

UC Davis

Agriculture and Resource Economics Working Papers

Title

Improving Estimates of Inequality and Poverty From Urban China's Household Income and Expenditure Survey

Permalink

<https://escholarship.org/uc/item/8q20g554>

Authors

Gibson, John
Huang, Jikun
Rozelle, Scott D.

Publication Date

2001-04-01

**Department of Agricultural and Resource Economics
University of California Davis**

Improving Estimates of Inequality and Poverty From Urban China's Household Income and Expenditure Survey

by

John Gibson, Jikun Huang and Scott Rozelle

Working Paper No. 01-018

April, 2001



Copyright © 2001 John Gibson, Jikun Huang and Scott Rozelle
All Rights Reserved. Readers May Make Verbatim Copies Of This Document For Non-Commercial Purposes By
Any Means, Provided That This Copyright Notice Appears On All Such Copies.

**California Agricultural Experiment Station
Giannini Foundation for Agricultural Economics**

April 2001

**Improving Estimates of Inequality and Poverty From
Urban China's Household Income and Expenditure Survey**

John Gibson*
University of Waikato

Jikun Huang
Center for Chinese Agricultural Policy
Institute of Geographical Sciences and Natural Resources Research,
Chinese Academy of Sciences

Scott Rozelle
University of California, Davis
and
Giannini Foundation

Acknowledgements

We are grateful to the World Bank for financial support under grant RPO 682-84 and to Shaohua Chen, Angus Deaton, Margaret Grosh, and Hu Dinghuan for helpful comments. We remain culpable for all errors and omissions.

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

Improving Estimates of Inequality and Poverty From Urban China's Household Income and Expenditure Survey

Abstract

In urban China the Household Income and Expenditure Survey requires respondents to keep a daily expenditure diary for a full 12-month period. This onerous reporting task makes it difficult to recruit households into the survey, compromising the representative nature of the sample. In this article we use data on the monthly expenditures of households from two urban areas of China to see if data collection short-cuts, such as extrapolating to annual totals from expenditure reports in only some months of the year, would harm the accuracy of annual expenditure, inequality and poverty estimates. Our results show that replacing 12-month diaries with simple extrapolations from either one, two, four or six months would cause a sharp increase in estimates of annual inequality and poverty. This finding also undermines international comparisons of inequality statistics because no country other than China uses such comprehensive 12-month expenditure records. But a corrected form of extrapolation, based on correlations between the same household's expenditures in different months of the year, gives much smaller errors in estimates of inequality and poverty.

Keywords: China, Income distribution, Inequality, Survey methods

JEL: D31; 015

1. Introduction

China's national statistical system, the largest in the world, is facing a crisis, especially in its Household Income and Expenditure Survey (HIES) activities. There are increasingly reasons to worry about the quality of the data for urban households, although existing analyses of quality problems are confined to the rural sector (Ravallion and Chen, 1999). In the urban survey, each member of a sampled household is required to keep a diary of all expenditures (disaggregated for hundreds of product categories) for each day for a full year. In theory, an enumerator visits the household once or twice each month, reviews the records for accuracy and completeness, assists the household with questions, and takes away the household records for data entry in the local State Statistical Bureau (SSB) office. The data are then aggregated into the annual expenditure estimates reported by the SSB for each province.¹

But in practice, the falling real value of payments made by local statistical bureaus to households makes it almost impossible to convince selected households to participate in some areas.² In response, some statistics officials rely on less random procedures, using personal, political, or patriotic pleas to solicit the support of work units—both firms and government agencies—to assist in finding respondents. Problems occur even once households are in the sample, with the rising value of time causing some respondents to ignore daily recording and just recall their previous month's expenditures when the enumerator visits, thus losing some of the perceived advantages of diaries.

¹ It is important to emphasise that the purpose of the survey is to estimate annual expenditures. The SSB does not report any monthly expenditure estimates, which highlights the unusual feature of the data used in the empirical section of the paper.

² The observations in this paragraph come from interviews with SSB personnel in four municipality districts. We also thank Shaohua Chen of the World Bank for verifying some of our observations and adding some of her own.

While additions to the survey budget could alleviate these problems, and also allow the survey to broaden in scope so as to cover emerging issues such as urban migration and reform of urban pensions, the reality is one of budget shortfalls. Hence, the government needs accurate survey estimates to help it track problems such as rising inequality and poverty in cities, while at the same time minimising the cost of data collection.

One potential solution to this dilemma would be to change to a system of collecting data from households for shorter periods. In contrast to China's detailed, year-long data collection, surveys in many other countries observe households for a week, a fortnight, or a month, and estimates of income and consumption from these periods are annualised by multiplying by 52, 26, or 12. The length of the recall period typically depends on the category of consumption, with long reference periods used for costly and/or infrequently consumed items and short reference periods for frequently consumed and minor items that would be easily forgotten (ILO, 1994). Hence, the scaling factors that convert these short duration observations into annual figures may vary, but the principle in all cases is the same; an estimate of annual expenditures can be made by simple extrapolation from shorter observation periods.

While certainly less costly, such data collection shortcuts may come at a price. Extrapolation to annual totals can give accurate estimates of mean annual expenditures for the population if the sample is spread evenly over the year, but dispersion across households may be overstated (Deaton, 1997). Random shocks that occur during the observation period, but are subsequently evened out over the rest of the year, get included along with the genuine between-households inequality in annual expenditures, causing measures of inequality to be overstated. In any setting where the poverty line is below the

modal value of per capita expenditure the overstated dispersion also will lead to an overstatement of the poverty head-count and other measures of poverty.³

So, China's statistical system finds itself facing difficult choices. Officials want to maintain and increase the quality of the data. The scope of the data, already narrow, is in great need of being broadened to provide analysts with more information on which to assess and create policies. Yet, at the same time, SSB budgets have shrunk in real terms, staffing has fallen, and salaries have lagged far behind. Local and national officials have begun debating the benefits and costs of their options.

The aim of this paper is to assess whether China's system of 12-month diaries can be modified into a system in which households report expenditures for only some months, while not incurring loss of accuracy in measuring annual expenditures, inequality and poverty. We consider several types of extrapolation from short observation periods to annual totals, including a 'corrected' method, suggested by Scott (1992) and based on correlations between the same household's expenditures in different months of the year – correlations implicitly assumed to be 1.0 by simple extrapolation. Hence, the results of our study have wider application – beyond the case of China – since they illustrate how well simple and corrected extrapolations work against the benchmark of annual expenditures measured with 12-month diaries.

2. China's Household Income and Expenditure Survey

In some respects China's HIES is one of the most comprehensive in the world. The combined urban and rural sample includes more than 100,000 households surveyed over the course of a year, every year. Local statistical bureaus collect China's urban HIES

³ Ravallion (1988) gives a proof of this proposition in the context of variability due to risk.

data in selected municipalities, cities, and towns. In theory, SSB officials work with a sample framework that is based on randomly selected “neighborhood committees” (*juweihui*) from throughout the city’s jurisdictional area, and randomly selected households from the comprehensive list that is kept by each committee head. Households are expected to keep records on a year-round basis and stay in the program for 3 years. Local enumerators, employees of the local SSB, train the households in maintaining a monthly diary form, assist households in completing their work each month, and collect the completed forms, returning them to their office for entering into a nationally designed coding program. Households receive a minimal payment from the SSB for their effort (between 3 to 5 yuan – around US\$0.50 per month).

The survey instrument focuses mainly on collecting household expenditures and incomes, but also contains sections covering some durable assets and household characteristics. The survey begins with a monthly accounting of the household’s demographic and employment situation, which is reconciled with a month-end income form that traces 25 types of income. The rest of the form is devoted exclusively to recording detailed information about household expenditures. The household must disaggregate its purchases of food into nearly 150 categories, recording both the quantity and total amount spent (information on prices is not asked). Separate forms, centered on categories such as clothing, household goods, health, transportation and communication, entertainment, and education are broken into more than 150 individual items.

The system places considerable emphasis on avoiding recall error by encouraging respondents to record their expenditures on a daily basis. Enumerators expect each member of each household to record information on each day’s transactions on a single

household form. If one member is temporarily away from the house, or if someone in the household does not know how to read or write, another household member records the information for that person. Because China's urban literacy rates are so high, enumerators in the cities that we visited claim that they almost never encounter a case where everyone in a household is illiterate.

Fiscal problems in recent years have, however, exacerbated China's HIES problems. Budgetary restrictions have kept payments to trivial levels. Randomly selected households increasingly have refused to accept the time-consuming task of keeping diaries and meeting with enumerators up to 20 times per year. In some areas local SSBs have begun to compromise the representative nature of the sample, relying on personal friendships and connections and making pleas to sympathetic leaders of government agencies and work units to find willing households. SSB officials at all levels are quite eager to discuss these matters, knowing that the statistical system that they had fought so hard to re-build after its destruction during the Cultural Revolution is facing its most serious crisis of the reform era.

3. Data and Methods

In this study, we use 1997 data on the expenditures of households in two urban areas in Hebei and Sichuan Provinces. The data were originally collected and aggregated into the statistics generated and reported by China's HIES. Rather unusually, for our study local SSBs provided us with the raw data files and we were able to reconstruct a set of expenditure variables for each of 232 households, for each month of the survey year.

We focus on non-durable consumption expenditures by excluding purchases of furniture, motor vehicles and other durable goods. The urban HIES lacks a full inventory of durable goods, so if durables expenditure was included there would be no way to estimate the value of services from durables for non-purchasing households. The components of our expenditure aggregate are food, clothing, rents, and other non-food daily necessities. In the context of China these daily necessities include expenditures on transportation, communication, travel, toiletries, stationary, education, and utilities. Two factors cause unusual within-year variability in a household's expenditure on these items. First, China's urban residents tend to pay far in advance for certain goods (e.g. semi-annual bus passes), and second, ambiguities in some categories may cause respondents to include larger expenditures, such as tuition payments, amongst daily necessities. This must be borne in mind when interpreting our results in the context of other countries.

We model the effect of replacing the current system of 12-month diaries with a system where expenditures from shorter observation periods are extrapolated to annual totals. Under this new survey design all of the components of non-durable consumption expenditures are assumed to be measured with a monthly reference period. However, in many surveys around the world a variety of reference periods are used; typically shorter for food items and longer for non-food, infrequent expenditures. Four types of simple extrapolation to annual totals are considered, where in each case the fieldwork is spread evenly over the year:

1. A single monthly observation of expenditures, annualised by multiplying by 12 (to spread evenly, 20 households are observed per month for the first 4 months and 19 per month for the last 8 months – $(20 \times 4) + (19 \times 8) = 232$).
2. Two monthly observations (six months apart) on each household, with expenditures for each month multiplied by six and then summed to give an annual total.

3. Four monthly observations on each household (with a gap of two months between each), with expenditures from each month multiplied by three to give an annual total.
4. Six monthly observations on each household (a one-month gap between each) with expenditures from each month multiplied by two to give an annual total.

To the extent that repeated observations capture within-year variation in expenditures by the same household, we would expect to see estimates of annual inequality fall as the extrapolated annual totals use more months of data.

To cover the four simulated fieldwork designs with their varying frequency of revisits to the household, eight months of data on 1997 expenditures were extracted from the records for each household in the sample (Table 1). If there was a natural ordering to the data we could do this just once, because the fieldwork designs that we mimic would always put the same households in, say, the January sample or the December sample. But the data were unordered so to guard against finding sample-specific bias, we draw 100 re-samples, each based on a different randomised ranking of the 232 households in our sample. We report summary measures from these 100 re-samples.

In addition to the four types of simple extrapolations, a *corrected* extrapolation to annual totals is made by multiplying the single monthly observation (scheme (a) in Table 1) by a factor that depends on the average correlation, \bar{r} between the same household's expenditures in all pairs of months in the year. This corrected extrapolation method is derived by Scott (1992) and is only briefly summarised here. Let \bar{x}_m refer to the average, and $V(x_m)$ the variance, of monthly expenditures across all i households and t months in the year. Extrapolating to annual expenditure totals by multiplying monthly expenditures by 12 gives an estimated variance of annual expenditures of $144 V(x_m)$. This is likely to overstate the variance in annual expenditures reported in 12-month diaries:

$$V(x_a) = \frac{1}{N} \sum_{i=1}^N (x_{i,a} - \bar{x}_a)^2 \quad (1)$$

where $x_{i,a}$ is annual diary expenditure by the i th household and \bar{x}_a is average annual expenditures. Equation (1) can be expressed as:

$$V(x_a) = \sum_{t,t'=1}^{12} r_{t,t'} \sigma_t \sigma_{t'} \quad (2)$$

where $r_{t,t'}$ is the correlation between expenditures in month t and month t' and σ_t is the standard deviation across households in month t . This follows because $x_{i,a} - \bar{x}_a$ in equation (1) can be expressed as the sum of the deviations of each household's monthly expenditure from the mean for that month, $d_{it} = x_{it} - \bar{x}_t$ and the d_{it} terms are components of the correlation coefficient:

$$r_{t,t'} = \frac{1}{N} \sum_{i=1}^N d_{it} d_{it'} / \sigma_t \sigma_{t'}. \quad (3)$$

Assuming that the dispersion across households does not vary from month to month, i.e., $\sigma_t = \sigma_{t'}$, equation (2) can be expressed as:

$$V(x_a) = [12 + 132 \bar{r}] V(x_m). \quad (4)$$

Hence the variance from simple extrapolation to annual totals, $144 V(x_m)$, equals $V(x_a)$ only in the special case of $\bar{r} = 1$.

The corrected extrapolation uses estimates of \bar{r} to scale the i th household's deviation from the overall monthly average, $(x_{it} - \bar{x}_m)$ up to an annual value, and adding this to the annual average across all households, $\bar{x}_a = 12 \bar{r} \bar{x}_m$, gives:

$$x_{i,A} = (x_{it} - \bar{x}_m) \sqrt{12 + 132 \bar{r}} + 12 \bar{r} \bar{x}_m. \quad (5)$$

For example, if the average correlation between the same household's expenditures in all pairs of months in the year is 0.25, the scaling factor is only 6.71 ($= 45$), rather than the scaling factor of 12 implied by simple extrapolation. Thus, the deviation of a household's one-month expenditures from \bar{x}_m has a smaller effect than under simple extrapolation, leading to a less dispersed distribution of annual expenditures.

The most reliable estimate of \bar{r} would use the 66 correlation coefficients, $r_{t,t'}$, between all $i-j$ pairs of months, but this requires observations on each household's expenditures in every month in the year, yielding no savings over the current fieldwork design. Instead, a sampling approach is taken, by estimating \bar{r} from only a few of the possible correlation coefficients for expenditures in the various $i-j$ pairs of months. This sampling approach reduces the cost of fieldwork but relies on the $r_{t,t'}$ having roughly the same value and varying little as the gap between t and t' increases. Existing evidence demonstrates that this may be a valid assumption, with at least one survey finding that $r_{t,t'}$ fell by just 0.0078 for each month that the gap between t and t' increased (CSO, 1995). In the survey designs simulated here, the $r_{t,t'}$ come from either a single revisit to the household six months after the first expenditure observation (i.e., scheme (d) in Table 1), three revisits to the household (i.e., scheme (c) in Table 1), or five revisits to the household (i.e., scheme (b) in Table 1).

Inequality and Poverty Measures

To compare the performance of the simple and corrected extrapolations with the benchmark of the expenditure estimates from the existing 12-month diary survey we use two inequality measures: the between-households standard deviation of expenditures and the Gini coefficient. We also use the P_α class of poverty measures: $P_\alpha = 1/n \sum_{i=1}^q (g_i/z)^\alpha$

where n is the total population, q is the number of poor individuals, z is the poverty line, g_i is the poverty gap, $z - y_i$, where y_i is expenditure per capita in the i th household and α is the poverty aversion parameter (Foster, Greer and Thorbecke, 1984). We use the three most popular members of the P_α class: the head-count index ($\alpha=0$), the poverty gap index ($\alpha=1$), and the poverty severity index ($\alpha=2$). The poverty gap index can show the resources needed to eliminate poverty through perfectly targeted transfers, a feature that we use below when comparing the costs and benefits of changing from 12-month diaries to extrapolations from monthly observations. In the simulations, the poverty line is set at Y2840 per person per year; applying this to the annual expenditures from the 12-month diaries gives a poverty headcount ratio of 30.5% -- a similar threshold used for setting relative poverty lines in other countries (Grootaert, 1994).

4. Results

Replacing urban China's current survey system of 12-month diaries with some form of extrapolation from monthly observations to annual totals would not affect the measurement of mean annual expenditures (Table 2, column 1). Extrapolation from one month's expenditures gives estimates of mean annual expenditure which have an average value of Y10,518 across our 100 re-samples. This average is almost exactly equal to the estimate obtained from using the full 12 months of observations on each household (Y10,524). Using more than one month's expenditures provides equally accurate point estimates (Table 2, rows 3 to 5). The corrected extrapolations give the same results as simple extrapolation from a single month, because the correction in equation (5) is designed only to affect measures of dispersion. Overall, the results in the first column of

Table 2 emphasise the well-known point that spreading a sample evenly over the year should lead to good estimates of mean annual expenditures (Deaton, 1997).

In contrast to the results for the mean, simple extrapolation to annual totals reduces the accuracy of measures of dispersion. When using all 12 months of data on each household, the standard deviation of annual expenditure is estimated as Y3,807 (Table 2, column 2). Multiplying one month's expenditures by 12 to get an estimate of annual expenditure increases the standard deviation by over 100%, to an average of Y7,718 over the 100 re-samples. Although the estimated standard deviation falls as more months of survey data are collected (rows 3 to 5), even extrapolating to annual totals from six months of data produces an estimate that is 12% larger than when the survey collects data for all 12 months of the year. There is also an upward bias, although not as pronounced, in the Gini coefficient: When using 12 months of data on each household, the Gini coefficient is estimated as 0.197 but when one month's expenditures are multiplied by 12 to give an annual total, the Gini is 64.6% higher, at 0.325.

The overstatement of the Gini coefficient falls as annual totals are extrapolated from more monthly expenditure observations but even when data are collected every second month, the average value of the Gini coefficient is 12% higher than when using the 12-month diaries. Moreover, there was no instance amongst the 100 re-samples of simple extrapolation producing an estimate of the Gini coefficient that either equalled or was lower than the Gini coefficient estimated on expenditures from the 12-month diaries. These results for the 100 re-samples are presented in the top panel of Figure 1 as smoothed densities, using a Gaussian kernel. The *relative* error is shown, with values greater than one implying exaggeration when extrapolated measures of annual

expenditures are used, and a value of 1.0 implying zero error. With simple extrapolation from one month's expenditures, the estimated Gini coefficient is between 1.39-2.10 times that found with the 12-month diaries, while with extrapolation from six month's data, the Gini coefficient is 1.03-1.19 times the estimate from 12-month data. Hence, increasing the intensity of fieldwork, by obtaining household's expenditures in more months of the year, reduces both the average level of error and the variance of that error.

The implication of these results is that replacing urban China's current survey system of 12-month diaries with simple extrapolations from either one, two, four or six months would cause a sharp increase in measured inequality. Because one of the concerns of the Chinese government is monitoring the trend in urban inequality, this may be an undesirable result of any cost-saving change to less frequent data collection. In addition to issues of temporal comparability, the results in Figure 1 have implications for the international comparability of inequality statistics. Household surveys in many countries use a design somewhat like the extrapolation from one month's expenditures in our simulations, but with results often reported on a monthly or shorter period basis rather than being annualised. Hence, it is likely that measured inequality in other countries is higher than in China due, in part, to the different periods over which expenditures are observed. If measured on the same basis as in other countries, and combined with other amendments suggested by Khan and Risken (1998), China's urban inequality levels might approach if not exceed levels found elsewhere in the world.

In contrast to these largely negative findings about the performance of simple extrapolation, the *corrected* extrapolations using equation (5) appear to perform well (Table 2, rows 6 to 8). Whereas the average value of the standard deviation of annual

expenditures using simple extrapolations from one, two, four or six months overstates the standard deviation from the 12-month diaries by 103%, 55%, 23%, and 12%, the overstatement using the corrected extrapolations is only between 0.8% and 5.5% (rows 6 to 8). Similarly, the estimated Gini coefficients are much closer when using the corrected extrapolations; if \bar{r} is estimated from a single revisit, the mean value of the Gini coefficient is just 6.4% higher than the value from the 12-month diaries and the error falls to 3.1 % when \bar{r} is estimated from five revisits.

The smoothed densities for the error when the Gini coefficient is calculated from the corrected extrapolations are shown in the bottom panel of Figure 1. In addition to a lower expected level of error, there is also less variance in the error when the corrected extrapolation method is used. The key to the correction is the average correlation between the same household's expenditures in different months of the year, and this differed little across the varying revisit frequencies studied, which accounts for the similarity of the smoothed densities for one, three and five revisits. It is notable that this correlation averaged 0.18; far from the 1.0 assumed by simple extrapolation. Moreover, because the only purpose of the revisits in our simulations is to measure \bar{r} , greater savings could come from revisiting just a subset of the households and estimating \bar{r} from this subset.

Estimating Poverty

The P_α poverty measures are fully characterised by the poverty line, the mean level of expenditures and the Lorenz curve representing the level of relative inequality. The same poverty line is used in all simulations and extrapolation gives accurate estimates of mean expenditures, consequently the greater inequality found with simple

extrapolation will also raise the measured level of poverty. How large is this bias? When one month's expenditures are multiplied by 12 to give annual totals, the average level of the poverty head-count index is 47% , rather than the expected 30% (Table 3, column 1). Even extrapolating from either every third or every second month produces substantial error, with the poverty head-count overstated by about 15%. The overstatement is even larger for the poverty gap and severity indices, because increasing the poverty aversion parameter, α makes the poverty index more sensitive to the greater dispersion in annual expenditures induced by simple extrapolation.

The improvement when using corrected extrapolation is at least as impressive as it was for the inequality measures (Table 3, rows 6 to 8). The head-count index is almost exactly the same when using corrected extrapolation, with a single revisit to estimate \bar{r} , as when using the 12-month diaries (30.52% versus 30.47%). There is more error in the poverty gap and, especially, poverty severity measures, but still much less than for any of the types of simple extrapolation considered. Because most studies of poverty concentrate on the head-count and poverty gap measures (despite their theoretical shortcomings), the results in Table 3 suggest that there may not be great loss of accuracy in replacing the system of 12-month diaries with corrected extrapolation to annual expenditure totals.

5. The Costs and Benefits of Reform

How costly to Chinese society is the overstatement of urban inequality and poverty when data collection shortcuts are used, and are these costs greater than the savings made by replacing 12-month diaries with some form of extrapolation from short observation periods? The answer depends on whether misguided policies are adopted on

the basis of the overstated inequality and poverty estimates and how much damage those policies do. It also depends on any offsetting advantages of a more representative sample, more diligent respondents and a broader topic coverage if the burden of using 12-month diaries is removed.

As a start in assessing these costs and benefits we assume that the only goal of the survey is to measure the total poverty gap, so that Chinese authorities can budget the correct amount for eliminating poverty (using costless and perfectly targeted transfers). We obtain this budgetary figure from the product of the poverty gap index, the value of the poverty line and the size of the population. The first column of Table 4 contains these estimates of the poverty gap *for the sampled households*, under each of the different survey designs that we study. The difference in the value of the total poverty gap between each potential survey design and the current system of expenditure reports for 12 months is used to calculate the excess size (or shortfall) of the poverty alleviation budget (Table 4, column 2). For example, if annual expenditures are simply extrapolated from one month's expenditures, the aggregate poverty gap is overstated for the sampled households by Y180,245, while if the corrected extrapolation is used (with \bar{r} from one revisit), the total gap is understated by Y6,028.

To move from these sample estimates to a population total for the miscalculated poverty gap we assume that sampling is at a 1:1000 rate. We also assume three levels of social loss from the payment of poverty alleviation funds to non poor households: 100% loss, 50% loss and 25% loss (Table 4, columns 3-5). Under the assumption of 100% loss, any money budgeted for poverty alleviation that is greater than the amount required to close annual poverty gaps is treated as being completely wasted, while the lower loss

rates assume that spending poverty funds on non-poor households still provides some social benefit. We also assume that the loss function is symmetric, so that not spending enough on poverty alleviation, due to an underestimation of the total poverty gap is equally as bad as spending too much. Under these assumptions, which are admittedly extreme, the social cost of overstating poverty by using simple extrapolation from one month's expenditures could be as high as 180 million yuan in our two cities. Even with the lowest rate of social loss and extrapolation from expenditure observations in every second month, the annual loss is 6 million yuan. But using the corrected extrapolations to calculate the aggregate poverty gap does improve these calculations considerably: social losses range from 6 million yuan (\bar{r} from one revisit, 100% loss rate) to only 0.3 million yuan (\bar{r} from five revisits, 25% loss rate).

How do these social losses compare with the possible savings from replacing 12-month diaries with some form of extrapolation? The local statistical bureaus in the two cities providing the data suggest that there are fixed costs of about 1 million yuan to collect the data, and marginal costs of 100,000 yuan per month of fieldwork. However, with this budget the local SSBs have to let survey personnel pursue commercial practices on the side as a way of raising their salaries to reasonable levels. So in addition to the actual marginal cost of 100,000 yuan per month, we also consider a doubling in the salaries, to the levels thought necessary to keep quality enumerators and professional statisticians, implying a marginal cost of 200,000 yuan per month of the survey activities. The last two columns of Table 4 contain estimates of the cost savings under these two salary schemes when less frequent observations on household expenditures are made.

Few of the combinations of social losses and reduced fieldwork costs in Table 4 make the move away from 12-month diaries appear attractive. With simple extrapolation, the minimum social loss is 5.8 million yuan but the reduction in fieldwork costs from observing households every second month is only 0.6 million yuan (or 1.2 million yuan when salaries are doubled to market levels). But with corrected extrapolation, and \bar{r} from three or five revisits, the loss to society is smaller than the additional cost of collecting data for all 12 months under the assumption that only 50% or 25% of misallocated poverty funds are considered a social loss.

One might infer from these results that other national statistical systems should follow China's lead and go to more comprehensive 12-month expenditure records. But few public expenditures are budgeted as exactly as the poverty alleviation fund that we model in our simulations. Moreover, the difficulties that Chinese authorities have in recruiting households into a 12-month diary survey show that this potentially more accurate data collection method may be, for all practicable purposes, infeasible. Our limited cost-benefit analysis has also excluded several potentially important benefits from reducing the burden on respondents, such as a more representative sample, more diligent respondents and a broader topic coverage. For example, questions that better capture the permanent living standards of respondent households (e.g., child anthropometrics, wealth) might be possible if less effort was spent collecting the expenditure data, and in this way the ability of the Chinese authorities to target the urban poor might even be improved. On the other hand, there may also be certain advantages in forcing households to record their expenditures continuously rather than starting and stopping (e.g., the elimination of "start-up bias" from all but one month's expenditure data).

6. Conclusion

Our work in this paper shows that replacing the current data collection method of 12-month diaries in urban China's Household Income and Expenditure Survey with simple extrapolation to annual totals from expenditures in only one, two, four or six months would cause a sharp increase in estimates of annual inequality and poverty. Under the most extreme form of data collection short-cut – collecting only one month's expenditure data on each household and annualising these estimates by multiplying by 12 – the standard deviation of annual household expenditures would be overstated by more than 100%, the Gini coefficient of inequality by 65% and the poverty headcount by 53%. Our limited cost-benefit analysis suggested that the loss to Chinese society from these inaccurate poverty statistics more than outweighs the cost savings from a reduced effort on data collection. However, we did not measure the potential benefits of a more representative sample, more diligent respondents and a broader topic coverage if respondents had to record only one month's expenditures.

In contrast to these negative findings about simple extrapolation, the corrected extrapolation method introduced by Scott (1992) performs well. If the correlation between the same household's expenditures in two different months of the year is used when extrapolating one month's expenditures to annual totals, the standard deviation is overstated by 5.5%, the Gini coefficient by 6.4% and the poverty headcount by just 0.1%. More frequent revisits to get a better measure of this average correlation between expenditures in the months of the year gives further small gains in accuracy. Hence, for many purposes it may be sufficient to replace urban China's current survey system of 12-month diaries with a system where household's expenditures are recorded in only two

months of the year. Moreover, further cost savings could be made by revisiting only a random sub-sample of households and using the \bar{r} estimated from this group to correct the expenditure estimates for the full sample.

One of the broader implications of our results is that international comparisons of inequality may be inconsistent, because no country other than China uses such comprehensive 12-month expenditure records. Instead, household surveys in many countries use a design somewhat like the extrapolation from one month's expenditures in our simulations, with results either annualised or reported on a monthly basis. Hence, it is likely that measured inequality in these other countries is higher than in China due, in part, to the shorter periods over which expenditures are observed. Our results also illustrate the importance of data collection methods in producing accurate estimates of variance-based measures. As the value of household's time rises with economic development, statistical agencies in many developing countries may find that they need to shorten the period over which they intrude upon households for observing expenditures or incomes. Such changes in the observation period need not affect estimates of the annual mean if the sample is spread over the year but estimates of the annual dispersion across households may be affected.

Table 1: Sampling Scheme for the Various Revisit Frequencies

Observation Number	Months that Household Expenditures are Observed Under Each Revisit Frequency											
	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
1-20	abcd		b	c	b		bcd		b	c	b	
21-40		abcd		b	c	b		bcd		b	c	b
41-60	b		abcd		b	c	b		bcd		b	c
61-80	c	b		abcd		b	c	b		bcd		b
81-99	b	c	b		abcd		b	c	b		bcd	
100-118		b	c	b		abcd		b	c	b		bcd
119-137	bcd		b	c	b		abcd		b	c	b	
138-156		bcd		b	c	b		abcd		b	c	b
157-175	b		bcd		b	c	b		abcd		b	c
176-194	c	b		bcd		b	c	b		abcd		b
195-213	b	c	b		bcd		b	c	b		abcd	
214-232		b	c	b		bcd		b	c	b		abcd

Notes:

- (a) Observe each household's expenditure in only one (staggered) month of the year.
- (b) Observe each household's expenditure every second (staggered) month of the year.
- (c) Observe each household's expenditure in one (staggered) month of each quarter.
- (d) Observe each household's expenditure in two (staggered) months of the year, with five month gap.

Table 2: Comparison of Statistics Calculated from Actual Annual Expenditures versus Extrapolations

	Mean ^a	Standard deviation	Gini coefficient	Percentage Error When Extrapolating		
				Mean	Standard deviation	Gini coefficient
Using 12 months of expenditure observations	10524	3807	0.1974			
Extrapolating to annual totals from:						
One months' expenditures	10518	7718	0.3250	0.1	102.7	64.6
Two months' expenditures	10531	5852	0.2692	0.1	53.7	36.4
Four months' expenditures	10538	4696	0.2323	0.1	23.4	17.7
Six months' expenditures	10534	4275	0.2202	0.1	12.3	11.6
Corrected extrapolations to annual totals with \bar{r} from:						
One revisit	10518	3599	0.1848	0.1	5.5	6.4
Three revisits	10518	3778	0.1893	0.1	0.8	4.1
Five revisits	10518	3854	0.1913	0.1	1.2	3.1

Note: Results are averages of 100 random draws of the various sampling schemes in Table 1.

^a Yuan per year.

Table 3: Comparison of Poverty Statistics Calculated from Actual Annual Expenditures, Naïve Extrapolations and Corrected Extrapolations

	Head- count (=0)	Poverty gap (=1)	Poverty severity (=2)	Percentage Error When Extrapolating		
				Head- count (=0)	Poverty gap (=1)	Poverty severity (=2)
Using 12 months of expenditure observations	30.47	6.02	2.01			
Extrapolating to annual totals from:						
One months' expenditures	46.66	15.04	7.01	53.1	149.8	248.5
Two months' expenditures	40.29	10.71	4.25	32.2	77.8	111.6
Four months' expenditures	34.73	8.08	2.84	14.0	34.2	41.1
Six months' expenditures	35.06	7.19	2.42	15.0	19.4	20.3
Corrected extrapolations to annual totals with \bar{r} from:						
One revisit	30.52	5.72	1.67	0.1	5.0	16.7
Three revisits	31.30	5.87	1.73	2.7	2.5	14.1
Five revisits	31.63	5.95	1.76	3.8	1.1	12.5

Note: Results are averages of 100 random draws of the various sampling schemes in Table 1 and are based on a poverty line set at Y2840 per person per year.

Table 4: Estimated Total Poverty Gaps, Misallocated Poverty Funds, and the Cost of Surveys

	Estimated Total Poverty Gap for Sampled Households	Excess (shortfall) in size of Poverty Budget for Sampled Households	Assumed Social Loss Due to Mis-allocating Poverty Funds			Savings From Reduced Fieldwork Costs	
			100%	50%	25%	Current	Post-Reform
Using 12 months of expenditure observations	120,340	0					
			<i>Millions Yuan^a</i>				
Extrapolating to annual totals from:							
One month's expenditures	300,585	180,245	180.2	90.1	45.1	1.1	2.2
Two months' expenditures	213,932	93,591	93.6	46.8	23.4	1.0	2.0
Four months' expenditures	161,549	41,209	41.2	20.6	10.3	0.8	1.6
Six months' expenditures	143,653	23,313	23.3	11.7	5.8	0.6	1.2
Corrected extrapolations to annual totals with \bar{r} from:							
One revisit	114,312	(6,028)	6.0	3.0	1.5	1.0	2.0
Three revisits	117,320	(3,021)	3.0	1.5	0.8	0.8	1.6
Five revisits	119,001	(1,339)	1.3	0.7	0.3	0.6	1.2

Note: Results are averages of 100 random draws of the various sampling schemes in Table 1.

^a The cost of mis-allocating (or Social Loss) is calculated by taking the "Excess Size of Poverty Budget" (column 2—which is the budget that would have to be allocated to sample households that were identified as being poor, when they really were not poor) and assigning the value to the entire population. Since the sample is 1 in 1000, the figure in column 3 is just 1000 times the figure in column 2. These are measured in millions of current yuan.

References

- Central Statistical Office (1995) *Republic of Zambia Household Budget Survey, 1993-94, Volume 1 – Main Tables and Report*, Central Statistical Office, Lusaka.
- Deaton, Angus. (1997) *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy* Johns Hopkins, Baltimore.
- Foster, J., Greer, J., and Thorbecke, E. (1984), “A Class of Decomposable Poverty Measures,” *Econometrica* 52, 761-765.
- Grootaert, C. (1994), “Poverty and Basic Needs Fulfilment in Africa During Structural Change: Evidence from Côte d’Ivoire,” *World Development* 22, 1521-1534.
- ILO (1994), “Household Income and Expenditure Surveys,” *Year Book of Labour Statistics, Sources and Methods*, Volume 6, International Labour Office, Geneva.
- Khan, A., and Riskin, C. (1998), “Income and Inequality in China: Composition, Distribution and Growth of Household Income, 1988 to 1995,” *China Quarterly* 157, 221-253.
- Ravallion, M. (1988), “Expected Poverty Under Risk-induced Welfare Variability,” *Economic Journal* 98(Dec): 1171-1182.

Ravallion, M., and Chen, S. (1999), "When Economic Reform is Faster Than Statistical Reform: Measuring and Explaining Income Inequality in Rural China," *Oxford Bulletin of Economics and Statistics* 61, 75-102.

Scott, C. (1992), "Estimation of Annual Expenditure From One-month Cross-sectional Data in a Household Survey," *Inter-Stat* 8, 57-65.