

Hybrid Survey to Improve the Reliability of Poverty Statistics in a Cost-Effective Manner

Faizuddin Ahmed

Cheku Dorji

Shinya Takamatsu

Nobuo Yoshida

The World Bank
Poverty Reduction and Economic Management Network
Poverty Reduction and Equity Unit
June 2014



Abstract

This paper studies the benefits, in terms of reliability and frequency of poverty statistics, of conducting a hybrid survey that collects non-consumption data from all surveyed households and consumption data from only a small subsample. Collecting detailed consumption or income data for the purpose of estimating poverty is costly and many low-income countries cannot afford to carry out such surveys on a regular basis. One option is to collect only non-consumption data and use consumption models developed from a previous round of household survey data to project poverty data. Although this approach is cost-effective because collection of non-consumption data is much cheaper than collection of consumption data, it is vulnerable to a structural change

between the current and previous household surveys and might produce poverty estimates that are not comparable with the previous ones. Instead, the hybrid approach creates consumption models from a subsample of the current survey and applies them to the entire survey to project consumption data for all households in the sample. This paper examines the hybrid approach with data from the Bangladesh Household Income Expenditure Surveys of 2000 and 2005. Improvements in accuracy are achieved even with subsamples of just 320 or 640 households. Budget simulations confirm that the additional cost of collecting consumption data for such small subsamples is minimal.

This paper is a product of the Poverty Reduction and Equity Unit, Poverty Reduction and Economic Management Network. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at <http://econ.worldbank.org>. The author may be contacted at nyoshida@worldbank.org.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

Hybrid Survey to Improve the Reliability of Poverty Statistics in a Cost-Effective Manner

Faizuddin Ahmed, World Bank
Cheku Dorji, National Statistics Bureau, Bhutan
Shinya Takamatsu, World Bank
Nobuo Yoshida, World Bank

JEL: I32.

Keywords: Survey to survey imputation, subsample, poverty measurement

Acknowledgments

The authors thank Peter Lanjouw for his thorough review of the previous draft and useful suggestions, which improved the quality of analyses in this paper; Kinnon Scott for her support in designing the pilot survey to test price data collection and the effect of switching from diary- to recall-based approaches in collecting consumption data; and the Bangladesh Bureau of Statistics for their help in designing the sampling of the pilot survey and for providing details of survey logistics to help estimate survey implementation costs. Finally, the authors are grateful for their discussions with Roy van der Weide and Kathleen Beegle, and insightful comments from Salman Zaidi.

I. Introduction and Literature Review

Following the recent economic crisis, which affected both developed and developing countries, policy makers and other stakeholders are demanding more frequent and timely poverty data. During the crisis, food and fuel prices rose sharply and employment opportunities deteriorated substantially. These impacts aggravated the situation of poor and vulnerable groups and likely increased the overall prevalence of poverty in many developing countries. Policy makers and stakeholders need to know to what extent poverty has increased, whether existing poverty alleviation programs are working, and if not, how they can be improved. However, for many developing countries, it is not yet (and may never be) known how much the prevalence of poverty increased simply because poverty statistics were not available during and in the immediate aftermath of the crisis.

Why do so many developing countries fail to publish poverty data regularly and frequently? Producing good quality poverty data requires collecting detailed household expenditure (or income) information at the household level. This is a resource-intensive and time-consuming process. To overcome these constraints, conducting rapid monitoring surveys with highly simplified survey questionnaires and/or shortened survey periods is sometimes suggested. However, many analysts are concerned about the quality of poverty data from such rapid monitoring surveys in terms of their comparability to official poverty data from large-scale full-fledged household surveys. If the data are not comparable, tracking poverty becomes extremely tricky and policy makers and other stakeholders cannot be sure how much (and sometimes even in what direction) poverty has changed.

Imputation procedures offer a potentially cost-effective solution for improving the frequency and comparability of poverty data.¹ The idea is simple: suppose consumption or income data are not collected in a particular year, but nonconsumption data from other surveys are available, then consumption or income data can be imputed into the nonconsumption/income data set using imputation models calibrated on consumption data collected previously. Because we are using data already available, there is no data collection cost (see, for example, Stifel and Christiaensen [2007] or Doudich et al. [2013]). Even if a survey is necessary to collect the nonconsumption/income data, its cost is likely to be low compared to that of a full-fledged consumption survey. Ravallion (1996), one of the early advocates of the imputation approach, experimented with how well household consumption expenditure can be estimated from information on housing conditions and subjective well-being, which can be collected quickly and cheaply.

For the imputation approach to work, it is important to assess whether imputation models can accurately predict household consumption/income and produce comparable poverty data. The imputation models are calibrated in a year when both household consumption data and nonconsumption data are available, and then applied to impute household consumption data for years when only the nonconsumption data are available.¹ A potentially serious problem is that there is no guarantee that imputation models will accurately reflect consumption patterns in those periods when only nonconsumption data are available. For example, ownership of cell phones is a good predictor of household wealth, but its relationship with total household expenditure has been changing rapidly in many Sub-Saharan African countries as cell phone coverage increases and cell phone fees decline. Ten years ago, ownership of cell phones was a good predictor of extreme wealth; today such phones are prevalent among the lower-middle-income classes and even among the poor. Using imputation models calibrated 10 years ago with cell phones as covariates would likely underestimate present-day poverty.

Because of these considerations, there is ongoing research and debate among analysts on methods to ensure that imputed poverty data are reliable and comparable. Stifel and Christiaensen (2007) provide

¹ The imputation method can be used for many other purposes as well. It is used for poverty mapping (e.g., Elbers, Lanjouw and Lanjouw 2003), restoring comparability of consumption data over time (e.g., Deaton and Dreze 2002, Kijima and Lanjouw 2003, Allwine et al. 2013, Revilla et al. 2010), and creating synthetic panel data for vulnerability analysis (Dang and Lanjouw 2013).

theoretical guidance regarding variables to be included in imputation models to ensure comparability and reliability of imputed poverty data. They recommend including covariates that change over time, but call for excluding variables whose rates of return are likely to change markedly in the face of evolving economic conditions. This argument makes sense in theory, but it is difficult to identify which variables would satisfy these conditions. Stifel and Christiaensen (2007) included several household durables in their imputation models, but Harttgen, Klasen and Vollmer (2012) criticized this decision because of the so-called “asset drift” effect—where the pace of improvement in asset ownership is much faster than the pace of income growth. As a result, the relationship between asset ownership and income or consumption could well be changing over time. In the end, which variables satisfy the recommendations of Stifel and Christiaensen (2007) is intrinsically an empirical question. At the same time, ascertaining empirically how best to proceed is difficult because actual household consumption data are needed. Remember, imputation methods are being employed precisely because actual household consumption data are not available.

Christiaensen et al. (2012) examined some of the arguments raised by Stifel and Christiaensen (2007) through an experiment. They collected multiple rounds of cross-sectional data including both household consumption and nonconsumption variables from four developing countries, namely rural China, Kenya, the Russian Federation, and Vietnam. They created imputation models from one round and imputed household consumption expenditure into other rounds that also had actual consumption data. They illustrated with their four case studies that imputation models can reproduce poverty statistics generated from directly collected consumption data. The results support the notion that imputation models constructed from past surveys can predict household consumption data of future rounds. Such results are encouraging and give us a certain level of confidence in the stability of imputation model coefficients over time. However, this exercise is still experimental. Validation is possible only in those settings where imputed estimates can be compared to actual estimates. In practice, the imputation approach needs to be used in those circumstances where actual consumption data are not available, and therefore, analysts must be willing to assume that the imputation model is accurately predicting household expenditure and poverty rates.

This paper explores a hybrid approach—implementation of a survey that collects a large sample of nonconsumption data, which can be collected easily and cheaply, but also collects a small subsample of both consumption and nonconsumption data. Collection of household consumption data in the subsample provides several advantages. First, the data can be used to estimate poverty rates directly. Because of the small size of the subsample, this direct estimation will not be precise: sampling errors can be significant. However, direct estimation of the poverty rate from the subsample might still provide a good ball park figure. Second, using the subsample, which includes both consumption and nonconsumption data, we can create imputation models and apply them to the full (and larger) sample to impute consumption for all sample households. This hybrid approach will result in an overall poverty estimate with the sampling error associated with the large household survey and a modeling error associated with our having imputed consumption into the large survey. Encouragingly, literature on small area estimation (for example, Elbers, Lanjouw, and Lanjouw [2003], hereafter ELL) suggests that model errors need not be excessive if proper care is taken in specifying and estimating the imputation model.

This hybrid approach can address the problems faced by both Stifel and Christiaensen (2007) and Christiaensen et al. (2012). Both papers create imputation models from past data and apply them to the latest household survey data. As discussed above, we cannot test whether the imputation models are still valid or reflect the relationship among household consumption and other correlates. In the hybrid approach, the imputation models are created from the subsample of the latest household survey, and so, by construction, the models reflect the latest relationship among household consumption and other correlates.

Despite these advantages, the hybrid approach is not perfect. First, collecting consumption data is costly, even if it is collected only for a subsample. We know that the larger the subsample size, the more accurate

the poverty estimates will be. But, we also know that as the subsample size increases, the cost of implementing the hybrid survey also rises. So, a key policy question is how much will a reduction in the size of the subsample increase the standard errors (SEs) of the poverty estimates. Furthermore, consumption data need to be comparable over time.

Second, the hybrid approach calls for regular re-estimation of the underlying imputation models. Developing reliable imputation models has become simpler since the World Bank Development Research Group created the PovMap 2 software for this purpose, but it remains a not insignificant task and a formidable challenge in those settings where statisticians have limited experience with these methods.

Third, the hybrid approach will produce the current value of imputed expenditure; therefore, some adjustments for intertemporal and spatial price differences need to be made. Such price adjustments often use unit values from household surveys, but the reliability of unit values will be limited since the size of subsample is small. In contrast, the approach proposed by Christiaensen et al. (2012) is attractive because (i) imputation models will be estimated only for the year in which a big household survey with a full sample of consumption data was carried out; (ii) imputed expenditure/income is already adjusted to prices of the big household survey year and, as a result, we do not have to make any additional price adjustments for the following years; and (iii) except for the big household survey year, we do not have to collect consumption data at all.

The objective of this paper is to compare the following three different poverty estimation approaches: (i) estimating poverty rates from consumption data in a subsample (“Subsample Direct Estimation”); (ii) estimating poverty rates using household expenditure imputed from models created from a subsample (“Subsample Imputation”); and (iii) estimating poverty rates using imputed household expenditure based on models created from a previous round of a large household survey (“Full Sample Imputation,” used by Christiaensen et al. [2012]).

This paper conducts the above analysis using the 2000 and 2005 Bangladesh Household Income and Expenditure Survey (HIES) data. Bangladesh was chosen because Bangladesh Bureau of Statistics expressed interest on the subsample imputation approach and also offered help to carry out additional surveys to see the effects of slight modification of survey modules on survey implementation costs. The 2000 and 2005 HIES data were chosen although a more recent round of HIES data (the 2010 HIES) is available because the poverty lines are slightly modified for the 2010 datasets. The modification of poverty lines in addition to price adjustments usually causes additional challenge to “Full Sample Imputation.” Therefore, the 2000 and 2010 HIES were selected. However, its robustness against changes in methodology is also can be seen as an advantage of the subsample approach.²

This paper also studies to what extent the limited sample size affects price adjustments, poverty lines, and poverty estimates. For situations where unit values from a subsample are too noisy to produce reliable poverty estimates, we also evaluated another way of conducting price adjustments using price data from an independent price survey. An issue here is that if unit values and prices are too different, poverty lines updated by price data would be very different and produce poverty estimates that are not comparable to the previous official poverty rates, which were estimated with price adjustments based on unit values. This paper examines whether differences between prices and unit values, both collected in a pilot survey, are statistically significant.

Finally, this paper examines the cost implications of collecting consumption data for a subsample, as well as the reliability of price data collected from price surveys. Collection of consumption expenditure data for a subsample might be useful to improve and ensure the statistical accuracy of imputed expenditure data; however, if it is costly, annual collection of such data might not be a realistic option for developing countries. We therefore examine the costs of collecting consumption data for a subsample after reviewing the current data collection process adopted by the Bangladesh Bureau of Statistics (BBS).

² This insight was drawn from discussions with Salman Zaidi.

This paper is not the first paper that evaluates the statistical precision and cost-effectiveness of the hybrid approach. A Fujii and van der Weide (2013) study examined the hybrid approach in a general environment. Their paper focused on the statistical properties of indicators derived from the hybrid approach and, after considering the implementation cost of the hybrid approach, derived the optimal sample sizes of both a large survey and its subsample given a total implementation cost. There are several differences between this paper and Fujii and van der Weide (2013). First, this paper not only estimates the statistical property of the hybrid approach, but also compares it with a simpler survey-to-survey imputation method, while Fujii and van der Weide (2013) focus on the statistical property of the hybrid approach. Second, this paper evaluates the accuracy of price adjustments—as mentioned, this can be a big obstacle in obtaining precise and reliable poverty statistics using the hybrid approach. Third, survey cost estimation requires detailed information on survey logistics, which can differ from country to country. By focusing on one country, this paper produces more precise estimates of survey implementation costs.

Finally, this paper focuses on

The following section presents the three proposed approaches to increase the frequency of poverty data in Bangladesh and the methods used to evaluate these approaches. Section III presents the key results of the performance of each approach. Section IV discusses the effects of updating inflation rates and poverty lines using a subsample or an independent price survey. Cost is an important factor in survey implementation, and section V assesses cost implications each approach. Section VI summarizes the study findings and their validation of the usefulness of the hybrid survey approach.

II. Three Approaches for High-Frequency Poverty Measurement in Bangladesh

This paper examines three low-cost approaches to increase the frequency of poverty data using Bangladesh 2000 and 2005 Household Income Expenditure Surveys (HIESs) data: Subsample Imputation, Subsample Direct Estimation, and Full Sample Imputation.

The Subsample Imputation approach estimates consumption models from a subsample of a hybrid survey, where both consumption and nonconsumption data are collected. It then uses the models to impute household expenditures for all households in the sample, even though there are consumption data for only the households in the subsample. Poverty rates are estimated from the imputed expenditures.

The advantage of this approach is that the consumption models are estimated from current year data. So, if a crisis hits a country and changes consumption patterns, the consumption models are also updated. The SE of this estimation includes both imputation errors and sampling errors, but the latter can be reduced significantly by ensuring a sufficiently large sample size of the overall hybrid survey. Because collection of the nonconsumption data is significantly cheaper, the cost of implementing such a survey is far lower than implementing a HIES. The imputation errors are expected to be non-negligible, although the literature on poverty mapping suggests that with proper care, imputation errors can be small, even when consumption models are estimated with a very small sample size.

The Subsample Direct Estimation approach simply estimates poverty statistics from the consumption data in the subsample of a hybrid survey. This approach is easier than the Subsample Imputation approach because it does not estimate imputation models and simulate poverty rates. However, there is a challenge, because the poverty estimates will likely come with higher SEs than the official poverty estimates due to the small sample size.

The Full Sample Imputation estimates consumption models using the previous round of household survey, like the HIES, where both consumption and nonconsumption data are collected, and imputes household expenditure in the future round where only nonconsumption data are collected. Like the Subsample Imputation, this approach will also face two types of errors—imputation and sampling errors.

The sampling errors can be reduced by increasing the size of the sample collected in the future round. The imputation errors can also be low since the models are estimated using a full sample. However, this approach requires the assumption that underlying parameter estimates in the imputation models remain stable over time. This could become controversial if a large structural change or economic shock occurs between a household survey with consumption data and a survey without consumption data. In the face of such shocks, the consumption models estimated from the household survey with consumption data might not reflect current consumption patterns accurately.

Data and “Full Sample Direct Estimation” Approach as a Benchmark

Bangladesh has been conducting a nationally representative multitopic household survey every five years since 1995; poverty statistics are available every five years. Both the 2000 and 2005 HIESs have a very similar questionnaire and follow the same sampling design. Food consumption data were collected using daily diary forms for 200 items and weekly for 30 items. Nonfood consumption data were collected using monthly and yearly recall forms for about 50 and 150 nonfood items, respectively. Enumerators stayed in primary sampling units (PSUs) for around two weeks, and visited each household every two days to collect information on daily food consumption for two days. Also, enumerators visited the same households for two more days to collect data on nondaily food and nonfood consumption. The 2000 HIES has 14 strata and includes 7,440 households, while the 2005 HIES has 16 strata and includes 10,800 households. The number of PSUs ranges across strata. Each PSU had 10 or 20 households in 2000, and exactly 20 households in 2005. In total, the 2000 and 2005 HIESs had 442 and 504 PSUs, respectively.² Both the 2000 and 2005 HIESs selected PSUs following the probability-proportional-to-size (PPS) sampling while selecting households using the simple random sampling (SRS) from the selected PSUs.

The hybrid survey is constructed from the 2005 HIES data. Subsamples are drawn from the 2005 HIES data and include both consumption and nonconsumption data—procedures for selection of subsamples are explained in the next section. The full sample is the same as the 2005 HIES data, and except for households selected for the subsample, has only nonconsumption data. For the Full Imputation approach, consumption models are developed from the full sample of the 2000 HIES data. The models are applied to impute household expenditures into the entire 2005 HIES sample using nonconsumption data.

All three approaches produce poverty headcount rates at the national level and for urban and rural areas in 2005. A benchmark for all three approaches is the poverty statistics estimated from actual consumption data of the 2005 HIES data. This benchmark is the Full Sample Direct Estimation approach. As shown in table 1, the poverty headcount rate is 40.0 percent for the entire country, and 43.8 and 28.4 percent in rural and urban areas, respectively. The SEs are 1.1 percent for the entire country, and 1.3 and 1.9 percent, respectively, in the rural and urban areas.

Table 1: Poverty Headcount Rates Using Full Sample of 2005 HIES

Full sample	Poverty rates	95% CI		Range of CI
National	40.0 (1.1)	37.8	42.2	4.5
Rural	43.8 (1.3)	41.3	46.3	4.9
Urban	28.4 (1.9)	24.6	32.1	7.5

Source: Authors’ estimations using 2005 HIES data.

Note: Parentheses are SEs. CI = confidence interval.

Subsample Selection in a Hybrid Survey

To preserve the sampling design in the 2005 HIES, we followed its sampling as much as possible. For example, like the 2005 HIES, PSUs of a subsample were selected following PPS sampling (see Lohr

[1999, 108] for more details). For selection of households in select PSUs, we used the SRS and selected 10 households per PSU. We reduced the 20 households in the 2005 HIES to 10 to reduce the sampling error. To see the effects of sample size, this analysis used two different sizes for the subsample: 320 and 640. Note that increasing the size of the subsample increases the cost of data collection, but likely improves the precision of poverty estimates as well.

Table 2: Sampling Schemes and Biases in Poverty Estimates

Sampling scheme	Poverty rate
1. PPS to select 4 PSUs from each stratum and PPS to select 10 households each	42.2 (2.4)
2. PPS to select 4 PSUs from each stratum and SRS to select 10 households each	40.4 (3.7)
3. PPS to select PSUs of different size from each stratum (64 PSUs in total) and SRS to select 10 households	40.2 (2.9)

Source: Authors' estimations from 640 households from the 2005 HIES.
Note: SEs are in parentheses.

The process of selecting subsamples impacts the results. If we use actual consumption data from the 2005 HIES, the national poverty headcount rate is 40 percent (table 1). If we select the same number of PSUs (or four PSUs) from each of all 16 strata and select 10 households by PPS, then the poverty rate from the subsample is biased upward (table 2). If we keep the same number of PSUs from each stratum, but select households by SRS, then the national poverty rate is close to the true rate. Finally, if we use a different number of PSUs from each stratum based on population, then the national poverty estimate is very close to the true estimate and the SE also declines. This analysis used the third approach in table 2 to select a subsample from the 2005 HIES.

Imputation Process for Subsample

The Subsample Imputation and Full Sample Imputation both involve imputation of household expenditures from an ancillary survey. The imputation process follows methodology developed in ELL (2003), which has two stages: first, a model of log per capita consumption expenditure of household h ($\ln y_h$) is estimated in a subsample of the hybrid survey (in the case of Subsample Imputation) or the 2000 HIES data (in case of Full Sample Imputation). The model can be defined as

$$\ln y_{ch} = \alpha + X'_{ch}\beta + Z'_c\gamma + u_{ch}$$

where α is an intercept, X'_{ch} is the vector of explanatory variables for household h and location c , and β is the vector of regression coefficients, Z' is the vector of location specific variables, γ is the vector of coefficients, and u_{ch} is a residual. This residual is decomposed into two independent components: $u_{ch} = \eta_c + \varepsilon_{ch}$ with a cluster-specific effect, η_c , and a household-specific effect, ε_{ch} . This structure allows both a location effect—common to all households in the same area—and heteroskedasticity in the household-specific errors. The location variables can be at any level—*zila*, *upazila*, union, *mauza*, and village (in the case of Bangladesh)—and can be drawn from any data sources that include all locations in the country. All parameters regarding the regression coefficients (α, β, γ) and distributions of the disturbance terms are estimated by feasible generalized least squares. In the second stage, poverty estimates and their SEs are computed. There are two sources of errors involved in the estimation process: errors in the estimated regression coefficients ($\hat{\alpha}, \hat{\beta}, \hat{\gamma}$) and in the estimated residuals, both of which affect poverty estimates and the level of their accuracy. ELL (2003) propose a way to properly calculate poverty estimates as well as their SEs while taking into account these sources of bias. A simulated value of expenditure for each census household is calculated with predicted log expenditure $\hat{\alpha} + X'_h\hat{\beta} + Z'\hat{\gamma}$ and

random draws from the estimated distributions of the residuals, η_c and ε_h . For each subsample, these simulations are repeated 100 times. The mean across the 100 simulations of a poverty statistic provides a point estimate of the statistic for the subsample. To respond to apparent differences in consumption patterns between urban and rural areas, separate consumption models for urban and rural areas are estimated.

Estimation of Poverty Headcount Rates and Their Standard Errors: Subsample Imputation and Subsample Direct Estimation

With either the Subsample Imputation or Subsample Direct Estimation, subsample selection can have a significant impact on poverty estimates. To see the statistical performance, we randomly drew subsamples and repeatedly estimated poverty rates with the subsamples. For the Subsample Direct Estimation approach, we estimated poverty numbers from each subsample and used the mean of poverty rates as the estimate of poverty rate and the standard deviation as the SE.

For the Subsample Imputation approach, we estimated consumption models from each subsample and estimated poverty rates using household expenditure imputed from the models. For the first 30 subsamples, we estimated the consumption models using PovMap2 software, and so were able to follow ELL's method exactly. However, although most results become stable after the exercise is repeated 30 times, there remains a concern on the robustness of the findings.

To see the robustness of the results from the 30 observations, we simplified the imputation method slightly and repeated it 1,000 times. Like ELL, we assumed the error structure consists of the cluster (η_c) and household errors (ε_{ch}). But, although ELL and PovMap2 allow for a flexible error structure, we assumed both cluster and household level errors are normally distributed with zero means.

The variances of the errors were estimated as follows: first, residuals from the ordinary least squares (OLS) estimation, e_{ch} , were calculated. Secondly, the empirical cluster effects, e_c , were calculated using the residuals within each cluster. The SE of η_c is equal to the standard deviations of the empirical cluster effects (e_c). Third, the empirical household effect was defined by the difference between the residuals and the empirical cluster errors, such that $\varepsilon_{ch} = e_{ch} - e_c$. Using the obtained two SEs of η_c and ε_{ch} , the two error terms were randomly produced from two normal distributions and the imputed consumption was calculated using this formula:

$$\ln y_{ch} = \alpha + X'_{ch} \hat{\beta} + Z'_c \hat{\gamma} + \eta_c + \varepsilon_{ch}$$

Such simplifications will increase the SEs of coefficients and model errors.

Finally, although the first 30 simulations trimmed extreme values, the following 1,000 simulations did not control them. Such trimming likely has a small effect on the estimation of poverty headcount rates, although it can have a significant impact on inequality measures, such as Gini coefficients. Because this paper focuses on poverty headcount rates, ignoring trimming is of less concern.

All in all, the 1,000 simulations following the first 30 simulations likely increased the SEs of poverty statistics estimated by the Subsample Imputation approach, but can be used to check whether the results of the first 30 simulations are vulnerable to the small sample errors.

SEs estimated by ELL's method need to be modified to incorporate sampling errors because the ELL method was originally used to impute consumption for all households in census data. As a result, estimated SEs consist only of modeling errors and do not include sampling errors. However, in both Subsample Imputation and Full Sample Imputation, we imputed consumption for all households in a household survey, and this means that the estimated SEs should include sampling errors that are not originally incorporated into the ELL method. To incorporate sampling errors, we used the SEs from the Full Sample Direct Estimation approach, which included only sampling errors, and added them to the SEs obtained from the ELL method. We assumed zero correlation between the modeling and sampling errors since we could not predict the sign or size of the correlation.

III. Results

This section presents the results of the Subsample Direct Estimation, Subsample Imputation, and Full Sample Imputation.

Table 3: Results from the First 30 Replications, Using Subsamples of 640 Households

Levels/Areas	Poverty rate		95% CI		Range of CI
<i>Subsample Imputation</i>					
National	39.4	(2.0)	35.4	43.4	7.9
Rural	43.1	(2.5)	38.3	47.9	9.7
Urban	28.2	(3.1)	22.2	34.2	12.0
<i>Subsample Direct Estimation</i>					
National	40.2	(2.7)	35.0	45.5	10.5
Rural	43.8	(3.2)	37.5	50.1	12.6
Urban	29.0	(4.5)	20.2	37.8	17.6

Source: Authors' estimations using 2005 HIES data.
Note: Cumulative means of poverty estimates and SEs (standard deviation of estimates) with 30 replications are shown. Estimates for 30 replications are shown in appendix table A2.2. Numbers in parentheses refer to SEs.³

Results Using Subsample Imputation and Subsample Direct Estimation

Table 3 shows the cumulative means of the poverty estimates after 30 replications with a subsample of 640 households, out of which 410 households are rural and 230 are urban. The cumulative means of the 30 replications are presented in the poverty rate column and the standard deviations of poverty rates in 30 simulations are reported as SEs in parentheses.

All cumulative means of both the Subsample Imputation and the Subsample Direct Estimation are close to the true or Full Sample Direct Estimation of poverty headcount rates at the national level as well as for urban and rural areas. For example, the cumulative means of the Subsample Imputation are 39.4 percent at the national level and 43.1 and 28.2 percent for rural and urban areas, respectively. The corresponding full sample direct estimates are 40.0 percent (national), 43.8 percent (rural), and 28.4 percent (urban) (table 1). The cumulative means of the Subsample Direct Estimation are 40.2 percent at the national level, and 43.8 and 29.0 percent for rural and urban areas, respectively. Furthermore, all cumulative means from the Subsample Imputation and the Subsample Direct Estimation are in the 95 percent confidence interval (CI) of the corresponding full sample direct estimates (table 1).

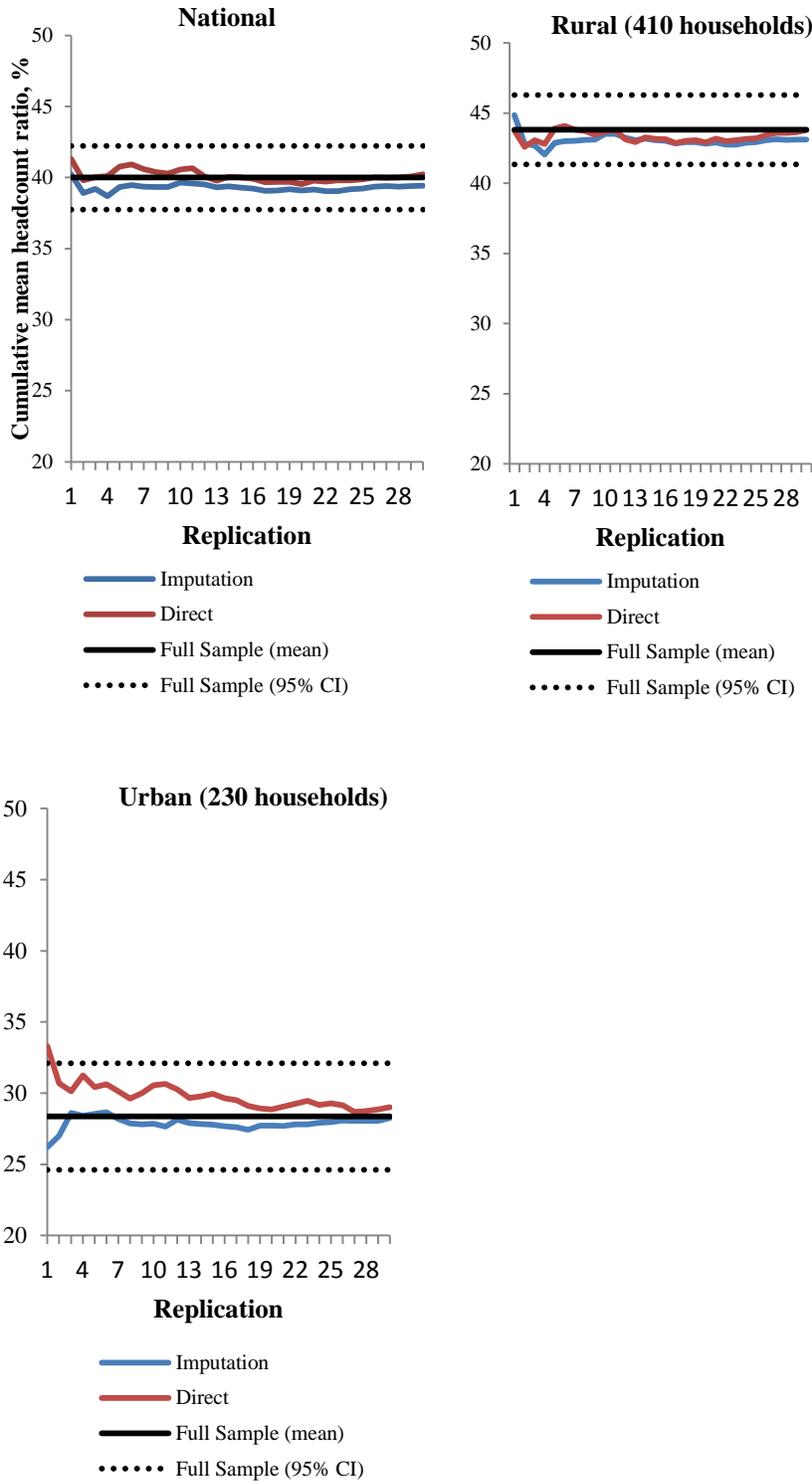
In terms of precision, Subsample Imputation outperforms Subsample Direct Estimation. For example, the SEs of poverty rates at all three levels/areas using Subsample Imputation are 2.0 percent (national), 2.5 percent (rural) and 3.1 percent (urban), which are substantially smaller than those of the Subsample Direct Estimation, which are 2.7 percent (national), 3.2 percent (rural), and 4.5 percent (urban).

Figure 1 plots the cumulative means of poverty estimates using the 640 households in the subsample (410 rural and 230 urban) for each replication. This indicates that both poverty estimates from the Subsample Imputation and Direct Estimation approaches very quickly converge to the “true” poverty rates, or

³ All standard errors of poverty estimates except for Full Sample Direct Estimation do not include the sampling error of poverty estimates in the full sample. For more details, see Fujii and van der Weide (2013) and Matloff (1981).

poverty rates based on the Full Sample Direct Estimation. But, estimates from Subsample Imputation appear to outperform those of Subsample Direct Estimation, particularly for urban areas.

Figure 1: First 30 Simulations for Subsample Imputation and Subsample Direct Estimation



Source: Authors' estimations using 2005 HIES data.

Reducing Subsample Size

Reducing the size of the subsample is expected to worsen the precision of the poverty estimates of both the Subsample Imputation and the Subsample Direct Estimation. Reducing the subsample size increases the sampling errors of poverty estimates from the Subsample Direct Estimation. Also, reducing the subsample size likely increases the SEs of poverty estimates from the Subsample Imputation, because the sample size that can be used to estimate consumption models is reduced. As a result, consumption models are more vulnerable to small sample bias, or the so-called “overfitting” problem. Restricting the number of variables is a typical way of minimizing the risk of small sample bias, but such a strategy tends to worsen the model’s fit.

However, reducing the size of the subsample is a cost-effective strategy for survey implementation. As discussed above, collection of consumption data is costly, particularly so in Bangladesh, where enumerators need to visit households seven times in two weeks during their stay in one village. Halving the size of the subsample with detailed consumption data will reduce survey implementation costs dramatically. This point is discussed more later in this section.

Table 4 shows that cumulative means from both the Subsample Imputation and the Subsample Direct Estimation remain close to those of the Full Sample Direct Estimation. It appears that the Subsample Direct Estimation slightly outperforms the Subsample Imputation in terms of the cumulative means as the size of the subsample declines. The SEs of the Subsample Direct Estimation deteriorate rapidly as the size of the subsample declines while SEs of the Subsample Imputation did not. For poverty estimates at the national level, the estimate from the Subsample Imputation remains reasonably accurate and its SE hardly increases. In this sense, if the government is interested in poverty rates at the national level only, reducing the subsample size to 320 households using the Subsample Imputation approach might be an option.

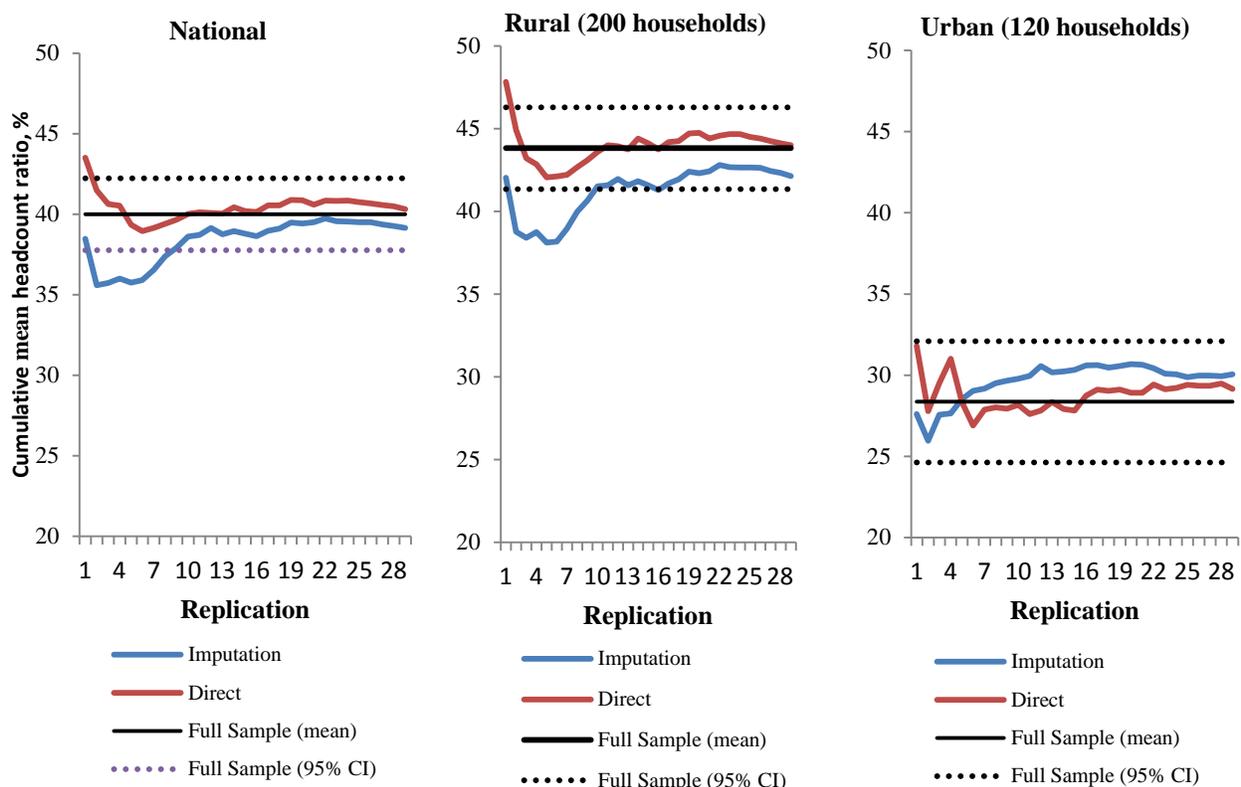
Table 4: Comparing Results from First 30 Replications

Approach	Level/Area	320 Households		640 Households	
Subsample Imputation	National	39.1	(2.1)	39.4	(2.0)
	Rural	42.1	(3.1)	43.1	(2.5)
	Urban	30.0	(3.6)	28.2	(3.1)
Subsample Direct Estimation	National	40.3	(3.5)	40.2	(2.7)
	Rural	44.0	(4.3)	43.8	(3.2)
	Urban	29.2	(5.9)	29.0	(4.5)
Full Sample Direct Estimation ^a	National		40.0		(1.1)
	Rural		43.8		(1.3)
	Urban		28.4		(1.9)

Source: Authors’ estimations using 2005 HIES data.

a. Full Sample Direct Estimation reports poverty rates by using household consumption per capita data from all households included in the 2005 HIES. So, the subsample size does not apply to the numbers. Estimates for 30 replications are available in appendix tables A2.2 and A2.3. Numbers in parentheses are SEs.

Figure 2: Convergence in Poverty Estimates Based on Subsample Imputation and Subsample Direct Estimation



Source: Authors' estimations using 2005 HIES data.

Figure 2 shows the convergence of the cumulative means of poverty rates as the number of replications increases. The cumulative means of both the Subsample Imputation and the Subsample Direct Estimation appear to converge even with only 30 replications. This gives us a certain level of confidence about the results.

Robustness Check for Subsample Imputation and Subsample Direct Estimation for 1,000 Replications

In the above analysis, even with only 30 replications, the cumulative means appear to converge. This is encouraging. Nevertheless, there is no full guarantee that the above results would hold even if the number of replications increased. As a robustness check, we examined whether results would change dramatically after increasing the number of replications.

However, as mentioned before, increasing the number of replications is not a simple task. Implementing ELL's methodology is simplified dramatically thanks to development of PovMap 2 software. But, each replication needs to be implemented manually. For example, users of PovMap2 need to select models and run simulations for each replication separately. It is not possible to repeat the selection of subsamples and estimation of poverty headcount rates 1,000 times. Therefore, we simplified the estimation and simulation process significantly so that we can repeat the process 1,000 times (section II provided an exact description of the simplified process).

Table 5: Results of 1,000 Replications—Cumulative Means and Standard Errors

Approaches	Areas	First 30 replications		Following 1,000 replications		95% CI from Full Sample Direct Estimation ^a		
						Lower	Upper	
		Subsample size = 640 households						
Subsample Imputation	National	39.4	(2.0)	38.4	(2.4)	37.8	42.2	
	Rural	43.1	(2.5)	41.8	(2.9)	41.3	46.3	
	Urban	28.2	(3.1)	27.9	(3.2)	24.6	32.1	
Subsample Direct Estimation	National	40.2	(2.7)	40.4	(2.9)			
	Rural	43.8	(3.2)	44.3	(3.5)			
	Urban	29.0	(4.5)	28.4	(4.6)			
		Subsample size = 320 households						
Subsample Imputation	National	39.1	(2.1)	38.5	(3.5)			
	Rural	42.1	(3.1)	42.6	(4.2)			
	Urban	30.0	(3.6)	29.0	(4.3)			
Subsample Direct Estimation	National	40.3	(3.5)	40.6	(4.1)			
	Rural	44.0	(4.3)	44.1	(5.1)			
	Urban	29.2	(5.9)	29.8	(6.9)			

Source: Authors' estimations using 2005 HIES data.

Notes: Numbers in parentheses are SEs. For each replication, a model is estimated using OLS and it keeps variables with significance levels higher than 0.001 and 0.0001 in rural and urban areas, respectively. More details are available in appendix tables A2.4–A2.6.

a. The 95% CIs are presented in the first three rows only because these numbers remain the same irrespective of approaches or subsample size.

For subsamples of 640 households

The cumulative means of poverty rates and SEs after 1,000 replications are both close to those based on the first 30 replications for both approaches. As shown in table 5, the cumulative means of the Subsample Imputation are 38.4, 41.8, and 27.9 percent for the national, rural and urban levels/areas, respectively. These are slightly underestimated compared with the poverty rates of the first 30 replications and the Full Sample Direct Estimation. However, all estimates from the Subsample Imputation are in the 95 percent CIs of the Full Sample Direct Estimation. The poverty estimates of the Subsample Direct Estimation are 40.4, 44.3, and 28.4 percent for the national, rural and urban levels/areas, respectively. They are very close to those of the Full Sample Direct Estimation. Like the results from the first 30 replications, the SEs of poverty estimates from the Subsample Direct Estimation are higher than those of the Subsample Imputation, and particularly higher for urban areas.

For subsamples of 320 households

The results on the cumulative means of poverty rates are similar to those of the subsample with 640 households. Also, they do not change much from the results of the first 30 replications. The SEs of poverty rates increased significantly by reducing the size of subsamples to 320 households. But, the simpler and less accurate estimation process used for the 1,000 replications resulted in much bigger SEs of the poverty rates when the sample size is reduced. For example, when the Subsample Imputation approach is used, the SE of the national poverty rate increases from 2.2 percent to 3.5 percent after switching from a more accurate estimation method used in the first 30 replications to a less accurate and

simpler estimation method used for the following 1,000 replications. This is the case for SEs of poverty estimates in both urban and rural areas.

Table 6: Comparison of Full Sample Direct Estimation between Only Households Included in both the 2000 and 2005 HIESs and All Households in the 2005 HIES

Area/Level	Poverty estimates from “new” full sample	Poverty estimates from all households in the 2005 HIES
National	39.8 (1.0)	40.0 (1.1)
Rural	43.6 (1.2)	43.8 (1.3)
Urban	28.4 (1.6)	28.4 (1.9)

Source: Authors’ estimations using data from the 2000 and 2005 HIESs.

Note: Numbers in parentheses are SEs. Because of the mismatch of location codes between 2000 and 2005, 860 samples, 60 (8.1 percent of samples in 2005) were dropped for the “new” full sample estimates. Cluster-robust SEs were used.

Full Sample Imputation

Using the ELL method as in Christiaensen et al. (2012), the parameters for consumption models and errors were estimated from the 2000 HIES data and then applied to the 2005 HIES to impute household expenditure, from which poverty headcount rates for 2005 were estimated. To reflect differences in consumption patterns between urban and rural areas, we prepared separate consumption models for urban and rural areas. As discussed above, a critical assumption of the Full Sample Imputation is that the estimated consumption models are stable between 2000 and 2005.

Because of changes in subdistrict boundaries between the two years, a small portion (8.1 percent) of samples in 2005 cannot be used, because we could not match some subdistricts in 2000 to those in 2005. The poverty rates and SEs in 2005 were obtained using the remaining 91.9 percent of the 2005 sample. However, as shown in table 6, the poverty rates in the “new” 2005 full sample are very close to those estimated from all households in the 2005 HIES.

Models with and without ownership of household durables were estimated to test “asset drift.” As explained earlier, the relationship between asset ownership and income or consumption likely changes over time, so Full Sample Imputation with household durables could be affected or biased if an asset drift exists. Household durables are important correlates of income or consumption, so not including these variables could reduce the model’s fit and thus increase the SEs of poverty estimates. To see the potential impact of the asset drift, we compared poverty estimates with and without household durables in consumption models.

Table 7: Poverty Estimates from Full Sample Imputation (%)

Area/Level	With durables		Without durables		Full sample direct estimation ^a
National ^b	38.2	(4.0) ^c	38.9	(3.7)	39.8
Rural	41.8	(5.0)	43.1	(4.6)	43.6
Urban	26.0	(3.1)	24.3	(3.6)	28.4

Source: Authors’ estimations using data from the 2000 and 2005 HIESs.

Notes: Numbers in parentheses are SEs. Povmap2 Software is used for the small area estimation.

a. Full sample refers to the number of households in the 2005 HIES that belong to subdistricts in both the 2000 and 2005 HIES data.

b. Weighted average of rural and urban estimators.

c. SEs are calculated assuming poverty rates of urban and rural areas are not correlated.

Table 7 summarizes the results of the Full Sample Imputation. Inclusion of ownership of durable goods has less impact on the national poverty rate, but it has a noticeable impact on urban and rural poverty rates. Compared with the true poverty rates (or Full Sample Direct Estimation), inclusion of durable

goods tends to underestimate poverty rates in rural areas, while estimating a more accurate poverty rate for urban areas. Dropping ownership of durable goods improves rural poverty rates, but underestimates urban poverty rates.

Finally, we compared the poverty estimates of all three approaches: Subsample Imputation, Subsample Direct Estimation, and Full Sample Imputation. Although poverty rates from the Full Sample Imputation were estimated from a slightly different set of households in the 2005 HIES, we directly compared the poverty rates of all three approaches; table 8 contains the results. For this comparison, the subsample size was 640 households, and the Full Sample Imputation approach excluded ownership of durable goods from both urban and rural consumption models. All three approaches estimated rural poverty rates very well, while the subsample approaches appeared to show better estimations of urban poverty rates.

The Full Sample Imputation approach was worst, underestimating the urban poverty rate by more than 4 percentage points. As a result, the urban poverty rate estimated by the Full Sample Imputation approach was outside the 95 percent CI of that of the Full Sample Direct Estimation. Finally, the Subsample Direct Estimation and Full Sample Imputation both suffered large SEs compared with those from the Subsample Imputation. In all levels, the SEs from Subsample Imputation were smaller than those from the Subsample Direct Estimation and Full Sample Imputation. In particular, the urban poverty estimate in Subsample Direct Estimation and the national and rural poverty estimates in Full Sample Imputation were large. All in all, the best performer was the Subsample Imputation approach using a subsample size of 640 households.

Table 8: Comparison of Three Approaches with the Full Sample Direct Estimation

Approach	Area	Poverty rate	SEs
Subsample Imputation ^a	National	39.4	(2.0)
	Rural	43.1	(2.5)
	Urban	28.2	(3.1)
Subsample Direct Estimation ^a	National	40.2	(2.7)
	Rural	43.8	(3.2)
	Urban	29.0	(4.5)
Full Sample Imputation ^b	National	38.9	(3.7) ^d
	Rural	43.1	(4.6)
	Urban	24.3	(3.6)
Full Sample Direct Estimation ^c	National	40.0	(1.1)
	Rural	43.8	(1.3)
	Urban	28.4	(1.9)

Source: Authors' estimations using data from the 2005 HIES.

a. The size of subsample is 640 households.

b. Consumption models without durables are shown here.

c. Full sample includes all households included in the 2005 HIES data.

d. This SE is estimated assuming there is no correlation between urban and rural poverty rates.

IV. Price Data Analysis

Although the Subsample Imputation approach performs well, there is one more challenge to overcome to implement it in Bangladesh, where poverty lines are estimated and updated for 16 strata separately. When updating the poverty lines, the government calculates an average of food price inflation and nonfood price inflation weighted by averages of the budget shares between two survey years. An important point here is

that food price inflation is estimated with unit values calculated from consumption data in the HIES. Unit values are used because there is a concern that price data collected for constructing the Consumer Price Index (CPI) have limited sample sizes at the stratum level, resulting in potentially large sampling errors. Unit values can fill the data gaps because they can be calculated from consumption data in the HIES, which has a large enough sample for each stratum. Two issues in the use of unit values are: (i) unit values are not available for most nonfood items because households do not have to report quantities for nonfood consumption, and (ii) unit values are known to be quite noisy. As a compromise, the Bangladesh Bureau of Statistics (BBS) uses an urban-rural breakdown of nonfood CPI data to estimate nonfood inflation, and estimates food inflation using unit values to construct a Tornqvist price index while carefully controlling for outliers.³

Price adjustment for food items is a challenge when using the Subsample Imputation approach. The hybrid survey does not collect consumption data, except for the subsample; as a result, unit values can be calculated only from food consumption data in the subsample. Since the size of the subsample is restricted to just 640 households, or less, the number of unit values for each stratum can become very small, resulting in a potentially large sampling error.

This section tests two approaches for conducting price adjustments. First, we followed the BBS's approach using the subsample, and quantified the extent of the potential bias. Second, we collected price data for food items and compared them with unit values. Compared to unit values, price data are often argued to be more reliable. For this experiment, we collected price data and consumption data in a few villages and compared them with unit values.

Using Unit Values in Subsamples

We tested the effect of updating poverty lines using subsamples of poverty estimates and SEs obtained by bootstrapping 100 subsamples. For each subsample, we calculated food inflation rates for 16 strata based on unit values available in the subsample. We then updated 16 poverty lines along with nonfood CPI, and implemented the Subsample Imputation. We compared the resulting poverty rates with those estimated from the Subsample Imputation using the official poverty lines of 2005 based on the full sample of 2005 HIES data. To determine the impact of sample size, we conducted this test with two subsample sizes: 640 and 320 households.

The food price inflation rates were calculated in four steps: first, we defined 13 food groups and calculated average budget shares of these groups for all 16 strata and both the 2000 and 2005 HIES rounds, say w_{0k}^j and w_{5k}^j . Second, we calculated unit values of a representative item (k) in each of 13 food groups and selected the median value across all 16 strata (j) and for both rounds (0 or 5), say p_{0k}^j and p_{5k}^j . Third, the stratum-specific food price inflation rates were calculated following Tornqvist:

$$\ln P_{50}^j = \sum_{k=1}^{13} \frac{w_{5k}^j + w_{0k}^j}{2} \ln\left(\frac{p_{5k}^j}{p_{0k}^j}\right)$$

Finally, stratum-specific inflation rates were calculated by taking averages of the food price inflation rates and the nonfood inflation rates (estimated from nonfood CPI) weighted with average budget shares of food and nonfood items.

In the simulations to follow, the 2005 budget shares and median unit values were estimated from a subsample, although those of 2000 were estimated from its full sample. For the official update of the poverty lines, the 2005 budget shares and median unit values were estimated from the full sample of the 2005 HIES data.

Limited observations of unit values might inflate the SEs of budget shares and median prices of 2005 substantially, or even worse, we might not be able to calculate inflation rates for some strata because no unit value is available for some food groups in some rounds of the simulation. If this happened, we excluded the groups, and reweighted the expenditure shares of the rest of food groups so that the revised weights totaled 1.

Table A3.1 shows the numbers of simulations without any unit value for each food group and each stratum. We can see unit values for cooking oil and soft drinks were missing for several rounds of simulation, but the impact would be negligible since the budget expenditure shares for these items are very small (table A3.2). On the other hand, it is worrisome that unit values of food grains were frequently missing for Rajshahi (SMA) and Sylhet (Municipality) because the budget shares of food grains are large.

Table 9 shows the effects of the limited observations of unit values on the inflation rate and poverty line estimates. The first column shows the inflation rates constructed using the full 2000 and 2005 HIES samples, while the second column shows the inflation rates constructed using the full sample of 2000 and a subsample from 2005. The differences are very limited. For most strata, the differences in inflation rates between the first and the second columns are minimal. As expected, Rajshahi (SMA) and Sylhet (Municipality) exhibit relatively large differences, but they are still below 10 percentage points. For rural and urban areas, the differences are less than 1 percentage point.

Table 9: Estimates of Price Index and 2005 Poverty Lines from subsamples

Stratum	Price index based on:			2005 poverty line based on:		
	Full sample	Subsample (640 HHs, 100 repetitions)		Full sample	Subsample (640 HHs, 100 repetitions)	
		Estimate	SE		Estimate	SE
Barisal (Rural)	1.283	1.277	0.014	915.6	911.2	9.8
Barisal (Municipality)	1.229	1.224	0.046	938.5	935	35.4
Chittagong (Rural)	1.206	1.208	0.018	884	885.3	13
Chittagong (Municipality)	1.153	1.154	0.021	954.8	955.2	17.3
Chittagong (SMA)	1.183	1.177	0.019	1,157.9	1151.5	18.6
Dhaka (Rural)	1.278	1.28	0.011	831.7	832.9	7
Dhaka (Municipality)	1.183	1.183	0.013	877.6	877.5	9.8
Dhaka (SMA)	1.173	1.17	0.023	1,003.9	1001.2	19.8
Khulna (Rural)	1.263	1.263	0.011	735.3	735.4	6.4
Khulna (Municipality)	1.177	1.172	0.016	812.8	809.5	11.1
Khulna (SMA)	1.193	1.197	0.029	922.2	925	22.7
Rajshahi (Rural)	1.270	1.273	0.018	758.8	760.8	10.7
Rajshahi (Municipality)	1.194	1.198	0.019	844.5	847.5	13.4
Rajshahi (SMA)	1.235	1.138	0.05	843	776.7	34.3
Sylhet (Rural)	1.237	1.213	0.04	817	801	26.5
Sylhet (Municipality)	1.201	1.182	0.046	1,012.3	996.4	39
Rural (population weighted)	1.257	1.257	0.016	815.2	814.8	10.7
Urban (population weighted)	1.182	1.178	0.023	965.2	961.6	18.9

Source: Authors' estimations using 2000 and 2005 HIES data.

Note: HH = household.

Another important observation is that SEs at stratum-level inflation rates are quite low. Despite extremely limited observations at the stratum level, SEs of inflation rates are all below 5 percentage points. SEs of rural and urban average inflation rates are just 1.6 percent and 2.3 percent, respectively.

As a result, stratum-specific poverty lines were also quite accurately estimated using subsample data from the 2005 HIES. The fourth column of table 9 shows poverty lines estimated using the full 2000 and 2005 HIES samples, while the fifth column shows those estimated using the full sample of the 2000 HIES and a subsample from the 2010 HIES. Most exhibit less than a 1 percent difference between the fourth and fifth columns. Like above, Rajshahi (SMA) and Sylhet (Municipality) show relatively large differences, but still they are less than 10 percent. If all poverty lines are aggregated for rural and urban areas, the differences become less than 1 percent. Furthermore, the SEs of poverty lines based on the subsamples of 2005 are also, in general, small. All of these observations suggest that poverty lines and inflation rates estimated from the subsamples of 2005 are robust and reliable.

Table 10: Estimated Poverty Rates

Updating poverty lines	Using subsamples of 2005 data				Using the full sample of 2005 data					
	Subsample Imputation ^a		Subsample Direct		Subsample Imputation ^a		Subsample Direct		Full Sample Direct Estimation	
National, 640 households	38.0	(2.5)	38.8	(2.7)	38.9	(2.5)	39.8	(2.7)	40.0	(1.1)
Rural, 410 households	41.2	(3.1)	42.9	(3.4)	42.1	(3.1)	43.9	(3.6)	43.8	(1.3)
Urban, 230 households	28.3	(3.6)	26.4	(4.2)	29.2	(3.6)	27.4	(4.2)	28.4	(1.9)
National, 320 households	38.2	(3.4)	39.6	(4.5)	39.2	(3.4)	40.9	(4.5)	40.0	(1.1)
Rural, 200 households	41.0	(4.3)	42.8	(5.5)	41.9	(4.3)	44.1	(5.5)	43.8	(1.3)
Urban, 120 households	29.7	(5.3)	29.7	(6.8)	31.0	(5.3)	31.2	(7.3)	28.4	(1.9)

Source: Authors' calculations using 2000 and 2005 HIES data.

a. A simplified method, which is described in detail in section II, is used.

Finally, we examined whether updating poverty lines using a subsample of the 2005 HIES data has a large impact on poverty estimates derived from the Subsample Imputation. To determine this impact, we compared poverty estimates from poverty lines and inflation rates estimated from subsamples of 2005 HIES data with those estimated from the full sample of 2005 HIES data.

In the Subsample Imputation described in section III, inflation rates and poverty lines are estimated using the full 2000 and 2005 HIES samples with only the consumption models developed using a subsample from the 2005 HIES. Here inflation rates, poverty lines, and consumption models are all estimated using a subsample from the 2005 HIES. To reduce the computational burden of this analysis, we adopted the simplified projection model (described in detail in section II) rather than a full-fledged ELL method when creating consumption models. Also, to see the effect of sample size, we conducted this analysis with subsamples of 320 and 640 households.

Table 10 summarizes estimated poverty rates. The first column shows poverty estimates from Subsample Imputation when a subsample of 2005 HIES data is used to estimate poverty lines and inflation rates,

while the fifth column shows poverty estimates using the full sample of 2005 HIES data to estimate poverty lines and inflation rates. Like before, the simplified projection method slightly underestimates the national and rural poverty rates, and SEs are slightly higher than when the full-fledged ELL is applied. The difference in poverty rates between the first and the fifth columns is minimal, no more than 1 percentage point for national, rural, and urban poverty rates. This is the case for both sample sizes, 320 and 640. This makes sense, because both inflation rates and poverty lines estimated from a subsample of 2005 HIES data are very similar to those from the full 2005 HIES sample. Also, it is interesting to see that SEs of all poverty rates do not change when a subsample of 2005 HIES data are used to estimate inflation rates and poverty lines. The results of Direct Estimation of poverty rates are also similar.

This analysis confirms that using data from the 2000 and 2005 HIESs, even if only a subsample of 2005 HIES data is used, does not change the inflation rates and poverty lines much. As a result, the results of the Subsample Imputation do not change much. This is very encouraging for the imputation approach.

Using Price Survey Data

Even though estimating inflation rates and poverty lines using a subsample of 2005 HIES data worked well, unit values are known to be noisy and vulnerable to outliers. Price data are often argued to be more reliable than unit values. One of the issues with using price data, which are currently collected for the CPI, is that these data do not have enough observations to create stratum-specific inflation rates. Therefore, the BBS is currently using unit values from HIES data to estimate stratum-specific food price inflation rates. However, the lack of observations can be overcome with relatively limited additional costs. If price data are collected in all PSUs of a hybrid survey, there will be enough observations to create inflation rates that are representative at the stratum level. Furthermore, collecting price data in this manner saves transportation and lodging costs because enumerators are in PSUs anyway to collect other data and collecting price data is not time consuming.

Even though price data can be collected at minimal additional cost, switching to price data can be problematic. Switching to price data can make new poverty estimates incomparable with past estimates if the price data are not comparable to unit values.

To determine the quality and comparability of price data, as well as the cost of collecting price data, the BBS and a World Bank team carried out a pilot survey, which collected price data and consumption data from 60 PSUs and 720 households (12 households in each PSU) in four strata: Dhaka urban, Dhaka rural, Khulna rural, and Rajshahi rural. The pilot was conducted in June and July 2011 with trained enumerators.

Table 11: Difference between Median Unit Values and Prices for Select Items

Items	Unit	Dhaka rural			Khulna rural			Rajshahi rural			Dhaka urban		
		Price	Unit value	% diff	Price	Unit value	% diff	Price	Unit value	% diff	Price	Unit value	% diff
Rice, medium	Kg	35	35	0	35	33	6	32	31	2	36	36	1
Rice, coarse	Kg	32	34	-6	31	31	0	29	30	-3	32	35	-8
Beef meat	Kg	250	250	0	250	250	0	250	250	0	250	259	-4
Chicken meat	Kg	145	150	-3	131	141	-7	140	143	-2	145	147	-1
Milk, liquid	L	40	40	0	30	30	0	30	30	0	50	50	-1
Soybean oil	L	118	120	-2	120	120	0	120	120	0	120	120	0

Source: Authors' estimations using pilot survey data.

Note: Bold numbers indicate the medians are statistically different between survey prices and unit values with p-values less than 5 percent. Shares of differences between survey and unit values = (survey price - unit value) / mean (survey price, unit value). All prices are reported in Bangladesh taka.

The consumption module was identical to that of the 2000 and 2005 HIESs, while the consumption data were collected via two week recall, as well as by diary. Price data were collected using a questionnaire created for this pilot by the World Bank team and a survey company, DATA. Price data include 32 major food and nonfood items.

Table 11 shows the median prices from the price survey and unit values from the consumption module for the six most important food items by four strata (Dhaka rural, Khulna rural, Rajshahi rural, and Dhaka urban). The histograms of these numbers are shown in the appendix (figures A2.1–A2.4). Table 11 shows the differences between the median price and unit value as proportions of their average so that the size of relative difference can be easily discerned. If the relative difference is statistically significant, percentage difference is in bold with p-values less than 5 percent. Positive (negative) differences indicate that prices (unit values) are larger than unit values (prices).

According to table 11, medians of prices and unit values are usually very similar, but reflecting small movements in both values, SEs are small, and even small differences in median prices and unit values become significant. More specifically, in Dhaka rural, the prices and unit values for coarse rice and chicken meat are statistically significantly different, but the relative differences of the two items are only 6 and 3 percent, respectively. In Khulna rural, the noticeable differences in medians between unit values and prices are found in medium size rice and chicken meat, although only the difference for chicken meat is statistically significant. In Rajshahi rural, medium rice is statistically significantly different, but the relative difference is only 2 percent. In Dhaka urban, median prices of coarse rice and beef meat are statistically significantly different, and the relative differences are 8 and 4 percent, respectively.

Overall, medians of unit values can be similar to those of prices if the price surveys were collected from each stratum. This implies that unit values from the subsample might be quite reliable even though the sample size is limited. Also, although we need to see whether the results hold when expanding the sample to the whole country, this analysis suggests that switching to price data might have minimal impacts on price adjustments and poverty estimates.

Bangladesh poverty measurement uses unit values from consumption data of the HIES to estimate inflation rates and poverty lines. Therefore, if consumption data are collected from only a subsample of the HIES data, estimation of inflation rates and poverty lines, and thus the resulting poverty estimates, can be noisy and unreliable. However, the above analysis shows, at least in the case of Bangladesh's 2000 and 2005 data, that using only a subsample of 2005 data does not increase noise in estimated inflation rates, poverty lines, or poverty headcount rates.

Price data collection is simple and less costly, but raises the possibility of incomparability between price data and unit values, which would make future poverty estimates incomparable with past poverty estimates. However, analysis based on the pilot survey reviewed here suggests that price data are comparable to unit values from consumption data. These results suggest that inflation rates and poverty lines can be accurately updated with a subsample of HIES data or additional price surveys. This is certainly encouraging for the Subsample Imputation approach, and even for the Subsample Direct Estimation.

V. Cost Analysis of Three Proposed Approaches and Effects of Switching from Diary- to Recall-Based Data Collection

Results here show that even if the Bangladesh economy did not experience a large economic crisis between 2000 and 2005, the subsample approaches provide better results for 2005 than the Full Sample Imputation that builds consumption models from the previous round of the HIES, that is, from the 2000 HIES data. But, as discussed earlier, collecting consumption data is costly, even for a subsample. Also,

the Subsample Imputation approach needs a large sample of nonconsumption data. The government of Bangladesh recently carried out a simpler version of the HIES to collect nonconsumption data, the Welfare Monitoring Survey (WMS). It simplified the HIES questionnaire significantly and does not include consumption and income data. According to the BBS, WMS implementation costs are only around one-tenth of the HIES.

How do the implementation costs of the HIES compare with those of the WMS and a collection of a subsample of consumption data? Although this cost estimation is admittedly rough, it was carried out in close collaboration with BBS survey experts to incorporate several aspects of the survey data collection process. In terms of the WMS sample size, we tentatively assumed that 12,240 households formed 612 PSUs, following the sampling frame of the latest round of the HIES (2010). Instead of estimating survey implementation costs in local currency, we estimated three key aspects of survey implementation, namely staff weeks, number of moves and days of lodging, which are all closely related to survey implementation costs. For example, the number of moves should be closely related to transportation costs, the days of lodging should be proportional to the total lodging cost, and the staff weeks should be proportional to the total staffing costs.

This cost comparison also included an experiment: switching the mode of consumption data collection from a diary-based approach to a recall-based approach. Consumption data in the HIESs are collected by a diary-based approach, or more precisely, seven consecutive visits by enumerators to a household to collect the last two day's consumption data. In this sense, the approach undertaken in the HIES could also be considered a recall-based approach, but since consumption data are collected in an almost continuous manner for two weeks, we call it a diary-based approach.

Obviously, the diary approach is very costly since a group of enumerators needs to stay in a village for 14 days to collect consumption data. If consumption data were collected in one visit with two weeks of recall, then the cost of implementation would be much cheaper. But this can cause comparability problems between the consumption data and data from previous rounds. Beegle et al. (2011) reviewed previous studies on this method of consumption measurement and noted that “well-implemented diary surveys might be expected to yield higher (and presumably closer to actual) levels of consumption, experimental evidence for this from developing countries is fragmentary.” The following analysis finds cost savings from switching from the diary approach to the two-week recall approach, and shows the impact of switching on comparability, based on analysis from pilot survey data.

Costs of WMS as Baseline and Assumptions for Estimation

WMS costs were used as a baseline. In the WMS, consumption data are not collected, and without the consumption module the survey is shorter than the HIES. Therefore, after consultation with BBS survey experts, we assumed that one enumerator could handle the survey in two PSUs for one week. So, the total staff weeks for 612 PSUs was equal to 306 ($612 \text{ PSUs} \times 0.5 \text{ week}$). In addition, the number of movements is assumed to be proportional to PSUs covered by the number of enumerators, and the days of lodging are assumed to be $(\text{the number of PSUs}) \times (\text{the number of weeks}) \times 7$. So, the number of movements is 612, and the days of lodging are 2,142 ($= 612 \text{ PSUs} \times 0.5 \text{ week} \times 7$) days.⁴

HIES Costs Compared to WMS Costs

Compared to the WMS, the HIES questionnaire is much more detailed and comprehensive. For example, even if the questionnaires of both the WMS and HIES include an education section, the HIES ask many more questions in the section than the WMS. Also, the HIES questionnaire includes questions regarding economic activities and agricultural and nonagricultural enterprises, which are not included in the WMS. Collection of consumption data for the HIES requires seven visits to each household in two weeks. After consultation with BBS survey experts, we assumed as reasonable the following regarding staff weeks, number of moves, and days of lodging:

- Five enumerators have to stay for two weeks to complete the survey in one PSU. Among the five enumerators, two persons engage in collecting (diary-based) consumption data, and three persons collect nonconsumption data for two weeks.
- Since the HIES has 612 PSUs, the total number of staff weeks in the HIES is equal to 6,120 staff weeks (3,672 for nonconsumption + 2,448 for consumption parts).
- For the number of movements and days of lodging, we assumed that the enumerators for consumption and nonconsumption parts move together and can share transportation, but not lodging. So, the number of movements is independent of the number of staff in a PSU, but the days of lodging depend on staff weeks in which all staff stay in a PSU. As a result, the number of movements is 612, and the days of lodging are 42,840 (= 5 enumerators × 2 weeks × 612 PSUs × 7 days).

Table 12 confirms that the HIES costs much more than the WMS. HIES data are attractive because the consumption data are available for all households, but the required staff work weeks and days of lodging for the HIES are 20 times more than those of the WMS.

Costs of Collecting Consumption Data from WMS Subsamples

Collecting consumption data from just small subsets of WMS households instead of from all households in the survey is obviously a cost-saving option. Depending on the size of subsamples, 64 and 32 PSUs were selected. Unlike the sampling of the HIES, 10 (not 20) households were chosen from each PSU. Another important assumption was that collection of consumption data was carried out separately from data collection of other modules. The first WMS data collection was conducted over one month, since most questions were not affected by seasonality. However, consumption data are likely to fluctuate in a year, so collection of the current WMS data was separated from the collection of consumption data from a WMS subsample.

Table 12: Comparison of Survey Implementation Costs

Options	Number of PSUs		Cost			Comparison to HIES 2010		
	Total	With subsample	Staff weeks	No. of moves	Days of lodging	Staff weeks (%)	No. of moves (%)	Days of lodging (%)
Full 2010 HIES	612	612	6,120	612	42,840	100	100	100
WMS only	612	NA	306	612	2,142	5	100	5
WMS with a subsample of 640 households	612	64	434	676	3,038	7	110	7
WMS with a subsample of 320 households	612	32	370	644	2,590	6	105	6
WMS with a subsample of 640 households (recall)	612	64	370	676	2,590	6	110	6

Source: Authors' projections after consultations with BBS.

Since the number of sample households in a PSU was reduced from 20 to 10, we assumed only half the staff weeks would be needed to collect consumption data from each PSU, which implies a total of 128 staff weeks needed for collecting consumption data from 640 sample households. If the size of the subsample was set at 320, then we would need a total of 64 staff weeks. As for the number of moves, enumerators needed to move among PSUs 64 times if the size of the subsample size is 640, or just 32 times for a subsample of 320. As for the days of lodging, 896 days were needed for the subsample of 640 households, or 448 days for the subsample of 320 households.

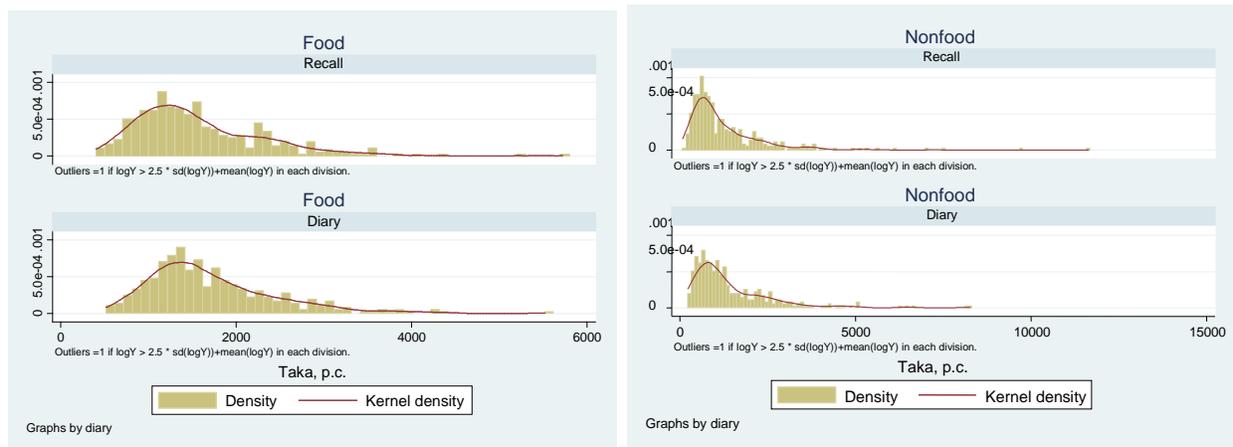
Even if the additional requirements in staff weeks, number of moves, and days of lodging are added, the cost savings from the WMS with a subsample of consumption data seem significant. For example, if the subsample size is set at 640, the combination needs only 7 percent of staff weeks and days of lodging required for HIES data collection, although the number of moves needed for the combination is projected to be slightly higher, 110 percent of the HIES (table 12). The latter is due to separate data collections for consumption data and nonconsumption data. If the size of the subsample is set at 320, the WMS with a subsample of consumption data needs even fewer staff weeks, moves, and days of lodging.

Diary-Based Versus Recall-Based Consumption Data Collection

As discussed above, there is potentially a large cost-saving effect by switching the current diary-based data collection to a recall-based collection. Currently, two enumerators need to visit each household seven times during their two-week stay in a PSU. This implies that if a two-week, recall-based collection is selected, in theory, an enumerator would need only two days to cover 20 households in a PSU. Since the number of households for a subsample of the WMS is reduced to 10 households, theoretically, one enumerator could collect consumption data from all 10 households in one day. In the cost projection above, we made a more conservative assumption—that one staff week would be needed to collect consumption data from all households in each PSU. As a result, if the subsample size is set at 640, then the WMS with a subsample of consumption data needs a total of 370 staff weeks, 676 moves among PSUs, and 2,590 days of lodging, which correspond to 6 percent, 110 percent, and 6 percent, respectively, of the HIES.

As mentioned earlier, how consumption data are collected affects how households respond. In general, we do not know whether a switch from a diary-based collection to a recall-based collection increases reporting of consumption expenditures (Beegle et al. 2011). Such a switch often reduces consumption aggregates simply because households remember their expenses better when entering them regularly in a diary.

Figure 3: Comparison of Distributions of Food Expenditures and Nonfood Expenditures between Diary- and Recall-Based Data Collections



Sources: Authors' estimations using pilot data.

To examine the effects of switching the method of collecting consumption data, the BBS and the World Bank conducted a pilot survey. To investigate the comparability of consumption aggregates when diary-based consumption data collection is replaced by recall-based collection, the two collection methods were randomly assigned to half of the households in each PSU. In four strata: Dhaka urban, Dhaka rural, Khulna rural and Rajshahi rural, 720 households from 60 PSUs (12 households from each PSU) in the

2010 HIES were surveyed. Within each PSU, six households were randomly allocated to each the diary and recall collections.

The experiment shows that per capita food consumption collected with the diary survey is higher than that of the recall survey. Figure 3 shows the histograms and kernel distribution of per capita food consumption for all households. The natural logarithm was taken and outliers removed. The mode of the distribution for the diary approach is to the right of that of the recall approach.

Table 13 contains the medians of per capita total, food, and nonfood consumption that were collected by the diary and recall methods for all four areas. Medians were used because they are less vulnerable to outliers. For all the areas combined, the median per capita food consumption from the diary collection was 10.1 percent larger from the recall collection. This difference is statistically significant. In all four areas, median per capita food consumption expenditure collected by diary was larger than that collected by recall although the difference is not statistically significant except for the Khulna rural stratum, possibly because of limited sample sizes. Results from the full sample suggest that collecting consumption data via the recall would result in the underestimation of per capita food consumption.

Table 13: Comparison between the Diary and Recall Based Collections for Food and Nonfood Expenditure per Capita (medians)

Areas	Category	Per capita, taka, median		Difference between diary and recall	
		(D)iary	(R)ecall	Difference	% of D - R
All	Food	1,557.0	1,406.8	150.3	10.1
	Nonfood	1,061.7	908.0	153.7	15.6
	Total	2,662.0	2,414.6	247.4	9.7
Dhaka rural	Food	1,570.3	1,455.7	114.7	7.6
	Nonfood	1,051.4	833.8	217.7	23.1
	Total	2,650.7	2,326.1	324.6	13.0
Khulna rural	Food	1,373.4	1,208.1	165.3	12.8
	Nonfood	950.9	750.1	200.8	23.6
	Total	2,448.9	2,097.6	351.2	15.4
Rajshahi rural	Food	1,454.0	1,369.5	84.5	6.0
	Nonfood	875.0	909.5	-34.5	-3.9
	Total	2,466.1	2,394.8	71.3	2.9
Dhaka urban	Food	2,115.2	1,913.5	201.7	10.0
	Nonfood	2,125.7	1,597.7	528.0	28.4
	Total	4,127.6	3,499.1	628.5	16.5

Source: Authors' calculations based on pilot survey.

Note: Outliers are defined as = 1 if $\log y > 2.5 \times \text{sd}(\log Y) + \text{mean}(\log Y)$ in each division. P-values are calculated based on bootstrap resampling in the median regression. Percentage of D - R = $(D - R) / \text{mean}(D, R)$.

An interesting observation is that even though nonfood consumption data were collected by the recall, irrespective of how food consumption data were collected, median per capita nonfood consumption expenditure was larger when food consumption was collected via the diary approach. This is puzzling, but a similar result in the nontreated part (nonfood) was also found in Beegle et al. (2011).

Switching the data collection methods from diary to recall likely underestimates household consumption data, overestimates poverty rates, and thus makes poverty estimates incomparable with those of previous

rounds in which consumption data were collected by diary. Therefore, even if the recall approach is cost saving, it is not recommended for the calculation of poverty estimates.

VI. Conclusion

Despite suggestive empirical evidence provided by Christiaensen et al. (2012), imputing household expenditure or income using models estimated from past data is risky because there is no guarantee that consumption models will not change over time. For example, even Christiaensen et al. (2012) show the consumption models estimated from past data could not predict future poverty data in the Russian Federation. They argued that the divergence in poverty rates between direct estimation and imputation can be attributed to questionable inflation adjustments, but it is difficult to prove that the inflation adjustment is the only source of the divergence.

The biggest challenge for the existing imputation approach is that if consumption or income data are not collected, there is no way to test whether the imputed poverty rate is correct until new consumption data are collected in the future. Christiaensen et al. (2012) could test the accuracy of their imputation model because they could compare the imputed consumption data with the actual consumption data. But, if we truly move to the imputation approach, except for big survey years, only nonconsumption data will be collected and poverty rates will be estimated using imputed consumption data. The accuracy of the poverty rates calculated using imputed consumption data cannot be tested in years when only nonconsumption data are available.

This paper provides empirical evidence that a hybrid approach—collecting consumption data for a subsample only—can be extremely useful in improving the statistical reliability of poverty data. For example, collecting consumption data from a subsample enables us to produce consistent but likely noisy poverty estimates by directly estimating them from the subsample. These estimates can be used to examine whether imputed poverty rates are accurate. Furthermore, this analysis shows that if consumption models are developed from the subsample, poverty rates can be estimated reliably by applying the consumption models to a large sample of the nonconsumption data. The estimates based on the Subsample Imputation approach are shown to be as precise as, or even better than, those based on models estimated from past data, even if the size of a subsample is as small as 320, at least as far as poverty estimates at the national level are concerned. Another advantage of the Subsample Imputation approach is that consumption models reflect the contemporaneous relationship between consumption data and nonconsumption data, which is not the case for the approach proposed and tested by Christiaensen et al. (2012).

Subsample Imputation, or even Subsample Direct Estimation, is very attractive, but there is one challenge to overcome, particularly in the case of Bangladesh, where poverty lines are updated with inflation rates estimated for each of the 16 strata separately. The official method constructs food inflation rates based on unit values calculated from HIES consumption data. If consumption data are collected only from a subsample, the number of unit values will be very limited, and thus, the reliability of unit values will be questionable. This is particularly true because unit values are known to be much noisier than price data.

However, evidence from the pilot survey suggests that unit values for key consumption items are not so noisy and quite similar to price data. Also, this analysis shows that even if unit values from a subsample are used to estimate inflation rates and poverty lines, the estimates are quite accurate. As a result, poverty rates estimated via either the Subsample Imputation or Subsample Direct Estimation are quite precise, even when the consumption models as well as the price adjustments are made using consumption data from a small subsample.

When determining which approach to use, implementation cost is also an important consideration. Subsample Imputation needs a large sample with nonconsumption data, and a small subsample with both consumption and nonconsumption data. If the subsample size is 640, then this combined survey needs

only 7 percent of the staff weeks and days of lodging required for typical HIES data collection although it requires a slightly higher transportation cost—110 percent of the transportation costs required for HIES data collection. The transportation cost is slightly higher than that of HIES because consumption and nonconsumption data need to be collected by different teams. Although the total cost saving depends on costs of lodging, staff time, and transportation, it is likely that the cost of data collection for Subsample Imputation is likely to be much lower than for HIES data collection.

These results are encouraging for the Subsample Imputation approach. But, it is also worth noting that all the analyses were conducted for a specific time period, for a specific country, that is, Bangladesh in 2000 and 2005. In a different time period, or for a different country, the superiority of the Subsample Imputation approach might disappear. However, this paper provides an analytical framework to determine the magnitude of cost savings and statistical advantages of Subsample Imputation over other approaches. Furthermore, to implement the Subsample Imputation approach requires a modeling exercise every year, the burden of which cannot be underestimated in some developing countries where technical and analytical capacity of statistics bureau is often severely limited.

Finally, this is not the first evaluation of the statistical precision and cost-effectiveness of the hybrid approach. Fujii and van der Weide (2013) studied the hybrid approach in a quite general environment. For the most part, these two papers are mainly complementary and reinforce the evidence of the hybrid approach's usefulness. However, there are differences: first, this paper compares the existing survey-to-survey imputation approach, which is identical to Full Sample Imputation, with the new hybrid approach. Fujii and van der Weide (2013) focus on derivation of properties of statistics estimated from the hybrid approach. Second, one potential shortcoming of the hybrid approach is that it needs to conduct price adjustments and update poverty lines using a limited number of unit values available in a subsample. Practically, this is an important challenge since it might expand the margin of error in poverty estimates. This paper studied the effect of limited sample size, while Fujii and van der Weide (2013) did not. Third, this paper conducts all analyses using data from Bangladesh 2000 and 2005, while the results from Fujii and van der Weide (2013) are general. All in all, these two papers fill each other's gaps and together build a strong case for a hybrid survey approach.

References

- Allwine, M., H. Uematsu, S. Takamatsu, and N. Yoshida. 2013. "A Technical Note on Lesotho Poverty Measurement," Mimeo. Poverty Reduction and Equity, World Bank.
- Beegle, K., De Weerd, J., Friedman, J. & Gibson, J. 2011. "Methods of household consumption measurement through surveys: Experimental results from Tanzania." *Journal of Development Economics*, 98(1): 3-18.
- Christiaensen, L., P. Lanjouw, J. Luoto, and D. Stifel. 2012. "Small Area Estimation-Based Prediction Methods to Track Poverty: Validation and Applications." *Journal of Economic Inequality* 10 (2): 267–97.
- Dang, H. and P. Lanjouw. 2013. "Measuring Poverty Dynamics with Synthetic Panels Based on Cross-Sections." *Measuring Poverty Dynamics with Synthetic Panels Based on Cross-Sections*, Policy Research Working Paper Series. No. 6504. World Bank.
- Deaton, A. and J.P.Dreze. 2002. "Poverty and Inequality in India: A Reexamination" *Economic and Political Weekly*, September 7, 2002.
- Deaton, A., and M. Grosh. 1999. "Chapter 17: Consumption." In *Designing Household Survey Questionnaires for Developing Countries: Lessons from Ten Years of LSMS Experience*, ed. Margaret Grosh and Paul Glewwe. Washington, DC: World Bank.
- Dorji, C., and N. Yoshida. 2011. "New Approaches to Increase Frequent Poverty Estimates." Unpublished manuscript.
- Doudich, M., A. Ezzrari, R. Van der Weide, and P. Verme. 2013. "Estimating Quarterly Poverty Rates Using Labor Force Surveys: A Primer." Policy Research Working Paper Series 6466, World Bank, Washington, DC.
- Elbers, C., J. O. Lanjouw, and P. Lanjouw. 2003. "Micro-Level Estimation of Poverty and Inequality." *Econometrica* 71 (1): 355–64.
- Fujii, T., and R. van der Weide. 2013. "Cost-Effective Estimation of the Population Mean Using Prediction Estimators." *Policy Research Working Paper* 6509.
- Harttgen, K., S. Klasen, and S. Vollmer. 2012. "An African Growth Miracle? Or: What Do Asset Indices Tell Us about Trends in Economic Performance?" *Poverty, Equity, and Growth Discussion Paper* 109, Courant Research Centre.
- Kijima, Y. and P. Lanjouw. 2003. "Poverty in India during the 1990s - a regional perspective," Policy Research Working Paper Series 3141, The World Bank.
- Lohr, S. 1999. *Sampling: Design and Analysis*, First Edition. Belmont, CA: Brooks Cole Publishing.
- Matloff, N. 1981. "Use of Regression Functions for Improved Estimation of Means." *Biometrika* 68: 685-689.
- Ravallion, M. 1996. "How Well Can Method Substitute for Data? Five Experiments in Poverty Analysis." *World Bank Research Observer* 11 (2): 199–221.
- Revilla, J., R. Katayama, N. Yoshida, and L. Fox. 2010. "Analysis of Poverty Trends using Poverty Mapping (Small Area Estimation-Based) Methods," available in *Zambia Poverty Assessment: Stagnant Poverty and Inequality in a Natural Resource-Based Economy*, 2012, World Bank, Washington DC, USA.
- Stifel, D. and L. Christiaensen. 2007. "Tracking Poverty Over Time in the Absence of Comparable Consumption Data." *World Bank Economic Review* 21 (2): 317–41.

Appendix 1.

Robustness Check 2: Different Model Specification in Subsample Imputation Approach

During the first 30 repetitions, district dummies were used in the consumption models, although all of the 64 districts in full samples were not chosen in subsamples (12 districts). This can be a problem because the reference district will change from 1 district from the 12 districts selected for a subsample to all districts not included in the subsample. As a result, households in all districts not included in the subsample receive zero from the district effect.

To examine the impact of this misspecification, using data from the 1,000 replications, poverty rates and their SEs are imputed using three different specifications of consumption models: (i) without division and district dummies, (ii) with division dummies, and (iii) with district dummies. Table A1.1 summarizes the results and shows that simulation results hardly change at all among the different specifications of consumption models. This analysis shows that including district dummies in the Subsample Imputation approach for the first 30 replications has a minimal effect on poverty estimates.

Table A1.1: Poverty Headcount Rates with Different Model Specifications

Size of subsamples	640 households		320 households	
<i>Without division and district dummies</i>				
National	38.7	(2.1)	38.8	(3.4)
Rural	42.1	(2.7)	41.9	(4.4)
Urban	28.1	(2.7)	29.3	(4.1)
<i>With division dummies</i>				
National	38.4	(2.1)	38.8	(3.3)
Rural	41.8	(2.6)	42.1	(4.2)
Urban	27.9	(2.6)	28.9	(4.0)
<i>With district dummies</i>				
National	38.4	(2.0)	38.8	(3.4)
Rural	41.8	(2.5)	41.9	(4.3)
Urban	28.1	(2.6)	29.2	(4.0)

Source: Authors' estimations using 2005 HIES data.

Note: The significance levels for variables to be included are 0.01 and 0.001 in rural and urban areas, respectively. The national poverty rates are the weighted averages of rural and urban poverty rates; 1,000 replications were used. Parentheses are SEs of poverty estimates.

Appendix 2. Data Tables and Figures

Table A2.1: Sampling Structure of the 2005 HIES

Stratum no.	Urban/rural	PSUs	Households in each PSU	Total households
1	Rural	28	20	560
2	Urban	13	20	260
3	Rural	58	20	1,160
4	Urban	23	20	460
5	Urban	9	20	180
6	Rural	86	20	1,720
7	Urban	37	20	740
8	Urban	24	20	480
9	Rural	44	20	880
10	Urban	22	20	440
11	Urban	7	20	140
12	Rural	85	20	1,700
13	Urban	36	20	720
14	Rural	5	20	100
15	Urban	19	20	380
16	Urban	8	20	160
Total		504		10,080

Source: Authors' estimations using 2005 HIES data.

Table A2.2: Poverty Rates from PovMap and Subsample Direct Estimation, Using Sample of 640 Households (230 urban, 410 rural)

Sample	PovMap						Direct Estimation					
	All	SE	Urban	SE	Rural	SE	All	SE	Urban	SE	Rural	SE
1	40.2		26.2		44.9		41.3		33.3		43.8	
2	37.6	1.9	27.8	1.2	40.8	2.9	38.3	2.1	28.1	3.7	41.4	1.7
3	39.8	1.4	31.8	2.9	42.4	2.0	40.5	1.5	29.0	2.8	44.1	1.5
4	37.1	1.6	27.8	2.4	40.2	2.1	40.3	1.3	34.6	3.2	42.0	1.3
5	41.9	2.0	29.0	2.1	46.2	2.6	43.5	1.9	27.1	3.3	48.4	2.7
6	40.1	1.8	29.3	1.9	43.7	2.3	41.7	1.7	31.6	3.0	44.9	2.5
7	38.7	1.7	25.3	2.2	43.1	2.1	38.5	1.8	27.3	3.0	42.2	2.4
8	39.2	1.6	25.7	2.2	43.7	2.0	39.1	1.8	25.9	3.2	43.1	2.2
9	39.4	1.5	27.3	2.1	43.3	1.9	39.3	1.7	33.1	3.2	41.4	2.2
10	42.5	1.7	28.2	1.9	47.1	2.2	43.1	1.8	35.7	4.5	45.6	2.2
11	39.0	1.6	25.7	2.0	43.4	2.1	41.6	1.8	31.4	3.3	44.7	2.1
12	38.8	1.6	33.6	2.5	40.5	2.1	33.8	2.6	26.1	3.4	36.2	3.0
13	36.8	1.7	25.0	2.6	40.7	2.2	36.0	2.7	22.4	3.9	40.4	3.0
14	40.4	1.6	26.8	2.5	44.8	2.1	43.7	2.8	31.3	3.8	47.7	3.1
15	37.8	1.6	27.3	2.4	41.3	2.1	39.4	2.7	32.4	3.7	41.6	3.0
16	38.4	1.6	26.1	2.4	42.4	2.1	38.5	2.7	25.2	3.8	42.9	2.9
17	36.6	1.7	26.4	2.3	39.9	2.1	35.8	2.8	27.4	3.7	38.4	3.1
18	39.6	1.6	24.6	2.4	44.5	2.1	40.0	2.7	22.1	4.0	45.6	3.0
19	40.5	1.6	32.7	2.6	43.0	2.0	39.6	2.6	25.8	4.0	44.2	3.0
20	37.6	1.6	27.9	2.5	40.8	2.0	36.7	2.6	27.7	3.9	39.5	3.0
21	40.4	1.6	27.4	2.5	44.6	2.0	44.9	2.8	32.9	3.9	48.7	3.2
22	36.8	1.6	29.9	2.4	39.1	2.1	37.9	2.8	33.5	3.9	39.3	3.2
23	39.2	1.6	27.7	2.4	43.0	2.1	42.2	2.8	33.8	3.9	44.8	3.2
24	42.1	1.7	30.9	2.4	45.8	2.1	39.5	2.7	22.6	4.1	44.8	3.1
25	40.5	1.7	28.7	2.4	44.3	2.1	41.7	2.7	31.8	4.0	44.8	3.1
26	42.4	1.7	30.7	2.4	46.2	2.2	43.2	2.7	26.0	4.0	49.0	3.2
27	40.7	1.7	27.6	2.3	44.9	2.2	39.9	2.6	16.7	4.6	47.7	3.3
28	38.6	1.7	28.1	2.3	42.0	2.1	40.0	2.6	30.0	4.5	43.1	3.2
29	40.2	1.7	28.4	2.3	44.0	2.1	42.1	2.6	32.5	4.5	45.1	3.2
30	40.4	1.7	33.6	2.4	42.7	2.1	44.9	2.7	33.6	4.5	48.5	3.2
Mean	39.4		28.3		43.1		40.2		29.0		43.8	

Source: Authors' estimations using 2005 HIES data.

Table A2.3: Poverty Rates from PovMap and Subsample Direct Estimation, Using Sample of 320 Households (120 urban, 200 rural)

Sample	PovMap						Direct Estimation					
	All	SE	Urban	SE	Rural	SE	All	SE	Urban	SE	Rural	SE
1	38.5		26.9		42.4		43.4		29.1		48.6	
2	38.4	4.1	28.3	0.2	41.7	1.1	43.6	0.2	31.8	1.9	47.8	0.6
3	29.8	2.9	22.7	0.8	32.2	0.9	37.4	3.5	23.8	4.1	42.0	3.6
4	36.1	2.6	32.4	3.4	37.3	1.8	38.1	3.3	32.9	4.1	39.8	4.3
5	37.1	2.5	27.9	3.8	40.1	1.9	40.1	2.9	35.6	4.5	41.7	4.0
6	34.4	2.5	32.9	4.0	34.9	1.7	33.5	3.9	17.7	6.6	38.9	4.1
7	36.9	2.5	32.1	3.7	38.5	2.1	36.5	3.7	19.6	6.9	42.4	3.8
8	41.0	2.4	30.2	3.4	44.6	2.8	40.6	3.5	33.7	6.8	42.8	3.5
9	44.2	2.3	32.1	3.3	48.2	2.9	41.5	3.3	29.0	6.4	45.8	3.4
10	42.8	2.3	31.1	3.4	46.6	2.8	41.9	3.2	27.4	6.0	46.5	3.3
11	45.4	2.4	30.9	3.2	50.2	2.7	43.5	3.3	30.1	5.7	48.2	3.5
12	39.8	2.3	31.9	3.2	42.4	2.7	41.2	3.1	22.1	5.7	48.0	3.5
13	44.4	2.3	38.0	3.1	46.5	2.6	39.5	3.0	30.0	5.5	43.3	3.4
14	33.6	2.2	25.0	3.1	36.4	2.5	39.7	2.9	34.7	5.6	41.3	3.3
15	41.9	2.2	31.0	3.1	45.5	2.6	45.8	3.2	22.4	5.6	53.0	3.9
16	36.3	2.1	32.0	3.0	37.8	2.5	36.7	3.2	26.4	5.5	40.1	4.0
17	36.1	2.2	34.9	2.9	36.6	2.4	39.2	3.1	42.5	6.4	38.1	4.1
18	44.4	2.2	31.1	2.9	48.8	2.8	47.7	3.5	35.1	6.3	51.5	4.4
19	41.8	2.1	27.5	2.8	46.4	2.8	40.5	3.4	27.6	6.2	45.2	4.3
20	46.6	2.1	32.6	2.9	51.2	2.7	47.4	3.6	30.7	6.0	52.8	4.5
21	38.3	2.1	32.9	3.1	40.1	2.7	40.3	3.5	25.1	5.9	45.5	4.4
22	41.2	2.1	29.9	3.2	44.9	2.6	35.1	3.7	28.7	5.8	37.5	4.6
23	45.0	2.0	25.6	3.1	51.4	2.6	46.2	3.8	40.3	6.1	48.4	4.6
24	35.1	2.0	22.2	3.0	39.4	2.6	40.7	3.7	22.9	6.2	46.8	4.5
25	39.1	1.9	29.5	3.0	42.2	2.5	40.8	3.6	30.7	6.0	44.5	4.4
26	38.5	1.9	25.5	3.0	42.8	2.5	38.8	3.6	34.2	6.0	40.3	4.4
27	39.7	1.9	32.4	3.0	42.2	2.5	38.6	3.5	27.7	5.9	42.2	4.3
28	35.4	1.9	30.0	3.0	37.2	2.6	37.6	3.5	29.4	5.8	40.3	4.3
29	36.5	1.9	28.7	3.1	39.0	2.8	38.5	3.5	33.1	5.7	40.6	4.3
30	35.8	1.8	33.2	3.0	36.7	2.8	35.2	3.5	20.0	5.9	40.5	4.3
Mean	39.1		30.0		42.1		40.3		29.1		44.1	

Source: Authors' estimations using HIES 2005 data.

Table A2.4: Subsample Imputation and Direct Estimation, 1,000 Replications (no division or district dummies)

P-values in stepwise	Subsample Imputation		Subsample Direct		Model	
	Estimate	SE	Estimate	SE	R ²	Mean no. coef
<i>Rural, 410 households</i>						
Full sample	43.8	1.3
0.0001	42.0	2.6	44.4	3.5	0.46	6.4
0.0005	42.0	2.6	44.2	3.5	0.49	8.2
0.001	41.8	2.6	44.3	3.6	0.50	9.1
0.005	42.0	2.7	44.3	3.6	0.55	12.4
0.01	42.1	2.7	44.2	3.4	0.57	14.6
0.05	42.5	3.1	44.2	3.4	0.62	23.6
<i>Urban, 230 households</i>						
Full sample	28.4	1.9
0.0001	27.9	2.8	28.2	4.6	0.64	4.8
0.0005	28.2	2.8	28.0	4.6	0.66	5.6
0.001	28.1	2.7	28.5	4.6	0.67	6.2
0.005	28.3	2.9	28.3	4.7	0.70	8.3
0.01	28.5	2.9	28.4	4.5	0.71	9.7
0.05	28.9	3.3	28.1	4.6	0.76	16.4
<i>Rural, 200 households</i>						
Full sample	43.8	1.3
0.0001	41.5	3.8	44.4	5.1	0.41	4.0
0.0005	41.4	3.9	44.2	5.1	0.46	5.2
0.001	41.6	4.0	44.0	5.0	0.48	5.8
0.005	41.7	4.2	44.0	5.0	0.54	8.2
0.01	41.9	4.4	44.1	5.1	0.56	9.6
0.05	42.3	4.8	43.9	5.1	0.65	16.1
<i>Urban, 120 households</i>						
Full sample	28.4	1.9
0.0001	29.0	3.9	30.1	6.9	0.62	3.5
0.0005	29.2	4.2	29.9	6.9	0.64	4.2
0.001	29.3	4.1	30.3	7.1	0.66	4.6
0.005	29.3	4.2	29.6	6.7	0.70	6.1
0.01	29.5	4.4	29.5	6.6	0.72	7.2
0.05	30.3	5.1	29.5	6.7	0.79	12.1

Source: Authors' estimations using 2005 HIES data.

Note: P-values used in forward stepwise function in Stata. R-squares and mean number of coefficients from the model used in 1,000 replications are shown.

Table A2.5: Subsample Imputation and Direct Estimation, 1,000 Replications (with division dummies)

P-values in stepwise	Subsample Imputation		Subsample Direct		Model	
	Estimate	SE	Estimate	SE	R ²	Mean no. coef
<i>Rural, 410 households</i>						
Full sample	43.8	1.2
0.0001	41.6	2.4	44.3	3.5	0.46	6.5
0.0005	41.7	2.4	44.3	3.3	0.49	8.3
0.001	41.7	2.5	44.1	3.5	0.51	9.2
0.005	41.7	2.7	44.1	3.5	0.55	12.5
0.01	41.8	2.6	44.3	3.5	0.57	14.7
0.05	42.1	2.9	44.1	3.5	0.62	23.5
<i>Urban, 230 households</i>						
Full sample	28.4	2.6
0.0001	27.8	2.6	28.1	4.6	0.64	4.7
0.0005	27.8	2.6	28.5	4.8	0.66	5.7
0.001	27.9	2.6	28.4	4.6	0.67	6.4
0.005	28.2	2.8	28.3	4.5	0.70	8.5
0.01	28.1	2.9	28.3	4.7	0.72	10.1
0.05	28.4	3.1	28.3	4.6	0.76	16.8
<i>Rural, 200 households</i>						
Full sample	43.8	1.2
0.0001	41.4	3.8	44.4	5.0	0.41	4.0
0.0005	41.6	3.8	44.1	5.0	0.46	5.2
0.001	41.7	4.0	44.0	4.9	0.48	5.8
0.005	41.8	4.0	44.1	5.1	0.54	8.3
0.01	42.1	4.2	44.1	5.1	0.57	9.9
0.05	42.4	4.5	44.2	5.1	0.65	16.1
<i>Urban, 120 households</i>						
Full sample	28.4	2.6
0.0001	28.9	4.0	30.1	7.0	0.61	3.4
0.0005	28.9	3.9	30.2	6.9	0.65	4.2
0.001	28.9	4.0	29.8	6.9	0.67	4.7
0.005	29.3	4.3	30.2	6.5	0.71	6.2
0.01	29.2	4.5	30.2	7.0	0.72	7.3
0.05	29.3	4.7	30.0	7.0	0.79	12.2

Source: Authors' estimations using 2005 HIES data.

Notes: P-values used in forward stepwise function in Stata. R-squares and mean number of coefficients from the model used in 1,000 replications are shown.

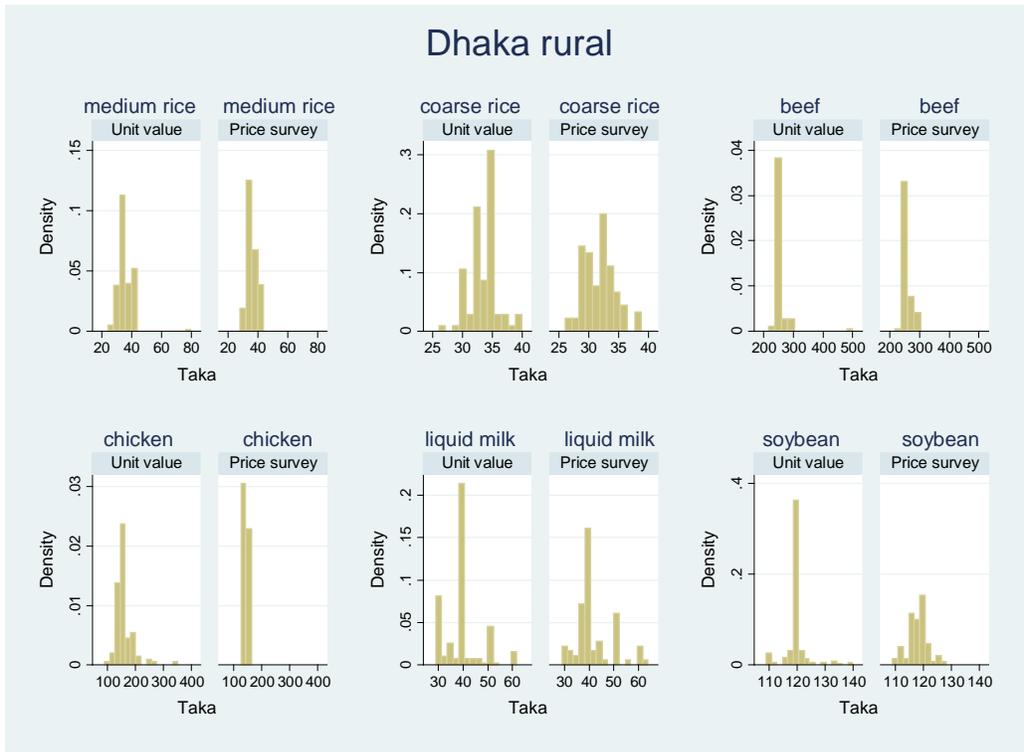
Table A2.6: Subsample Imputation and Direct Estimation, 1,000 Replications (with district dummies)

P-values in stepwise	Subsample Imputation		Subsample Direct		Model	
	Estimate	SE	Estimate	SE	R ²	Mean no. coef
<i>Rural, 410 households</i>						
Full sample	43.8	1.2
0.0001	41.4	2.4	44.4	3.5	0.48	7.0
0.0005	41.4	2.5	44.2	3.5	0.51	9.1
0.001	41.6	2.3	44.3	3.6	0.53	10.2
0.005	41.7	2.5	44.3	3.6	0.57	13.9
0.01	41.8	2.5	44.2	3.4	0.59	16.3
0.05	42.1	2.7	44.2	3.4	0.64	26.0
<i>Urban, 230 households</i>						
Full sample	28.4	2.6
0.0001	27.8	2.6	28.2	4.6	0.64	4.8
0.0005	28.1	2.7	28.0	4.6	0.66	5.7
0.001	28.1	2.6	28.5	4.6	0.67	6.4
0.005	28.2	2.7	28.3	4.7	0.70	8.7
0.01	28.5	2.7	28.4	4.5	0.72	10.4
0.05	28.7	3.0	28.1	4.6	0.77	17.4
<i>Rural, 200 households</i>						
Full sample	43.8	1.2
0.0001	41.5	3.8	44.4	5.1	0.41	4.1
0.0005	41.4	3.9	44.2	5.1	0.47	5.5
0.001	41.7	4.0	44.0	5.0	0.49	6.1
0.005	41.8	4.0	44.0	5.0	0.55	8.7
0.01	41.9	4.3	44.1	5.1	0.58	10.2
0.05	42.5	4.8	43.9	5.1	0.66	16.8
<i>Urban, 120 households</i>						
Full sample	28.4	2.6
0.0001	28.9	3.9	30.1	6.9	0.62	3.5
0.0005	29.1	4.2	29.9	6.9	0.65	4.2
0.001	29.2	4.0	30.3	7.1	0.66	4.6
0.005	29.2	4.2	29.6	6.7	0.71	6.2
0.01	29.6	4.5	29.5	6.6	0.73	7.4
0.05	30.1	5.0	29.5	6.7	0.79	12.6

Source: Authors' estimations using 2005 HIES data.

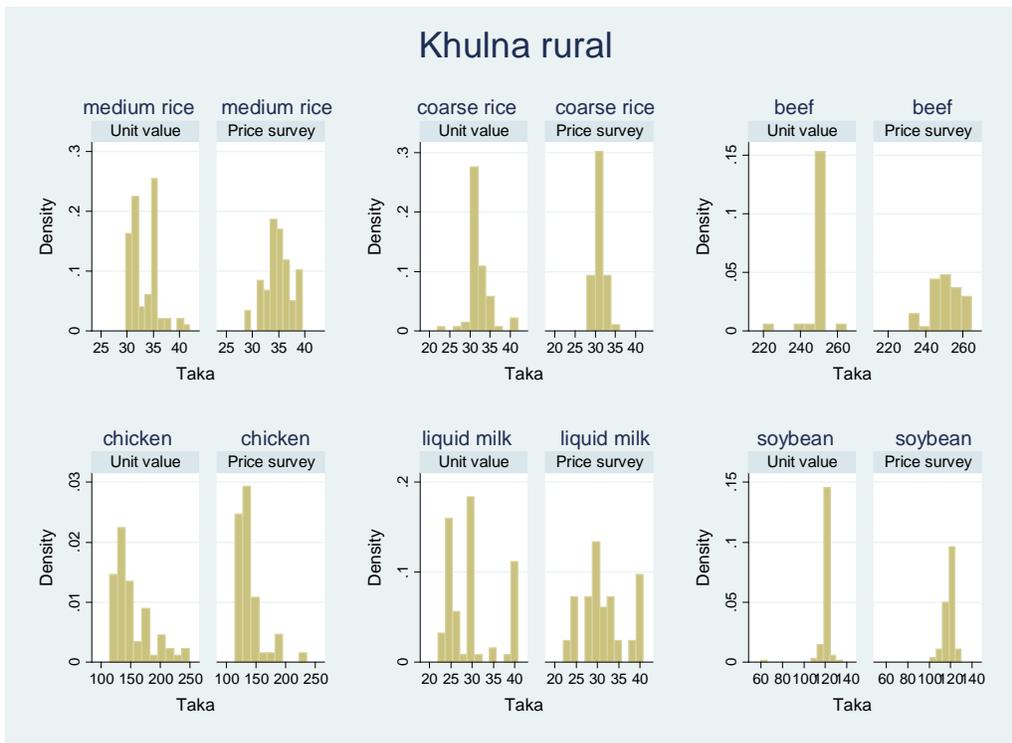
Notes: P-values used in forward stepwise function in Stata. R-squares and mean number of coefficients from the model used in 1,000 replications are shown.

Figure A2.1: Unit Values and Prices from Price Survey, Dhaka Rural



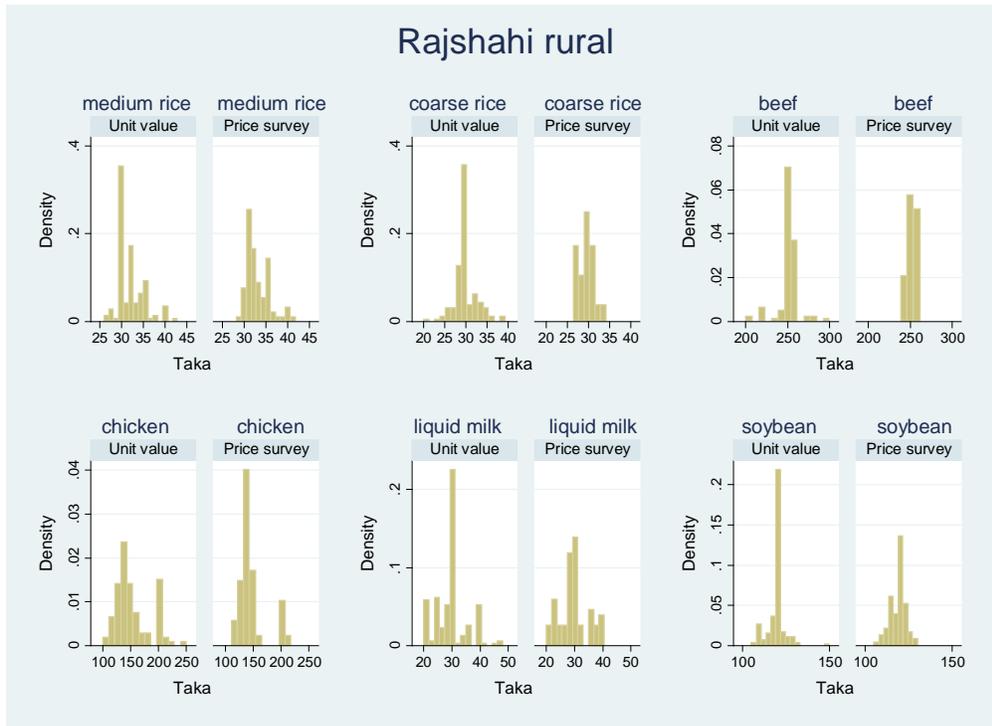
Source: Authors' estimations using pilot survey data.

Figure A2.2: Unit Values and Prices from Price Survey, Khulna, Rural



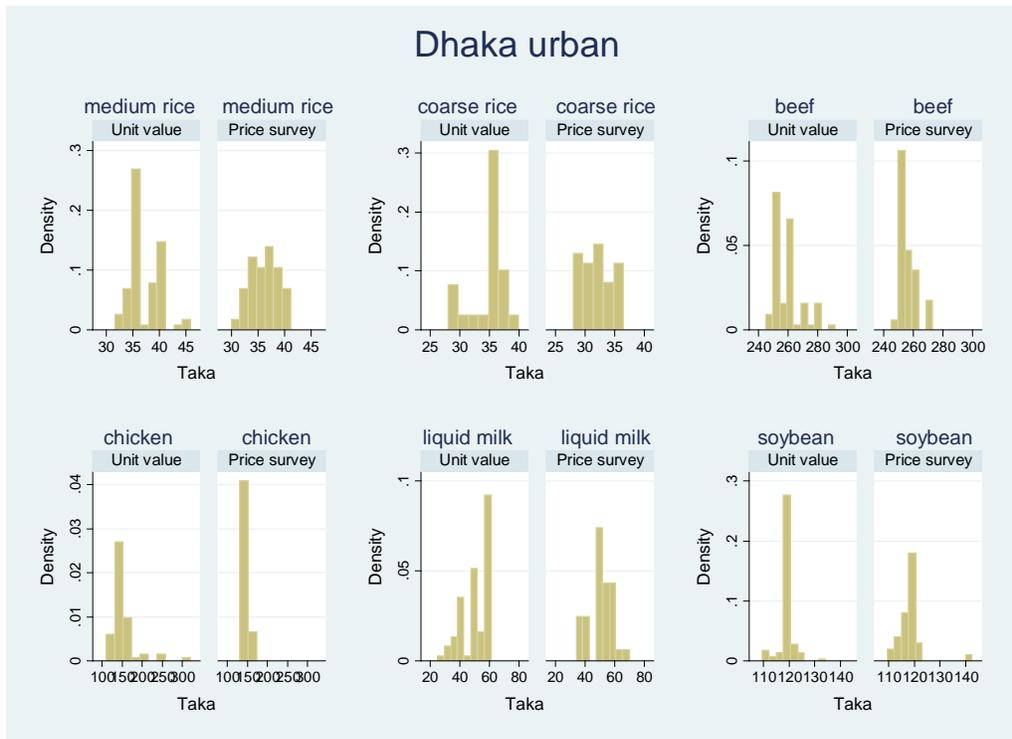
Source: Authors' estimations using pilot survey data.

Figure A2.3: Unit Values and Prices from Price Survey, Rajshahi, Rural



Source: Authors' estimations using pilot survey data.

Figure A2.4: Unit Values and Prices from Price Survey, Dhaka, Urban



Source: Authors' estimations using pilot survey data.

Appendix 3. Price Adjustments Using Unit Values in Subsamples

Table A3.1: Frequency of Missing Price Data for Food Categories, 100 Replications by Stratum (times)

Food group	Food grains	Pulses	Fish	Eggs	Meat	Vegetables	Milk	Sugar	Cooking oils	Salt, spices	Fruits	Soft drinks	Betel and cigarette
Barisal (Rural)	0	0	0	0	4	0	0	0	0	0	0	73	12
Barisal (Muni.)	0	0	0	5	2	0	0	0	9	0	5	66	0
Chittagong (Rural)	0	0	0	0	0	0	0	0	0	0	0	21	0
Chittagong (Muni.)	0	0	0	0	0	0	0	0	1	0	0	16	0
Chittagong (SMA)	17	0	1	2	8	0	9	0	64	0	1	34	2
Dhaka (Rural)	0	0	0	0	0	0	0	0	0	0	0	8	0
Dhaka (Muni.)	0	0	0	0	0	0	0	0	0	0	0	7	0
Dhaka (SMA)	0	0	0	0	0	0	0	0	10	0	0	22	0
Khulna (Rural)	0	0	0	0	2	0	0	0	0	0	0	34	0
Khulna (Muni.)	0	0	0	0	0	0	0	0	3	0	0	5	0
Khulna (SMA)	0	0	0	0	9	0	2	3	0	0	10	53	4
Rajshahi (Rural)	0	0	0	0	0	0	0	0	0	0	0	16	0
Rajshahi (Muni.)	0	0	0	0	0	0	0	0	0	0	0	47	0
Rajshahi (SMA)	89	0	30	0	8	0	3	0	5	0	0	61	1
Sylhet (Rural)	12	0	0	3	17	0	6	1	0	0	3	78	1
Sylhet (Muni.)	48	0	0	1	6	0	4	0	9	0	14	41	0

Source: Authors' estimations using 2000 and 2005 HIES data.

Table A3.2: Mean Weights for Food Categories in 2005 (%)

Food group	Food grains	Pulses	Fish	Eggs	Meat	Vegetables	Milk	Sugar	Cooking oils	Salt, spices	Fruits	Soft drinks	Betel & cigarette
Barisal (Rural)	44.3	3.3	16.6	1.3	5.8	9.4	2.4	1.3	3.7	7.1	2.1	0.5	2.2
Barisal (Municipality)	36.8	3.9	19.0	1.6	9.2	9.8	2.3	1.4	3.5	6.9	2.0	0.9	2.7
Chittagong (Rural)	40.5	3.2	13.1	1.2	5.5	10.2	2.4	2.0	4.6	7.9	3.6	0.6	5.2
Chittagong (Municipality)	35.8	3.0	13.7	1.6	7.6	9.1	2.9	2.8	4.7	7.9	4.8	1.6	4.5
Chittagong (SMA)	34.5	2.8	14.5	1.9	6.1	10.3	4.1	3.3	4.4	7.2	5.0	2.2	3.8
Dhaka (Rural)	45.0	3.0	12.9	1.1	4.9	8.4	3.3	2.0	4.5	8.4	2.9	0.2	3.4
Dhaka (Municipality)	39.5	3.8	13.2	1.9	6.4	8.5	4.5	2.3	4.7	8.0	3.5	0.7	3.0
Dhaka (SMA)	31.4	4.5	16.0	1.7	8.8	9.5	3.4	1.7	6.0	11.3	3.4	1.0	1.4
Khulna (Rural)	49.5	2.1	10.1	1.1	4.4	11.4	2.0	2.1	5.2	7.1	2.8	0.1	1.9
Khulna (Municipality)	40.5	2.7	12.4	1.8	6.9	11.1	2.4	2.6	5.8	7.6	3.4	0.7	2.1
Khulna (SMA)	40.3	2.4	16.3	1.9	7.7	9.6	2.2	1.4	5.9	7.9	1.3	0.5	2.6
Rajshahi (Rural)	54.6	1.5	7.8	1.1	6.1	9.1	2.1	1.2	4.5	7.1	2.0	0.3	2.6
Rajshahi (Municipality)	48.9	2.1	9.0	1.6	6.1	9.6	2.9	1.9	4.8	7.7	2.1	0.9	2.4
Rajshahi (SMA)	41.8	4.1	7.6	2.5	9.4	8.8	2.2	2.0	5.3	8.6	4.5	0.9	2.4
Sylhet (Rural)	40.0	2.1	15.1	0.8	4.5	9.4	1.8	2.1	4.7	8.8	2.6	1.0	7.3
Sylhet (Municipality)	30.9	2.8	17.7	1.0	6.4	10.3	3.8	2.9	5.8	9.5	3.5	1.3	4.2

Source: Authors' estimations using 2000 and 2005 HIES data.

Notes

¹ For the sake of exposition, hereafter household income will not be mentioned, although some countries do use income for measuring poverty.

² The structure of the full data sample of the 2005 HIES is included in appendix table A2.1.

³ Ahmed (2004, appendix A) describes the more detailed procedure to update poverty lines from 1991–92 to 1995–96 and to 2000. The same procedure was used for updating the poverty line in 2005, except the number of food item categories decreased by 1 from 14 and the number of strata in which price indexes were calculated.

⁴ The cost of supervisors is not included in this assessment.