****

BASIC INFORMATION

****

PAPUA NEW GUINEA HOUSEHOLD SURVEY 1996

A GUIDE TO THE CONSTRUCTED VARIABLES

(AVAILABLE IN THE FILE *CONSTRUCTED DATA FILES.ZIP*)

JOHN GIBSON
DEPARTMENT OF ECONOMICS
UNIVERSITY OF WAIKATO

FEBRUARY 2000

Introduction

The Papua New Guinea Household Survey of 1996 (PNGHS96) is a comprehensive and multipurpose survey. In addition to estimating the value of household consumption, the survey also collected data on demographics, education levels and costs, income sources, the use of health facilities, birth details for young children, diet, housing conditions, agricultural assets, agricultural input spending, and participants' perception of their quality of life. Young children and their parents were weighed and measured. This document describes the constructed variables that are available in the file *Constructed Data Files.zip*. The full (raw) data set and data dictionaries are also available, in the file *PNGH96.zip* for researchers who want variables that are not included in the constructed data files or for those who wish to replicate or vary the methods used to construct the variables described here.

There are four files that will be extracted when *Constructed Data Files.zip* is unzipped and these are named:

1. *Annual household consumption value (87 items and total).xls*
2. *Demographic data and short-term to annual extrapolation.xls*
3. *Census unit weights.xls*
4. *Survey design variables for N1200 cross-sectional sample.xls*

These files are all in Excel 4.0 format and so should be easily read by (or pasted into) statistical software such as *Stata* as well as being in a convenient format for users just wanting basic descriptive statistics.

A note on identification

In PNGHS96 households are uniquely identified by the combination of the codes for:

Province
Census Division
Census Unit
Dwelling
Month.

It is particularly important that all five of these variables be used when identifying a household because the survey (unusually) included a longitudinal component where households in 20 of the 120 surveyed Census Units ("clusters") were revisited several months after the initial interviews. For these households, the codes for the first four identifying variables are the same for both sets of observations, and it is only the month of the interview that distinguishes between the initial interviews and the revisit. When calculating summary statistics or carrying out regressions users may want each household to occur in the data only once. One cross-sectional sample which has been obtained with only one observation on each household in the longitudinal sub-sample is described in the file *"Survey design variables for N1200 cross-sectional sample.xls"*. Appendix 1 provides more details for choosing which observation to use on a household with observations available in two different months.

The Constructed Consumption Variables (Brief Version)

The survey used a closed interval recall method, with households interviewed twice so that the start of the recall period was signalled by the first interview. These two interviews were usually two weeks apart, which is the length of the pay period in Papua New Guinea.[1] 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,[2] net gifts received, and stock changes, so they should be a good measure of consumption during the recall period. These short-period consumption estimates were then extrapolated to annual totals and combined with the annual recall estimates for 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. A detailed description of the procedures used to form the consumption estimates is in Appendix 2.

The file *Annual household consumption value (87 items and total).xls* has estimates of the annual value of household consumption of 87 items, in PNG Kina (nb: K1.30=US$1.00 in 1996), for a sample of 1396 households. Some of these households lack full information on their consumption.[3] Moreover, household sampling weights normally need to be applied before descriptive statistics are calculated or multivariate analysis carried out with these data.

Some users of the consumption estimates may require deflators, to correct for both differences in price levels between regions and differences in prices through time (because the survey was spread over a 12 month period). One spatial price index that is available for researchers is that which is based on the regional poverty lines. These poverty lines were formed for five regions of Papua New Guinea: the National Capital District (NCD), Papuan/South Coast (which includes Oro Province), the Highlands, Momase/North Coast, and the New Guinea Islands. These poverty lines (and hence the price index) combine rural and urban areas within each region because usually there was only one urban Census Unit per region (there were no rural CUs in the sample for the NCD). Appendix 3 describes the price data used in the construction of the poverty lines. One advantage of basing a price index on poverty lines is that these lines can still be constructed when non-food prices are missing (see Appendix 4 which describes the formation of the poverty lines) but a disadvantage is that the price index is based on the consumption of the "poor" rather than of an "average" household.[4] Therefore, some alternative spatial price indexes are reported in Appendix 5. Table 1 reports the spatial price index derived

---

[1] 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. See Appendix 2 for details.

[2] 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).

[3] Typically because the second, consumption recall, interview was not completed. The file *Demographic data and short-term to annual extrapolation.xls* includes a variable "2nd visit" which indicates this.

[4] Just under 38 percent of the PNG population are poor, at the upper poverty line, so there is likely to be some difference between the average budget shares of this group and the average budget shares of the total population.

from the regional poverty lines, with the base set as either the NCD or the national average price level.

Table 1: Spatial Price Index Derived From Poverty Lines

	All Consumption		Food Consumption	
National Capital District	100.0	195.3	100.0	180.9
Papuan/South Coast	63.6	124.3	71.9	130.0
Highlands	50.1	97.9	52.9	95.7
Momase/North Coast	36.0	70.4	40.0	72.4
New Guinea Islands	54.4	106.3	60.0	108.5
Papua New Guinea[a]	**51.2**	**100.0**	**55.3**	**100.0**

[a]Population weighted average.

Ideally temporal price variation should be controlled for at the same time that spatial price variation is controlled for. However, the composition of the sample changed from month to month, tending in some months to be more remote (and therefore likely to have higher prices due to transport costs). Thus, the comparison of average monthly prices may pick up more than just temporal variation. The only temporal comparisons of prices that avoid this problem are for the 20 clusters revisited as part of the longitudinal sub-sample. However, these revisits did not start until after June, so a March quarter to June quarter price comparison is unavailable. Instead of using the prices collected by the survey, variation in nominal price levels over the 12 month period of the field work was controlled for by using the Consumer Price Index as a deflator (Table 2). The deflator was applied uniformly to each region. It should be stressed that the CPI is actually calculated just for urban areas so it may not be an ideal guide to intra-year changes in the cost of living for the whole of PNG in 1996.

Table 2: Temporal Price Deflator

	March	June	Sept	Dec	Annual
Index	344.4	345.9	351.4	350.8	348.1
Deflator	98.9	99.4	100.9	100.8	100.0

Sampling Weights and Survey Design Effects

PNGHS96 was not a self-weighting survey. Instead, a set of household weights were derived from (i) the variation between the 1990 Census estimates of the size of each Census Unit (which along with the sampling interval and the systematic ordering of Cus determined the selection of that particular CU) and the actual size found during the survey, (ii) the deviation of the actual number of households surveyed in each cluster from the target number and (iii) the unequal sampling rates between strata (the NCD was over-sampled to ensure a sufficiently large sample for comparisons with an earlier urban household survey). In regard to reason (ii) for the use of weights, it was a feature of the survey that non-respondent households were replaced only at the first interview stage, which differs from the practice in some surveys. If a household did not respond to the consumption recall interview (e.g., they had left the village for an extended period) a replacement household was not used. The rate of progress of survey teams would have been impeded if they had to stay behind in a village to re-do the survey (with a 14 day bounded recall period) for one replacement household, while the other alternative of using an unbounded recall for the replacement household would defeat the efforts made to accurately measure consumption (particularly the recording of own-food production). Non-replacement means that there are more first visit interviews than second visit interviews, and so a different set of weights could be used for variables that are collected on the first visit. The file *Census unit weights.xls* contains these different estimates of the weights, while Appendix 6 has a fuller and more technical discussion of the weighting variables.

In addition to weights, analysts interested in robust statistical inference will also need to take account of the stratification and clustering effects. The sample was divided into two main strata (i) the NCD and (ii) the rest of the country, with the non-NCD strata then divided into 18 sub-strata based on elevation (3 classes), rainfall (2 classes) and an indicator of economic accessibility (3 classes). The cluster effects occur because sampling took place at two stages: initially Census Units were selected, with probability proportional to estimated size, and in the second stage 12 households (6 in the NCD) were selected from within each sampled CU. The file *Survey design variables for N1200 cross-sectional sample.xls* reports the stratum number, the cluster number, and the household weight for each household in the cross-sectional sample which is drawn by using only one of the two possible observations on the households in the longitudinal sub-sample. Note that if population weights are required, they can be calculated by multiplying the household weight by the number of people in the household.

Appendix 1
The Longitudinal Sub-sample

1. A subset of one-sixth of the selected CUs were chosen as a longitudinal sub-sample. Households were interviewed in these CUs at two different times of the year, roughly eight months apart, so that the correlation between expenditures in different months of the year for the same household could be measured. This correlation is needed for unbiased estimation of the between-households variance of annual expenditures when the data come from short period recall expenditures (Scott, 1992). The usual method of calculating the variance of annual expenditures, which is based on extrapolating annual expenditure from short period recall data, overstates the true variance. This can bias estimates of the incidence of poverty because the overstated variance increases the proportion of households whose expenditures fall below a certain absolute cut-off, e.g., a poverty line.

2. The Census Units included in this longitudinal subsample are (Province No., Census Division No., and Census Unit No.):

1	27	501	11	5	507
3	19	13	11	10	10
4	84	4	12	50	26
4	84	24	12	51	15
4	84	40	13	33	27
4	84	59	13	81	8
7	32	7	14	31	15
7	82	3	14	83	621
9	8	22	17	6	5
11	4	43	17	83	1

3. The revisit interviews for the longitudinal sub-sample also provide an opportunity to make the sample more representative of an entire year. The survey had initially been planned to take place only in the early months of 1996 because an October deadline had been set for availability of the results, to coincide with preparations for the annual PNG Budget. This deadline was subsequently relaxed after fieldwork had started but by then a greater number of interviews had been carried out in January and February than would have taken place if interviewing was spread evenly over the 12 months of the year (Figure 1a). However, by using revisit interviews, which took place from August onwards, to replace first visit interviews, the over-representation of January and February in the sample is reduced, and interviews are spread more evenly over the year (Figure 1b). This replacement is possible because the CUs to be revisited were chosen from a randomly selected group and the complete survey was repeated when the household was revisited. In fact, the interviews are possibly more accurate during the revisits than the initial visits because of the greater experience of the interview teams.

Figure 1a: Distribution of Interviews - Excluding Revisits

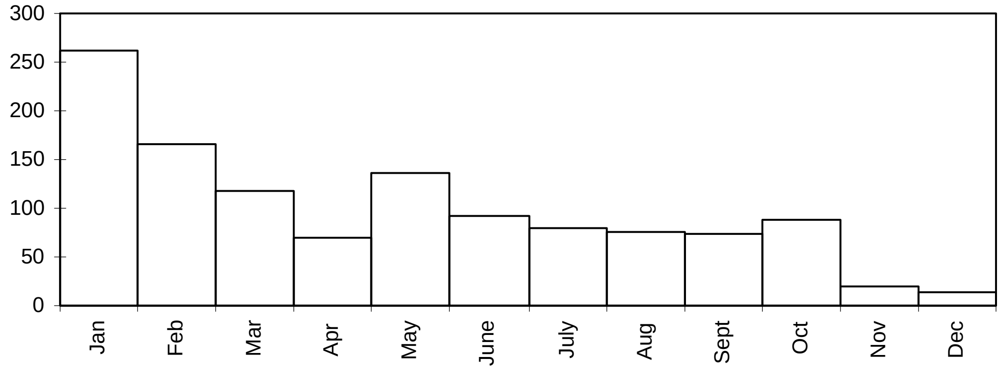

Figure 1b: Distribution of Interviews - Revisits Replacing First Visits

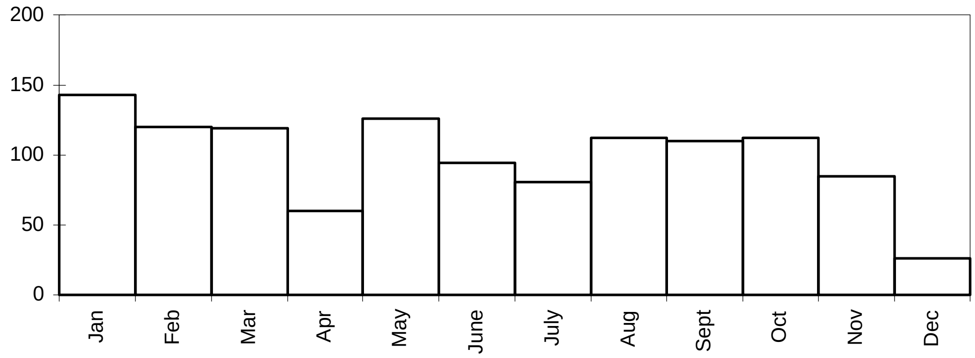

Appendix 2
Measuring the Value of Household Consumption

1. Five components of household consumption were measured by the survey: food consumption; other frequent expenses; annual expenses; consumption of non-housing durables; and consumption of housing services. The main components of consumption that were not measured are consumption of leisure, consumption of environmental services, and consumption of freely provided (or subsidised) public services. These items were excluded because of the conceptual and practical difficulties in measuring them.

2. The reference period for food consumption and other frequent purchases was set by the time between the first and second interviews. This bounded period between the two interviews was designed to be two weeks, but often differed from this, so the dates of each interview were recorded. The reference period for annual expenses was the past 12 months. The estimates of total consumption are made on an annual basis by extrapolating food consumption and other frequent expenses from the bounded recall period up to a full year. The inter-household variance of these estimates of annual consumption will be greater than if consumption was continuously measured over a full 12 month period because some of the shocks that occur in the short reference period would be evened out over the course of a year. The extent of this upward bias in the measures of dispersion has been demonstrated by Gibson (1999),[5] using results from the longitudinal sub-sample.

3. One problem in extrapolating short-term consumption from the bounded recall reference period, and then adding it to annual expenses, is that the size of the household may differ between the short and the long time periods. For example, a household may usually have five residents but during the bounded recall period two members leave to visit friends in another household for six days. This will cause the consumption of food and other short-term expenses to fall and when this is extrapolated to the full year it will lower the estimated annual value of the household's consumption. If total annual consumption was divided by the number of usual residents (five), the result would understate the true level of per capita consumption, because for the rest of the year the two people who were absent during part of the recall period are living in the household and contributing to its consumption. Therefore some adjustment needs to be made for the effect of absent residents, and for the opposite effect of short-term guests (and these corrections need to offset so that there is no double counting).

4. The specific adjustment used was as follows: the potential number of person-days during the recall period was estimated as the number of calendar days between the two interviews multiplied by the number of persons listed in the household roster. The actual number of person-days was measured from a schedule of absences and guests (present for at least seven days). An allowance was made for the smaller impact of absent or additional young children by estimating person-days in adult-equivalent terms, where each child aged six and below had a weight of 0.5. The ratio of the actual number of person-days to the potential number was

[5] http://www.mngt.waikato.ac.nz/depts/econ/staff/johng/PDF_Files/WorkingPapers/ExtrapolatedPoverty.pdf

formed and used to estimate the annualised equivalent of the bounded recall period consumption as:

$$\left(Food + Other\,Frequent\;Expenses\right) \times \left(\frac{365}{\left(Bounded\;Period\right)} \times \frac{Potential\;days}{Actual\,days} \right).$$

The extrapolation factors used to convert short-reference period expenditures to annual totals, are contained in the file *Demographic data and short-term to annual extrapolation.xls*.

5. If someone was listed as a usual resident on the roster but was substantially absent during the bounded recall period (absent at least 10 days), they were not counted when estimating household size, and therefore the above correction factor for absences was not applied to get the annual estimate of the recall period consumption. An exception to this rule was children who were away at school during the bounded recall period but whose school fees and other costs are still being met by the household. These children were considered to be part of the household and they therefore counted when estimating household size.

Food Consumption

6. Food consumption was measured by a residual approach using the equation:

	Purchases
+	Own-production
+	Gifts received
-	Sales
-	Gifts given
-	Stock increases
=	Consumption.

This approach was chosen for two reasons. First, it was the method that had been used previously in PNG in the Urban Household Survey and the Household Expenditure Survey. Second, there is substantial interest in some of the components of consumption in their own right, especially the quantity and value of domestic food production and sales.

7. The inclusion of sales of items that were either self-produced or purchased is an important part of the formula for calculating consumption. Previous evidence from urban areas shows that ignoring this can substantially overstate the consumption of items like beer that are bought to be then resold on the black market, and also overstates the consumption of foods where a substantial amount of production is sold, e.g., betel nut (Gibson, 1995). The importance of small-scale selling activity might be overstated if the purchases and sales of more formal business entities, like trade stores, were recorded as household transactions. Therefore, food purchases that were expenses of a formal business were ignored, as were the corresponding sales by these same household-business units. However, the expenses of an informal business, like flour purchases by a scone seller, were counted as food purchases, and sales (valued at input costs) were then subtracted to give the estimate of consumption.

8. The residual approach to measuring consumption does suffer some drawbacks. If forms of food disappearance that are not included in the formula -- such as plate waste and feeding household food to domestic animals -- are important, apparent food consumption levels will be overstated. The residual approach is also less intuitive than directly asking households what they eat, and there is international evidence that estimates of calorie *intake* based on what households eat are less extreme (higher for poor households and lower for rich households) than are measures of calorie *availability* based on an expenditure approach (Strauss and Thomas, 1995). Because the residual approach depends on six separate components there is also more scope for error. For example, having higher stocks of a food at the end of the recall period than at the start, but no record of purchases, production or gifts received to account for the higher stocks, could be due to error in any one of four possible parts of the interview.

9. The most difficult part of the food consumption equation to obtain reliable measures of was food production. Locally produced foods, which include root crops, bananas, vegetables and sago, are not sold by weight in Papua New Guinea. Instead, they are sold in fixed price bunches, bundles, heaps, or piles. Although some rural people are familiar with kilogram weights from sales of coffee or purchases of flour, rice or stock food, the densities of these items are much different to the densities of root crops. A further difficulty is the high volume of production, due to the bulky nature of the local staples (excluding sago); an average sized rural household can easily produce over 100 kg of root crops per week. In these circumstances it was considered unwise to ask respondents: "how many kilograms of food *x* did your household produce since my last visit?" Such answers reported in kilogram units would give the appearance of accuracy but this would just disguise the problem of measuring food production and force it onto the interviewers and the respondents.

10. Instead of relying exclusively on kilogram units, the question on food production allowed flexibility in the choice of units. Five units were available. The first was an empty 25 kg rice bag, which had three graduations ("¼", "½", "¾") marked on the outside. This bag was given to the household at the start of the survey and they were asked to put their produce into it during the recall period. This was the recommended unit for bulky crops. The other units were bunches and heaps; kilograms; singles, which were recommended for items like coconut and betelnut, and livestock (e.g., one pig, three chickens); and "other". The 25 kg rice bag was the most frequently used measuring unit for the staple root crops and sago. The other important staple -- bananas – was measured with equal frequency in terms of either rice bags or bunches. Other foods where bunches or heaps were frequent units of measure are aibika (a leafy vegetable), a residual category for 'all other vegetables and nuts', and betelnut. Coconuts, fruit, sugar cane, pork, fish and other meat were all commonly measured by counting the number of single units produced.

11. Three of the measuring units -- rice bags, singles, and kilograms – were considered "reliable" because in principle factors to convert them into kilogram weights could be estimated and applied. The remaining two units – bunches/heaps and other – were considered uninformative as to the actual weight of food produced by the household. The problems that would have been encountered in assuming an average weight for a bunch or heap are best illustrated by the example of bananas: a bunch could refer to a 'hand' of approximately

10 single bananas weighing between one and two kilograms or else it could refer to a ‘rope’ or a ‘branch’ that is made up of several hands of bananas and can sometimes weigh 20 kilograms.

12. The first step in transforming bunches/heaps and other unspecified measures into kilogram quantities was to convert the “reliable” measures -- rice bags and singles -- into kilograms. The conversion factors used are reported in Table 1. The unit value was then calculated for each household that produced the food and reported production using one of the “reliable” measures. The median of these production unit values was then estimated for each food in every Census Unit in the non-NCD sample, with the exception of meat and fish, where the median was calculated over all households in the region. The final step in imputing kilogram quantities was to take the monetary value that households had assigned to the particular food that was measured with an uninformative quantity unit, and divide that value by the median unit value. If no median unit value had been calculated for that Census Unit – either because no other households produced the food or because the quantity measurement for all households who produced the food was an uninformative one – the imputation was based on the median unit value from the full sample of households in all non-NCD Census Units. In the NCD sample, the monetary values assigned to production of foods with uninformative quantity units were divided by market prices, in order to impute the quantity of production in kilograms.

13. The distinction made above between “reliable” and uninformative quantity units is blurred somewhat by the imprecision of the conversion factors estimated for rice bags and singles. The three foods having the largest number of measurements of both weight *and* rice bags were sweet potato, taro and bananas. The coefficient of variation (standard deviation relative to the mean) of the weight of a rice bag full of each of these foods was estimated as 0.3, 0.2, and 0.2. A further imprecision is that the relationship between how full the rice bag was and the weight is not linear. Instead a full bag would contain roughly one kilogram more than two half full bags, but the survey form just records the total number of bags produced and not whether that total is made up of very many fractions of bags or just a few full bags. The situation is even worse for the variability in the weight of single items: the average weight of a single sweet potato assumed by the unit value (and calorie) calculations was 0.45 kg, but weights found during the survey ranged from 2.5 kg to 0.1 kg. The estimated coefficients of variation for single sweet potato and cassava were 0.9, for taro 1.0, for yam 1.4, and for banana (single ‘fingers’) 0.7. Even items where individual units have a relatively regular size, like coconuts and betelnuts, had estimated coefficients of variation of about 0.4. The problems are bigger when assigning an average weight to an individual animal such as a pig, because whole pigs were not able to be weighed during the course of the survey. For items that are aggregations or residual categories (e.g., “other fresh fruit”, “other greens, vegetables, and nuts”) the average weight used for both rice bags and singles will be a very noisy measure of the actual weight of production of any particular household because we don’t know if, for example, the item produced is a pineapple, whose average weight is about 2 kg, or a mango, whose average weight is just 0.3 kg.

14. The imprecision of the conversion factors in Table 1 means that for any given household, the estimated quantity of food production will contain a substantial element of random noise. If this noise increases the measured between-household dispersion, the number of households

below a poverty line might be overstated. But sample summary statistics such as means and total should not be affected if the conversion factors are right, on average. This also applies to any measures derived from food quantities, such as per capita calorie availability.

Table 1: Weight Conversion Factors Used for Measuring Food Production

	---------- Kilograms per ----------		
	25 kg rice bag	Single	Bunch/heap/other
Sweet potato	19	0.45	2.0
Cassava	16	0.7	1.5
Taro and Chinese Taro	16	0.6	1.8
Yams	16	0.7	1.4
English potato	20	0.1	0.6
Bananas	12	0.16	1.6
Sago	30	n.a.	1.0
Sugar cane	20	2.5	5.0
Other fresh fruit	12	0.8	1.0
Coconuts	18	1.3	2.6
Peanuts	8	0.01	0.2
Aibika	4	0.06	0.2
Other greens, vegetables, nuts	10	0.4	1.0
Rice	20	n.a.	1.0
Lamb and mutton	15	20	n.a.
Chicken	15	1.5	n.a.
Pork	15	30	n.a.
Other meat (incl. bush meat)	15	3.0	n.a.
Fish (fresh, dried, shellfish)	15	0.5	1.0
Eggs	15	0.06	0.5
Betel nut	19	0.03	0.15

15. Although considered an uninformative unit of measure, there was one circumstance when weights had to be assumed for bunches/heaps and "others". This occurred when a household either gave away or sold food that it had produced, and the gift or sale was measured in units of bunches or heaps, while the original production had been measured in another unit. The final column of Table 1 contains the conversion factors that were assumed for this purpose. This should not cause large errors because the most common way of reporting gifts and sales was in the same units as production was reported, which made calculating the fraction given or sold easy, regardless of whether production was recorded using uninformative quantity units.

16. The computer programs that calculated food consumption used the median production unit value in three other ways, additional to converting uninformative quantity units into kilograms. These uses were:

- Imputing the quantity of gifts received and the quantity of purchases, from the value of gifts received and value of purchases, if the quantities had not been recorded. Although in principle, a Census Unit median could be calculated for purchase unit values and gift unit values, for the locally produced foods there was much more information on production than on gifts received or purchases.

- Assigning a value to beginning food stocks, if the household reported no production, purchases or gifts received of the food concerned. Usually there would be either production or purchases or gifts received, so the ratio of sub-total values and quantities could be used to implicitly give food stock changes the same household-specific unit value as was applied to production, purchases and gifts.
- Creating an alternative measure of the value of each household's production of each food, by multiplying the household's production quantity by the Census Unit median of the production unit value for that food. This alternative measure can only be created if "reliable" quantity units were used.

For all three of these uses, the regional average market price was used for foods that are not produced in PNG (including rice and sheepmeat, whose local production is almost zero). For the food consumption of households in the NCD sample, the average market price was used for all foods instead of the median production unit value.

17. Using the median unit value to adjust the value of production is appropriate if (i) it is assumed that the law of one price holds within CUs, and (ii) it is believed that quantities measured with "reliable" units are more accurate than reports of how much it would cost to buy or sell the quantity of a certain food produced by the household. The reason for doubting the reports of values is that very few households who are producing a food also buy that same food, so there may be a lot of misinformation about prices. Sales activity is also low so many households would not be able to infer purchase values from the prices they receive when selling the food. In these circumstances there may be a lot of random noise in the values attributed to self-produced foods.

18. The unit values for production of sweet potato were examined to see how important this random noise in the values attributed to self-produced foods may be. Attention was restricted just to households who used the 25 kg rice bag as the unit of measure and just to CUs where unit values for the rice bag were available for at least four households (n=694). The *within* Census Unit coefficient of variation in the value that each household placed on a rice bag full of self-produced sweet potato was 0.45, which is twice as high as the variation in the measured weight of rice bags of sweet potato. These households were ranked within their CUs, according to the value they placed on a rice bag of sweet potato. The ratio of the highest unit value to the lowest unit value was formed, and the median value of this ratio was estimated to be four. Hence, there appears to be large variation in the value that households within the same Census Unit placed on a given volume of production. Some of this variation will be caused by the differing weight of sweet potato that each household's rice bag contained, but this source cannot explain it all. Some more of the variation may be due to differences in the quality of sweet potato that each household was producing. However, at least some of the remaining variation would seem to be due to differences in opinions about prices.

19. The variation in sweet potato unit values found within CUs is likely to occur with other self-produced foods as well. It may seem unreasonable that two households, who produce the same quantity of a food in the same location, can have that production valued differently. A

household might fall below the poverty line just by being too pessimistic when valuing their own food production because they think prices are lower than they truly are. Using the adjusted values of food production, based on the CU median unit value, would seem to be an improvement. However, two factors caution against too much reliance on this adjustment: (i) it cannot be applied to food production measured by uninformative units, and (ii) it places a lot of faith in quantity measurements, which, especially for animal products, are based on little more than values taken from the PNG literature and informed guess work regarding conversion factors.

20. The effect of adjusting the reported value of food production by basing it on the Census Unit median unit value is to cause a slight reduction in measured inequality in the value of food consumption. The Gini index for the value of food consumption is 0.44 when the value placed on self-produced food is not adjusted. The Gini index falls to 0.42 if self-produced foods are valued at the CU median unit values. Therefore, variation in unit values within CUs, which seems to be due at least partly to differing opinions about prices, is not a big contributor to measured inequality.

Other Frequent Expenses

22. Measuring consumption of non-food frequent expenses (e.g., gambling, P.M.V. fares, kerosene, tobacco) was easier than measuring food consumption because, with the exception of leaf tobacco and firewood, consumption was entirely from purchases and gifts, rather than from own-production. The other simplification was that quantities were not required.

23. The only problem encountered was assigning values to firewood, because in some areas firewood is not bought and sold, so respondents had difficulty assigning a value to the firewood they harvested or received as gifts. A regression equation was used to impute firewood values if they were missing, using three assumptions to guide the specification of the equation. More firewood, and hence a higher value, is needed if:

- there is more food to be cooked,
- there are more people in the household to cook for and keep warm, and
- the outside temperature is colder.

The equation used was:

Firewood Value = 39.1 (Number in household)
(4.0)
+ 0.017 (Value of Food Purchased & Produced)
(2.8)
+ 61.5 (Number in household * Highlands dummy)
(4.6)
+ 0.027 (Value of Food * Highlands dummy)
(4.0)

N=821 R^2=0.30 (*t*-statistics).

Annual Expenses

24. Annual expenses were estimated from an unbounded 12 month recall, with consumption derived from three components:

Value of Purchases - Value of Gifts Given + Value of Gifts Received.

Spending on some of the items on the annual expenses schedule was also recorded in some sections of the first interview questionnaire (e.g., school fees, charges for water) and these data were used to check the consistency of results. Any purchases of an item on the annual expenses schedule during the bounded recall period were also recorded, and these short-term data were used to assess the consistency of the reported values given in the long term recall.

25. One of the categories of annual expense items that had data collected – loan payments and finance fees -- is not included in measured consumption. Many of the data reported for this category appeared to be for repayments of principle as well as payments of interest. The appropriate component for measuring consumption is just the interest payment because this is the cost to the household of bringing forward in time the acquisition of assets. If principal repayments were included, measured consumption would have been overstated.

26. Two other notable items of the annual expenses are spending on bride price ceremonies, and spending on death feasts. Consumption of these two items was defined in the conventional way, so for example, contributing to the bride price expenses of someone who lives outside the household subtracts from the household's measured consumption. It could be argued that in Papua New Guinea contributing to the bride price and death feast expenses of kinsmen is a requirement of participating in the everyday life of the community. Hence, these expenses on people living outside the household might be a legitimate component of consumption and of a consumption-based poverty line (World Bank, 1990). However, that line of argument was not followed when constructing the annual expenses variable. Thus spending on bride price ceremonies and death feasts for people outside the household was treated in the same way as other items purchased and then given away as gifts, by subtracting it from the households measured consumption.

Consumption of Durables

27. An inventory of 16 durable assets was used to collect data on the purchase price (or value if it was a gift) and date of acquisition of each asset, and the price that it would realise if sold. Purchase prices were inflated to 1996 terms using the movement in the Consumer Price Index. The straight-line annual depreciation was calculated as:

(Purchase Value - Estimated Sales Value) / (1996 - Year of Acquisition).

If components of the formula were missing, a default value of annual depreciation was assigned based on the average depreciation for that asset calculated from the sample of households with complete data. Table 2 contains these default depreciation values along with estimates of the ownership rates for each durable in each region and the country as a whole.

28. The questions in the inventory of durable assets used *ownership* rather than *use* as the criteria for including or not including a particular asset. This may result in underestimates of consumption for people living in households where chattels are provided (e.g., chairs, refrigerators) but it is hoped that the rent paid on the house would include a component due to the provision of these assets.

Table 2: Default Depreciation Values for Durables

Chairs and tables	*20*	Metal/fibreglass dinghies	*280*
Primus or portable stoves	*12*	Outboard motors	*310*
Kerosene lamps	*3*	Bicycles	*20*
Refrigerators or freezers	*120*	Motorcycles	*650*
Sewing machines	*23*	Cars or pickup trucks	*1200*
Generators	*240*	Cameras	*10*
Guns	*60*	Radios/cassette players/stereo	*25*
Traditional canoes	*18*	Television/video equipment	*95*

Note: Default depreciation is the value of annual depreciation imputed for an item if the actual depreciation was not able to be calculated due to missing data on acquisition dates or prices.

The Consumption of Housing Services

29. Houses were divided into two classes: "mainly traditional" and "mainly non-traditional" based on the building materials used in the floor, roof, and walls. All houses in the National Capital District sample were allocated to the "mainly non-traditional" class, and a further 115 houses from the sample outside the NCD were also allocated to this class. The remaining 839 "mainly traditional" houses were further divided into those with a traditional roof (*n*=731) and those with an iron or metal roof (*n*=108). The "mainly non-traditional" houses were divided into a high/medium-cost group (*n*=248) and a low-cost/village/urban settlement group (*n*=116).

30. For owner-occupied dwellings, the economic life of the house is needed so that capital costs can be depreciated over a certain number of years in order to estimate the annual flow of housing services. The life expectancy of a mainly traditional house with a traditional roof was estimated as 3.5 years, based on the age structure of the houses in the sample. The life expectancy of a traditional house with a metal roof was estimated as 4.8 years. An upward adjustment in the estimate of the economic life was made to reflect the fact that some housing materials are salvaged and reused in the next dwelling (e.g., sound roofing iron can be transferred from a dilapidated house to its replacement). The economic life of a traditional house with a traditional roof was set at four years, and a traditional house with a metal roof was set at six years. The economic life of a high-cost non-traditional house was estimated as 20 years and a low-cost non-traditional house was estimated to have a life of 10 years. For all house types, if a dwelling was observed to have a greater age than the expected economic life, the actual age was used when calculating annual depreciation.

31. The capital cost of traditional houses was established from two questions. The first asked how much it cost to build or buy the house. These costs were inflated to 1996 terms, using the

movement in the Consumer Price Index since the year of construction of the house. The second question asked about the number of days of unpaid labour that were used to build the house. Ideally, person-days spent on house building would be converted into kina terms using an opportunity cost of labour that was specific to each household. However the survey did not include questions on incomes and earnings, so external data on the opportunity cost of labour were used. Specifically, estimates of mean annual cash income per household from agricultural activities, for the agricultural system that each census unit was located in, were used (see Allen, Bourke, and Hide, 1995). These estimates ranged from K25 per household to K1200 per household, and were divided by 200 to get estimates of the opportunity cost of labour per person per day. For CUs that were not assigned to agricultural systems (urban areas, rural non-village and plantations), the opportunity cost of labour was set at K6 per day. Obviously, these estimates of the opportunity cost of labour are only rough approximations. A further source of underestimation of capital costs is that unpaid for materials used in construction (e.g., bamboo, sago thatch) did not have a value imputed for them (but see the sensitivity analyses).

32. The capital cost of some traditional houses had to be imputed because the data on unpaid labour used constructing the house were not reliable. For example, 20 people may have spent three days building the house, and the data reported are that the house was built with three days of unpaid labour. The answer that was really wanted was 60 person-days of unpaid labour but the survey question was not adequately worded to get this answer in all cases. Unreliable data were identified from scatter plots of the capital cost of the house (which included unpaid labour days valued at opportunity cost) against floor area, plotted separately for each census unit, and also taking account of construction materials. Data for 170 houses were deemed unreliable using this method. A cost function was estimated using the data from the remaining 669 traditional houses:

$$\begin{aligned}\ln(\text{Capital Cost}) = {} & \underset{(6.08)}{1.29} + \underset{(8.66)}{0.024}\,(\text{Floor Area}) - \underset{(6.13)}{0.000069}\,(\text{Floor Area})^2 \\ & + \underset{(2.41)}{0.062}\,(\text{Number of Rooms}) + \underset{(8.15)}{0.877}\,(\text{Iron Roof}) \\ & - \underset{(2.17)}{0.746}\,(\text{Temporary Walls}) - \underset{(0.14)}{0.012}\,(\text{Dirt Floor}) \\ & + \text{cluster dummy variables} \qquad \bar{R}^2 = 0.79\end{aligned}$$

where ln is the natural logarithm and *t*-statistics are in ( ). This equation was used to impute the capital cost of the 170 houses where the original data were unreliable.

33. Estimating the capital cost of non-traditional houses was more difficult. Fewer of the occupiers were owners, especially in urban areas, so they did not know details about the age of the dwelling or construction costs. In principle, this should not matter because if these occupiers are paying rent, their rental payments serve as an estimate of the consumption of housing services. However, in Papua New Guinea the rental market is distorted by shortages of housing (due partly to collective clan ownership of land) and employer provision of housing. Survey results reported by Dahanayake (1991) suggest that two-thirds of middle-class urban dwellers live in employer-provided housing and, on average, have 94 percent of their rental costs subsidised. These subsidies can cause enormous variations in measured consumption

levels. For example, in one high-cost Census Unit in the NCD sample, the occupiers of one house were paying K1700 per month rent while the occupiers of another house of similar age, size and quality were paying only K40 per month rent. Therefore it is important to have estimates of the capital cost of the house so that the estimated annual depreciation could be compared with the annual rent paid. If the rent paid was much less than the estimated depreciation, it was assumed that a rental subsidy was being provided and the estimated value of annual depreciation was used instead.

34. The capital cost of non-traditional houses was established by two different methods. For low-cost/settlement houses, the rules used for traditional houses (described above) were applied because occupiers had usually built the house themselves or purchased it on an open market, and this made them more informed about costs. But there were 15 houses where cost data were unavailable so imputation was based on the following cost function:

$$\begin{aligned}
\ln(\text{Capital Cost}) = {} & \underset{(17.8)}{6.39} + \underset{(2.54)}{0.030}\,(\text{Floor Area}) - \underset{(2.04)}{0.00016}\,(\text{Floor Area})^2 \\
& + \underset{(1.23)}{0.100}\,(\text{Number of Rooms}) + \underset{(3.20)}{0.677}\,(\text{Cement walls}) \\
& + \text{cluster dummy variables} \qquad \bar{R}^2 = 0.33
\end{aligned}$$

where *t*-statistics are in ( ).

35. For medium/high-cost houses where the occupiers were unable to give estimates of costs, an imputed cost was based on the construction costs of a house of similar size, built in a similar era, by the National Housing Commission. The Housing Commission has been the main provider of urban housing and it has used a set of basic designs which are usually identifiable by the floor area of the house. Housing Commission costs were used rather than estimating a cost function from the group of medium/high-cost houses in the sample with complete cost data because there was a much higher ratio of houses with missing data to those with complete data than in the traditional and low-cost non-traditional sub-samples.

36. Specifically, for the medium/high-cost houses, data on the construction costs of standard Housing Commission houses in several different years were obtained and used to estimate a regression of real construction costs, in 1996 terms, on floor area (an intercept and quadratic term were found to be statistically insignificant):

$$(\text{Construction Cost})_t \times \left(\frac{CPI_{1996}}{CPI_t}\right) = \underset{(27.4)}{816.2}\,(\text{Floor Area}) \qquad R^2 = 0.86$$

where the *t*-statistic is in ( ). This equation was used to impute the construction cost of any medium/high-cost houses that lacked cost data but had estimates of floor area. The cost that is being imputed is for a basic design, so can be considered as a lower bound estimate of the actual construction cost.

37. The annual depreciation for all houses was calculated using a straight-line method, with the capital cost in 1996 terms being divided by which ever was the greater of either the expected

economic life for the type of house or the actual life of the house. If the calculated annual depreciation exceeded, by more than 150 percent, any annual rent payments that were reported, the estimated depreciation was used as a proxy for the true economic rental of the house. The annual depreciation was also used as the estimate of the consumption of housing services for all owner-occupied dwellings.

Appendix 3
The Price Data Used for the Poverty Lines and Spatial Price Index

1. During the course of the household survey, price data were collected from each of the 120 CUs that were in the sample. The price survey was also repeated when CUs that were part of the longitudinal sub-sample were revisited. Prices were collected for food and non-food items. Two types of price survey were carried out by the interview teams. The prices of packaged food items (e.g., rice, sugar, tin meat, beverages) and non-food items (e.g., soap, kerosene, cigarettes) were collected from the two main trade stores or supermarkets used by the households in the Census Unit. The prices of locally produced foods were collected from the nearest local market, with the price and weight of up to six different lots of each food being recorded. The market price survey was carried out on two different days, to help improve the precision of average prices, so potentially the weights and prices of 12 lots of each food were recorded for each Census Unit. However in some areas markets were held only infrequently so the price survey could be carried out only once. In a few cases, interviewers made special return visits to an area to coincide with market day.

2. The average price of each item was calculated for each Census Unit. For items from the trade store survey, this usually meant taking an average of two prices, but if the item was only available in one store the single price was used. Sometimes the items available were not of the desired specification (e.g., the specification for tinned meat was a 340g can of “Ox and Palm” brand, but sometimes the price collected may have been for a 200g can, or for a 340g can of a different brand). When the price for the non-specification good was the only one available in a particular Census Unit it was used to predict the missing price of the good of desired specification, if price relativities between that specification and the desired specification were able to be estimated from the rest of the sample. Otherwise, the missing price was predicted from a cross-sectional regression of the price of the desired item on the prices of all other trade store goods, for the sample of CUs where the price was not missing. The logic of this regression is that the spatial pattern of prices for trade store goods reflects transport costs.

3. An exception to these procedures was made for the CUs in the NCD sample, where a simple average was taken of the prices for each item collected in all 40 CUs. The reason for this exception was that people who live in one part of the capital city can easily buy in other parts of the city. Therefore the spatial distribution of prices inside Port Moresby is less relevant than is the spatial distribution of prices within the other, larger, regions.

4. If a locally produced food was not able to be priced in a particular Census Unit during the survey period, the price was treated as missing data and was not predicted. There are two reasons for treating the prices of locally grown foods in a different manner to the treatment for packaged food items. First, the reasons for not observing a price are more complex: the Census Unit may lie outside the environmental range of production for the food, or the crop may be available in only certain seasons of the year, or in more remote areas the food could be so widely grown that no market exists because demand is meet by own-production. Second, there is no single factor like transport costs that could be used for predicting the price.

5. The collection of prices spanned 12 months so there will be temporal as well as spatial variation. The temporal variation should not affect the average prices calculated for the four regions outside of the NCD because surveying progressed at approximately the same pace in those four regions. However, the NCD survey did not start until April, three months after the first prices were collected in the other regions. The quarterly Consumer Price Index, calculated by the NSO, was used to correct for the different distribution across time of the samples. Specifically, a weighted average of the CPI was taken for the non-NCD sample, where the weights were the number of CUs surveyed in each quarter. A similar average was taken for the NCD sample and the ratio of the two was estimated as 1.0051. Therefore, all prices collected in the NCD survey were reduced by a factor of 1/1.0051, so as to be directly comparable, timewise, when comparisons were being made with the prices collected in the other regions.

6. Some items were priced at regional level because market prices for these items are not widely available in all Census Units. This included formally marketed items that are sold mainly in urban areas (such as bread, lamb and mutton) and informally marketed items produced in only a few areas but sold more widely throughout urban markets (e.g., potatoes). Appendix Table 1 lists the regional average prices for all food items excepts meals eaten out of the home.

Appendix Table 1: Regional Average Food Prices

	NCD	Papuan	Highlands	Momase	Islands
			toea per kilogram		
Sweet potato	69	60	20	24	38
Cassava	54	43	17	17	25
Taro	101	53	43	34	67
Yam	73	69	63	26	52
Banana	84	44	42	23	42
Sago	120	55	50	39	62
Coconut	40	39	43	19	13
Rice	93	129	129	121	111
Lamb and Mutton	340	320	282	364	333
Pork	400	350	300	310	370
Chicken	460	550	400	520	560
Other meat (incl. bush meat)	380	170	100	60	200
Fish	295	202	402	312	306
Sugar cane	47	42	29	19	24
Other fresh fruit	132	69	61	67	58
Peanut	326	444	246	202	331
Aibika	131	62	44	29	46
Other greens, vegetables, nuts	80	75	40	40	50
Potato	120	225	55	114	155
Betel nut	623	414	654	150	213
Flour	97	169	182	155	124
Tinned meat	647	777	766	758	784
Tinned fish	414	497	408	439	443
Milk	1049	1181	1102	1084	1139
Sugar	158	213	203	189	185
Bread	221	198	217	198	198
Biscuits	381	504	475	458	458
Dripping	500	534	490	493	490
Other diary, cereals, and eggs	392	508	452	430	458
Tea, coffee, milo	556	773	849	734	739
Snack food	1069	1646	1293	1267	1350
Salt, spices, sauces	375	379	423	426	378
Soft drink	232	306	300	278	280
Beer	454	566	786	541	578
Other alcohol	3254	3429	3140	3471	3343

Appendix 4
The Formation of the Regional Poverty Lines

1. The initial poverty line food baskets were formed from the food consumption budgets of households where nominal total consumption per adult equivalent was below K570 per year. This group of 437 households represents the poorest 50 percent of the population. he weighted average consumption quantities were estimated in units of grams per adult equivalent per day for 35 foods in each of five regions (National Capital District, Papuan/South Coast, Highlands, Momase/North Coast, and New Guinea Islands). The average calorie availability per adult equivalent was then calculated from these food quantities using data on the calorie content and edible fraction of each food . An allowance was made for calories obtained from meals consumed away from home, by assuming that this source of calories was twice as expensive as the average 'price' of all other calories that the household obtained. These calorie 'prices' (strictly speaking, 'unit values') were calculated separately for each household. A scaling factor was calculated for each region by dividing the average daily calorie availability by 2200. This scaling factor was applied to all quantities in the food basket (values in the case of meals away from home) to ensure that the scaled basket supplied exactly 2200 calories.

2. An analysis of covariance (dummy variable) model was used to statistically test whether regional differences in the composition of the food baskets were statistically significant. The set of regression equations estimated by weighted least squares were of the form:

$$w_i = \alpha + \beta_1 \text{(Papuan)} + \beta_2 \text{(Highlands)} + \beta_3 \text{(Momase)} + \beta_4 \text{(New Guinea Islands)}$$

where w_i is the share of calories provided by the *i*th food and the right-hand-side variables are of 0-1 form, indicating which region the observation (i.e., household) came from. Shares are used as the dependent variable so that differences in mean levels of food consumption between regions do not affect the results (the same effect is achieved in the formation of the food basket by scaling quantities). Also, the regression is in terms of calories rather than food quantities, so that foods without quantity data (e.g., meals consumed away from home) can be included. The hypothesis that $\beta_1=\beta_2=\beta_3=\beta_4=0$ was then tested (because this is a linear hypothesis the Wald test results are not affected by the choice of the NCD as the base region). This equation was estimated separately for each of the 36 foods in the basket because software for estimating systems of equations with clustered standard errors is unavailable. Equation-by-equation estimation is inefficient, because it ignores the correlations between the residuals in each equation, so the hypothesis of equal shares between regions is less likely to be rejected. Nevertheless, the hypothesis of equal shares was rejected for 27 out of the 36 foods (at p=0.05), and these 27 foods account for 97.4 percent of the average calorie supply for households below the cut-off. A major contributor to inequality in regional average shares was sweet potato, which provides 64 percent of the calories in the highlands, only 1.7 percent of the calories in the NCD, and between 10 and 16 percent of the calories in the other regions.

3. Although there is a dissimilarity of diets across the five regions as a whole, there is considerable similarity between the Papuan, Momase and New Guinea Islands regions. The null hypothesis that the share of calories contributed by a particular food to each regions diet is the same in all three of the regions (i.e., $\beta_1 = \beta_3 = \beta_4$) was rejected for only six out of the 36 foods. This merging of food preferences in these three regions is more defensible than the use of a single

national basket of foods because the dietary patterns in these three lowlands regions are the most closely alike. In contrast, diets in the Highlands are dictated partly by environmental constraints (e.g., altitude and climate preclude any significant coconut and sago production), while diets in the National Capital District are conditioned by the arid climate, the poor links to the rest of the economy and the abundance of imported and commercially produced foods.

4. The poverty line food baskets for the NCD, the Highlands and the combined Papuan/Momase and New Guinea Islands regions were priced using regional averages of local market prices. The only food in the baskets without price information was meals consumed away from home. An allowance for this cost was made by scaling the sub-total price of all of the other foods in the basket up by 1/((2200-*k*)/2200), where *k* is the number of calories in the basket provided by meals consumed away from home.

5. The region with the highest cost of buying 2200 calories per day was the NCD, where the price was K543 per year. Households in the Highlands and Momase regions faced the lowest cost, at K307 and K259 per year. Note that these cost differences do not seem to imply differences in the quality of the diets supplied. Consumers in the NCD would have to pay K696 per year to buy the basket of foods for the Highlands region and K640 per year to buy the Momase basket. Both of these are more expensive than the actual basket bought by NCD households, and therefore the NCD basket is not revealed to be preferred to the other baskets.

6. An allowance for essential non-food items was calculated from the typical value of non-food spending by households whose total expenditure just 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 these households (in region *j*) is found from the following Engel curve (dropping household subscripts):

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

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 *k*th demographic category, and D_j is an intercept dummy for region *j*. When total expenditure exactly 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 lower poverty line z_j^L is given by the sum of the food and non-food components, $z_j^L = z_j^F + z_j^F (1-\alpha_j) = z_j^F (2-\alpha_j)$.

7. The resulting poverty lines suggested that the highest cost of living for poor people was in the NCD, followed by the Papuan and New Guinea Islands regions. The lowest cost of living was in Momase. Setting the poverty line in the NCD equal to 100, the poverty lines in the other regions were: Papuan 72.5; Highlands 53.5; Momase 43.1; and New Guinea Islands 61.4. These poverty lines, in index form, were then used as deflators to convert nominal consumption

into real consumption. This gave a group of 185 households (21 percent of the population) with real consumption per adult equivalent less than the cut-off value of K570 per year. A new set of regional food baskets was formed from the budgets of these households and the resulting set of poverty lines was calculated. These poverty lines, in index form, showed some changes from the results of the previous iteration so the process was repeated. Figure 1 shows that after the fifth iteration, the only poverty line changing was for the Highlands where the index number oscillated between 49.9 and 50.1. These changes were small enough to accept the convergence. The iterations were also started from the values of a Törnqvist regional price index (calculated with the available prices) and the poverty lines still converged to the same levels.

Figure 1: Convergence of Price Index Implied by Lower Poverty Line

(combined basket of foods for Papuan, Momase and Islands)

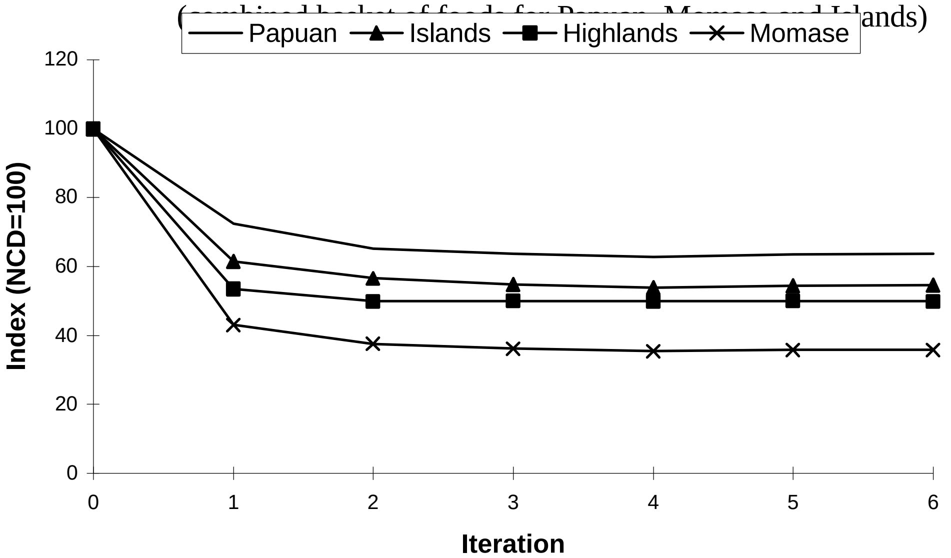

8. The values of the food poverty lines in each region are reported in Appendix Table 4.1. This table also reports whether the cost in region *j* of buying the region *i* basket of foods is less than the cost of buying the region *j* basket of foods, at region *j* prices. If the cost is less, households in region *j* could buy the region *i* basket of foods but seeing as they do not they must prefer the region *j* basket, implying that the region *i* basket is of lower quality. None of the regional poverty line diets were revealed to be inferior to the diets of the other regions.

Appendix Table 4.1: Cost of Region *i* Poverty Line Basket of Foods at Region *j* Prices (combined basket for Papuan, Momase, and Islands regions)

	region *j* prices				
region *i* basket	NCD	Papuan	Highlands	Momase	NGI
			(kina per year)		
National Capital Dist.	**543**	578	544	510	525
Papuan + Momase + Islands	594	**391**	330	**218**	**326**
Highlands	708	580	**288**	271	408

9. The results of estimating the Engel curve for calculating the typical value of non-food spending by poor households were:

$$w_f = 0.732 - 0.052 \ln\left(\frac{x}{n \cdot z^F}\right) - 0.008 n_a - 0.001 n_{c1} - 0.004 n_{c2} - 0.133 \text{ NCD} + 0.033 \text{ Papuan}$$

(2.94) (2.19) (0.22) (0.63) (5.48) (0.89)

$$- 0.054 \text{ Highlands} + 0.014 \text{ Momase} \qquad R^2=0.13.$$

(2.12) (0.55)

The regression results show that the average food share for households who could just afford the food poverty line, if they devoted all expenditure to food, is 0.70 in the base region (the New Guinea Islands), taking account of the effect of the demographic variables. The food shares in the NCD and Highlands are lower than in the rest of the country. Consequently, the non-food allowance will be proportionately larger for these two regions than it is for the rest of the country. The resulting poverty lines, in terms of annual expenditure per adult equivalent, are: NCD K779; Papuan K496; Highlands K390; Momase K280; and New Guinea Islands K424 (with a population weighted average for all of PNG of K399).

Appendix 5
Alternative Spatial Price Indexes

1. The price index reported in Table 1 measures the relative cost of a "basic needs" standard of living in each region. A 'general-purpose' spatial price index may be more appropriate for other uses, so results from four other price index formulas are reported here. The first general-purpose price index considered is the Laspeyre's index, which is widely used. This index calculates the relative cost in each region of buying the base region's basket of goods:

$$L = \frac{\sum_{j=1}^{J} Q_{kj} P_{ij}}{\sum_{j=1}^{J} Q_{kj} P_{kj}},$$

where k is the base region, i=every other region, and j is each item in the consumption basket. The Laspeyre's index overstates the cost-of-living in high price areas because it does not let households make economising substitutions away from items that are more expensive in their home region than they are in the base region. For example, coconut is more expensive in the Highlands than in Port Moresby, so the quantity of coconut consumed per household is lower in the Highlands than in Port Moresby. The Laspeyre's index calculates the cost of purchasing a Port Moresby quantity of coconut at Highlands prices. A true cost-of-living index would calculate the cost of obtaining the Port Moresby level of utility when facing the high prices for coconut that prevail in the Highlands, letting the household rearrange its consumption bundle to minimise cost. Technically, the Laspeyre's index assumes a Leontief utility function.

2. The second price index considered – the Paasche index - also does not allow substitution (i.e., it is exact only for a Leontief utility function). But in contrast to the Laspeyre's index, the Paasche index understates the cost of living in high price areas because it evaluates relative prices using a basket of goods that varies for each of the *i* regions:

$$P = \frac{\sum_{j=1}^{J} Q_{ij} P_{ij}}{\sum_{j=1}^{J} Q_{ij} P_{kj}}.$$

In other words, the Paasche index takes a weighted average of relative prices, where the weights reflect prior economising substitutions by households. Continuing the above example, the Paasche index assumes that Port Moresby households consume the same (low) quantity of coconut as is consumed by Highlands households. This understates the cost of living disadvantage of living in the Highlands compared with Port Moresby because Port Moresby households will reap greater savings by consuming more low-priced coconut.

3. The problem of the Laspeyre's index overstating, and the Paasche index understating, the true cost of living in high price areas has led to the use of *superlative* price indexes in some applications. (Using the Laspeyre's and Paasche indexes to bound the true cost of living index is appropriate only in the unlikely case of income elasticities equal to one for all commodities, i.e., homothetic utility.) A superlative price index will closely approximately a true cost-of-

living index for any utility function (Diewert, 1976). One superlative price index is the Fisher index, which is a geometric average of the Laspeyre's and Paasche indexes: $F = (L \times P)^{1/2}$. Another superlative price index, which is the final one considered here, is the Törnqvist index:

$$T = \exp\left[\sum_{j=1}^{J}\left(\frac{w_{kj} + w_{ij}}{2}\right)\ln\left(\frac{P_{ij}}{P_{kj}}\right)\right]$$

where w_{ij} is the average share that item j has in the consumption basket in region i, and region k is the base region.

4. In principle, all four of these price index formulas require prices for the full set of items in the consumption basket. However, items that comprise 28 percent of the average household's consumption basket were not priced by the survey. Therefore, the results are calculated with just the budget share, quantity and price information for items that comprise the remaining 72 percent of the average household's consumption budget. The results can be interpreted as deflators for total consumption only if it is assumed that spatial variation in prices is the same, on average, for the 'no price' items as it is for the items with prices.

5. The results for the four price index formulas are reported in Appendix Table 5.1.The National Capital District was chosen as the base region. The quantities and budget shares used were population averages for each region, while the prices were averages across the Census Units in each region. The NCD appears to be the region with the highest cost of living, according to every index except the Laspeyre's, which shows a higher cost of living in the Highlands. The basket of goods for the Laspeyre's index is set by consumers in the NCD. Hence, imported items and items produced in the urban import-competing sector have a high weight, and these items are expensive in the Highlands due to internal transport costs. All of the deflators rank Momase as the region with the lowest cost of living. Overall, the results in Appendix Table 5.1 show that the choice of a particular price index formula can have a big impact on the size of deflators, although results for the two superlative indexes are reasonably close. A modification to the Laspeyre's index, with national average quantities and prices used as the base, also gives results that are close to the values for the superlative indexes (100.0, 78.5, 83.7, 58.7, and 70.0).

Appendix Table 5.1: Results for Alternative Spatial Price Deflators

	Fisher Index	Laspeyre's Index	Paasche Index	Törnqvist Index
National Capital District	100.0	100.0	100.0	100.0
Papuan/South Coast	77.7	83.0	72.8	79.9
Highlands	85.7	119.1	61.7	80.7
Momase/North Coast	65.7	81.3	53.1	59.9
New Guinea Islands	79.9	84.9	75.1	80.9

Notes: Quantities and budget shares used in the calculations are population averages.
Results based on the items in the consumption basket (78 percent of total) with prices available.

6. How do the deflators in Appendix Table 5.1 compare with the spatial price index derived from the poverty lines and reported in Table 1? The regional patterns are similar but the values of the deflators are much different. The poverty lines imply a cost of living outside of the NCD that is only 36-64 percent of the cost in the NCD. In contrast, the deflators in Appendix Table 5.1 give a smaller cost advantage to living outside the NCD (e.g., 60-81 percent of the NCD cost according to the Törnqvist index). One likely cause is the assumption that spatial variation in the prices of the 'no price' items is the same as it is for the items with prices observed. The most important 'no price' item is firewood (budget share of 6.2 percent). A small survey of firewood prices was carried out in the Highlands, and in the National Capital District. Average firewood prices were K0.08/kg in the Highlands and K0.14/kg in Port Moresby. This implies a firewood price index of 57 for the Highlands (NCD=100), instead of the implicit index of 80.7 which is calculated from the items with price data and extrapolated across the items without price data. Thus, the assumption that the spatial variation in prices for the 'no price' items is the same as for the items with prices is not a very good one. In fact, if the index value of 57 is taken as a proxy for the relative price of all non-food items with missing prices, the Törnqvist index for the Highlands would fall from 80.7 to 71.5.

7. The problem of missing prices can be overcome by restricting the comparison of the various deflators just to food items. Prices were collected for all 35 foods, so the same information is available for both the food poverty line and the 'general purpose' price indexes. Results for the four price index formulas are reported in Appendix Table 5.2.

Appendix Table 5.2: Price Indexes for Food[a]

	Fisher Index	Laspeyre's Index	Paasche Index	Törnqvist Index
National Capital District	100.0	100.0	100.0	100.0
Papuan/South Coast	92.8	108.4	79.5	91.1
Highlands	78.8	111.5	55.7	74.2
Momase/North Coast	71.2	98.4	51.6	62.6
New Guinea Islands	87.5	102.4	74.8	86.0

Notes: Quantities and budget shares used in the calculations are population averages.
[a]Excludes meals consumed away from home.

8. The food price indexes in Appendix Table 5.2 show a smaller cost advantage to living outside the NCD than do the food poverty lines. For example, the Törnqvist index in Appendix Table 5.2 suggests that the cost of food consumption outside of the NCD is just 63-91 percent of the cost in the NCD. But, the food poverty lines outside the NCD are only 40-72 percent of the NCD food poverty line. A major cause is the difference between the diets of poor households – used for the food poverty line – and the diets of average households – used for the other deflators (Figure 1).

Figure 1: The High Cost in the NCD of Items Consumed by the Poor

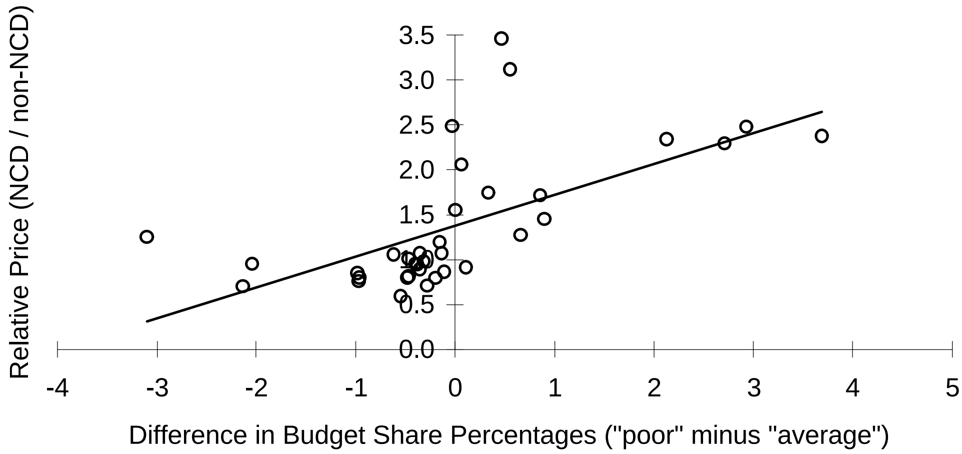

9. Figure 1 shows the relationship between the relative price of each food (NCD/non-NCD) and the importance of that food to the poor relative to its average importance. Specifically, the difference in budget shares between the consumption baskets used for the food poverty lines and the consumption baskets used for the Appendix Table 5.2 deflators is used to indicate the importance of the food to poor households. The relationship in Figure 21shows the high relative price in the NCD of foods that are important to the poor. In other words, the cost of food consumption in the NCD, relative to the rest of the country, is higher for a poor household than for an average household. The high cost of internal transport, which raises the urban price of the local staples consumed by the poor, may cause this. Another reason is that more expensive food items, which are important in the diets of average households but do not have much weight in the poverty line basket, are relatively cheap in the NCD because these foods are either imported or produced in the protected import-competing sector located close to urban areas.

10. In summary, the general-purpose deflators do not match the poverty line deflator used for the main analyses because they are evaluated at different points in the income distribution. This matters because budget shares change as households become richer (i.e., the utility function is not homothetic). Hence, the poverty line deflator and the general-purpose deflators use different weights. The other causal factor is that relative prices are not uniform: some foods are cheap in one area and expensive in another, while for other foods the pattern is reversed. Therefore, the choice of appropriate deflator will depend very much on the purpose at hand, and on the answer to the question: ‘whose cost of living is being measured?’

Appendix 6
Sampling Weights

1. A set of weights, which potentially can vary for each Census Unit in the sample, are needed when estimating statistics from the household data. These weights are needed to correct for the varying selection probabilities of areas and households in different parts of the sample. They can be regarded as made up of three components:

- Correction for the differing sampling rates used in the strata at the area stage of sampling. This affects the two principal strata NCD and non-NCD. In addition, since areas are selected with probability proportional to "size", the size variable (census households reported) has to be included in the weighting. Essentially the weights here are obtained as the sampling interval divided by the census size variable.

- Correction for varying numbers of households selected in the CU. This is made up of two components: (1) the planned variation between 6 (in NCD) or 12 (in non-NCD), and (2) non-response, when the number of interviewed households within the CU differs from the desired number (6 or 12). These selected households have to be blown up to the number existing, in each CU, that is to the number listed during the household listing operation. The weight in each CU is the ratio of the number listed to the number interviewed. Separate weights are needed for the first visit (household data) and the second visit (consumption data).

- A third weight is introduced in a few special cases, namely where the CU is subdivided into segments. If $\mathbf{s_i}$ is the number of segments created in the i-th CU, and one of these is selected at random, and if the listing covers only the selected segment, the weight must be $\mathbf{s_i}$, the listing total $\mathbf{M_i'}$ being the segment listing total. However, if the household listing is extended to cover the whole CU with a CU listing total of $\mathbf{M_i''}$, then the weight should be $\mathbf{M_i''/M_i'}$. In the latter case the $\mathbf{M_i'}$ in the preceding term will of course cancel the $\mathbf{M_i'}$ in the final term.

2. Summarising, the weights to use for each Census Unit are given by the formula:

$$\mathbf{W_{hi} = [\, I_h / b_{hi} \,] \,.\, [\, M_{hi}' / M_{hi} \,]}$$

where the main strata (NCD and non-NCD) are indexed by *h* and Census Units by *i*, and

- I_h is the interval used in the first stage sampling to select CUs (given by the ratio of the number of households in that strata of the sampling frame to the number of CUs selected)
- b_{hi} is the number of households interviewed in the Census Unit (where this can vary between first interviews (household data) and second interviews (consumption data))
- M_{hi} is the number of households in the CU recorded at the 1990 Census (and this number, in conjunction with the interval I_h and the systematic ordering of the CUs, lead to the selection of the particular Census Unit into the sample)
- M_{hi}' is the number of households listed in the Census Unit or segment in 1996 (and this determines the selection probability of any given household in the CU).

3. The sample of CUs for the National Capital District was drawn from two sampling frames that were put together. The first sampling frame was the 1990 census, which had 468 CUs. The second frame was 50 new CUs that had been recently recognised (and mapped) by the NSO as part of the work requirements for the DHS. A systematic sample of 40 CUs was drawn from this combined frame.

4. The count of households done in 1996 by the interview teams was often substantially different from the 1990 Census count, ranging from being over four times larger, to being only one-quarter as large. Many of the CUs with apparently large changes in household numbers were in the Highlands, where the 1990 Census operated by 'line-up' enumeration. To ensure that the apparent size changes were not caused by the method of splitting Census Unit into segments, return visits were made to completely list several of the large Highlands CUs. Even after re-listing, several CUs with apparently large shifts in household numbers, and therefore potentially large or small weights, remained in the sample. Some of the apparent declines in household numbers were able to be explained by factors such as people moving out of areas with severe tribal fighting or crime problems (e.g., a 70 percent decrease in Buk in Western Highlands). Instances of apparently rapid population might be more easily explained as errors of undercoverage in the 1990 Census

5. This possible error in the Census need not cause undue concern because although the aim was to select CUs with probability proportional to size, there is no assumption that the measures of size $\mathbf{M_i}$ are accurate. If they are, the sampling (random) error is likely to be smaller, which is an advantage, but whether they are accurate or not there will be no bias, provided the above weights are used. Therefore the inaccuracy of the 1990 census is not a cause for serious concern. Statisticians refer to the method used here as "probability proportional to estimated size" and it is just as "respectable" as the more ideal case of true PPS. Above all it is important not to try to "correct" the values $\mathbf{M_i}$ after the sample has been used. The value $\mathbf{M_i}$ which appears in the denominator of the weighting formula is introduced in order to cancel the varying probabilities which were used in selecting the CUs. This cancelling will only occur if we use identically the same values $\mathbf{M_i}$ in the weighting formula as were used in selecting the sample. That the size measures are inaccurate does not matter because they are going to be updated by introducing the values $\mathbf{M_i'}$, coming from the very latest count.