Testing the Infrequent Purchases Model Using Direct Measurement of Hidden Consumption from Food Stocks

John Gibson† and Bonggeun Kim‡

Abstract

Reports of zero expenditure on individual commodities during the reference period of a household survey are a frequent but awkward feature of applied demand analysis. These zeros may come from two distinct phenomena: (i) genuine non-consumption, whether for economic or non-economic reasons, and (ii) purchases that happen too infrequently to be captured within the survey reference period, with hidden consumption out of stocks. Distinguishing between these types of zeros is difficult in the disaggregated survey data increasingly used by demand analysts. Hence econometric models for infrequent purchases rely on untested hypotheses. In this paper we test such models, using data from an unusual household survey where food stocks are measured at the start and end of the survey reference period. Parameter estimates using these direct measures of hidden consumption out of stocks are compared with estimates from infrequent purchase models that attempt to recover this hidden consumption. The results suggest considerable bias when using the infrequent purchase models.

Keywords: Engel curves, food consumption, infrequent purchase, storage

JEL Codes: C81, D12

† Department of Economics, Private Bag 3105, University of Waikato, Hamilton 3240 New Zealand (e-mail: jkgibson@waikato.ac.nz)

‡ Department of Economics, Seoul National University, Gwanangno 599, Gwanak-gu, Seoul, Korea (e-mail: bgkim07@snu.ac.kr)

I. Introduction

Understanding how consumer demand shifts in response to income changes is fundamental to many areas of economics. The surveys used to estimate Engel curves and income elasticities typically observe households for only a short period; ca. 7-14 days for foods. Because it is possible to eat from food stocks, a household may be observed to spend nothing on a food that it nonetheless eats, causing bias in estimated demand parameters (Keen, 1986). This issue also arises in developing countries, where consumption is not only from the market but also from own-production. Eating is more frequent than harvesting, and in the extreme of seasonal crops that are then stored, the harvest occurs only once while consumption out of that harvest lasts for many months. But surveys that indirectly derive consumption from questions on purchases, own-production, and inter-household transfers typically miss the hidden consumption from food stocks.[1]

Econometricians have developed infrequent purchase models that attempt to recover this hidden consumption, starting with the *p*-Tobit model of Deaton and Irish (1984). These models adopt a two-equation structure, with a latent purchase equation for whether spending occurs in the survey reference period followed by a Tobit of the spending level (including zeros). Consumption is estimated as the product of the odds of purchase and the observed spending, and with a long enough reference period is observationally equivalent to spending. Thus, for households with positive purchases of a food during the survey, the reported value is $1/p$ times consumption during the survey period, where p is the ratio of the length of the survey period to the length of the purchase period. For example, if a household buys a sack of rice every four weeks and the survey reference period is one week, then p=0.25 and for

[1] Some surveys directly ask about food consumption, along the lines of "what did you eat yesterday?" but these tend to be used more by nutritionalists than economists, in part because they have limited information on total household resources, which are needed to model demand behaviour (Deaton, 1997). We do not consider these types of direct consumption surveys in our analysis below and instead focus on surveys that are structured according to the various means of acquisition of commodities (purchases, own-production, transfers), which includes all household budget surveys in developing countries and many living standards surveys.

households buying rice during the survey (with this probability being *p*, if no household has a corner solution), the reported expenditure is four times the rate of weekly consumption.

Clearly this problem is wider than just one of zero reported purchases. A household may supplement foods purchased during the survey period with takings from stocks, so that reported spending is not zero but still understates actual consumption in the reference period. Conversely, spending may overstate consumption if stocks are being built up during the survey period. Allowing other forms of commodity acquisition, such as from own-production and inter-household transfers also does not alter the problem. As long as food has some durability and there are transactions costs (e.g. the food gardens in the setting providing the data used below are often an hour travel from dwellings), consumption will occur more frequently than own-production and there is potentially hidden consumption out of stocks.

The *p*-Tobit and infrequent purchase models (IPMs) described by Blundell and Meghir (1987) are increasingly used because many demand studies now utilise disaggregated household survey data, so zeros in dependent variables are more common.[2] For example, Nordstrom and Thunstrom (2009) use IPMs to study demand for 32 separate grain products in household budgets, with the occurrence of zero expenditure on individual products as high as 92 percent. But despite the growing relevance of these models, there is yet to be a critical test of their identifying assumptions, in part because of lack of suitable data.

In this paper, we use unusual data from a survey where food stocks are observed at the start and end of the reference period, so that hidden consumption out of stocks can be measured. Despite the attention of econometricians to infrequent purchase and storability, many surveys assume that food is non-storable and do not measure food stocks.[3] With the extra information on the estimated consumption from stock changes we can distinguish

[2] Of the 71 citations in the ISI Web of Science to Deaton and Irish (1984), 14 are for articles published since 2005, while 19 of the 74 citations to Blundell and Meghir (1987) are for articles since 2005.

[3] Even in papers that are motivated by the infrequent purchase problem, the examples that are given of durability beyond the short reference period of the survey exclude food at home (e.g., Kay et al, 1984, p.170), presumably because of the implicit assumption that food has no storability.

between non-consumption and purchase infrequency as reasons for zero observed expenditure on particular foods. Moreover, we can form a measure of consumption for each food that has verisimilitude (that is, is closer to the truth) over other measures based just on purchases (and other acquisitions), and compare parameter estimates that use the direct measures of hidden consumption out of stocks with estimates from IPMs that attempt to recover this hidden consumption. A further advantage of our data is that the survey reference period varied across households, so it is possible to test if the IPMs perform better when the reference period is longer. Such a test would not normally be possible because in most surveys, while the reference period may vary across commodities (longest for semi-durables and shortest for foods), for a given commodity all households are observed for the same time.

This testing of IPMs is needed since these models rely on identifying assumptions to distinguish them from other models of zero expenditures that are derived from quite different household behavior. Specifically, zero expenditures also may occur because the household is a genuine non-consumer of the good, due either to abstention (for non-economic reasons such as health or religious preferences) or to a corner solution where the household contains potential consumers of the good who cannot afford it at current income and prices. In the case of abstention, the double-hurdle model of Cragg (1971) is typically used, where a Probit models the participation decision and a Tobit models the expenditures, and economic variables should be excluded from the participation equation (Jones, 1992). With corner solutions, the assumption is that demand for the good is censored from below at zero and a Tobit model is conventionally used. However, without prior information, it is not possible to be certain of the cause of the observed zeros (Meghir and Robin, 1992) and so the choice of empirical specification is just based on the identifying assumptions.

The remainder of the paper is as follows. In Section II, we provide a simple measurement error model and relate our model to the estimation models for non-consumption and infrequent purchase. In Section III, we describe our data and report the empirical results. Section IV concludes with a summary of the main findings.

II. Models for Zeros

We present a simple model to illustrate how errors in measuring consumption affect demand analysis. It helps at the outset to distinguish between three consumption variables: latent consumption, y^*; actual consumption during the survey period, y; and observed acquisitions (purchases) during the survey period, y'. Latent consumption is observable in principal, but only if actual consumption occurs, since y^* is assumed to be normally distributed and hence may be negative. It is this latency which gives rise to the classic Tobit model, with its assumption of censoring at zero (corner solutions), and nothing done in this paper to directly measure hidden consumption from stocks alters that feature of demand data.

Observed acquisitions during the survey period are the sum of actual consumption, which includes hidden consumption out of stocks, and an error term:

$$y_i' = y_i + \varepsilon_i \tag{1}$$

where ε_i represents the difference between the two measures. When $y_i' = 0$ and $y_i > 0$ the classic case of zeros due to infrequent purchases occurs. However, equation (1) also covers other measurement errors, such as acquiring a small quantity during the survey period with remaining consumption coming from stocks, or conversely acquiring a much larger quantity than is eaten during the survey period and building up stocks.

When the classic Tobit model is applied to food demand data, it is based on the assumption that $\varepsilon_i = 0$, so:

$$y_i' = y_i = \max(0, y_i^*) \quad . \tag{2}$$

and in the regression model that underlies much applied demand analysis, with latent consumption as the dependent variable, a set of x_i consumption determinants (e.g., income) and $u_i \sim N(0,\sigma^2)$ representing preference heterogeneity in demand:

$$y_i^* = x_i\beta + u_i\,, \tag{3}$$

the sample likelihood function is:

$$\ln L = \Sigma_+(-\ln\sigma + \ln\phi((y_i - x_i\beta)/\sigma)) + \Sigma_0 \ln(1-\Phi(x_i\beta/\sigma)) \tag{4}$$

where Σ_+ is the summation for households with positive consumption and Σ_0 is the summation for households with zeros, while $\Phi(\cdot)$ and $\phi(\cdot)$ represent the standard normal cumulative and density functions.

But when a household survey is not able to observe all consumption during the reference period, because of hidden consumption from unmeasured food stocks, $\varepsilon_i \neq 0,$ and cases of $y_i' = 0$ do not necessarily correspond to $\max(0, y_i^*) = 0.$ Thus, these are not genuine corner solutions. To deal with this, the literature on IPMs defines D_i as a latent variable for whether an acquisition (purchase) occurs during the survey period ($D_i > 0$ if and only if $y_i' > 0$) and P_i is the probability of purchase. This yields:

$$E(y_i') = E(y_i' \mid D_i > 0)P_i + E(y_i' \mid D_i \le 0)(1-P_i) = E(y_i' \mid D_i > 0)P_i \tag{5}$$

and to derive the new censoring rule the literature assumes that $E(y_i') = E(y_i)$ or $E(y_i' \mid D_i > 0)P_i = E(y_i)$ according to equation (5). By using this assumption, the new censoring rule is

$$y_i' = \begin{cases} y_i^*/\Phi(z_i\theta) & if\ \ y_i' > 0 \\ 0 & otherwise \end{cases} \tag{6}$$

where $\Phi(z_i\theta) = P_i$ is the probability of purchase from a model where $D_i = z_i\theta + w_i$ and the z_i are determinants of purchase probability and $w_i \sim N(0,1)$. Under independence,[4] the probability that positive consumption occurs with no purchase occurring over the short survey period is $(1-\Phi(z_i\theta))\Phi(x_i\beta/\sigma)$ and the corresponding sample log-likelihood function is

$$\begin{aligned}\ln L = \Sigma_+(-\ln\sigma + 2\ln\Phi(x_i\beta/\sigma) + \ln\phi((\Phi(z_i\theta)y_i' - x_i\beta)/\sigma))) \\ + \Sigma_0 \ln(1-\Phi(z_i\theta)\Phi(x_i\beta/\sigma)),\end{aligned} \tag{7}$$

When $\Phi(z_i\theta) = 1$ or a good is non-storable, equation (7) becomes the Tobit model of equation (4).

The discrepancy between acquisitions, y' and consumption, y_i that is caused by storability is corrected by the infrequent purchase model by the assumption, $E(y_i') = E(y_i)$. However, if some non-consumption zeros are incorrectly treated as infrequent purchase zeros, then the estimate of $\Phi(z_i\theta) = P_i$ is upwardly biased and the infrequent purchase model overestimates consumption. The failure of this assumption may also cause biased estimates of income elasticities of demand and it is that empirical implication that we focus on here.

III. Empirical Analysis

A. Survey Context

Data used in this paper come from the 1996 Papua New Guinea Household Survey (PNGHS), based on a random sample of 1200 households for whom 1144 have complete information on all variables used here. Some results are reported separately for rural (n=830) and urban (n=314) sub-samples since the transactions costs of food acquisition and resulting incentives to store food may vary by sector. The PNGHS is a multi-topic survey similar to the Living

[4] Many subsequent papers usually focus on distributional assumptions of the estimation techniques developed by Deaton and Irish (1984) and Blundell and Meghir (1987) in their applications to tobacco (Kimhi, 1999), clothing demand (Majima, 2008), savings and remittances (Sinning, 2007).

Standards Measurement Surveys of the World Bank.[5] Consequently, the food consumption module is less detailed than in household budget surveys, with only 36 separate food groups distinguished, rather than the 100 or more typically used for a budget survey. But even with these less detailed groups, almost 55 percent of households reported zero expenditures (including zero non-market acquisitions) on the typical food group during the survey period and the food with the fewest reported zeros still had 13 percent report zeros.

A key feature of the survey for our purposes is that it used bounded recall, where this design was originally chosen to restrict the scope for telescoping errors.[6] Specifically, the start of the recall period was signaled by the first interview and in the second interview households were asked about purchases, own-production and net transfers received in the time elapsed since the first interview. This design gives two potential advantages for testing IPMs; variation in the period over which consumption is observed, and the possibility of measuring starting and ending food stocks so that what is typically hidden consumption out of food stocks can be directly observed.

In terms of the variation in recall length, the second visit was designed to be approximately two weeks after the first interview. However the very difficult topography and poor transport infrastructure in PNG meant these return visits were sometimes sooner and other times later, depending on the logistics of moving survey teams around. While the median bounded recall period was 14 days, the standard deviation is 3.65 days. Almost eight percent of households had a recall period of only one week and one percent had a three-week recall. Hence there is some variability in the length of the observation period, which we can use to test the hypothesis that the performance of IPMs improves when the observation period lengthens. This would not be possible in other surveys, which either use a diary for a fixed

---

[5] Details on the survey and data downloads are available at: http://go.worldbank.org/CJ9LIGVJ00

[6] With telescoping, a respondent compresses purchase occasions and other forms of acquisition that occurred over a longer period of time into the reference period and thus the reported value of acquisitions exceeds the true value over that period.

period (typically two weeks) or else use a single visit (unbounded) recall where the respondent is asked about acquisitions in the last week, fortnight or month, with the period depending on the type of commodity, but with no variation across households in the recall period for the same commodity.

In terms of observing what typically is hidden consumption from stocks, during the first interview the enumerators weighed the beginning food stocks for 18 foods. The ending stocks of these same foods were reweighed at the second interview so that stock change could be measured. Several of these foods are not produced by households (e.g., rice and flour are imported, and sugar, canned fish and canned meat are commercially produced). Stocking of these foods reflects infrequent purchase behavior, possibly to exploit pecuniary economies from buying larger volumes (e.g. a 25-kg rice sack). In this respect, these foods are no different to the situation in previous uses of IPMs in developed countries.

In terms of the own-produced foods, are all harvested non-seasonally since PNG is in the humid tropics and lacks both a regular monsoon and a dry season. In this regard, the setting is similar to equatorial Africa (e.g. Rwanda, Burundi and Uganda). Despite the lack of seasonality, harvesting is not as frequent as consumption due to transactions costs, since food gardens can be more than an hour walk from the dwelling, especially in the mountainous highlands areas. Also, religious and social commitments mean that people do not go to their food gardens every day. The locally produced crops are mainly plantains and root crops, which are storable for about a month. These locally grown foods are also widely purchased, but infrequently due to the transactions costs of going to market, since the rural population is poorly served by roads and many markets only open for limited times each week.[7]

---

[7] Gibson and Rozelle (2003) note that the average household in PNG is 2.5 hours from the nearest road, and 3 hours from the nearest government station (the equivalent of a market town). Hence there are large transactions costs in going to market.

These purchase and production characteristics mean that household food stocks may be large, but not so impossibly large as to prevent accurate measurement. For example, the foods with the largest stocks recorded were cassava and yams, with average holdings of about nine kilograms (kg) amongst stockholders, and a maximum of 90 kg. For a family of five, this maximum stock could feed them for 20 days if the stocked food provided half of daily calories. Hence this is far shorter storage than may occur in highly seasonal environments but still exceeds the reference period used by almost all surveys and so allows a role for hidden consumption from stocks. Amongst the store-bought foods, rice was most heavily stocked, with average holdings of 7 kg amongst stockholders and a maximum of 50 kg. Stocks of other store-bought foods averaged about 2 kg (flour and sugar) or less than one kg (canned foods), with maxima of 25 kg (flour and sugar) or up to 10 kg (canned foods).

This unusual environment, where it is feasible to measure food stocks at the start and end of the survey, is of little inherent interest since few other countries have the same characteristics. But it provides a good opportunity to test the infrequent purchases model since we are unaware of any survey elsewhere in the world where hidden consumption from food stocks has been directly measured in this manner.

B. Descriptive Results

Table 1 reports descriptive statistics on the importance of hidden consumption from stocks. Although stock measurements were made for 18 foods, we focus here only on 10 of them, ignoring the fresh meats, sago and sweet potato. The fresh meats are the one group of foods where corner solutions may represent non-economic motivations (e.g., PNG has many Adventist households with religious prohibitions on eating pork). To allow the cleanest comparisons between the Tobit on y_i and the IPM on y_i' we exclude the meats, which may have non-economic reasons for zeros. Sago is ignored because it is only grown and consumed in a small area of PNG while sweet potato is ignored because its IPM estimates would not converge. Nevertheless, the 10 foods considered comprise just under one-quarter of the total household budget, ranging from banana at 5.6 percent to flour at 0.7 percent (Table 1). The budget shares are similar, whether based on consumption, y_i or observed acquisitions, y_i'. This is to be expected in a non-seasonal environment, since households may be adding to, as well as running down, food stocks.[8]

How important is hidden consumption from food stocks? The third column of Table 1 reports the proportion of households who recorded zero acquisitions of each food during the survey period but for whom consumption actually occurred out of stocks. One-sixth of households reported no acquisition of bananas (plantains), but actually consumed them from stocks. While the proportions are lower for the other foods, averaging just over six percent, there is still a non-trivial fraction relying entirely on storage to provide their consumption of each food during the survey period.

(Table 1 about here)

Even though the proportions in column 3 relate to the zero expenditures (acquisitions) that motivated the development of IPMs, this is not the full extent of hidden consumption. In column 4, the proportion of households acquiring less food during the survey than their consumption out of stock is reported. This situation can be thought of as partially-hidden consumption and is surprisingly frequent. Around one-quarter of households consumed more coconut, sugar, yams and rice from stocks than their recorded acquisitions, while for the other foods the proportions are at least one-tenth. That both store-bought and own-produced items are in the group where a high proportion of households have partially hidden consumption indicates that the 'infrequent purchase' title is a misnomer and that the problem of hidden consumption from food storage is more widespread than previously thought.

[8] In a country where an annual harvest is stored and then consumed over many months, it would be expected that $y_i > y_i'$ for any survey that occurs outside the harvest period.

To provide an estimate of the overall importance of hidden consumption for each food, column 5 reports the weighted average of the proportions in columns 3 and 4. According to this average, hidden consumption from storage is most important for banana, coconut, rice and sugar. These foods offer the most favorable conditions for the IPMs because when the degree of storability of a food increases more of the zeros are due to infrequent purchase. Hence, fewer non-consumption zeros (corner solutions) will incorrectly be assumed to be infrequent purchase zeros. Thus, a negative relationship between the degree of bias in the IPMs and the storability of foods is expected.

C. Econometric Results

Recall from Section 2 that the aim of the IPMs is to give parameter estimates that rely not just on observed acquisition of the food but also on the hidden consumption recovered by the latent purchase equation. To see how well this procedure works in practice, we estimate two sets of demand models, which are both in budget share form. The first applies the classic Tobit model to y_i, our data on actual consumption during the survey period that uses the direct measurements of what typically would be hidden consumption out of stocks. We use a Tobit model since observing this hidden consumption does nothing to obviate the censoring at zero due to corner solutions. The second model is an infrequent purchases model, applied to the survey data on y_i' (acquisitions, as measured from purchases, own-production and net transfers) which ignores the hidden consumption from stock changes. Since equations (4) and (7) have the same parameters from the underlying linear model of latent demand, we should get the same income elasticities of demand, from the two sets of estimates.

The budget shares for the demand model are based on shares of total expenditure,[9] which in addition to the 36 foods, includes 20 types of frequently purchased non-foods, an

[9] Therefore we do not assume any multi-stage budgeting, and the estimated expenditure elasticities can be directly interpreted as income elasticities, rather than needing to be first combined with parameters from a multi-stage budget allocation equation.

annual recall of 31 infrequent expenses and the value of the flow of services from durables and owner-occupied dwellings. The measure of household income used is the logarithm of real total expenditure per capita, which is deflated for price differences over both time (the survey was staged over 12 months) and space (deflators for five regions). This is the best estimate of permanent income available from the survey and captures the total economic resources available to households. We also use household size and the share of household members in various age and gender groups, the characteristics of the household head (under the assumption that this person has some influence over family diets) and the region in which the household is located as control variables in the Engel curve estimation.

How well do the IPMs replicate the results of using Tobit on the data where the hidden consumption out of stocks is directly measured? The results in Table 2 suggest that the IPMs provide substantially biased estimates of the elasticities, compared with what is calculated from direct measurement of consumption from stocks. For example, although flour, sugar, and canned meat are all luxuries according to the Tobit estimates, the income elasticities estimated by the IPMs are only 0.45, 0.72 and 0.49.[10] The pattern of lower income elasticities coming from the IPMs holds for eight of the ten foods studied, and for five of the eight there is no overlap of the 95 percent confidence intervals for the income elasticities for the same food from the two estimation procedures. Amongst these five with statistically significant differences in the income elasticities, the degree of bias in the IPM estimates ranges from 30.2 percent for rice to 81.4 percent for flour.

(Table 2 about here)

The income elasticities from budget share models are calculated as: $[(\beta/w)+1]$, where β is the coefficient on income in the budget share regression and w is the budget share. The

[10] Readers in rich countries may be surprised that income elasticities of food demand could be so high. But in a poor country (headcount poverty from PNGHS was 38 percent) with significant undernutrition (42 percent of the PNG population did not meet food energy requirements), food overall has a high income elasticity of demand, and these particular foods have been introduced to the diet and are effectively luxuries.

estimates reported in Table 2 are evaluated at the mean budget shares. Hence, the income elasticities could differ either because average budget shares differ, or because of different estimates of β. In fact, the average budget shares are very similar when using either acquisitions or consumption (as seen previously in columns 1 and 2, Table 1). However, the estimates of β are statistically significantly different between the IPM results on acquisitions and the Tobit results on consumption so the bias in the income elasticities is coming from the parameter estimates of the IPMs.[11]

In addition to showing the large biases in the IPM estimates of the income elasticities, the results in Table 2 also illustrate the negative correlation between the bias and the degree of storability of the food. When the bias estimates in column 5 are compared with either the share of households recording zero acquisitions of each food but for whom consumption actually occurred out of stocks (column 3, Table 1) or the share with smaller acquisitions than the consumption from stocks (column 4, Table 1), significant negative relationships are apparent. The Spearman (rank) correlation coefficients are -0.58 and -0.56, and both are statistically significant at the p<0.09 level. Thus, the foods that have the highest degree of storability have the least bias in the IPM estimates, because with higher storability more of the zeros are actually due to infrequent purchase and so there is less scope for non-consumption zeros (corner solutions) to be incorrectly treated as infrequent purchase zeros.

The IPM elasticity estimates in Table 2 are from equations without exclusion restrictions, so that $x_i = z_i$. From our reading of the literature, this is common practice in the applied

[11] Another concern with the elasticities in Table 2, and with the unreported parameter estimates from the Tobits and IPMs underlying them, may be that there is a logical inconsistency in using a dependent variable that either does or does not measure hidden consumption from stocks, while the empirical proxy for the main independent variable – income – always uses the data on consumption out of stocks for the 18 foods where stocks were measured. We believe that the more comprehensive the measure of consumption, which here includes stock change, the better it will proxy permanent income of the household. However, we also redefined total per capita consumption to exclude the contribution of stock changes and used this (worse) proxy for permanent income in a replica of the regressions that contribute to Table 2. The result of a significant downward bias in the estimates of the income elasticities from the IPMs was unchanged. For the five foods with significant differences between the IPMs and the Tobits, the new estimates of the income elasticities were (IPM, Tobit): sugar (.759, 1.477), tinned meat (.537, 1.300), flour (.467, 2.345), tinned fish (.494, .949), rice (.701, .958).

demand studies that use IPMs; the same set of covariates used to explain demand are used to predict whether a purchase occurs in the survey period. Relying on functional form assumptions to identify parameters may not produce robust results, as the literature on exclusion restrictions in Heckman models has shown.[12] We therefore re-estimated the IPMs with various exclusions imposed. To reduce clutter, the results shown in Table 3 are for just two of the foods where the IPM income elasticities showed significant differences from the Tobit elasticities; rice and sugar.

Imposing a variety of exclusion restrictions makes no difference to our conclusion that income elasticities from the IPMs substantially understate the income elasticities coming from the more correct procedure of the Tobit model applied to data where hidden consumption out of stocks is directly measured (Table 3). The exclusion restrictions that we used are for variables that relate to the transactions costs of going to stores or markets, which could make households shop infrequently and stock foods, and the scope for food storage in the home. Specifically, for transactions costs we use the number of hours needed to reach the nearest transportation network (road, airstrip or boat launching place), the number of trade stores and markets in the respondent's village, and an index of market development (the total number of stores, markets and transport businesses in the village). For storage opportunities we use the floor area of the dwelling and whether it has an iron roof. These variables almost always have the predicted sign in the first stage model for whether a purchase is reported (negative for remoteness and dwelling attributes, positive for the others) and are statistically significant. But including these variables in the purchase probability equation but not in the demand equation makes almost no difference to the income elasticity estimated by the IPMs. The income elasticity estimates for sugar range from 0.720-0.744, compared with 0.719 when no exclusion restrictions are used, and for rice from 0.691-0.722 compared with 0.711 when

[12] See for example XXX (YYYY). Bonggeun – can you find a reference in labour literature to a paper criticising lack of exclusion restrictions. Maybe even to Deaton 1997 if nothing comes to hand.

just relying on the functional form assumptions. Thus the bias in the IPM estimates of the elasticities does not appear to be coming from incorrect exclusion restrictions.

Next we turn to explain why the IPM provides systematically biased estimates.

Although the recall period averaged 14 days, it varied from about 8 to 25. With this variance in the observation period, we can devise another test for the assumption of $E(y_i') = E(y_i)$. We hypothesize that the IPM should work better for households who had a longer recall period than those with shorter, thus we could split the sample in half and test for better performance of the IPM with the longer observation period (greater than or equal to 14 days). Table 4 confirms our prediction and the discrepancy between the income elasticities of purchase and that of consumptions becomes much smaller with the longer observation period. The pattern is similar for nine foods as the income elasticities of consumption substantially decrease from those of the full sample while those of purchase do not change from those of the full sample.

One way to investigate the causes of bias is to see whether the basic assumption of infrequent purchase models is valid by directly comparing two measures of expenditures: the self-reported purchase expenditure, y_i' and the true consumption, y_i. Since storability of goods is assumed to cause the discrepancy between self-reported expenditures and true consumptions, the IPM assumes that observed and latent expenditures are the same on average over the long survey or observation period as the identifying assumption, $E(y_i') = E(y_i)$. As the IPM argues, the discrepancy between y_i' and y_i^* could be mainly caused by storability of purchased goods (e.g. storable purchased 20 kg rice sack in urban areas of PNG). The assumption of $E(y_i') = E(y_i)$ will be safe for the case of purchased storable goods, but it may not be valid for the case of seasonally produced goods. For the storability of purchased goods, the type of measurement error (ε_i) would be random. It is because some households purchased 20 kg rice sack just before the survey started and so using typical survey would have lower than true expenditure of goods, and others who purchased during the survey and had not consumed all by the end of survey period would have reported expenditures higher than true consumptions. This type would lead to that observed and latent expenditures are the same on average over the long survey or observation period as the identifying assumption, $E(y_i') = E(y_i)$. However, the measurement errors in the seasonally produced storable goods are more likely to be non-random because outside of the harvest period, the households are only drawing down their stocks, not increasing the stocks. There is not counterpart to buying a 20 kg rice sack during the survey period. This type of measurement errors would lead to $E(y_i') = E(y_i)$. This prediction is confirmed by our data and Column 1 of Table 3 shows that the assumption of $E(y_i') = E(y_i)$ is rejected at conventional level of significance for four foods under $H_1 : E(y_i') < E(y_i)$. Especially, the assumption is rejected for three foods (sugar, flour, and rice) where there is no overlap with the 95 percent confidence intervals surrounding the income elasticity with true consumptions.

Furthermore, if the type of errors is dependent on where people live and its corresponding purchase and self-production patterns, then the test results for the assumption of $E(y_i') = E(y_i)$ may be different between urban and rural areas where there is difference in the structure of rural and urban diets. For example, on any given day, almost 90 percent of urban residents may be found eating rice while the rate for rural residents is only one-quarter. The error types due to location specific consumption patterns could lead to different test results across areas and Column 2 and 3 of Table 5 support this prediction. For five foods, the assumption is rejected in one area (either urban or rural) but not in the other area. Table 4 provides the alternative test results for the assumption of $E(y_i')=E(y_i)$ with $H_1 : E(y_i') \neq E(y_i)$.

IV. Conclusion

Reports of zero expenditure on individual commodities during the reference period of a household survey are a frequent but awkward feature of applied demand analysis. Infrequent purchase models have been developed to deal with this problem of hidden consumption out of household stocks but have not been able to be tested because of lack of suitable data. In this paper we use data from an unusual household survey with direct observation of food stock changes during the period of the survey. Parameter estimates using these direct measures of hidden consumption out of stocks are compared with estimates from infrequent purchase models that attempt to recover this hidden consumption.

For five out of the ten foods used in the test, the infrequent purchase model gave a significantly different estimate of the income elasticity of demand for the food than did the Tobit model applied to the data that directly measured hidden consumption from stocks. Specifically, the income elasticity was biased downwards for all five of these foods. The bias is much smaller for the sub-sample of households for whom expenditures were measured over a longer period. Thus, the IPM works least well when applied to demand data observed over a short observation period, which is the type of situation that the model is designed for. These results suggest that further effort should be made to test the infrequent purchases models that have been developed in the literature, and to come up with robust procedures for uncovering hidden food consumption from stocks.

References

Boizot, C., Robin, J-M., and M. Visser. (2001), "The demand for food products: An analysis of interpurchase times and purchased quantities", *The Economic Journal*, 111(April): 391-419.

Blisard, N. and Blaylock, J. (1993), "Distinguishing between market participation and infrequency of purchase models of butter demand", *American Journal of Agricultural Economics,* 75(2):314-320.

Blundell, R. and Meghir, C. (1987), "Bivariate alternatives to the Tobit model", *Journal of Econometrics,* 34: 179-200.

Deaton, A. and Irish, M. (1984), "Statistical models for zero expenditures in household budgets", *Journal of Public Economics,* 23: 59-80.

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

Gibson, J. and Rozelle, S. (2003), "Poverty and access to roads in Papua New Guinea", *Economic Development and Cultural Change,* 52: 159-185.

Jones, A. (1992), "A note on computation of the double-hurdle model with dependence with an application to tobacco expenditure", *Bulletin of Economic Research,* 44: 67-74.

Kay, J., Keen, M., and Morris, C. (1984), "Estimating consumption from expenditure data", *Journal of Public Economics*, 23: 169-181.

Keen, M. (1986), "Zero expenditures and the estimation of Engel curves", *Journal of Applied Econometrics,* 1: 277-86.

Kimhi, A., (1999), "Double-hurdle and purchase-infrequency demand analysis: A feasible integrated approach", *European Review of Agricultural Economics*, 26(4): 425-442.

Majima, S. (2008), "Fashion and frequency of purchase: womenswear consumption in Britain, 1961-2001", *Journal of Fashion Marketing and Management*, 12(4): 502-517.

Meghir, C. and Robin, J. (1992). "Frequency of purchase and the estimation of demand Systems", *Journal of Econometrics*, 53: 53-85.

Nordström, J. and Thunström, L. (2009) “The impact of tax reforms design to encourage healthier grain consumption”, *Journal of Health Economics*, 28: 622-634.

Newman, C., Hencion, M., and Matthews, A. (2001). “Infrequency of purchase and double-hurdle models of Irish households' meat expenditure” *European Review of Agricultural Economics*, 28(4): 393-419.

Shinning, M. (2007), “Determinants of savings and remittances: Empirical evidence from immigrants to Germany”, *IZA Discussion Paper* No. 2996.

Table 1. Sample Statistics and Storability Parameters, PNG data, N=1144

Variable	Budget Shares y_i'	y_i	Proportion of zeros with hidden consumption	Proportion of all households with hidden consumption	Average Storability
	Mean(SE)	Mean(SE)	Proportion	Proportion	Proportion
Sugar	.0089(.0144)	.0093(.0170)	.073	.276	.187
Tinned meat	.0165(.0007)	.0167(.0008)	.024	.170	.096
Flour	.0066(.0184)	.0082(.0299)	.049	.218	.103
Tinned fish	.0173(.0243)	.0173(.0243)	.029	.113	.081
Banana	.0561(.0852)	.0549(.0835)	.167	.246	.233
Coconut	.0126(.0262)	.0132(.0260)	.083	.303	.206
Yam	.0189(.0579)	.0189(.0560)	.035	.245	.089
Cassava	.0112(.0318)	.0111(.0306)	.034	.140	.071
Rice	.0380(.0490)	.0410(.0510)	.076	.243	.200
Taro	.0397(.0797)	.0389(.0763)	.061	.196	.131

Note: average storability is a weighted average of the shares with completely hidden consumption from stocks (column 3) and the shares with partially hidden consumption from stocks (column 4).

Table 2. Expenditure Elasticities, PNG data, N=1144

Variable	Acquisitions (1)	SE	Consumption (2)	SE	Bias (=1-(1)/(2))
Sugar[a]	.719	.074	1.499	.128	0.520
Tinned meat[a]	.489	.105	1.341	.124	0.635
Flour[a]	.450	.203	2.428	.411	0.814
Tinned fish[a]	.532	.074	.949	.092	0.439
Banana	.612	.070	.636	.070	0.037
Coconut	.663	.087	.428	.136	-0.549
Yam	.899	.152	.655	.381	-0.372
Cassava	.470	.151	1.015	.276	0.5366
Rice[a]	.711	.058	1.019	.069	0.302
Taro	.668	.087	.719	.136	0.070

Note: Superscript [a] indicates that there is no overlap of the 95 percent confidence intervals for the income elasticity estimates in column (1) and column (2).

Table 3. Sensitivity Analysis: Use of Various Exclusion Restrictions

Variable	Sugar			Rice		
	IPM Income Elasticity	Bias	Acquisition Decision Coefficient (SE)	IPM Income Elasticity	Bias	Acquisition Decision Coefficient (SE)
No Exclusion Restriction IVs	.719	.520		.711	.302	
IV1: Hours to Nearest Transportation	.744	.503	-.0025 (.0003)[c]	.691	.321	-.0025 (.0005)[c]
IV2: # of Trade Stores in Village	.735	.509	.0441 (.0091)[c]	.710	.303	.0599 (.0210)[c]
IV3: # of Markets in Village	.720	.519	-.0194 (.0410)	.692	.320	.2262 (.0694)[c]
IV4: Index of Market Development	.733	.511	.0388 (.0073)[c]	.692	.320	.0912 (.0216)[c]
IV5: Floor Area of Dwelling	.723	.517	.0023 (.0019)	.708	.305	-.0046 (.0020)[b]
IV6: Dwelling has Iron Roof	.723	.517	.3510 (.1108)[c]	.722	.291	-.2902 (.1181)[b]

Note: Superscript [a] , [b] and [c] represent the levels of statistical significance of 10%, 5% and 1% respectively.

Table 4. Expenditure Elasticities with Different Observation Periods

Variable	Bias Total Sample N=1144	Bias Longer Observation Period, N=633	Bias Shorter Observation Period, N=511
Sugar	.520	.350	.666
Tinned meat	.635	.437	.767
Flour	.814	N/A	.914
Tinned fish	.439	.146	.613
Rice	.302	.001	.347

Table 5. Tests for the Assumption of the IPM, PNG data, N=1144

Variable	Total N=1144		Rural N=830	Urban N=314
	$D = E(y')-E(y)$	t-Test $H_0{:}E(y')=E(y)$ $H_1{:}E(y')\neq E(y)$	t-Test	t-Test
Sugar	-.0004	p=.136	p=.201	p=.402
Tinned meat	-.0002	p=.256	p=.421	p=.424
Flour	-.0016	p=.028[b]	p=.035[b]	p=.534
Tinned fish	.00002	p=.879	p=.855	p=.980
Banana	.0011	p=.244	p=.150	p=.132
Coconut	-.0006	p=.021[b]	p=.010[a]	p=.883
Yam	.0006	p=.924	p=.961	p=.767
Cassava	.00009	p=.791	p=.517	p=.154
Rice	-.0022	p=.018[b]	p=.148	p=.010[b]
Taro	.0007	p=.361	p=.338	p=.847

Note: Superscript [a], [b] and [c] represent the levels of statistical significance of 10%, 5% and 1% respectively.

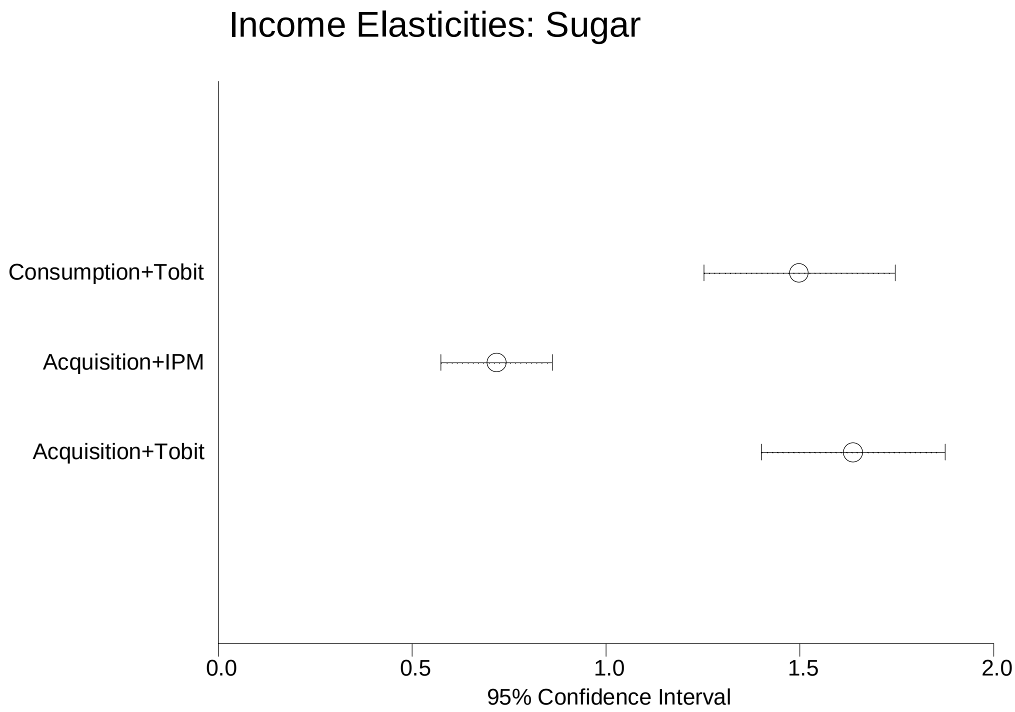
Income Elasticities: Sugar
Consumption+Tobit
Acquisition+IPM
Acquisition+Tobit
0.0
0.5
1.0
1.5
2.0
95% Confidence Interval

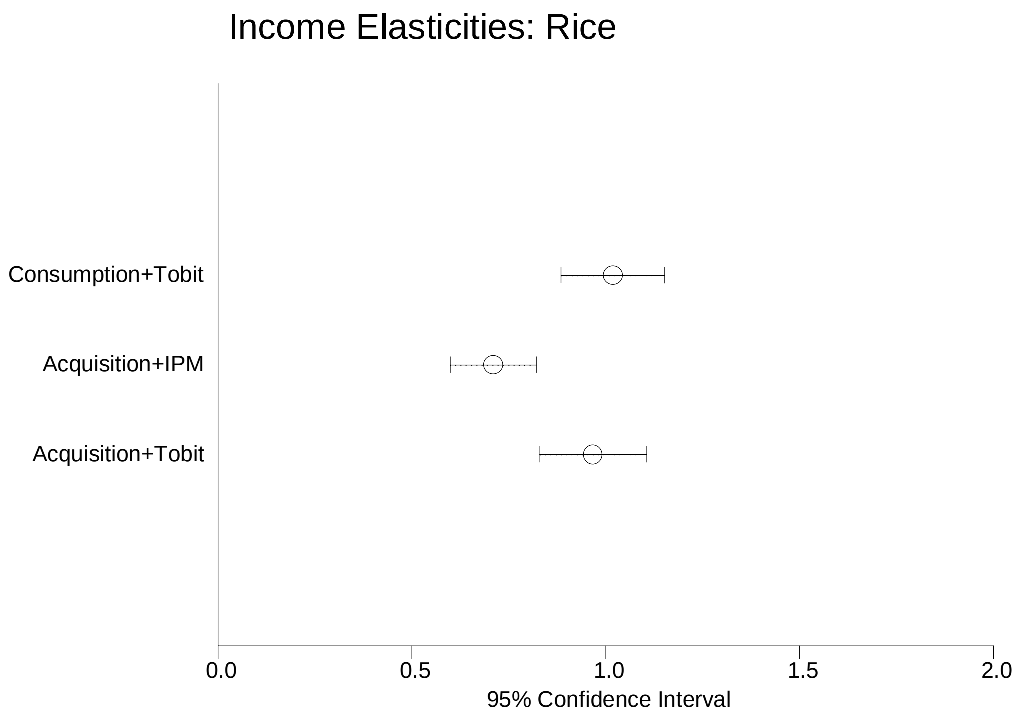
Income Elasticities: Rice
Consumption+Tobit
Acquisition+IPM
Acquisition+Tobit
0.0
0.5
1.0
1.5
2.0
95% Confidence Interval