not be accurate. This would be the case for procedural or overall inconsistencies, but is not for response inconsistencies—any subsequent inconsistencies resulting from a response error are, by construction, procedural and are not counted in the response discrepancy rate.

While all this is interesting and reassuring, it may not be entirely obvious that the fit of the unconditional distribution is particularly relevant in developing a multivariate model. As it turns out however, if the mean number of inconsistencies (λ(i)) given by individual i over a number of independent trials is related to a set of individual characteristics X(i) according to:

$$\lambda_i = Q_i \exp(X_i\beta) \qquad 3.2)$$

and if n_i follows a Poisson distribution, then

$$E(n_i \mid X_i) = \lambda_i \qquad 3.3)$$

Expressions 3.1) – 3.3) form the basis of what is sometimes referred to as Poisson Regression (see Maddala, 1984). The likelihood of observing a sample of N cases

$$L(n_1, ..n_N|X_1, .., X_N; \beta) = \prod_{i=1}^{N} P(n_i) \qquad 3.4)$$

$P(n_i)$ can be obtained by substituting 3.2) into 3.1). That is:

$$P(n_i) = [\exp(-Q_i \exp(X_i\beta))]\, Q_i^{n_i} \exp(n_i X_i \beta)/n! \qquad 3.5)$$

Substituting 3.5) into 3.4), taking logs, and collecting terms yields the following log likelihood function:

$$\mathbf{L} = \sum_{i=1}^{N} [-\exp(X_i\beta + \ln(Q_i)) + n_i(X_i\beta + \ln(Q_i)) - \ln(n_i!)] \qquad 3.6)$$

It can be shown that so long as the X's are not perfectly colinear (and so long as $\exp(X\beta)) > 0$ for some i) this log-likelihood function is globally concave in the β's.[4] In the present analysis we employ the Davidson-Fletcher-Goldfarb-Shanno algorithm to maximize 3.6) and obtain our estimates of β.[5] Estimated standard errors are constructed from the diagonal elements of the inverse-Hessian matrix.

Figure 2
Actual and Theoretical Probabilities
of Counts of Response Discrepancies

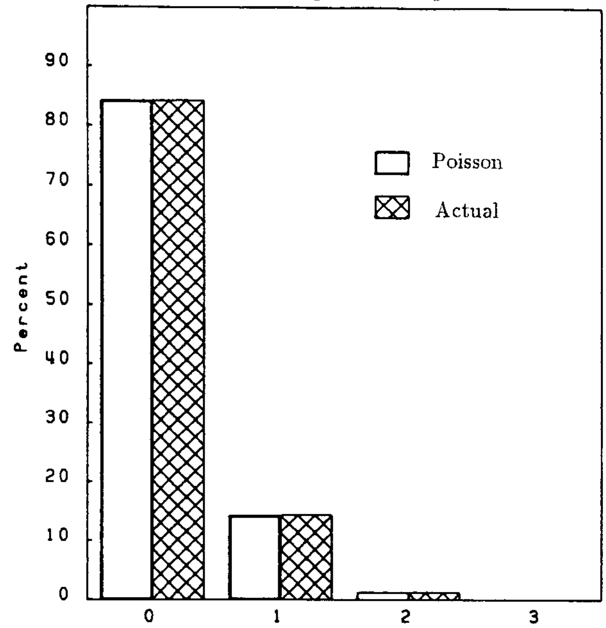

The independent variables we employ in our analysis are intended to capture (at least some of) the effects variability in interviewing process, and in characteristics of respondents which might affect response quality. The interviewing process variables are 1) the calendar month in which the original interview was taken, 2) a scale based on the overall performance of interviewers in the various Regional Offices, 3) whether a proxy informant was used in the original interview, and 4) whether the individual is some one other than the reference person or his/her spouse (*e.g.* child, aunt, *etc.*).

The individual characteristics included in our empirical specification are the same ones thought to affect market productivity in the human-capital model of earnings. These consist of age (and its square), education, race, and gender. We also include income itself in some of our specifications.

Table 3 presents both bivariate and multivariate estimates of the Poisson regression model for response discrepancies obtained by maximizing 3.6) with respect to the β. The first column of figures, labeled 'Bivariate Parameters', are obtained when the Poisson Regression model is estimated with only a constant and the variable listed to the left of the coefficient included as predictors. As hypothesized, response inconsistencies decline significantly with interview month. Since the month is included as a proxy for interviewer and respondent experience with the SIPP, and since the logarithm of month is used, the coefficient of –.275 is interpretable as the experience elasticity of response inconsistencies—a one percent increase in experience is associated with a .275 decrease in response inconsistency rates. This result is encouraging because it indicates that progress was being made in improving response quality early in the SIPP program.

The fact that the coefficient on the log on the Regional Office inconsistency rate is so close to unity, and is highly

Table 3
Maximum Likelihood Poisson Regression
Estimates of Response Inconsistencies
(Asymptotic SRS Standard Errors in Parentheses)

	Bivariate	Multivariate	
	Parameter	without Income	with income
Constant	−3.609** (.037)	−1.455* (.623)	−1.724** (.603)
Interview Month	−.275* (.132)	−.251+ (.132)	−.235 (.130)
Regional Office Discrepancy Rate	.935** (.322)	.980** (.313)	.962** (.313)
Proxy Respondent	.175 (.132)	.107 (.146)	.113 (.146)
Odd Relationship to Reference Person	.383* (.161)	.109 (.199)	.030 (.203)
Age (decades)	−.485** (.176)	−.369+ (.207)	−.215 (.197)
Age-squared (decades-squared)	.470** (.175)	.345+ (.202)	.205 (.197)
Education	−.044* (.019)	−.042* (.021)	−.020 (.022)
Whthr Female	.098 (.123)	.080 (.128)	−.071 (.137)
Whthr Black	.162 (.203)	.140 (.209)	.128 (.203)
Income ($100's)	−.827** (.211)		−.701** (.241)
ln(likelihood) (d.f.)		−761.5 (10)	−757.8 (11)

+significant at the 10% level.
*significant at the 5% level.
**significant at the 1% level.

significant means that differences in something at the regional level are important, but the bivariate results can provide no clue as to what it might be. The positive coefficient for the relationship to reference person dummy variable indicates that the response consistency for reference persons and their spouses is higher (by about 38.3 percent) than that obtained from other persons in the household.

The effects of age on response inconsistency rates is highly non-linear. The coefficients of −.485 and .47 on age (in 100's years) and age square, respectively, suggest that response quality increases with age at a decreasing rate until age 51 where it attains its maximum. For respondents much older or younger than this, response quality is significantly lower.

The final two variables with significant bivariate associations with response inconsistencies are education and income. Each one-year increase in educational attainment is associated with a 4.4 percent decrease in the response inconsistency rate. The extremely significant coefficient of −.827 on income, similarly, is interpreted as indicating that a dollar increase in monthly personal income is associated with a .83 percent decrease in the response inconsistency rate. Monthly personal income is the most powerful predictor of response inconsistencies included in our analyses. Conceivably some of this effect may be a reflection a tendency for fewer imputations being made for relatively complete interviewers and these interviewers tend to be interviews with people who have some income to report.

The bivariate results just discussed are analogous to simple correlations in linear models. The multivariate results presented in the last two columns of Table 3, in contrast, are analogous to multiple correlation coefficients. These coefficients are, therefore, interpretable as the net effects of the various factors on response inconsistency one obtains when the effects of other factors are controlled. Thus, it is not surprising that these multivariate effects are, in general, weaker than their bivariate counterparts. Indeed, with the single exception of the Regional Office inconsistency index, all the coefficients in column 2—the specification which includes everything but income—are of the same sign as those in column 1, but are smaller in absolute value. The estimated standard errors are also, in general, larger in the multivariate analyses—a second indication that the various predictors are correlated with each other. The decreased size of the estimated effects and their increased estimated variance combine to decrease the significance of almost all predictors in the multivariate analysis which excludes income. The combined effect of age and age-squared remains significant even though the individual coefficients are not.

When income is added to the multivariate specification of the response inconsistency Poisson regression, every other individual characteristic becomes insignificant. Taken literally, this result would suggest that all of the effects of age and education on response quality discussed up to this point are the result of the correlation of these factors with income. We find this result hard to believe. Why income, itself, should have a positive effect on response quality is a mystery.

Two further aspects of the multivariate Poisson regression estimates of response inconsistencies should be noted. First, the overall goodness of fit of both versions of the multivariate model is highly significant. The χ^2 under the null hypothesis of no association for the model presented in column 2 is 27.8 with 10 degrees of freedom and that for the model in column 3 is 35.2 with 11 degrees of freedom. Second, and of more substantial interest, the coefficient on the Regional Office inconsistency index was unaffected by the inclusion of respondent characteristics. In fact, this coefficient increased slightly when the other factors were controlled. This suggests that the source of the regional differences in response inconsistencies is something other than regional differences in the characteristics of respondents. One possibility is that the quality of interviewer training or selection varies by region. Alternatively, it may be that the care given to the reinterview program varies from one Regional Office to the next. In either event, future analysis of the reinterview data with data on interviewer characteristics, would seem worthwhile.

Total Inconsistency Rates

Response inconsistencies are relevant when one is trying to understand the response process itself, but in many respects a better measure of the reliability of survey items is the total inconsistency rate. Unlike the response inconsistency rate, the Poisson distribution is not a good choice for describing or modeling total inconsistencies. Figure 3 presents the a histogram of the actual inconsistency counts from the SIPP reinterview data, along side those implied by Poisson and Negative-binomial distributions constructed using sample moments. The probabilities predicted by the Poisson, based on the sample mean of .572 per reinterview, grossly under estimate the fraction of clean cases (n = 0) as well as of very

Figure 3
Actual and Theoretical Probabilities
of Counts of Total Discrepancies

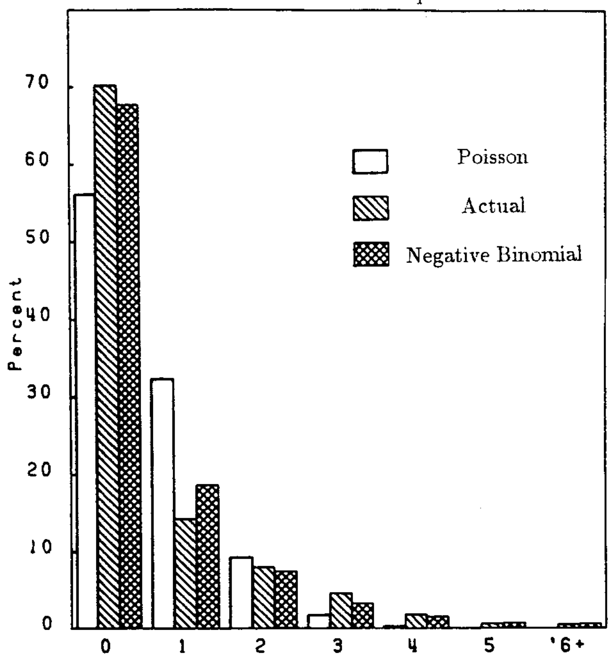

dirty cases ($n \geq 3$). The problem is that there is more variability in the data than is implied by a Poisson process. If total inconsistencies were following a Poisson process, then their variance should equal their mean. In fact, it is more than twice (1.16/.572) as large.

Such problems of excessive variability are often encountered in fitting data to counting distributions. In the Poisson, all of the variability is due to the fact that the n(iq) are determined by a Poisson process—the $\lambda(i)$ are deterministic functions of the X(i). If we assume instead that the $\lambda(i)$ are themselves random variables, and that they follow a Gamma distribution with parameters $\exp(X\beta)$ and δ then it can be shown that:

$$p(n_i) = \prod_{j=1}^{n_i}(\exp(X_i\beta) + j)\, p^{\exp(X_i\beta)}\, q^{n_i}/n! \qquad 3.7)$$

where $p \equiv \delta/(1+\delta)$ and $q \equiv 1/(1+\delta)$. The mean and variance of n_i for such a negative binomial distribution are, respectively:

$$\exp(X_i\beta)/\delta, \quad \text{and} \quad \exp(X_i\beta)(1+\delta)/\delta^2. \qquad 3.8)$$

Figure 3 includes the predicted probabilities for this negative binomial distribution with p set equal to the sample mean divided by the variance ($\bar{n}/v(n)$), and $\exp(X\beta)$ set to the square of the sample mean divided by the variance minus the mean ($\bar{n}^2/(v(n)-\bar{n})$.[6] Clearly the negative binomial fits the unconditional distribution significantly better than does the Poisson. The log-likelihood function can be obtained by substituting equation 3.8) into 3.4) and taking logs. This yields, for a sample of size N:

$$L = \sum_{i=1}^{N}[\exp(X_i\beta)\ln(p_i)+n_i\ln(q_i)+\{\sum_{j=1}^{n_i}\ln(\exp(X_i\beta)+j)\}-\ln(n_i!)]$$

Maximization of this with respect to the β was accomplished using the same algorithm employed in our earlier estimation of the Poisson regression model. The results of this estimation are presented in Table 4.

The results of the maximum likelihood negative-binomial analysis of total inconsistencies (Table 4) look very much like those obtained for response inconsistencies using Poisson regression (Table 3). The interpretation of these coefficients is the same as that of the Poisson regression coefficients—for those variables entering linearly (e.g. education), a one unit increase is associated with a proportionate change in the inconsistency rate of β (a 5.2% decrease for education in the bivariate model). The only real difference between the Poisson regression coefficients for response inconsistencies and those of the negative-binomial for total inconsistencies is that the latter are generally larger in absolute value and have lower estimated variances. The same substantive results hold.

Table 4
Maximum Likelihood Negative-Binomial
Regression Estimates Total Inconsistencies
(Asymptotic Standard Errors in Parentheses)

	Bivariate	Multivariate	
	Parameter	without Income	with income
Constant	−.872** (.094)	.807+ (.434)	.592 (.455)
δ	.726** (.076)	.779** (.065)	.797** (.062)
Interview Month	−.198* (.097)	−.191* (.090)	−.180* (.090)
Regional Office Discrepancy Rate	.985** (.236)	1.043** (.209)	1.022** (.221)
Proxy Respondent	.149 (.108)	.262+ (.142)	.197 (.143)
Odd Relationship to Reference Person	.382** (.121):	.109 (.100)	.121 (.103)
Age (decades)	−.511** (.125)	−.326** (.140)	−.189 (.143)
Age-squared (decades-squared)	.5745** (.123)	.400** (.136)	.271* (.138)
Education	−.052** (.013)	−.028* (.014)	−.011 (.015)
Whthr Female	.011 (.094)	.019 (.085)	−.149+ (.091)
Whthr Black	.284* (.143)	.275* (.135)	.258+ (.134)
Income ($100's)	−.698** (.135)		−.588** (.162)
ln(likelihood) (d.f.)		−1546.4 (11)	−1541.7 (12)

The bivariate results are obtained by estimating the model with the variable interest and the constant and shape parameter (δ) only.

Given the strong similarity of the results of the Poisson regression model of response inconsistencies and the negative-binomial model of total inconsistencies, we are lead to suspect that response and procedural inconsistencies share a common causal structure. Whatever this structure is, it evidently involves characteristics of both the respondent and the interviewer (or at least of the Regional Office).

Before closing out our discussion of the negative-binomial regression results it is useful to explore briefly the implications of the fact that response errors are well described as a Poisson process whereas procedural errors are not. What it means is that, abstracting from skip sequence effects, the occurrence of a response error in one question has no effect on the probability of a response error in a subsequent question.

That the inclusion of procedural errors destroys the fit of the Poisson model to the data suggests that the sequencing processes itself acts as a correlating influence on the inconsistency probabilities from one question to the next. This raises the possibility that more sequencing is being done in studies like the SIPP than is optimal. This potential problem is analogous to the problem of optimal interviewer workloads when the interviewer acts as a correlating influence for response errors. The trade-off in that case is that training costs decrease with work load while response variance increases. In the present case, the overall interview length can be reduced by skipping entire classes of respondents around questions based on their responses to earlier questions. The resulting interviewing time savings come at a cost of increased response (broadly defined) variance and therefore decreased question reliability. As is the case with interviewer workloads, this cost is generally unknown and is often ignored in the survey design process, with the result that sequencing may be over utilized just as work loads are often too high.

4. Conclusions

In this paper we have analyzed data from the SIPP reinterview program to see if it can be of value in understanding nonsampling error issues. We conclude that it can, indeed, be very valuable in several ways. First, it allows us to appreciate the fact that not all inconsistencies in the data are due to respondents providing unreliable reports. A goodly portion of the discrepancies between interview and reinterview reports is due to inconsistencies in the interview procedures. The skip sequences used in the SIPP are complex and are not always successfully followed by the interviewers. Second, the reinterview data has proven valuable in identifying particular questions with unusually high response variances. This is important not just for analyst who may wish to correct for question reliability, but for future redesigns of the SIPP questionnaire. Third, we have shown with the reinterview data that data quality does vary systematically from one type of respondent to the next. Data quality appears to be significantly lower for low income, Black, and either very young or very old respondents. Finally, while there are significant effects of things which can only be attributed to the interviewing procedure or the interviewer her or himself, the quality of SIPP data apparently improved significantly between February and August of 1984.

Finally, the results of the present analysis lead to one recommendation for the future redesign of the main SIPP instrument itself. This is that the rather complicated skip sequences currently being used be simplified—they are causing relatively minor response errors to be amplified into much more serious problems.

References

Hausman, J., B.H. Hall, and Z. Griliches, 1984, "Econometric Models for Count Data with an Application to the Patents-R&D Relationship", *Econometrica*, 52:4, pages 909–938.

Maddala, G.S., 1983, *Limited-Dependent and Qualitative Variables in Econometrics*, (Cambridge University Press). pages 51–55.

O'Muircheartaigh C. A. and R.D. Wiggins, 1981, "The Impact of Interviewer Variability in an Epidemiological Survey", *Psychological Medicine*, 11, pages 817–824.

Press, W.H., B.P. Flannery, S.A. Teukolsky and W.T. Vetterling, 1986, *Numerical Recipes*, (Cambridge University Press).

Smith, R., 1987. "The SIPP Results of the Reinterview Program for 1985", (Washington, D.C.: Bureau of the Census, June 18, 1987).

Footnotes

[1]I would like to thank Dan Kasprzyk, Fred Cavanaugh and Chet Bowie of the Census Bureau for making the data available and Laura Klem of the Survey Research Center for merging the reinterview data with the public release files. I would also like to thank Graham Kalton, Dan Kasprzyk, Jim Lepkowski, Jeff Moore, Gary Shapiro, Irv Schriener and Vicki Stout for their helpful comments on earlier versions of this paper. Finally, I would especially like to thank Dan Kasprzyk for presenting the paper in my unavoidable absence from the meetings. The work was sponsored, in part, by a Joint Statistical Agreement (JSA-87-5) between the Census Bureau and the University of Michigan. Any errors are my responsibility.

[2]This is not to say that SIPP data are in any sense more error prone than other survey data. The error that exists, however, is more easily seen because of the longitudinal nature of the data.

[3]If the respondent's reinterview response is affected by his memory of the interview response, errors in the two will tend to be positively correlated rather than independent. Thus, to this extent, the estimated response variances presented in the present analysis will tend to be conservative.

[4]See Hausman, Hall, and Griliches (1984).

[5]The algorithm we employ is written in Pascal by the author using sub-routines described in Press, Flannery, Teukolsky and Vetterling (1986).

[6]This is the method of moments technique for fitting the data to the distribution.