Measuring Chronic Poverty Without a Panel

John Gibson[†]

Abstract

A new method of decomposing cross-sectional poverty estimates into chronic and transient components is demonstrated using data from a recent household survey in Papua New Guinea. This method is simpler than previously used panel methods because it does not require data on every household in every period. The only requirement is that a subset of the surveyed households have a repeat observation made on their welfare indicator some time after the initial observation. In the setting considered, the chronic and transient components of headcount poverty are roughly equal, while three-quarters of the mean poverty gap is transient – in the sense of being due to within-year fluctuations in expenditures.

Acknowledgements:

I am grateful to two anonymous referees, Scott Rozelle, Christopher Scott, Jeffrey Williams and Fred Zimmerman for helpful advice. Data used in this paper were originally collected as part of a World Bank poverty assessment for Papua New Guinea, for which financial support from the governments of Australia (TF-032753), Japan (TF-029460), and New Zealand (TF-033936) is gratefully acknowledged. All views in this paper are those of the author and should not be attributed to the World Bank.

JEL: C81, I32

Keywords: Chronic poverty, Household expenditures, Survey design, Transient poverty

[†] Department of Economics, University of Waikato, Private Bag 3105, Hamilton, New Zealand.
Fax: (64-7) 838-4331. E-mail: jkgibson@waikato.ac.nz.

I. Introduction

In a given time period, people may be poor either because their mean consumption over time is below the poverty line or because they have suffered a temporary shortfall in consumption. The poverty that is due to intertemporal variability in consumption can be considered as "transient", while that which persists in mean consumption over time can be considered as "chronic poverty" (Jalan and Ravallion, 1998).[1] Distinguishing between these two types of poverty may be useful because the policies for dealing with chronic poverty – increasing the level of, and returns to, assets owned by the poor – may not be successful in dealing with transient poverty. Policies designed to even out fluctuations in welfare over time, such as insurance and income stabilisation schemes, may be more appropriate ways of dealing with transient poverty (Lipton and Ravallion, 1995).[2]

Even if investigators are just interested in chronic poverty, they may want to remove the transient component because it is unlikely to be a constant fraction of the total poverty measure either over time or across population sub-groups. In particular, greater environmental variability, restricted access to smoothing mechanisms, and shorter reference periods over which consumption is observed will tend to raise the fraction of measured poverty which is of a transient nature.[3] If the transient component is not removed it will also reduce the sharpness of "poverty profiles", which seek to identify the characteristics of the poor for targeting purposes. For example, a well-educated person who is not chronically poor but has bad luck in the particular month that her consumption is observed may be counted as poor, making it harder to see that low education is a structural characteristic of the poor.

The usual method of distinguishing chronic from transient poverty is to identify those households whose welfare is below the poverty line in every period, and then separate this group out from each period's poverty measure, leaving the transient component as the residual (Ravallion, 1988). Hence, this approach requires longitudinal observations, in which the same households are tracked to form a panel. Because there are few panel surveys available in developing countries (Deaton, 1997), most studies do not distinguish chronic and transient poverty.[4] Consequently, existing cross-sectional estimates of poverty are mixtures of chronic and transient poverty, with differing (and unknown) weights on each component.

The aim of this paper is to measure chronic poverty without using a panel survey. This measurement also allows the decomposition of a cross-sectional estimate of poverty into transient and chronic components. In the empirical application, chronic poverty is defined in terms of annual expenditures, as distinct from transient poverty due to short term fluctuations in expenditures within a year. However, the method could be easily adapted so that chronic poverty was defined in terms of a multi-year average of some welfare indicator. The only data requirement is that a subset of the surveyed households have a repeat observation made on their welfare indicator some time after the initial observation. Hence, the method used here is potentially cheaper and simpler than a full panel study, because it does not require data on every household in every period.

The decomposition used here is adapted from the statistical literature on household surveys. In this literature, the problem is one of obtaining estimates of the variance of annual

expenditures for surveys that observe each household's expenditure for only a short period – a week, a fortnight, or a month – and extrapolate to annual totals.[5] This survey design can give good estimates of annual means if the sample is spread evenly over a year but it may give exaggerated estimates of the between-households dispersion in annual expenditures. This exaggeration occurs because the variance calculated from extrapolated expenditures includes the effect of shocks that happen during the short observation period but are evened out over the rest of the year (Scott, 1980). In the survey literature, these within-year fluctuations for a given household are treated as noise and an adjustment method has been devised by Scott (1992) to remove this component from the calculated variance.

In the context of poverty measurement, the within-year fluctuations in household expenditures represent transient poverty rather than random noise. Hence, the adjusted estimates of annual expenditures that remove this component are suitable for the measurement of chronic poverty. The main requirement for the calculation is an estimate of the average correlation between the same household's expenditures in all periods of the year. At a minimum, observations on the same household in two different time periods are needed to estimate this correlation. The required data are available from a recent survey in Papua New Guinea where observations on a subset of households were repeated during the year (Gibson and Rozelle, 1998). In this setting, the chronic and transient components of headcount poverty appear roughly equal, while three-quarters of the mean poverty gap is transient – in the sense of being directly attributable to within-year fluctuations in expenditures. The sensitivity of these results to using just two time periods to measure the average correlation in expenditures over time is also assessed, using additional survey evidence from urban China.

The paper proceeds as follows: Section II describes the adjustment method used to remove the effect of within-year expenditure fluctuations on the variance of extrapolated annual expenditures. Section III discusses the implications of this adjustment for poverty measurements. The household survey data used in the empirical application are described in Section IV. The results using the adjusted expenditures to decompose poverty into chronic and transient components are reported in Section V, while conclusions are in Section VI.

II. An Adjustment Method for the Variance of Extrapolated Annual Expenditures

Assume that the goal of a household survey is to estimate the average value of household annual expenditures, $\bar{x}_A$. Observing each household for 12 months (e.g., using expenditure diaries),[6] would give a measure of their annual expenditures, $x_{i,A}$ allowing the estimate:

$$\bar{x}_A = \frac{1}{N}\sum_{i=1}^{N} x_{i,A}\,. \tag{1}$$

The problem with the approach embodied in equation (1) is that the high cost of observing each household for so long may restrict the sample size and the number of other surveys fielded by the statistical agency while lack of accuracy may bedevil other methods of estimating $x_{i,A}$, such as an annual recall of expenditures from a single visit to the household. Statistical agencies have responded to these problems by observing each household's expenditure for only a short period – a week, a fortnight, or a month – staggering data collection so that the observations are spread evenly over the whole year. Assume that the observation period for each household is 14 days and that the sample is spread evenly (and randomly) over the 26 non-overlapping fortnights in the year. Further assume that all types of expenditure are measured by the 14-day observation

(i.e., there are no infrequent expenses). If x_{it} is the expenditure by household *i* in fortnight *t*, then an estimate of average expenditure in fortnight *t* for the whole country, $\bar{x}_t$ is given by:

$$\bar{x}_t = \frac{1}{N/26} \sum_{i=1}^{N/26} x_{it} \,. \tag{2}$$

Average annual expenditures are given by the sum of average expenditures for each fortnight, so:

$$\bar{x}_A = \sum_{t=1}^{26} \bar{x}_t \,. \tag{3}$$

Another way to get the same estimate of average annual expenditures as is given by equations (2) and (3) is to extrapolate from the 14-day observation on each household,

$$x^e_{i,A} = 26 \cdot x_{it} \tag{4}$$

and then take the average of these extrapolated estimates of annual expenditures:

$$\bar{x}_A = \frac{1}{N} \sum_{i=1}^{N} x^e_{i,A} \,. \tag{5}$$

The estimates given by either equations (2) and (3) or equations (4) and (5) can be used instead of the more expensive estimates obtained from equation (1).

When attention shifts from estimating means to estimating measures of dispersion, the above results no longer hold (Scott, 1992). Annual expenditure estimates formed by extrapolating from a short observation period will have a variance, $V(x^e_{i,A})$ that exceeds the variance of true annual expenditures, $V(x_{i,A})$. The variance is raised because some of the shocks that occur in the short reference period are evened out over the course of a year, and this would be observed if households were monitored for 12 months. Let $V(x_{i,14})$ be the estimate of the between-households variance in fortnightly expenditures:

$$V(x_{i,14}) = \frac{1}{N}\sum_{i=1}^{N}\left(x_{it} - \bar{x}_{14}\right)^2 \tag{6}$$

where $\bar{x}_{14}$ is the average value of fortnightly expenditures across all i households and t fortnights. The variance of extrapolated annual expenditures, which are estimated by multiplying fortnightly expenditure by 26 is:

$$V(x^e_{i,A}) = \frac{1}{N}\sum_{i=1}^{N}\left(x^e_{i,A} - \bar{x}_A\right)^2 . \tag{7}$$

If equation (4) is substituted into equation (7), and noting that $\bar{x}_A = 26\cdot\bar{x}_{14}$ and 26^2=676:

$$V(x^e_{i,A}) = 676\cdot V(x_{i,14}) . \tag{8}$$

The variance of $x_{i,A}$, $V(x_{i,A})$, which is calculated with annual expenditure estimates from data collected by observing each household for a full year, is substantially less than the variance implied by equation (8). To find an expression for $V(x_{i,A})$, note that the deviation of each household's annual expenditures from the annual mean is the sum of deviations of each household's fortnightly expenditure from the mean for that fortnight, $d_{it} = x_{it} - \bar{x}_t$:

$$x_{i,A} - \bar{x}_A = \sum_{t=1}^{26} x_{it} - \sum_{t=1}^{26}\bar{x}_t = \sum_{t=1}^{26} d_{it} . \tag{9}$$

By definition, the variance of $x_{i,A}$ is:

$$V(x_{i,A}) = \frac{1}{N}\sum_{i=1}^{N}\left(x_{i,A} - \bar{x}_A\right)^2 . \tag{10}$$

Equation (9) can be substituted into equation (10) to give,

$$
\begin{aligned}
V(x_{i,A}) &= \frac{1}{N}\sum_{i=1}^{N}\left[\sum_{t=1}^{26} d_{it}\right]^2 \\
&= \frac{1}{N}\sum_{i=1}^{N}\left[\sum_{t,t'=1}^{26} d_{it}d_{it'}\right] \\
&= \frac{1}{N}\sum_{t,t'=1}^{26}\left[\sum_{i=1}^{N} d_{it}d_{it'}\right].
\end{aligned}
\tag{11}
$$

Equation (11) is useful because it allows us to relate $V(x_{i,A})$ to the Pearson correlation coefficient for the expenditures by the same household in a given pair of fortnights t and $t\phi$, estimated over all households. This correlation coefficient is defined as:

$$
r_{t,t'} = \frac{1}{N}\sum_{i=1}^{N} d_{it}d_{it'} \Big/ s_t s_{t'}, \tag{12}
$$

where s_t is the standard deviation across households in the tth fortnight. If equation (12) is substituted into equation (11), the variance of annual expenditures can be expressed as:

$$
V(x_{i,A}) = \sum_{t,t'=1}^{26} r_{t,t'} s_t s_{t'}. \tag{13}
$$

Splitting the summation in equation (13) into two parts – one part where $t=t\phi$ and the other where $t\neq t\phi$, equation (13) can be rewritten as:

$$
V(x_{i,A}) = \sum_{t=1}^{26} V(x_{i,14}) + 2\sum_{t\neq t'} r_{t,t'} s_t s_{t'} \tag{14}
$$

because $r_{t,t}=1$ and s_t^2 is just $V(x_{i,14})$, as already defined in equation (6).

The expression for the variance of annual expenditures, $V(x_{i,A})$, can be simplified if it is assumed that dispersion across households does not vary from fortnight to fortnight.[7] With this assumption, we can define $s_t s_{t\phi} = V(x_{i,14})$, so that equation (14) can be simplified to:

$$V(x_{i,A}) = \left[26 + 2 \sum_{t \neq t'} r_{t,t'} \right] V(x_{i,14}) \,. \tag{15}$$

Note that the average correlation between the same household's expenditures in all pairs of fortnights in the year, $\bar{r}$ is given by:

$$\bar{r} = \frac{1}{325} \sum_{t \neq t'} r_{t,t'} \,. \tag{16}$$

If equation (16) is inserted into equation (15), the result is:

$$V(x_{i,A}) = \left[26 + 2 \cdot 325 \cdot \bar{r} \right] V(x_{i,14}) \,. \tag{17}$$

The proportion by which the variance of annual expenditure estimates that are formed by extrapolating from a short observation period, $V(x^e_{i,A})$ exceeds the true variance of annual expenditure, $V(x_{i,A})$ is seen by dividing equation (8) by equation (17):

$$\frac{V(x^e_{i,A})}{V(x_{i,A})} = \frac{676}{\left[26 + 650 \cdot \bar{r} \right]} \,. \tag{18}$$

It is clear from equation (18) that the variance of extrapolated annual expenditure, $V(x^e_{i,A})$ will equal the true variance only if expenditures by the same household in every pair of fortnights in the year are perfectly correlated (i.e., $\bar{r} = 1$).

Figure 1 shows the extent to which $V(x^e_{i,A})$ exceeds the true variance when there is less than perfect correlation between expenditures in different fortnights of the year. For example, if the average correlation between the same household's expenditures in all pairs of fortnights in the year is only 0.60, the extrapolation method will overstate the variance by a factor of 1.63. There is only a slight improvement if the extrapolation is made from an observation period of

one month rather than one fortnight, with the variance exaggerated by a factor of 1.58. To the extent that some types of expenditure are measured with an annual reference period, the exaggeration in the variance due to extrapolation will be somewhat less than that predicted by equation (18).[8] The second curve in Figure 1 shows the situation when one-fifth of expenditures are measured with an annual reference period, with this percentage being the relevant one for the PNG survey used in the empirical part of the paper.

Previous Evidence on $\bar{r}$ in Developing Countries

Scott (1980) reports that a 1967-68 survey in rural areas of Lesotho, where households were re-interviewed every month for a year, found an average correlation in expenditures between months for the same households of 0.5. A similar correlation was found in a 1964-65 survey in an urban area of Cameroun, where households were observed in just two months of the year, separated by a five-month gap (Scott, 1980). More recently, the 1993-94 Household Budget Survey in Zambia observed households for four one-month periods spread over a year (CSO, 1995). The estimated $\bar{r}$ was 0.55 for rural households, 0.66 for low-income urban households, and 0.69 for high-income urban households. These correlations suggest that the variance will be raised when household surveys extrapolate annual expenditure estimates from short observation periods, with the largest effect likely in the rural sector.

Adjusting the Variance of Annual Expenditure Estimates

Equation (17) suggests that the variance of extrapolated annual expenditures can be adjusted to equal the variance that would result from observing each household for a full year. The two components needed for an adjustment are (i) the variance in fortnightly expenditures,

$V(x_{i,14})$ and (ii) the average correlation between expenditures in all pairs of fortnights in the year, $\bar{r}$. Although the best estimate of $\bar{r}$ would use the 325 correlation coefficients, $r_{t,t\ell}$ obtained from observing each household in all fortnights, a good estimate of $\bar{r}$ may be obtained from observations in just a few fortnights – possibly even just two (Scott, 1995). This economical approach to estimating $\bar{r}$ will be valid if the correlations among periods have roughly the same value and vary little as the gap between observations increases, as was found by the 1993-94 Household Budget Survey in Zambia where $r_{t,t\ell}$ fell by just 0.0078 for each month that the gap between t and $t\ell$ increased.[9] The 1967-68 survey in Lesotho also found that the period between observations had only a slight effect on the estimated correlations (Scott, 1992).

In addition to the saving made by observing households in only a few periods of the year, further savings can be made by restricting the repeated observations to a random subset of the sampled households. This random sub-sample should be large enough to allow $\bar{r}$ to be calculated separately for major groups of the population (e.g., rural and urban, rich and poor), because the extent to which expenditures fluctuate within the year may differ between these groups. It would also be advisable to carry out sensitivity analyses with different values of $\bar{r}$ because the use of sub-samples will increase the sampling error of $\bar{r}$ in an unknown way. In fact, a sub-sample approach was used in the 1993-94 survey in Zambia, where one-fifth of households were observed repeatedly and $\bar{r}$ was estimated from this sub-sample and then used to correct the variance for the full sample (although sensitivity analyses on $\bar{r}$ were not carried out).

III. Implications for Poverty Measurement

The adjustment to the variance of extrapolated annual expenditure estimates, described by equation (17), removes the component of the calculated variance that is due to within-year fluctuations in expenditures. In the context of poverty measurement, these expenditure fluctuations contribute to transient poverty. If poverty measures based on a single fortnight's expenditures are viewed as the sum of transient and chronic components, then the adjustment method can be used to isolate the chronic component of poverty. Rather than adjusting the variance, as in equation (17), the expenditure estimates for each household can be adjusted to remove the effect of within-year fluctuations, and these adjusted expenditures can be used to measure chronic poverty.

The estimate of the *i*th household's extrapolated annual expenditures, $x^e_{i,A}$ can be broken into two parts: the household's deviation from the overall fortnightly average, $(x_{it} - \bar{x}_{14})$ scaled up to an annual value, and the annual average across all households, $\bar{x}_A\,(= 26 \cdot \bar{x}_{14})$:

$$x^e_{i,A}{}^* = (x_{it} - \bar{x}_{14})\sqrt{V(x_{i,A})/V(x_{i,14})} + 26 \cdot \bar{x}_{14}\,. \tag{19}$$

The key to the adjustment is the scaling factor. When annual expenditure is extrapolated from 14-day observations the scaling factor is 26 (see equation (8)) and the effect of transient shocks occurring during the 14-day observation period is fully included. To remove the transient component, adjusted estimates of annual expenditures, $x^e_{i,A}{}^*$ can be estimated by altering the scaling factor, with equation (17) substituted into equation (19):

$$x^e_{i,A}{}^* = (x_{it} - \bar{x}_{14})\sqrt{26 + 650 \cdot \bar{r}} + 26 \cdot \bar{x}_{14}\,. \tag{20}$$

So, for example, if the average correlation between the same household's expenditures in all pairs of fortnights in the year is 0.50, the scaling factor is only 18.7 (=√351).

Short-term fluctuations in expenditures may result in some poor households temporarily exiting poverty, while others who are not normally poor temporarily experience poverty. Therefore the measure of chronic poverty, based on the adjusted expenditure estimates from equation (20), may be either higher or lower than the total poverty measure that includes the transient component. Following the approach of Ravallion (1988), let z be the fixed poverty line, set in terms of the welfare indicator y, and $f(y)$ is the density function of y. The poverty of each individual is represented by a function $p(y, z)$, with $\partial p/\partial y \leq 0$ and $\partial p/\partial z > 0$. Amongst the general class of additive poverty measures, the value of aggregate poverty is given by:

$$P(z) = \int_0^z p(y, z)\ f(y)\, dy. \qquad (21)$$

When the headcount index is used as the poverty measure, $p(y,z)=1$ for all $y \leq z$, and equation (21) is just the cumulative distribution function, $F(y)$. As long as $f(y)$ is unimodal, $F(y)$ has a unique inflexion point at the mode of y, being convex below this point and concave above. The curvature in these two parts of the distribution function is greater for the less dispersed welfare indicator. Below (above) the mode, the distribution function of the less variable welfare indicator will be below (above) the distribution function of the more variable indicator. Thus, if the poverty line is below the mode of the welfare indicator, the removal of the transient component will cause the headcount index of chronic poverty to be lower than the headcount index for the total poverty measure, while the reverse applies if it is above the mode.[10]

IV. Data and Methods for the Case Study

The Papua New Guinea Household Survey (PNGHS) is a random sample of 1200 households, residing in 120 rural and urban communities ("clusters"), who were interviewed between January and December 1996 as part of a World Bank poverty assessment. The clusters were selected from the enumeration areas of the 1990 Census, using a stratified sample (15 strata). Household weights were derived from (i) the unequal sampling rates between strata, (ii) the variation between the 1990 Census estimates of the size of each cluster and the actual size found during the survey, and (iii) the deviation of the actual number of households surveyed in each cluster from the target number. All results presented below take account of the clustered, stratified and weighted nature of the sample. The results reported are estimated from the 1144 households who had complete information on their expenditures.

The survey used a closed interval recall method, with households interviewed twice so that the start of the recall period was signaled by the first interview. These two interviews were usually two weeks apart, which is the length of the pay period in Papua New Guinea.[11] Expenditure data were collected on all food (36 categories) and other frequent expenses (20 categories) during the recall period. The expenditure estimates include the imputed value of own-production,[12] net gifts received, and stock changes, so they should be a good measure of consumption during the recall period. An annual recall covered 31 categories of infrequent expenses. An inventory of durable assets was used to estimate the value of the flow of services from these assets, including rental services from owner-occupied dwellings. The expenditures by households of different size and demographic composition were standardized by dividing by the

number of adult-equivalents, where children aged 0-6 years count as one-half of an adult and everyone else counts as an adult.[13] There was no allowance for economies of household size.[14]

Twenty clusters were chosen, randomly, as a "longitudinal sub-sample". Expenditures by households in these clusters were observed for two periods of the year, roughly seven months apart. These two sets of observations on expenditures by the same households are needed to estimate the correlation coefficient, $r_{t,t\ell}$ (one of 325 possible correlations between different fortnights of the year). If the results from previous studies hold, this single correlation may give a reasonable estimate of $\bar{r}$, the average correlation between the same household's expenditures in all pairs of fortnights in the year. Furthermore, if the households in the longitudinal sub-sample are representative of the whole sample, the estimate of $\bar{r}$ can be used to adjust the estimates of annual expenditures for the whole sample (using equation (20)) so as to derive measures of chronic poverty. The assumption that intra-year correlations do not vary much as the gap between observations increases is clearly a crucial one, and one that was not able to be tested given the survey design. However, it is possible to conduct sensitivity analyses using different values of $\bar{r}$.

There is no statistically significant difference between the average values of household characteristics for the longitudinal sub-sample and the average for the households in the other 100 clusters (the "cross-sectional sub-sample"), thus the longitudinal sub-sample appears to be representative of the whole sample (Table 1). In particular, households in the two sub-samples have the same average level of expenditure per adult equivalent and the same pattern of economic activity. Households in the two sub-samples have similar access to transport

infrastructure, which affects price (and hence, expenditure) fluctuations during the year by allowing movement of goods from surplus to shortage areas. Finally, a similar proportion of clusters in the two sub-samples are in dry areas, which tend to have more variable food production over the course of the year.

The Poverty Lines

Poverty lines were set for five regions of Papua New Guinea – the National Capital District (NCD), the South Coast, the Highlands, the North Coast, and the Islands – using methods outlined by Ravallion (1994). These poverty lines were based on baskets of locally consumed foods that provide 2200 calories per day. A comparison of the food budgets of poor households in each region showed that a single national basket of foods was inappropriate, so separate baskets of foods were used for the NCD, the Highlands, and the lowland regions. To ensure that these baskets provided diets of the same quality, they were formed from the food budgets of households living below the same level of real expenditure per adult-equivalent, rather than from the poorest x percent of households in each region. Thus, differences in real consumption levels between regions did not translate into differences in the poverty line diets. Furthermore, revealed preference tests were carried out by checking if the cost in region j of buying the region i basket of foods was less than the cost of buying the region j basket of foods, at region j prices.[15] These tests provided no evidence that any region had a poverty line diet that was inferior to that of the other regions. The annual cost of the poverty line diets varied from K543 in the NCD to K218 in the North Coast region, with a national average of K300.[16]

The final value of the poverty lines was set by adding an allowance for non-food items. This allowance is based on the typical value of non-food spending by households whose total expenditure equals the cost of the food poverty line. Consuming these non-food items means that some food needs are ignored, so the non-food items can be considered as essentials (Ravallion, 1994). The average food share for households whose total expenditure equals the food poverty line is found from the following Engel curve:

$$w = a + b \ln\left(\frac{x}{n \cdot z_j^F}\right) + \sum_{k=1}^{K} g_k n_k + \sum_{j=1}^{J-1} f_j D_j + e$$

where w is the food budget share, x is total expenditure, n is the number of persons, z_j^F is the food poverty line for an adult-equivalent in region j, n_k is the number of people in the kth demographic category, and D_j is an intercept dummy for region j. If total expenditure equals the cost of the food poverty line, $\ln\left(x/\left(n \cdot z_j^F\right)\right) = 0$, so $a_j = \hat{a} + \sum_{k=1}^{K} \hat{g}_k \bar{n}_k + \hat{f}_j$ gives the average food share in region j, where $\bar{n}_k$ is the mean of the demographic variables for the households used to form the poverty line basket of foods. The poverty line z_j is given by the sum of the food and non-food components, $z_j = z_j^F + z_j^F (1-a_j) = z_j^F (2-a_j)$. This ranged from K779 in the NCD to K280 in the North Coast, with a national average of K400.

The Poverty Measures

The P_a class of poverty measures developed by Foster, Greer and Thorbecke (1984) is used. This class contains several commonly used poverty measures as special cases, including the head-count index and the poverty gap index. The head-count index indicates the incidence, but not the depth, of poverty. The poverty gap index – the overall shortfall between the poverty

line and the expenditure level of the poor, as a ratio to the product of the poverty line and the population size – indicates the average depth of poverty. However, the poverty gap index is insensitive to the distribution of expenditures amongst the poor, so an improvement is to weight each of the poverty gaps so that the poorest people have the biggest weight. With the P_a class of measures, the higher the value of the poverty aversion parameter, a the higher the weight on the poverty gaps of the poorest people. The general formula is:

$$P_a = \frac{1}{n}\sum_{i=1}^{q}\left(\frac{g_i}{z}\right)^a ,$$

where n is the total population of equivalent adults, q is the number of equivalent adults in poor households, z is the poverty line, and g_i is the poverty gap for the ith adult-equivalent, $g_i = z - y_i$, (where y_i is expenditure per adult-equivalent in the ith household). When a=0, the P_0 measure is the head-count index, when a=1 the P_1 measure is the poverty gap index. These two, plus the P_2 measure, are used in this study.

In keeping with most studies of poverty, the poverty measurements reported are based on the number of persons living in poor households, rather than on the number of poor households.[17] The shift in focus to individuals means that the equations in Section III – which are for *household* expenditures – may need to be modified to get adjusted estimates of annual expenditures *per adult-equivalent,* $y^e_{i,A}$*. Expenditure per adult-equivalent can fluctuate during a year because either household total expenditure fluctuates or household size fluctuates. The equations in Section III deal only with the first source of instability. Only in the special case where household composition is stable over the year can $y^e_{i,A}$* be constructed by first adjusting

estimates of annual household expenditures, and then dividing these by the number of adult-equivalents. Otherwise, $y^{e}_{i,A}$* has to be constructed by forming estimates of expenditure per adult-equivalent in each fortnight, using these to estimate $\bar{r}$, and then proceeding with the other steps in the calculation of adjusted annual estimates (following equation (20)).

V. Results

Table 2 contains the estimates of the between-fortnights correlations for the households in the longitudinal sub-sample of the PNGHS. The correlations are presented for three variables of interest: expenditure per household, household size, and expenditure per adult-equivalent. The correlations are estimated for the full longitudinal sub-sample, which is nationally representative, and separately for the households in rural and urban clusters. This split reflects the policy interest in these two sectors and the evidence from other countries that urban households may have more stable expenditures over the year (and, thus, a smaller fraction of poverty that is of a transient nature).

The correlation between each household's expenditures in two different fortnights, seven months apart, is 0.65. The correlation is 0.60 for rural households and 0.53 for urban households. These values are similar to the range of estimates reported above for other developing countries. The greater instability of expenditures by urban households is seemingly because these households have a less stable demographic composition. The correlation between the two reports of household size, for the same urban households seven months apart, is only 0.65. The composition of rural households is more stable, with a correlation of 0.75. One reason for the

greater instability in household size in urban areas is that urban households often host extended family members – who may be migrants from rural areas or urban residents without their own dwelling – for visits that can last several months. This is less common in rural areas because families have their own dwellings, and they can quickly construct new dwellings from bush materials if they move to a new location. Since there is instability in household composition over the course of the year, the correlations for expenditure per adult-equivalent are the appropriate ones to use. These correlations have the expected pattern of urban dwellers having more stable expenditures over the year than do rural dwellers, with correlation coefficients of 0.65 and 0.51.

A modified version of equation (20) is used to construct adjusted estimates of annual expenditure per adult-equivalent:

$$y_{i,A}^{e}* = \left(y_{it} - \bar{y}_{14}\right)\sqrt{26 + 650 \cdot \bar{r}} + 26 \cdot \bar{y}_{14} + AP_i / n_i \,, \qquad (22)$$

where $y_{it} = x_{it}/n_i$ is the value of expenditure per adult-equivalent in household *i* in fortnight *t*, $\bar{y}_{14}$ is the average of y_{it} across all *i* households and *t* fortnights, AP_i is the *i*th household's annual expenditures on items that were not surveyed with a 14-day reference period (e.g., infrequent expenses, annual services provided by durable assets), and n_i is the number of adult-equivalents in household *i*. Summary statistics for the extrapolated (i.e., unadjusted) and the adjusted estimates of annual expenditure per adult-equivalent are reported in Table 3.

The adjusted estimates of annual expenditure per adult-equivalent have the same average but a lower variance than the unadjusted estimates. Note that $V(y_{i,A}^{e}) \big/ V(y_{i,A}^{e}*)$ for the rural sector in Table 3 is 1.53 (=[733/592]^2), while equation (18) would predict that the variance of the

unadjusted extrapolation overstates $V(y_{i,A})$ by a factor of 1.89. The discrepancy is because equation (18) assumes that all expenditures are covered by the short-period observations (i.e., AP_i=0). In fact, the expenses measured with an annual reference period are common to both the adjusted and unadjusted estimates of annual expenditures, which reduces the difference.

The overall level of inequality in annual expenditures appears to be 17-30 percent higher when the unadjusted estimates are used, depending on the particular inequality measure (Table 3). The greatest difference is for the Theil index, which is sensitive to income differences at the top of the distribution, while there is less difference in the Gini coefficient, which is sensitive to income differences about the mode of the distribution. The adjustment to the annual expenditure estimates has the biggest effect on inequality measures in the rural sector, where the inclusion of within-year fluctuations raises the measures of annual inequality by 27-55 percent. These within-year fluctuations are sufficient to raise measured inequality in the rural sector above that in the urban sector, while the adjusted estimates imply the opposite ranking.

Further evidence on the effect of the adjustment method on measured dispersion comes from a comparison of the smoothed densities (Figure 2). These densities have been calculated using a kernel density estimator with a Gaussian kernel and a bandwidth of $0.9 \cdot n^{-0.2} \cdot \min\{sd, iqr/1.34\}$, where *sd* and *iqr* are the standard deviation and interquartile range of the estimates and n=1144 is the number of observations (see Salgado-Ugarte, Shimizu and Taniuchi, 1993) and are estimated from population-weighted data. A logarithmic transformation has been applied so that the effect of the adjustment can be seen more clearly for the lower tail of the distribution. Although the use of logarithms partly explains why the adjustment appears to affect the density

function mainly in the bottom half, there is a further reason for this pattern. The adjustment method is designed to be mean-preserving, so the positive skew in the expenditure distribution means that more households have their adjusted estimates of annual expenditures raised (57 percent) than lowered.

How much chronic, transient and total poverty is there in Papua New Guinea? Using the poverty estimates that include within-year expenditure fluctuations to get the sum of chronic plus transient poverty, it appears that the headcount poverty rate is 30.4 percent (Table 4). But the estimates based on adjusted annual expenditure, which should remove the effect of (within-year) transient poverty, suggest that only 15.0 percent of Papua New Guineans live in chronically poor households. Thus, the poverty headcount based on a fortnight's observation (extrapolated to an annual total) contains a transient component that is as large as the chronic component. This importance of transient poverty is similar to what has been found elsewhere using panel surveys,[18] although a complete test of the decomposition method used here would require applying it in a setting where the traditional method based on panel surveys has also been used.

The almost exact split of headcount poverty into chronic and transient components does not hold across the rural and urban sectors (Table 4). In the urban sector only one-third of the headcount poverty is of a transient nature. In contrast, slightly over one-half of the headcount poverty in the rural sector is of the transient variety, which follows from the greater within-year instability of expenditures per adult-equivalent for rural households. The lower fraction of poverty that is of a transient nature in the urban sector also holds for the poverty gap and poverty severity measures, reflecting the advantages that urban households have in accessing

consumption smoothing mechanisms.[19] The variation across sectors in the fraction of poverty which is transient means that poverty comparisons that are just interested in chronic poverty could be biased because cross-sectional estimates of poverty rates may have differing degrees of transient poverty included.

The transient component makes an even larger contribution to the total poverty measure for the higher order poverty statistics, P_1 and P_2. Three-quarters of the mean poverty gap is transient – in the sense of being directly attributable to within-year fluctuations in expenditures. One use of the poverty gap index is to indicate the minimum size of perfectly targeted transfers needed to eliminate poverty, given by the product of the poverty gap index (P_1), the poverty line, and the population size. Hence, the results in Table 4 show how the inclusion of a transient component affects the estimated cost of reducing chronic poverty. With a population of 4.4 million (in adult-equivalent terms) and a poverty line of K400, P_1=9.1 implies that at least K160m per year would be needed to eliminate poverty, as compared to a cost of only K40m per year, when the adjusted annual expenditure estimates are used for calculating the poverty gap index (P_1=2.3). This comparison is not meant to imply that eliminating chronic poverty should be the sole goal of policy. But chronic poverty may be more amenable to being reduced by resource transfers, while such transfers are less likely to help the transiently poor who face a problem of smoothing the consumption derived from their resources rather than an overall shortage of resources. Indeed, by causing the cost of reducing chronic poverty to be overstated, the inclusion of a transient component in cross-sectional estimates of the poverty gap may divert attention away from the interventions required to improve consumption smoothing amongst households who are normally not poor.

How sensitive are the results in Table 4 to the particular values of $\bar{r}$ used? Recalling that the estimate of $\bar{r}$ used to form adjusted annual expenditures is just one of 325 possible pairwise correlations, and is estimated from sub-samples for the urban and rural sectors, it is worth exploring what the results would look like with different values of $\bar{r}$. Similarly, methods of forming poverty lines vary, causing uncertainty about the exact location of the poverty line, so evaluating the results at different poverty lines may be valuable. In this regard, Table 5 reports estimates of the headcount index at $\bar{r} = 1$ (i.e., the unadjusted estimates), at $\bar{r} = \hat{r}$ (i.e., the estimated values reported in Table 3), and at $\bar{r} = \hat{r} \pm 25\%$ and for three different poverty lines: the estimate of *z* described above and $z \pm 25\%$.[20]

Three inferences can be drawn from Table 5, the first of which is that even at alternative values of *z* and $\bar{r}$ there is a substantial gap between the total poverty measure that includes the transient component (i.e., $\bar{r} = 1$) and the measures of chronic poverty based on the adjusted estimates of annual expenditures. But because this gap varies with the level of the poverty line and the degree of within-year expenditure fluctuations, the potential for cross-sectional estimates of poverty to give misleading inferences about chronic poverty rankings is underscored. Secondly, and unsurprisingly, the contribution of chronic poverty to the total poverty measure falls as the inter-temporal variability in consumption rises, as seen from the reduced level of chronic poverty as the estimate of $\bar{r}$ is lowered. Third, as the poverty line increases towards the mode of the expenditure distribution, the chronic component of the total poverty measure increases. This reduced contribution of transient poverty can be understood with the framework used by Ravallion (1988); for a unimodal distribution, both a more variable and a less variable

indicator will have similar cumulative values at the point of inflexion corresponding to the mode. Thus, the contribution of the transient component to cross-sectional poverty measures is likely to be greatest in settings with a low to moderate level of poverty and a high level of within-year consumption variability.

VI. Conclusions

Distinguishing between chronic and transient poverty is important because different types of policies are likely to be needed for dealing with each kind of poverty. Even analysts who are mainly interested in chronic poverty may need ways of measuring (and then removing) the transient component to ensure that their cross-sectional poverty comparisons do not give misleading inferences about chronic poverty rankings. The usual method of distinguishing chronic from transient poverty is to separate out those households whose welfare is below the poverty line in every period, leaving the transient component as the residual. Because this approach requires panel data, few empirical studies distinguish chronic from transient poverty in developing countries, and existing cross-sectional estimates of poverty are unknown mixtures of chronic and transient components.

The goal of this paper has been to decompose cross-sectional poverty estimates for Papua New Guinea into chronic and transient components without using panel data. The only data requirement of the method used here is that a subset of the surveyed households have a repeat observation made on their welfare indicator some time after the initial observation. Hence, only a small change in the usual design of cross-sectional household surveys is needed, adding about 10 percent to the fieldwork cost in the PNG environment. This is likely to be far cheaper than a

full panel survey and so the method may offer a simple and practical alternative to the usual methods of decomposing poverty estimates into chronic and transient components. Although the empirical application has involved adjusting annual expenditure estimates to remove the effect of within-year expenditure fluctuations, with the adjusted annual expenditure estimates used to measure chronic poverty, it could be easily adapted so that chronic poverty was defined in terms of a multi-year average of some welfare indicator.

Once the within-year fluctuations are removed from the expenditure data in Papua New Guinea, the headcount index of poverty falls from 30 percent to 15 percent. Hence, if chronic poverty is defined in terms of annual expenditures, the chronic and transient components appear roughly equal. This importance of transient poverty is similar to what has been found by studies in other countries that have used panel data. However the importance of transient poverty differs greatly between rural and urban sectors of Papua New Guinea because of the advantages that urban households have in accessing consumption smoothing mechanisms. The presence of transient poverty also raises the mean poverty gap and therefore also raises the estimated cost of eliminating chronic poverty with targeted transfers.

The major caveat to these results is that they are based on a survey with only a single revisit used to estimate the average correlation between the same household's expenditures over time. This design relies heavily on the assumption that intra-year correlations do not vary much as the gap between observations increases. Recent evidence from urban China supports this use of a single revisit for adjusting annual expenditure estimates in the manner described here (Gibson, Huang and Rozelle, 1999). This Chinese evidence comes from a survey where

respondents keep a 12-month expenditure diary, with a poverty line set to define 30 percent of the population as poor according to their 12-month expenditures. In contrast to this measure, which can be considered to reflect chronic poverty, using a single month's expenditure data for each household gives a measure of chronic plus transient poverty – for which the headcount index is 47 percent. If expenditures in a second month (six months after the first) are used to estimate the average correlation, so that an adjusted estimate of annual expenditures can be made following the methods described in this paper, headcount poverty is estimated as 30.5 percent – almost identical to the measure of chronic poverty obtained from the full 12-month data.

References

Behrman, Jere R., Foster, Andrew D., and Rosenzweig, Mark R. (1997), 'The dynamics of agricultural production and the calorie-income relationship: Evidence from Pakistan', *Journal of Econometrics* 77(1): 187-207.

Central Statistical Office (1995), *Republic of Zambia Household Budget Survey, 1993-94, Volume 1 – Main Tables and Report* Lusaka.

Deaton, Angus (1997), *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy* Johns Hopkins, Baltimore.

Deaton, Angus, and Muellbauer, John (1986), 'On measuring child costs: with applications to poor countries', *Journal of Political Economy* 96(4): 720-743.

Fields, Gary (1994), 'Data for measuring poverty and inequality changes in the developing countries' *Journal of Development Economics* 44(1): 87-102.

Foster, James, Greer, J., and Thorbecke, E. (1984), 'A class of decomposable poverty measures', *Econometrica* 52(4): 761-765.

Gibson, John, Boe-Gibson, Geua, and Scrimgeour, Frank (1998), "Are voluntary transfers an effective safety net in urban Papua New Guinea?' *Pacific Economic Bulletin* 13(2): 40-53.

Gibson, John, and Rozelle, Scott (1998), 'Results of the Household Survey Component of the 1996 Poverty Assessment for Papua New Guinea' *mimeo* Population and Human Resources Division, The World Bank.

Gibson, John, Huang, Jikun, and Rozelle, Scott (1999), 'Measuring poverty and inequality in urban China: Can we do it as accurately for less cost?' *mimeo* Department of Economics, University of Waikato.

Handa, Sudhanshu, and King, Damien, (1997), 'Structural adjustment policies, income distribution and poverty: a review of the Jamaican experience', *World Development* 25(6): 915-930.

Jalan, Jyotsna, and Ravallion, Martin (1998), 'Transient poverty in post-reform rural China', *Journal of Comparative Economics* 26(2): 338-357.

Lanjouw, Peter, and Ravallion, Martin (1995), 'Poverty and household size', *Economic Journal* 105(Dec): 1415-1434.

Lipton, Machael, and Ravallion, Martin (1995), 'Poverty and policy', in *Handbook of Development Economics, Volume III* (J. Behrman and T.N. Srinivasin, editors) Elsevier, pp. 2551-2657.

McCulloch, Neil, and Baulch, Bob (1999), 'Distinguishing the chronically from the transitorily poor: evidence from Pakistan', *IDS Working Paper*, No. 97, Institute of Development Studies, Sussex.

Morduch, Jonathan (1994) 'Poverty and vulnerability', *American Economic Review* 84(2): 221-225.

Morduch, Jonathan (1995) 'Income smoothing and consumption smoothing', *Journal of Economic Perspectives* 9(3): 103-114.

Ravallion, Martin (1988), 'Expected poverty under risk-induced welfare variability', *Economic Journal* 98(Dec): 1171-1182.

Ravallion, Martin (1994), *Poverty Comparisons* Harwood Academic Publishers, Chur.

Ravallion, Martin (1996), 'Issues in measuring and modeling poverty', *Economic Journal* 106(Sept): 1328-1343.

Salgado-Ugarte, I.H., Shimizu, M., and Taniuchi, T. (1993), 'Exploring the shape of univariate data using kernel density estimators', *Stata Technical Bulletin Reprints* vol. 3, pp. 155-173.

Scott, Chris (1980), 'Practical problems in conducting surveys on living standards' in *Conducting Surveys in Developing Countries: Practical Problems and Experiences in Brazil, Malaysia and the Philippines* (C. Scott, P. de Andre, and R. Chander, editors), LSMS Working Essay 10, World Bank, Washington D.C.

Scott, Chris (1992), 'Estimation of annual expenditure from one-month cross-sectional data in a household survey', *Inter-Stat* 8(Mar):57-65.

Scott, Chris (1995), 'Recommended Sample Design for the Papua New Guinea Household Survey, 1996' *mimeo* East Asia and Pacific Division, The World Bank.

Scott, Chris, and Ben Amenuvegbe (1991), 'Recall loss and recall duration: an experimental study in Ghana', *Inter-Stat* 4(1): 31-55.

STATIN (1994), *Jamaica Survey of Living Conditions Report* Statistical Institute of Jamaica and Planning Institute of Jamaica.

World Bank (n.d.), *The 1993 Nicaraguan Living Standards Measurement Survey: Documentation*, Poverty and Human Resources Division, Policy and Research Department, The World Bank, Washington.

World Bank (1993), *Basic Information: Peru Living Standards Measurement Survey (PLSS) 1991*, Poverty and Human Resources Division, The World Bank, Washington.

World Bank (1994), *Guyana: Strategies for Reducing Poverty* Report no. 12861, Latin American and Caribbean Region, Country Department III, Population and Human Resources Division, The World Bank, Washington.

World Bank (1996), *Guyana Living Standards Measurement Study, 1992-93: Basic Information Document* Country Operations Division II, Department III, Latin American and the Caribbean Region, The World Bank, Washington.

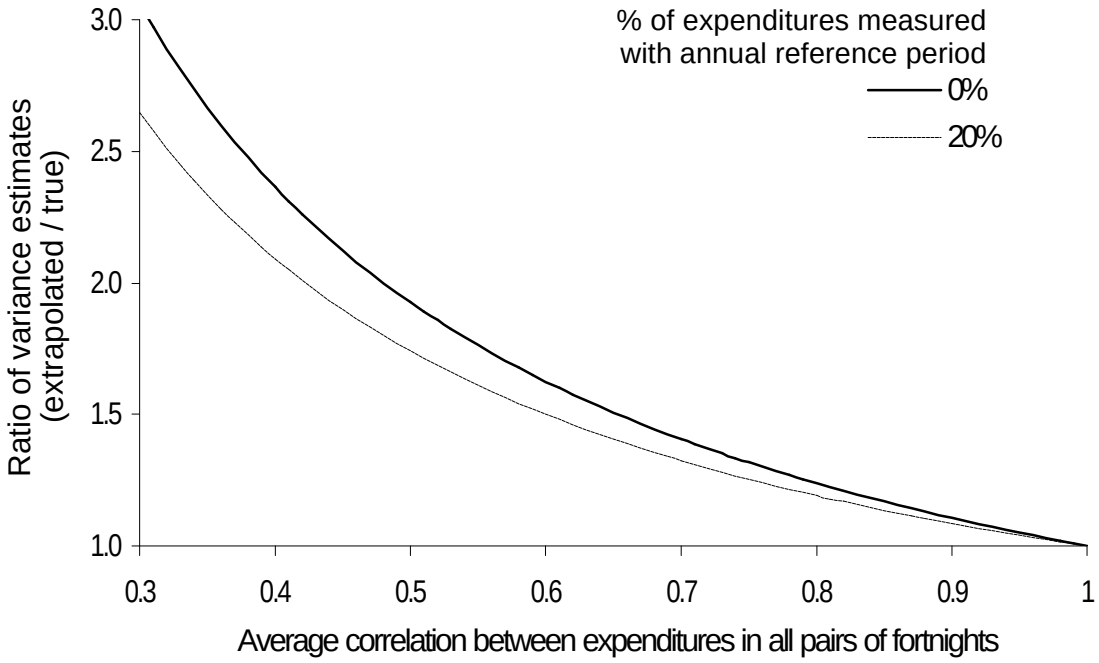

Figure 1: Variance of Extrapolated Annual Expenditures Compared With True Variance

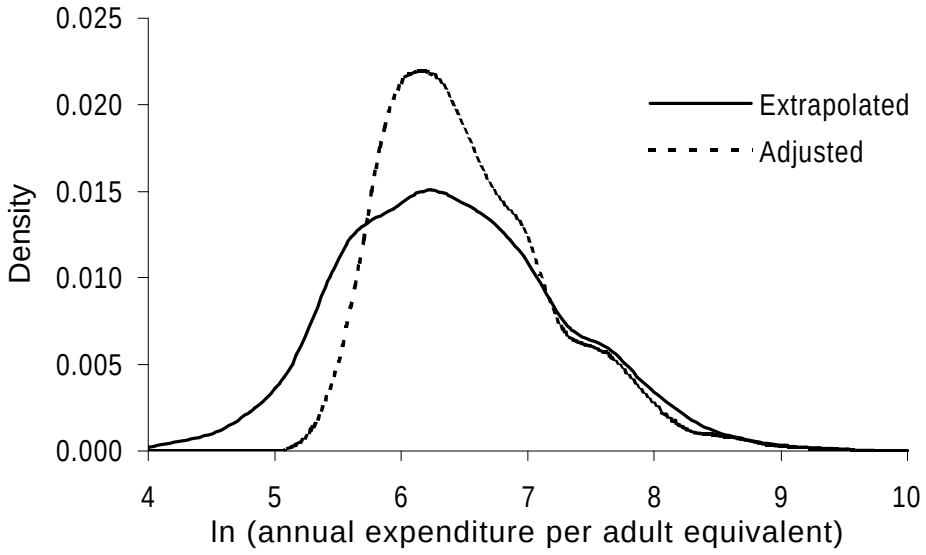

Figure 2: Extrapolated and Adjusted Annual Expenditures: Smoothed Densities

Table 1: Comparison of the Cross-sectional and Longitudinal Sub-samples

	Cross-sectional sub-sample[a] *N*=982	Longitudinal sub-sample[a] *N*=162	*t*-test for equal means[b]
Total expenditure per adult equivalent[c]	901.3 (83.2)	924.9 (130.7)	0.15 [0.89]
Household size	5.8 (0.2)	6.1 (0.3)	0.80 [0.43]
Years of school of household head	4.0 (0.3)	4.8 (0.8)	0.94 [0.36]
Age of household head	40.1 (0.6)	40.8 (1.7)	0.41 [0.69]
Female head (%)	7.9 (1.3)	8.4 (2.5)	0.16 [0.88]
Head's main income is wage job (%)	20.1 (2.4)	26.2 (8.9)	0.66 [0.52]
Head's main income is tree crop agriculture (%)	38.0 (4.4)	42.7 (9.1)	0.46 [0.65]
Minutes walk to nearest road, airstrip or port[d]	60.8 (15.4)	43.5 (21.9)	0.65 [0.52]
Dry climate (<2500mm rainfall/year) (%)[d]	39.5 (6.9)	42.6 (14.0)	0.20 [0.85]

Notes:

[a] Standard error of the mean in ( ) corrected for clustering, sampling weights and stratification.

[b] *p*-level for two-sided hypothesis test in [ ].

[c] Kina per year, in 1996 national average prices, where the value of the poverty line is used as the spatial price deflator and K1.3=US$1 in 1996.

[d] Data collected at cluster level, and weighted by the population in each cluster.

Table 2: Correlation Coefficients for Between-Fortnights Comparison of Expenditure and Household Size for the Same Household

	Full Sub-sample N=162	Rural N=127	Urban N=35
Expenditure per household	0.65	0.60	0.53
Household size (adult equivalents)	0.75	0.75	0.65
Expenditure per adult equivalent	0.57	0.51	0.65

Notes:
Correlation coefficients are calculated from data collected on the same household for two separate fortnights, approximately seven months apart.
Expenditure estimates are for food and other frequently purchased items and exclude annual expenses, the consumption of services from durable goods, and actual and imputed dwelling rent.
Expenditure estimates are corrected for inflation.

Table 3: Summary Statistics for Annual Expenditures per Adult-equivalent

	Papua New Guinea	Rural	Urban
Extrapolated			
Mean	911	743	2218
Standard deviation	1086	733	2064
Gini index	48.3	43.7	42.3
Theil index	42.9	33.6	31.2
Adjusted			
Mean	911	743	2218
Standard deviation	972	592	1911
Gini index	41.4	34.5	39.2
Theil index	33.1	21.7	26.8
Ratio of extrapolated to adjusted			
Mean	1.00	1.00	1.00
Standard deviation	1.12	1.24	1.08
Gini index	1.17	1.27	1.08
Theil index	1.30	1.55	1.16

Notes:
Means and standard deviations are population-weighted.
The "extrapolated" estimates of annual expenditure are constructed by multiplying fortnightly expenditure by 26 and adding infrequent expenses.
The "adjusted" estimates are calculated using equation (22) in the text, with separate estimates of $\bar{r}$ for urban and rural areas based on the values reported in Table 2.

Table 4: Poverty Measures for Papua New Guinea, 1996

Poverty measure (P_a) and sector	Total Poverty[a]	Chronic Poverty[b]	Percentage of total poverty that is: Chronic	Transient
Head-count index ($a=0$) (percent)				
Rural	33.2 (2.8)	16.2 (2.5)	48.9	51.1
Urban	8.2 (2.9)	5.3 (2.2)	64.7	35.3
Total	30.4 (2.6)	15.0 (2.2)	49.4	50.6
Poverty gap index ($a=1$) (percent)				
Rural	10.1 (1.2)	2.5 (0.5)	24.8	75.2
Urban	1.9 (0.8)	0.8 (0.4)	41.1	58.9
Total	9.1 (1.1)	2.3 (0.4)	25.2	74.8
Poverty severity index ($a=2$) (´100)				
Rural	4.4 (0.7)	0.6 (0.1)	13.9	86.1
Urban	0.7 (0.4)	0.2 (0.1)	30.2	69.8
Total	3.9 (0.6)	0.6 (0.1)	14.2	85.8

Note:

The poverty index is $P_a = \frac{1}{n}\sum_{i=1}^{q}\left(\frac{g_i}{z}\right)^a$ where n is the total population of equivalent adults, q is the number of equivalent adults living in poor households, z is the poverty line, and g_i is the poverty gap given by z minus expenditure per adult-equivalent in the ith household, y_i. The poverty index is more sensitive to lower expenditures by the poor, the higher is a.

Standard errors in ( ) corrected for clustering, sampling weights and stratification.

[a]The total poverty measure is based on extrapolated annual expenditures per adult-equivalent (obtained by multiplying fortnightly expenditure by 26 and adding infrequent expenses) and includes the effect of within-year expenditure fluctuations.

[b]The chronic poverty measure is based on adjusted annual expenditure per adult-equivalent (estimated using equation (22) in the text) and removes the effect of within-year expenditure fluctuations.

Table 5: Head-count Index of Annual Poverty in Papua New Guinea Under Differing Assumptions about Poverty Lines and Within-year Consumption Variability

	Poverty Line		
Annual expenditure estimates	$z-25\%$	z	$z+25\%$
Extrapolated ($\bar{r}=1$)	16.8	30.4	41.5
Adjusted ($\bar{r}=\hat{r}+25\%$)	5.9	19.9	34.8
	(35.0)	(65.6)	(83.9)
Adjusted ($\bar{r}=\hat{r}$)	3.4	15.0	30.8
	(20.4)	(49.4)	(74.3)
Adjusted ($\bar{r}=\hat{r}-25\%$)	0.9	8.7	24.8
	(5.5)	(28.5)	(59.9)

Notes: The values reported are the estimated percentage of the population (in adult-equivalent terms) living in poverty.

Values in ( ) are the percentage of total poverty that is due to the chronic poverty component, derived from the comparison of the poverty measures using adjusted annual expenditures with those using extrapolated annual expenditures.

The value of $\hat{r}$ is 0.65 for urban areas and 0.51 for rural areas, while the value of z varies by region but has a national average value of K400 per adult equivalent per year.

Notes

[1] A similar distinction is made by Morduch (1994) between *structural poverty*, where permanent income is less than the poverty line, and *stochastic poverty*, where current consumption falls below the poverty line even though permanent income exceeds the poverty line.

[2] The welfare implications of intertemporal variability may be substantial: Morduch (1995) calculates that poor households might be willing to give up 16 percent of their income to perfectly smooth consumption.

[3] Fields (1994) points out that long reference periods are more likely to capture permanent rather than transitory living standards. However, few empirical studies of poverty discuss the factors behind their choice of reference period and a recent survey is also agnostic on this issue, noting that poverty analysis usually begins with data on household expenditures or income *over some period* (Ravallion, 1996, p. 1328).

[4] Exceptions include Jalan and Ravallion (1998) and McCulloch and Baulch (1999).

[5] Developing countries which have such surveys include: Jamaica (STATIN, 1994, and Handa and King, 1997) where annualized expenditure was based on one week's food expenses multiplied by 52 and one month's recurring non-food expenses multiplied by 12; Guyana (World Bank, 1994 and World Bank, 1996) where annualized expenditure was based on one month's food expenses, and non-food expenses for either one month, three months, or 12 months; and Peru (World Bank, 1993) where annualized expenditure was based on a fortnight's food expenses, and non-food expenses for either a fortnight, three months, or 12 months.

[6] This is the approach used by the Chinese State Statistical Bureau for the national Household Income and Expenditure Survey but it is not used on a regular basis in any other country.

[7] One objection to this would be if the pay-period was monthly and if large purchases are made immediately after a pay-day, dispersion across households may vary between those surveyed in a fortnight that includes a pay-day, and those surveyed in a fortnight that excludes a pay-day. Another objection is that heterogeneity in preferences may affect food consumption in harvest season, compared with planting season when preferences matter less because food availability is close to subsistence level (Behrman, Foster, and Rosenzweig, 1997). However, these objections should not apply in Papua New Guinea because the pay-period is fortnightly, and most farmers continuously plant and harvest root crops throughout the year, rather than having defined planting and harvesting seasons.

[8] Of course, other types of errors may be introduced by these long term recalls (Scott and Amenuvegbe, 1991).

[9] I am grateful to Chris Scott for this information.

[10] Ravallion (1988) also shows that an increase in welfare variability will increase expected poverty, as measured by any member of the Atkinson class of poverty measures, if the welfare function is concave in the risky variable.

[11] The expenditures by different households were put on a consistent time period by using the ratio of the actual recall period to the target period of 14 days. Similarly, the effect that short-term guests and short-term absences by usual household members had on measured expenditures during the recall period was accounted for by using the ratio of actual to expected person-days.

[12] The monetary values for self-produced foods were the values used by respondents. Estimates of average expenditure and the poverty rate are unchanged if these respondent-reported unit values are replaced by either cluster medians of the unit values or cluster averages of market prices (Gibson and Rozelle, 1998).

[13] The adult-equivalence scale is based on estimates of child costs, made using the Engel and Rothbarth methods outlined by Deaton and Muellbauer (1986), and on a comparison of the dietary requirements of adults and children of various age groups. Details are provided by Gibson and Rozelle (1998).

[14] There was some evidence of household size economies, according to the Engel method used by Lanjouw and Ravallion (1995). However, doubts have been expressed about this method (Deaton, 1997), so this paper follows the usual practice for developing countries, of assuming that economies of size are absent. This assumption is unlikely to affect the main point of the paper because economies of size adjustments are usually normalized so that aggregate measures of poverty do not change.

[15] If the cost was less it would suggest that the region *i* basket of foods was less preferred than the region *j* basket, which would indicate that the region *i* poverty line gave a lower standard of living.

[16] The food poverty lines are slightly higher, and more variable between regions, if a single national basket of foods is used. Note that K1.30=US$1.00 in 1996.

[17] The two approaches do not give the same answer because poor households tend to be larger than average, so household-level calculations understate the poverty rate amongst persons.

[18] The transient component of headcount poverty, using income as the welfare indicator, averaged 60 percent in the ICRISAT panel of households (Ravallion, 1988). In rural China, the transient component of the headcount index (for consumption-poverty) ranged from 30 percent to 46 percent across survey years (Jalan and Ravallion, 1998).

[19] Gibson, Boe-Gibson and Scrimgeour (1998) report that four-fifths of households in a poor urban area had access to savings and loan accounts, and one half of these households had access to life or medical insurance. Few rural households have access to such formal consumption smoothing mechanisms.

[20] The sensitivity analyses for the poverty gap and poverty severity index are available from the author.