Preliminary Draft: February 12, 2009

Constructing a Consumption Aggregate for the Purpose of Welfare Analysis: Principles, Issues and Recommendations Arising from the Case of Brazil

Peter Lanjouw (World Bank)[1]

Introduction

The analysis of poverty and inequality is a well-established field of research that serves as an important input into policy-making. While the very meaning of poverty and inequality remains the subject of debate, and differences of opinion persist in how to best study such themes, certain basic steps in the empirical analysis of household welfare have become reasonably standard.

While the ultimate shape and scope of the analysis can vary dramatically, a nearly universal requirement for any empirical study of wellbeing is that individuals (or households) must be ranked on the basis of one or more indicators of living standards - usually income or consumption expenditures (but sometimes other indicators such as nutritional status, access to basic services, or even a composite measure). The choice and definition of an appropriate indicator might seem a fairly straightforward task. However, a person embarking on such an exercise is quickly confronted by a whole range of issues, many of which will require some kind of decision-making and on which guidance, in the form of best-practice conventions or theoretically derived results, is still rather scarce. As the basic welfare indicator serves as the foundation upon which most of the subsequent, detailed, analysis of welfare is based, it is important to select an indicator that can command broad endorsement and that will hopefully not require substantive revision.

[1] This paper has been prepared for the OECD/Univeristy of Maryland Conference entitled "Measuring Poverty, Income Inequality and Social Exclusion: Lessons from Europe", Paris, March 16/17, 2009. I am grateful to Fang Lai, Phillippe Leite and Emmanuel Skoufias for their many contributions to this project, and to Kathleen Beegle, Gero Carletto, Francisco Ferreira, Jed Friedman, Jesko Hentschel, and Kinnon Scott, for helpful discussions. The views in this paper are my own and should not be taken to reflect those of the World Bank or affiliated institutions.

In this note we are concerned with deriving a welfare indicator for households that captures the economic dimensions of wellbeing. We focus specifically on consumption data. There are both conceptual and pragmatic reasons why consumption expenditures available from household surveys might be preferred for the purpose of poverty and inequality analysis to an indicator such as household income. It is argued, for example, that consumption represents the *achievement* of a particular welfare level, while income reflects *opportunity* to achieve a certain level. Consumption expenditures reflect not only what a household is able to command based on its current income, but also whether that household can access credit markets or household savings at times when current incomes are low or even negative (due perhaps to seasonal variation or a harvest failure). For this reason consumption is often argued to provide a better picture of a household's longer run standard of living than a measure of current income. Further, calculating consumption expenditures is often easier than calculating household incomes, particularly for the poor. While poor households are probably purchasing and consuming only a relatively narrow range of goods and services, their total income may derive from multiple different activities with strong seasonal variation and with associated costs that are not always easily assigned. Getting an accurate net income figure for such households can be frustratingly difficult. Where consumption information is collected, an additional advantage is that not only are consumption expenditures available, but a poverty line can often be derived from the same survey, thereby strengthening the link between the welfare indicator used in the analysis and the threshold determined to separate the poor from the non-poor.

This note examines the detailed information on household consumption expenditures that has been collected in Brazil in the most recent *Pesquisa de Orçamentos Familiares* (POF) of 2002/3 fielded by the Instituto Brasileiro de Geografia e Estatística (IBGE). The goal is to consider how best to construct a consumption aggregate from these data for the purpose of analyzing poverty and inequality. The note first considers the available building blocks for producing a consumption aggregate in the POF, and reviews some of the principles and issues that can guide decisions as to whether, and provides a flavor of the details involved in deciding how specific items should be included in the aggregate. The paper points to a few limitations in the way certain components of

---

All errors are my own.

consumption are collected, and raises some concerns with respect to the observed presence of extreme values in the data set. Such outliers can have a very strong impact on poverty and inequality measures and to the extent that one might have doubts that the outliers are actually conveying accurate information, there are arguments in favor of trimming these extreme observations from the data. The note explores some of these issues via simple sensitivity analysis and proposes a moderate degree of trimming of the data prior to subsequent analysis.

Constructing the Benchmark Consumption Aggregate

Brazil's POF is a nationally representative expenditure survey that was fielded by IBGE in 2002/3. The principal objective of the POF is to provide the detailed information on household expenditures required to produce cost of living indices such as the Consumer Price Index. In previous waves of the POF (1987/8 and 1995/6) the survey was fielded only in 9 metropolitan areas (plus Goiânia and Distrito Federal). However, in 2002/3, IBGE extended the sample to the country as a whole - it is representative at the state level for urban and all-state totals; for rural areas the sample is designed be representative at the region level only. This latter feature makes the POF extremely interesting for the purpose of welfare analysis, because it marks the first time in several decades that a nationally representative survey yielding consumption information is available in Brazil.[2] The overall sample size of the POF 2002/3 is just under 50,000 households.

The process of creating a consumption aggregate is guided by a number of considerations. In this note we provide an overview of some of the issues that arise and the principles that can be applied. Our treatment is informal and far from exhaustive. A more complete reference document on this whole topic can be found in Deaton and Zaidi (2002). An important initial consideration is that, as our measure is supposed to proxy welfare, there is an interest in having as *comprehensive* a measure of consumption as possible. This is because a consumption measure that is narrowly

[2] Analysis of wellbeing in Brazil has tended to employ income data from the *Pesquisa Nacional por Amostro de Domicilios* (PNAD). Ferreira, Lanjouw and Neri (2003) suggest that the PNAD might not be suitable for at least some aspects of welfare analysis in Brazil. This appears to be especially the case for analyzing rural welfare (see also Elbers, Lanjouw, Lanjouw and Leite ,2001). One other survey fielded by IBGE in recent years has included consumption information – the *Pesquisa Sobre Padores de Vida* (PPV) of 1996. However, this survey had a relatively abbreviated consumption questionnaire, was fairly small in sample size (about 5,000 households) and covered only the Northeast and Southeast of the country (see Ferreira, Lanjouw and Neri, 2003, for further discussion).

defined would imply, when comparing welfare levels of households or individuals, that omitted components do not contribute in any way to welfare. Or alternatively, that while certain consumption components are omitted and are important to welfare, they would be distributed across members of the population in such a way that they would not affect rankings were they to be included. The extent to which these implicit assumptions seem reasonable varies with the specific components of consumption in question, but as a general rule one would want to include as many components of consumption as is feasible.

However, it is often not possible to include all components of consumption in an equally straightforward manner. For several components it becomes necessary to introduce additional assumptions in order to be able to add these to the consumption aggregate. This can quickly add to the complexity of the exercise and can threaten the *transparency* of the process.

Some additional complexity need not in and of itself justify abandoning the exercise, except insofar as the *credibility* of the entire undertaking is thrown into doubt. The value of the entire enterprise of welfare analysis rests crucially on the degree to which the conclusions are widely endorsed. As subsequent strategies aimed at income redistribution or poverty alleviation all rest on the credibility of the underlying consumption aggregate, it is vitally important not to sacrifice credibility in the process of adding some particularly tricky consumption component to the consumption aggregate.

In deriving the "preferred" consumption measure from the POF data, the multiple objectives of comprehensiveness, transparency and credibility were retained as central focus. The exercise was approached in a series of steps. Each step was approached in a tentative manner, as one would not want to force the components into the consumption aggregate at the cost of unacceptably speculative assumptions or convoluted argumentation.

Constructing the POF consumption aggregate

The POF survey collects information on household acquisitions of goods (purchased for own use or for other households, received as gift, and self-produced) in the previous periods of 7, 30 and 90 days, and 12 months. The 7 days recall includes acquisitions of food, both inside and outside the home, and transport expenses. The 30 days recall was applied to a range of nonfood consumption goods, such as pharmaceutical products, and also leisure and entertainment. The 90

days reference period include clothing and a variety of services, among others. The expenses made throughout the 12 month period comprise the acquisition of durable goods (like houses, cars and electronic appliances). The overall consumption measure can be aggregated up from 10 broad categories of items: "Food consumption (including consumption inside and outside the home)"; "Housing"; Health"; "Schooling/Education"; "Transport"; "Clothes"; "Culture/Leisure"; "Personal Services"; "Hygiene and personal care"; and "Others".

Within these broad categories, there are specific items which deserve special consideration. In general, care should be taken to avoid including in the consumption aggregate the following: a) "lumpy" items purchased sporadically; b) items that serve as inputs into production, or investments; c) items with low elasticity with respect to total expenditure; d) items acquired for other households. We discuss briefly below the criteria for not including such items in a consumption aggregate that is intended for welfare analysis.

a) *"Lumpy" and infrequent acquisitions* - Consumer durable purchases are typically large expenditures that occur very infrequently. A classic example is the purchase of a car or motorcycle. A particular household is likely to purchase a car only once every number of years. With a 12 month recall period, there will be a certain subset of households in the data who do indeed report purchasing a car. They will report spending a considerable sum of money for this item. Other households in the dataset will, in fact, own a car but will have purchased it in some preceding period, and will thus report zero expenditure in a car. Attributing a consumption value of zero to households that own but did not purchase a car in the specific recall period, will understate their welfare because they will in fact be consuming the services of a car. Attributing the purchase value of the car to those households in the data that happened to purchase a car during the reference period will overstate their welfare because they will not be consuming all of the services provided by a car in this one-year reference period. The car's services will be consumed over a period of several years. The attributes of a consumer durable imply that it is unappealing to simply add expenditures over the reference period directly to the consumption aggregate. Where possible, a flow of consumption from consumer durables can be added to the consumption aggregate, imputed

from the available information on ownership, age and replacement value of consumer durables. Deaton and Zaidi (2002) provide a good discussion of the available methods. In POF, although there is a section on the inventory of durable stocks for households owning goods in a longer period span, the questionnaire does not include information on value (either original purchase value or current replacement value), so it is not possible to calculate the flow of services from durables.

b) *Items that serve as inputs into production, or investments* - One key concern throughout the process must be to avoid treating spending for production or investment purposes, as consumption. If one includes expenditures on inputs into household production, and the income from household production is in turn devoted (at least in part) to consumption expenditure, then double counting occurs, and the consumption aggregate is overstating the actual welfare levels achieved by the household. In most circumstances, the distinction between productive inputs and consumption is rather obvious. For example, it is clear that fertilizer expenditures should not be reflected in the consumption aggregate for farming households. In some cases the distinction is less clear (see below).

c) Items with low elasticity with respect to total expenditure - In some cases, it is difficult to determine the effect on welfare of expenditure in items like health products and services. The analysis of whether to include health expenditures warrants an assessment of the elasticity of health expenses with respect to total expenditure. For instance, it is difficult to measure the extent to which health expenditures increase welfare, since it is not possible to measure the loss of welfare from illness and the increase in welfare from its alleviation. Including only the expenditure is incorrect, though excluding health expenditures altogether means that one may miss important differences between two people, both of whom are sick, but only one of whom is paying for treatment. Moreover, there are other considerations related to whether health expenditures may also be discretionary and welfare enhancing, but it is difficult to discriminate "necessary" from "unnecessary" expenditures. Deaton and Zaidi (2002) recommend analyzing the elasticity of the expenditure in health items with respect to total expenditure. The higher the elasticity, the stronger the case for inclusion. We analyze the elasticity of health and education expenditures in POF when explaining the components of

the consumption aggregate.

d) *Items acquired for other households* - Goods acquired for gifts to other households should be excluded from the consumption aggregate, since their inclusion would involve double-counting if, as one would expect, the transfers show up in the consumption of other households. Therefore, it is recommended to include only the goods acquired as a gift from others, which increase the well-being of that household, but not the expenses made in that household for increasing consumption of other households.

Food consumption

The food consumption module in the POF compises a diary left with each household for a period of seven days. Households are requested to provide a detailed description for each day of the week, of all items of food that are purchased or otherwise acquired. The questionnaire requires the household's key informant to note each specific food item acquired, the quantity obtained (including specifying the unit of measurement in which the quantity is recorded), the value of the acquisition, the location of acquisition, and the form of acquisition (e.g. purchased, received as gift, produced by the household itself, etc.). In this way very detailed information is collected from each household about its food acquisition during a period of seven days.

The food component of the consumption aggregate comprises the value of expenditures and acquisitions of food items for consumption both inside and outside the home. Aggregating across all items, over the whole week, yields a measure of household weekly food acquisition. Multiplying this by the number of weeks in a month or in a year yields a measure of monthly or annual food "expenditures". While it may not be strictly the case that all food acquired in a given week is consumed that week, the general assumption is that at the monthly or annual level, total food expenditures indicate the value of total food consumed by the household. This procedure provides the first component of the overall POF-based household consumption

aggregate.

Figure 1 relates household level food shares in one state, Ceara, to total consumption expenditures (based on the total expenditure figure that has been produced by IBGE to accompany the POF data). A non-parametric regression curve traces out the Engel-Curve summarizing the relationship, on average, between total household spending and the share of spending devoted to food. Contrary to "Engel's Law" which states that the food share declines with total spending, there is clear evidence here that the Engel curve in Ceara first rises before it starts to decline.[3] Why is this happening? Consider a simple example. Suppose that households in Brazil actually purchase food on a fortnightly basis (with households uniformly spread across weeks) but that the recording period in the POF is one week. Let F = *average* weekly food expenditure and Y = average weekly total expenditure. Suppose further that non-food expenditure = (Y-F) is correctly measured due to a longer recording period. If *n* households are sampled from a group with identical {F,Y}, *n*/2 will have food purchases of f=2F, with y = Y+F, and *n*/2 will have f=0, with y=Y-F. The mean food expenditure is correctly estimated as (1/2)*2F + (1/2)*0 = F. However, the distribution of consumption and income, and hence the Engel curve, are incorrectly estimated. The true food share is F/Y while the empirical food shares are zero, with probability one half, and 2F/(Y+F) with probability one half. The foodshare is increasing in y.

Figure 1: Share of Food Expenditure vs. Total expenditure: Ceara

[3] This relationship is observed not only in Ceara. A similar graph produced at the all Brazil level reveals a similar pattern. Thomas (1986) documents this phenomenon in many different datasets. It should be noted, however, that the range over which the Engel Curve is rising in these data is higher than one normally sees.

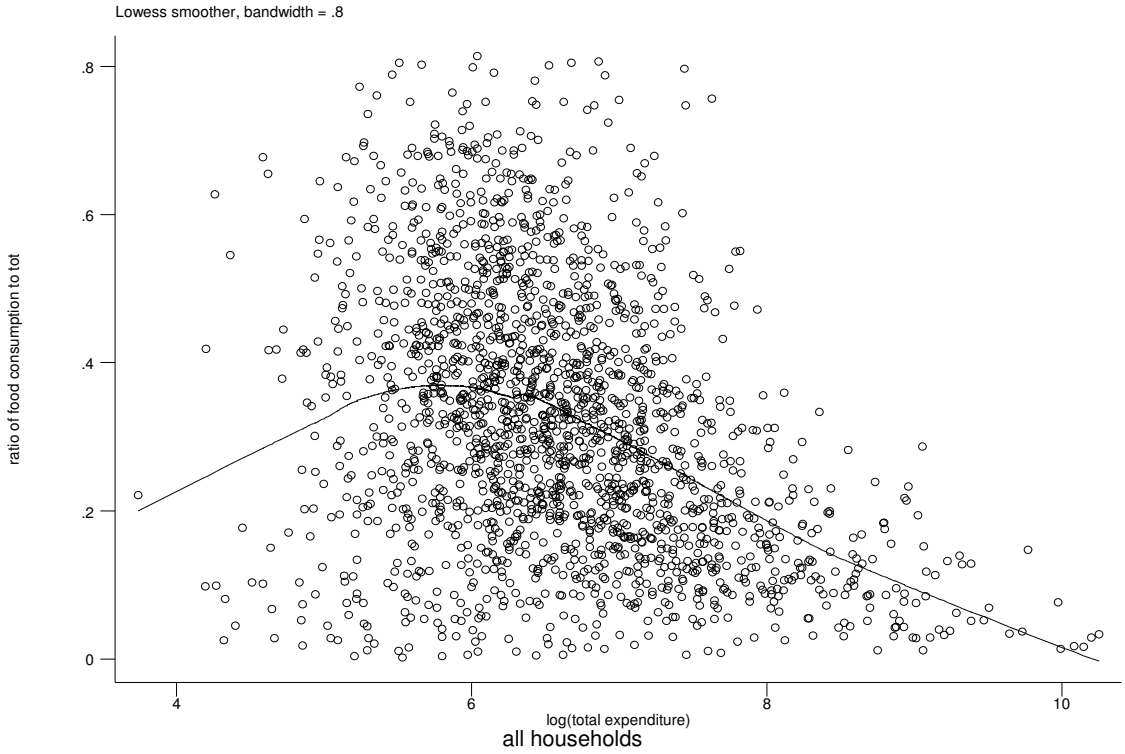

From Figure 1 it appears that one of the reasons why the Engel curve first rises is due to the presence of a number of very low food shares amongst households with low total expenditures. Indeed, in the raw data of the POF 2002/3 there are 1636 households (3.4% of the total) with no reports on food consumption whatsoever. It is difficult to imagine that those observations provide an accurate depiction of consumption patterns – that the very poor should be devoting all, or the bulk, of their budget to non-food items. Our conjecture here is that for those households food expenditures may simply have been inaccurately collected due to an inappropriate reference period in the consumption questionnaire.[4]

It is not clear what should be done about this. The presence of noise in the food consumption measure need not affect our calculation of average food consumption in the state as a whole. However, given that welfare analysis focusing on poverty or inequality is particularly

[4] Anecdotal evidence suggests that indeed, in Brazil, it is not uncommon for households to purchase many of their food items in bulk. Such behaviour is understandable in a setting where historically, high rates of inflation provided households with a real incentive to immediately convert their monthly salary into purchases of real goods and

interested in the tails of the consumption distribution the discussion here suggests that at least some of those who would be counted as the poor might be measured with error. To the extent that this problem occurs only with respect to food consumption, one might hope that for those households with significant non-food expenditures, their overall ranking in the welfare distribution may not be affected too badly by this problem. As a result this issue may be of less concern when trying to identify the rich (in an analysis of inequality, for example). However, amongst those with low incomes, for whom food expenditures are typically particularly important, the presence of noise in the food consumption data is likely to lead to an over-estimate of overall poverty and to make less sharp the distinction between the "poor" and the "non-poor" in terms of household and individual characteristics.[5]

To address this problem in the process of constructing the POF-based consumption aggregate, the following procedure was implemented. Food expenditure of the households with missing reports was predicted, based on a model for food expenditure as a function of a set of households' housing and demographic characteristics and area of residence, estimated on the subset of households with non-zero food expenditures. The parameters estimates derived from this model were used to impute food expenditures for all households that recorded zero food spending, (except for 97 households with per capita income below the political indigence line of R$50, who were expected to have no reports on food consumption because of real difficulties, rather than because of the short recall period).[6] Testing for the sensitivity of poverty and inequality measures to this imputation in the final consumption aggregate revealed that measures showed little sensitivity to the imputation in food expenditure.

Housing

The housing component of the consumption aggregate comprises rents; basic services;

---

services.

[5] Ravallion (1988) and Deaton and Zaidi (2002) discuss issues surrounding poverty measurement in the presence of noisy data in greater detail. See also Lanjouw and Lanjouw (2001).

[6] We also tried this procedure with Propensity Score Matching and found very similar results for the imputation.

small-scale home renovations; furniture and household items; appliances and appliance repairs; and cleaning materials. For those households that are renting their home, monthly rental payments can be straightforwardly included as a measure of the consumption of services that derive from housing. Households owning their dwellings do not pay rent, but are clearly consuming housing services, and so it is important to also include an expenditure figure for such households. In the POF, owner-occupiers are asked to provide an estimate of the rental value of their home. In a setting where there is an active rental market, owner-occupiers are likely to be well informed about the value of their home and the kind of rent they would have to pay for a home with similar quality and location attributes. Upon close examination and comparison also against predicted rental payments, it was found that household responses to such hypothetical questions in the POF were generally quite satisfactory and could be used in the consumption aggregate. In other settings it is sometimes necessary to consider imputation models (as applied with Food, above) in order to include housing consumption in the overall aggregate.

Expenditures on basic services (water, sewage, etc.) were included in the POF consumption aggregate. This is not an uncontroversial decision. Although such expenditures represent a large share of total expenditure for only a few households, Deaton and Zaidi (2002) generally recommend against the incorporation of expenditures on publicly provided services in the consumption aggregate. This is because finding the proper set of prices with which to value these goods is difficult. Including expenditures on networked water and sanitation, for example, while not being able to properly take account of the fact that some households are not connected to a water network at all, that some households do not receive bills although they are connected, and that some households receive only sporadic supply of water and supplement their publicly provided water with purchases from private vendors, could introduce important biases in rankings of households.[7] If there is any reason to think that expenditures on networked water, electricity and gas are only weakly linked to the welfare that is associated with the actual consumption of those services the general recommendation would be to exclude these expenditures from the consumption aggregate. Other services, such as internet access, telephone

expenditures and television subscriptions, are more straightforwardly added to the consumption aggregate.

Expenditures on home renovations in the consumption aggregate include the more frequent expenditures on housing maintenance, such as: upkeep, gardening and home repairs. These were collected as expenditures within the 90 days reference period. The POF survey also collected expenditures on renovations over the 12 month reference period. In this case, the renovations are less frequent and lumpy since they include spending on construction activities. These last expenditures were not included as part of the consumption aggregate. As commented in point a) above, this type of occasional and high expenditures can introduce a wedge between the welfare levels of households which incurred in this type of expenditure in the reference period and the households who spent on them in a previous period. The same consideration was taken for deciding on the inclusion of durable items as furniture, appliances (fridge, televisions) as well as repair of these appliances. Each item was scrutinized in turn in order to decide if the purchase of a given durable good was to be considered an occasional and lumpy expenditure. Only those items which were considered as frequent and less lumpy were included in the POF consumption aggregate.

Health and education

If one were to include expenditure on health then one should also take into consideration the implicit loss of welfare due to illness. This is very difficult to do. However, some items related to prevention and care can be considered as more discretionary and welfare enhancing. These can more reasonably be included in the consumption aggregate. The decision to include or exclude these expenditures, according to Deaton and Zaidi (2002), should be based on the analysis of the income elasticity of the health expenditures. If the income elasticity is high, then overall rankings of households may well change depending on whether such expenditures are included. When the elasticity is low, rankings are likely to be more robust. Deaton and Zaidi

[7] Hentschel and Lanjouw (1996) discuss these issues in some detail.

(2002) show that in developing countries, this elasticity is typically quite low (varying between 0.74 and 0.86), which suggests that non-inclusion in the consumption aggregate has some justification.

In a similar way, elasticities for education can be computed. Here too there can be reasons to be concerned about inclusion. On the one hand, education directly adds to welfare. On the other, education is an investment. Inclusion of education expenditures can clearly introduce a wedge in welfare levels between households without children going to school and those with children at school. Again, the higher the elasticity, the stronger the case for inclusion in the consumption aggregate.

The results presented in Tables 1, 2 and 3 compare the elasticity of health and education with respect to the total expenditures and to family income for Brazil using the POF. Table 1 indicates that the elasticity of education expenditures is larger than the elasticity of health expenditures, providing some justification to the inclusion of all of the education expenditures, but not of health. The elasticity of health is 0.97, which is lower than the elasticity of education expenditures, although greater than the elasticity found in the countries analyzed by Deaton and Zaidi (2002). The elasticity of the health and education expenditures was also estimated by income deciles (Tables 2 and 3). We can see that the elasticities are always higher for education expenditures than they are for health expenditures. In the case of health, the elasticity is highest in deciles four and six. For the bottom deciles, this elasticity is lower.

In light of these results – elasticities in Brazil that were not particularly low for health expenditures and rather high for education expenditures, the following procedure was adopted. Expenditures in health and dental insurance plans were added to the consumption aggregate because they provide insurance - which can be related to a higher level of welfare, and there is no indication of decrease in welfare from illness in insurance plans. Moreover, these expenses represent a fairly sizeable component of total expenditures incurred by the families. Other types

of health expenditures, such as the purchase of pharmaceutical products and medical attention, were excluded, since in this case it is not possible to capture the welfare loss from the diseases they are supposed to alleviate. Expenditures in education were included, since expenditures in private school fees can be directly related to a higher level of welfare of households paying for educational services. Although education can also be considered an investment instead of consumption, the inclusion of education expenditures in the consumption aggregate is unlikely to lead to double counting as the returns from this particular investment will probably not be reflected in current consumption levels. It is common to treat education as a consumption item, but it is obviously a matter of judgment.

Table 1 Elasticities of health and education expenditures

Variable	Elasticity	Standard deviation	t	P Value
Health * Income	0,81	0,0136	59,64	<.0001
Health * Expenditure	0,97	0,0100	69,80	<.0001
Education * Income	1,13	0,0200	54,88	<.0001
Education * Expenditure	1,30	0,0200	62,59	<.0001

Source: 2002-03 POF.

Note: The sample design of the survey was considered for the calculation.

Table 2: Elasticities of health expenditures by deciles of income distribution

Income decile	Elasticity	Standard deviation	T	P Value	Observations
1	0,037	0,061	0,60	0,548	2762
2	0,567	0,222	2,56	0,011	3116
3	0,550	0,276	2,00	0,046	3421
4	**1,589**	0,324	4,91	0,000	3655
5	0,572	0,316	1,81	0,071	3782
6	**1,214**	0,289	4,20	0,000	3968
7	0,953	0,248	3,85	0,000	4122
8	0,921	0,190	4,86	0,000	4338
9	0,964	0,123	7,86	0,000	4502
10	0,655	0,033	19,64	0,000	4633

Source: 2002-03 POF.

Note: The sample design of the survey was considered for the calculation.

Table 3: Elasticities of education expenditures by deciles of income distribution

Income decile	Elasticity	Standard deviation	T	P Value	Observations
1º	0,027	0,076	0,36	0,718	2067
2º	0,830	0,278	2,98	0,003	2190
3º	0,730	0,357	2,05	0,041	2397
4º	1,018	0,454	2,24	0,025	2491
5º	1,053	0,424	2,49	0,013	2755
6º	0,907	0,395	2,29	0,022	2852
7º	1,688	0,355	4,75	0,000	3070
8º	1,567	0,289	5,43	0,000	3283
9º	1,382	0,190	7,29	0,000	3619
10º	0,835	0,053	15,79	0,000	3954

Source: 2002-03 POF.

Note: The sample design of the survey was considered for the calculation

Transport services

Expenses in transport services were included as part of the consumption aggregate. Although some of these expenditures can also be considered as “regrettable necessities” for

getting to the work place, it was not possible to distinguish them from transportation expenses for other purposes.

Clothing, culture and leisure, personal services and personal hygiene and care

These components of the consumption aggregate comprise all types of expenditures in clothing, leisure (tickets to cinema, etc), personal services (haircuts, beauty, etc) and personal care; likely to increase welfare of the households without introducing biases in the comparability of households' welfare levels. Notwithstanding the fact that expenditures in clothing and shoes can be considered infrequent purchases, the value of these purchases is rather modest, so they were included in the aggregate.

Considerations for other expenditures

The remaining components of the consumption aggregate comprise professional services (such as notaries, lawyers); expenditures in ceremonies, celebrations and anniversaries (that are collected for the 12 month reference period); and expenses related to taxes, contributions, banking fees, among others. The procedure followed was to include all such items except for occasional expenditures (such as occasional ceremonies). As with consumer durables these are infrequent expenditures that can become very costly and ideally we would like to have some smoothed value rather than actual, total expenditure on the event. The sole exception was made with respect to birthday parties and wedding anniversaries – events that occur on an annual basis. For such items the 12 month reference period is the appropriate one and one could thus justify including these items in the consumption aggregate.

Following Deaton and Zaidi (2002) expenditures on taxes, contributions and levies are not part of consumption, but a deduction from income, and should therefore not be included in the consumption aggregate. Consequently the POF consumption aggregate does not include such payments. Deaton and Zaidi (2002) recommend including property taxes only when there is evidence that they could be linked to the provision of a specific service to the households. In the

Brazil case, there were no clear grounds for relating property taxes (IPTU and ITR)[8] to specific services and a better level of well-being. However payments that could be linked to service provision, like insurance payments were included. Taxes related to the acquisition of goods already excluded (e.g. purchases of cars) were excluded as well. Expenses related to financial transactions, such as the paying off of debts were not included as part of the aggregate.

As recommended by Deaton and Zaidi (2002) expenditures on gifts and transfers were excluded from the aggregate. Including them would involve double counting if the transfers show up in the consumption of other households. Large expenditures that may be considered investments, such as the purchase of real estate, gold bars, etc. were excluded from the consumption aggregate. They can also introduce bias in the comparison with households already owning these assets.

Trimming of Overall Consumption

Our examination of the components of the POF consumption questionnaire leads us to suggest that the POF consumption data may suffer more than perhaps other similar datasets for other countries, from measurement error. In the case of food expenditures we have suggested that there may be grounds for concern associated with an inappropriate recall period of one week – which resulted in a large percentage of zero, or very low, expenditures in the data. Further, we have noted above that we are unable to impute a stream of consumption services from the very long list of consumer durables included in the data set. We are compelled in this case to include actual expenditures on such items (abstracting away from those highly infrequent and costly items that we suggest should be excluded altogether). This means that, once again, there will be many households which record zero expenditures on specific items and other households that record expenditures that are probably in excess of the value of the stream of consumption that they derive from the item during the reference period. The overall effect, again, is the same as if we had measurement error in the data.

[8] IPTU: Imposto sobre a Propriedade Predial e Territorial Urbana. ITR: Imposto Territorial Rural.

The effect of measurement error in the analysis of welfare can be quite significant. As has been shown by Ravallion (1988), Lanjouw and Lanjouw (2001) and Deaton and Zaidi (2002), measured poverty is likely to be higher than it should in the presence of measurement error. Similarly, measured inequality will be biased upward.

These considerations often lead to the suggestion that some protocol for trimming extreme values from the consumption aggregate be considered prior to utilization of the consumption aggregate in applied welfare analysis. To assess the impact of trimming we conduct sensitivity analysis with reference to the most "naïve" possible consumption aggregate that simply brings together all of the consumption components without following any of the specific recommendations outlined above. Table 4 shows that sensitivity of measured inequality at the level of each Brazilian state is indeed quite significant. At the level of Brazil as a whole, untrimmed per capita consumption inequality yields a Gini coefficient of 0.559.[9] When half of a percentage point of all observations is dropped from both the bottom and from the top of the per capita consumption distribution the Gini declines to 0.531. A more draconian trimming protocol, cutting 2.5 percentage points from both ends of the distribution yields a Gini of 0.507.

At the all-Brazil level our measure of inequality is likely to be overstated due to the fact that spatial price differences have not been accommodated.[10] However, Table 4 indicates that high inequality occurs also at the level of each state, but that, again, measured inequality is quite sensitive to the question of whether extreme value observations are to be trimmed or not. In Ceara, for example, measured inequality is even more sensitive to trimming than at the all-Brazil level. The Gini coefficient in Ceara declines from 0.577 to 0.505 when 2.5 percentage points of observations are dropped from both tails of the income distribution.

What specific trimming protocol to adopt is not an easy question to answer. It is likely that a consumption aggregate constructed following the recommendations outlined above would be less sensitive to trimming than Table 4 suggests. After all, the consumption aggregate prescribed above incorporates imputed food consumption for some households and excludes some of the larger

[9] This compares with a Gini of 0.507 reported in World Bank (2006) based on the consumption aggregate definition outlined above.

[10] World Bank (2006) analyzes poverty and inequality in Brazil on the basis of the POF survey and includes as well a correction for spatial price variation. The overall Gini reported in World Bank (2006) following adjustment for spatial price variation is 0.479-0.481, depending on the specific price index used.

durable good expenditures and this would presumably remove some of the extreme values at the top end of the consumption distribution. It remains, however, that measured inequality, and to some extent poverty, will be quite sensitive to decisions regarding the definition of the consumption aggregate, to the adoption of trimming protocols, as well as to other adjustments such as those for spatial price variation.

Conclusion and Recommendations for the Future

In this note we have briefly scrutinized the Brazil POF 2002/3 questionnaire and dataset with a view towards identifying and discussing some of the principles and issues associated with construction of a consumption aggregate. We have emphasized that a consumption aggregate compiled for the purpose of welfare analysis (poverty and inequality, for example) may require different treatment than one compiled for some other purpose. We have emphasized that the objective here is to be able to produce reliable and credible comparisons of welfare across households and individuals. We have suggested that such a measure of consumption would exclude expenditures that are better seen as investments or inputs into production. Moreover we have described the desirability of capturing not simply expenditure levels, but rather a monetary value of the stream of services that is enjoyed by an individual or household from the consumption of a particular good or service. We have underscored that there can be a tension between the level of comprehensiveness of a consumption aggregate on the one hand (the more comprehensive the better – in principle) and the transparency and interpretation of the aggregate, on the other.

Our examination has led to various suggestions for treatment of specific consumption items. We have also indicated that there may be reasons to worry about measurement error in the final consumption aggregate that cannot be avoided. At least some of that measurement error may be associated with certain design features of the POF questionnaire, and in that light we can conclude with two principal recommendations for revision of the questionnaire fur future reference.

1. Recall period for food expenditures: we noted that the POF data appears to include a relatively large proportion of households that report zero, or very low, spending on food based on the recall period of one week in the questionnaire. It is unlikely that such households are, in fact, not consuming food. Rather it seems possible that many households in Brazil purchase food on a fortnightly, or even monthly, basis. Future experimentation with the design of the POF consumption questionnaire may wish to consider alternative

recall periods to the one-week recall. Experiments along such lines could reveal whether the conjecture above has any basis in fact.

2. Consumer Durables Consumption: We have noted above that the POF consumption questionnaire includes a very exhaustive listing of expenditures on a variety of infrequently purchased goods and services – notably consumer durables. However, information of stocks of durables owned is far less exhaustive and moreover, there is no information collected on households' estimation of the current value of the durables that they own. The absence of such information makes it very difficult to include in the consumption aggregate a calculation of the stream of services consumed by households of all durables that they own. In a relatively rich country such as Brazil, the ownership of consumer durables is likely to be quite widespread and as such there is good reason to expend the additional effort to produce a consumption aggregate that reflects well the contribution of that ownership to wellbeing.

In sum, the construction of a consumption aggregate is part science, part art. In most practical settings there will be many issues and difficulties encountered during the process of producing a consumption aggregate. Many of these cannot be resolved conclusively and to everyone's complete satisfaction. Judgment calls are required, and it is likely that not everyone will agree to the choices and judgments made. It is therefore very important that the exercise be approached systematically and that there is a clear documentation of how each step in the process has been followed. Welfare analysis that builds on the consumption aggregate can have far-reaching implications for policy debate and design. Sensitivity analysis that gauges the degree to which conclusions are robust to alternative definitions of the consumption aggregate should therefore be undertaken and reported. Comparisons across countries, time periods, and settings, in which it has not been possible to ascertain that the definition of consumption is identical should beware that at least some of the differences observed may be driven by non-comparability of the underlying welfare definitions.

Table 4: Sensitivity of Measured Inequality to Trimming of Top and Bottom Extreme Values

State	Untrimmed	Trimmed (both sides)		
	Gini	0.5 percent	2.5 percent	5.0 percent
Rondonia	0.513	0.476	0.455	0.428
Acre	0.568	0.537	0.518	0.489
Amazonas	0.550	0.509	0.481	0.450
Roraima	0.499	0.474	0.454	0.430
Para	0.471	0.445	0.425	0.402
Amapa	0.481	0.464	0.450	0.432
Tocantins	0.555	0.504	0.478	0.450
Maranhao	0.477	0.456	0.433	0.405
Piaui	0.534	0.498	0.470	0.438
Ceara	0.577	0.535	0.505	0.467
Rio Grande do Norte	0.567	0.528	0.497	0.461
Paraiba	0.548	0.518	0.491	0.451
Pernambuco	0.533	0.504	0.480	0.449
Alagoas	0.591	0.553	0.520	0.481
Sergipe	0.550	0.527	0.506	0.475
Bahia	0.557	0.520	0.498	0.468
Minas Gerais	0.528	0.503	0.479	0.451
Espiritu Santo	0.552	0.529	0.507	0.481
Rio de Janeiro	0.572	0.550	0.528	0.500
Sao Paulo	0.501	0.476	0.454	0.432
Parana	0.527	0.501	0.481	0.452
Santa Caterina	0.480	0.455	0.436	0.408
Rio Grande do Sul	0.522	0.490	0.468	0.440
Mato Grosso do Sul	0.519	0.491	0.466	0.436
Mato Grosso	0.523	0.490	0.461	0.428
Goias	0.504	0.474	0.453	0.427
Distrito Federal	0.572	0.541	0.521	0.497
All Brazil	0.559	0.531	0.507	0.480

References:

Deaton, A. and Zaidi, S. (2002) 'Guidelines for Constructing Consumption Aggregates for Welfare Analysis', Living Standards Measurement Survey Working Paper 135, The World Bank, Washington D.C.

Elbers, C., Lanjouw, J.O., Lanjouw, P. and Leite, P.G. (2001) *Poverty and Inequality in Brazil: New Estimates from Combined PPV-PNAD Data*, mimeo, DECRG, the World Bank.

Ferreira, F., Lanjouw, P. and Neri, M. (2003) *A new Poverty Profile for Brazil using PPV, PNAD and Census Data*, Revista Brasileira de Economia, Vol 57, No. 1.

Hentschel, J. and Lanjouw, P. (1996): *Constructing an Indicator of Consumption for the Analysis of Poverty: Principles and Illustrations with Reference to Ecuador*, Living Standard Measurment Survey Working Paper No. 124., The World Bank.

Lanjouw, J. O. and Lanjouw, P. (2001): *How to Compare Apples and Oranges: Poverty Measurement Based on Different Definitions of Consumption*, Review of Income and Wealth, 47(1), 25-42..

Ravallion, M. (1998) *Expected Poverty Under Risk-Induced Welfare Variability* Economic Journal, 98, 1171-82.

Thomas, D. (1986) The Food Share as a Welfare Measure, Ph.D. Dissertation, Princeton University, 1986.

World Bank (2006) 'Brazil: Measuring Poverty Using Household Consumption' Report No. 36358-BR, the World Bank, Washington D.C.