



9

Rapid Consumption Surveys

Utz Pape and Johan Mistiaen

1 The Data Demand and Challenge

Poverty is the paramount indicator used to gauge the socioeconomic well-being of a population. Particularly after a shock or in a volatile context, poverty estimates can identify who was affected, and how severely. This is particularly relevant in fragile countries where monitoring poverty dynamics help measure the country's progress toward stability, or increased risk of relapsing into conflict. As one of the main indicators for poverty, monetary poverty is measured by a welfare

This chapter is a summary of Pape, Utz Johann, and Johan Mistiaen. "Household Expenditure and Poverty Measures in 60 Minutes: A New Approach with Results from Mogadishu." Policy Research Working Paper Series. The World Bank, 2018. <https://ideas.repec.org/p/wbk/wbrwps/8430.html>.

U. Pape (✉) · J. Mistiaen
World Bank, Washington, DC, USA
e-mail: upape@worldbank.org

J. Mistiaen
e-mail: jmistiaen@worldbank.org

aggregate, usually based on consumption in developing countries and a poverty line. The poverty line indicates the minimum level of welfare required for healthy living.

Consumption aggregates are traditionally estimated based on time-consuming household consumption surveys. A household consumption questionnaire records consumption (how much was consumed) and expenditure (how much was purchased, or obtained in other ways like gifts or aid) for a comprehensive list of food and non-food items. Covering between 300 and 400 items, the questionnaire often exceeds 120 minutes to administer. In addition to the longer administering time leading to higher costs, response fatigue can increase measurement error, especially for items at the end of the questionnaire. In a fragile country context, a face-to-face time of 90–120 minutes can be prohibitively high. In the case of Somalia, security concerns restricted the duration of a survey visit in Mogadishu to about 60 minutes.

The extensive nature of household consumption surveys makes it difficult to obtain updated poverty estimates, especially when they are needed the most, such as after a shock and in fragile countries. Approaches have therefore been developed to reduce administering times to allow for the collection of consumption data. The most straightforward approach to minimize administering time is to reduce the number of items surveyed, either by asking for aggregates, or by skipping less frequently consumed items, which is called the reduced consumption methodology. However, both approaches—using aggregates, and skipping less common items—have been shown to underestimate consumption, which in turn overestimates poverty.¹ Splitting the questionnaire to allow for multiple visits is another solution, but potential attrition issues especially in fragile contexts increases the required sample size and may be costlier. In addition, multiple visits to the same household can increase security concerns.

The second class of approaches utilizes a full consumption baseline survey and updates poverty estimates based on a small subset of collected

¹Beegle et al. (2012).

indicators.² These approaches estimate a welfare model based on the baseline survey using a small number of easy-to-collect indicators. This allows poverty estimates to be updated by collecting only the set of indicators instead of the direct consumption data. While this approach is cost-effective and easy to implement in normal circumstances, it has two major drawbacks in the context of fragility and shocks. First, the approach requires a baseline survey, which is sometimes not available, as in the case of Mogadishu. Second, the approach relies on a structural model estimated from the baseline survey.³ In the case of shocks, structural assumptions that cannot be tested are often violated. Thus, poverty updates based on the violated assumptions tend to underestimate the impact of the shock on poverty. Therefore, cross-survey imputation methodologies are not applicable in the context of shocks and fragility.

2 The Innovation

To assess poverty in Mogadishu, we tested a new methodology combining an innovative questionnaire design with standard imputation techniques. This substantially reduces the administering time of a consumption survey from multiple hours or even days to about 60 minutes, while still resulting incredible poverty estimates. The gain in shorter administering time, however, is offset by the need to impute missing consumption values. Given the design of the questionnaire, this method circumvents the systematic biases identified for alternative methodologies.

2.1 Overview

The rapid consumption survey methodology involves five main steps (Fig. 1). First, core items are selected based on their importance for consumption. Second, the remaining items are partitioned into optional

²Doudich et al. (2013); SWIFT.

³Christiaensen et al. (2011).

The consumption module is partitioned into core and optional modules, which in turn are assigned to households. Consumption is imputed utilizing the sub-sample information of the optional modules either by single or multiple imputation methods.

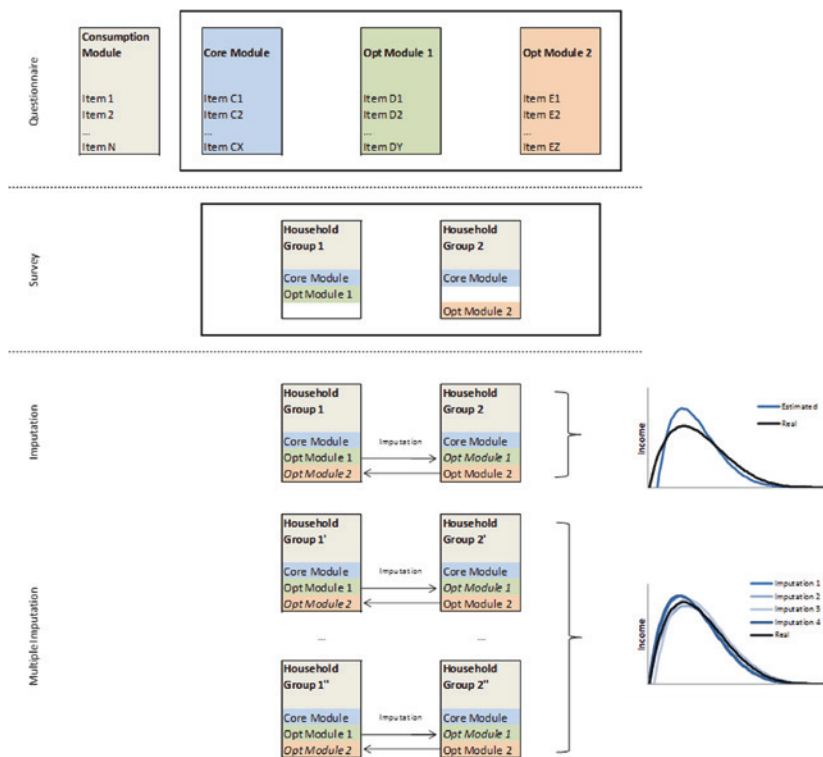


Fig. 1 Illustration of the rapid consumption survey methodology (using illustrative data only)

modules. Third, optional modules are assigned to groups of households. Fourth, after data collection, consumption of optional modules is imputed for all households. Fifth, the resulting consumption aggregate is used to estimate poverty indicators.

First, core consumption items are selected. Consumption in a country bears some variability, but usually a small number of a few dozen items captures the majority of consumption. These items are assigned to the core module, which will be administered to all households. Important items can be identified by their average consumption share

per household or across households. Previous consumption surveys in the same country, or consumption shares of neighboring or similar countries can be used to estimate consumption shares.

Second, non-core items are partitioned into optional modules. Different methods can be used for this partitioning. In the simplest case, the remaining items are ordered according to their consumption share and assigned one by one while iterating the optional module in each step. A more sophisticated method takes into account the correlation between items, and partition them in a way so that all items within a module explain consumption as best as possible, while the information between modules should be highly correlated. The partitioning influences the standard error of the estimation, but does not introduce bias. Thus, even in the absence of a previous survey, this methodology can be applied. More complicated partition patterns can result in a set of very different items in each module. However, the modular structure should not influence the layout of the questionnaire. Instead, all items should be grouped into categories of consumption (e.g. cereals) and different recall periods. It is therefore recommended to use CAPI technology, which allows the structure of the consumption module to be hidden from the enumerator.

Third, optional modules should be assigned to groups of households. Optional modules should be assigned randomly, stratified by clusters to ensure appropriate representation of optional modules in each cluster. This means that each cluster should include about the same number of households assigned to each optional module. This step is followed by the actual data collection.

Fourth, household consumption should be estimated by imputation. The average consumption of each optional module can be estimated based on the subsample of households assigned to the optional module. In the most straightforward case, a simple average can be estimated. More sophisticated techniques can employ a welfare model based on household characteristics and consumption of the core items. The next section presents six techniques and demonstrates their performance on the dataset from Hargeisa.

Single imputation of the consumption aggregate underestimates the variance of household consumption. Depending on the location of the

poverty line relative to the consumption distribution, this may either consistently under- or overestimate poverty. Multiple imputations based on bootstrapping can mitigate the problem but will render analysis more complicated. We use single as well as multiple imputation techniques for the evaluation of the methodology.

3 Key Results

In this section, the rapid consumption methodology will first be applied to a dataset including a full consumption module from Hargeisa, Somaliland. This will be used to assess the performance of the rapid consumption methodology compared to the traditional full consumption methodology. The results of the High Frequency Survey in Mogadishu are then presented. Security risks in Mogadishu restrict face-to-face interview time to less than one hour; therefore, the rapid consumption methodology was used to derive the first ever consumption estimates for Mogadishu. We present the resulting consumption aggregate, and perform consistency checks for its validation.

3.1 Ex Post Simulation

The rapid consumption methodology is applied ex post to household budget data collected in Hargeisa, Somaliland. Hargeisa was chosen as it is very similar to Mogadishu. Using the full consumption dataset from Hargeisa allows a full assessment of the new methodology. Based on selected indicators, we compare the results of the estimated consumption based on the rapid consumption methodology with the results from using the traditional full consumption module. We add a comparison with the results for a reduced consumption module.

The simulation assigns each household to one optional module. The consumption data for the modules not assigned to the household is deleted. Multiple simulations are performed, with various modules being assigned to households. Across the simulations, we calculate three consumption indicators and four poverty and inequality indicators. The

Table 1 Number of items and consumption share captured per module

	Number of food items	Share of consumption (%)	Number of non-food items	Share of consumption (%)
Core	33	92	25	88
Module 1	17	3	15	3
Module 2	17	2	15	3
Module 3	15	2	15	4
Module 4	17	2	15	3

consumption indicators capture the accuracy of the estimation at three different levels: the household level, the cluster level (consisting of about nine households), and the level of the dataset. In addition, we calculate the poverty headcount (FGT0), the poverty depth (FGT1), the poverty severity (FGT2), and the Gini coefficient to capture inequality.

Six estimation techniques are compared with respect to their relative bias and relative standard error, based on 20 simulations. All simulations used the same item assignment to modules using the algorithm as described (see Table 1 for the resulting consumption shares per module).⁴ The estimation techniques differ considerably in terms of performance. We also compare the techniques to using a reduced consumption module where the same consumption items are collected for all households. The number of items is equal to the size of the core module and one optional module, implying a comparable face-to-face interview time to the rapid consumption methodology.

Comparing the reduced consumption approach with the full consumption as a reference, the reduced consumption approach suffers from an underestimation of consumption. This is not surprising because the approach only collects information on the consumption of a subset of items. Applying the median as a summary statistic also results in an underestimation of consumption. As consumption distributions have a long right tail, the median consumption belongs to a poorer household than the average household. In the case of Hargeisa, several optional

⁴We performed robustness checks with different item assignment to modules, including setting the parameter $d=1$ and $d=2$. The estimation results are extremely robust to changes in the item assignment to modules.




Fig. 2 Average relative bias and standard error

modules have a median of zero consumption. Thus, the median underestimates the consumption in a similar way to the reduced consumption approach. In contrast, the average consumption of households is larger than the consumption of the median household. Thus, it is not surprising that the technique using the average as a summary statistic overestimates total consumption at the household and cluster levels.

The regression techniques have a similar performance, with a considerable upward bias at all levels. The Tobit regression performs slightly better at the household and cluster levels. As known from literature about small area estimates, the regression approaches do not model the error distribution correctly and, thus, underestimate the tails of the distribution. Depending on the value of the poverty line relative to the mode of the distribution, this results in an over- or under-estimation of the poverty rate. In contrast, both imputation techniques perform exceptionally well, with a bias below 1% at all levels (Fig. 2).

While the bias is important in order to understand the systematic deviation of the estimation, the relative standard error helps to understand the variation of the estimation. Other than in a simulation setting, the standard error of the estimation cannot be calculated, as only one assignment of households to optional modules is available. Thus, it is important that the estimation technique delivers a small relative standard error.

Generally, the relative standard error reduces when moving from the household level over the cluster level to the simulation level. The relative standard error for the reduced consumption methodology is smaller than for the summary statistic techniques because the reduced consumption is not subject to variation from the module assignment to households. The regression techniques have large relative standard errors of around 20% at the household level, while the multiple imputation techniques vary between 15 and 20%. At the cluster level, the relative standard error drops to 7% for regression techniques and 5% for multiple imputation techniques. At the simulation level, the relative standard error is around 3% for regression techniques and 1% for multiple imputation techniques.

The distributional shape of the estimated household consumption level can be compared to the reference household consumption by employing standard poverty and inequality indicators. The poverty headcount (FGT0) is 57.4% for the reference distribution.⁵ Not surprisingly, the reduced consumption technique and the median summary statistic overestimate poverty by several percentage points due to the underestimation of consumption, while the average summary statistic and the regression techniques underestimate poverty, since they overestimate consumption. The multiple imputation techniques overestimate poverty, but only by 0.5 percentage points (or about 1%), performing significantly better than the reduced consumption approach, which has a bias that is more than two times larger. The reduced consumption technique and the median summary statistic as well as the multiple imputation techniques deliver good results for FGT1 and FGT2, emphasizing that not only can the headcount be estimated reasonably well, but the distributional shape is also conserved. With the exception of the median summary statistic, these techniques also perform well estimating the Gini coefficient, with a bias of less than 0.5 percentage points. The relative standard errors show similar results as for the estimation of the consumption. The relative standard error of the reduced

⁵The FGT0 is calculated based on the US\$1.90 PPP (2011) international poverty line, converted into local currency in 2013.

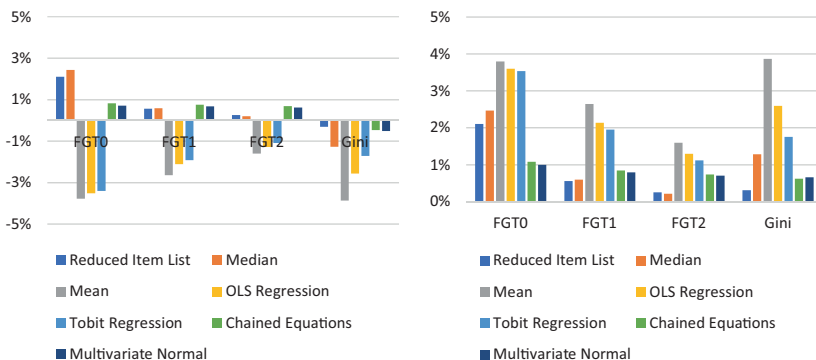


Fig. 3 Bias and standard errors

consumption for FGT0 is double that of the multiple imputation techniques. The relative standard errors for the multiple imputation techniques for FGT1 are comparable but larger than for FGT2 and Gini (Fig. 3).

In conclusion, the average summary statistic and the regression approaches cannot deliver convincing estimates. While the reduced consumption technique and the median summary statistic perform considerably better, they both overestimate poverty. Only the multiple imputation techniques are convincing in all estimation exercises. In terms of the estimation of the important poverty headcount (FGT0), the multiple imputation techniques are virtually unbiased.

4 Implementation Challenges, Lessons Learned, and Next Steps

In late 2014, consumption data using the proposed rapid consumption methodology was collected in Mogadishu using CAPI. The rapid consumption questionnaire reduced face-to-face interview time considerably. A household visit took about 40 minutes on average (with a median of 35 minutes), including greetings, household characteristics,

consumption modules, and a number of perception questions. Nine out of ten interviews took less than 65 minutes.

After data cleaning and quality assurance procedures, 675 households with consumption data were retained.⁶ A welfare model was built to predict missing consumption in optional modules. The welfare model was tested on the core consumption, after removing the core consumption as an explanatory variable. The model for food consumption retrieved an R^2 of 0.24, while non-food consumption was modeled with an R^2 of 0.16. It is important to emphasize that these models give a lower bound of the R^2 compared to the models used in the prediction, as the prediction models include the core consumption as an explanatory variable. Given the assessment of the different estimation techniques in the previous section, the multivariate normal approximation using multiple imputations is applied to the Mogadishu dataset.

For the Mogadishu dataset, the assignment of items to modules had to be manually refined.⁷ The refinement had a minor impact on the share of consumption per module. It is curious, though, that the share of consumption per module is different for Hargeisa and Mogadishu. Using the Hargeisa dataset, 91% of food consumption (and 76% of non-food consumption) is captured in the core module. In contrast, the core food consumption share is only 64% (and 62% of non-food consumption) in Mogadishu before imputing the consumption of non-assigned modules. Thus, employing a reduced consumption module based on consumption shares identified in Hargeisa would have crudely underestimated consumption in Mogadishu, without being able to evaluate the inaccuracy. In contrast, the rapid consumption methodology allows the estimation of shares for each module, while the consumption

⁶While the survey also covered IDP camps, the analysis presented is restricted to households in residential areas, excluding IDP camps.

⁷Manual refinement is necessary to ensure that items like ‘other fruits’ do not double-count types of fruits not assigned to the household. This is implemented by relabeling and manually assigning modules. In addition, some item groups items were split into individual items, which is generally preferable for recall and recording, as well as calculation of unit values.

estimation procedure implicitly takes into account the ‘missing’ consumption shares for each household (Table 2).

The cumulative consumption distribution can be compared for the consumption captured in the core module, the assigned optional modules, and the imputed consumption. By construction, the core consumption shows the lowest consumption per household. Adding the consumption from the assigned optional modules shifts the cumulative consumption curve slightly. The imputed consumption is shifted even further as the estimated consumption shares from the non-assigned modules are added (Fig. 4).

Without full consumption aggregate values for Mogadishu, we can only show the consistency of the retrieved consumption aggregate with other household characteristics to validate the estimates. Consumption per capita usually reduces with increasing household size. Indeed, we find that household size is significantly negatively correlated with estimated per capita consumption.⁸ Per capita consumption also decreases with a larger share of children among the household members. The proportion of employed members of the household significantly increases consumption per capita. Thus, the retrieved consumption estimate is consistent and using the evidence from the ex post simulations, highly accurate.

The results of the ex post simulation indicate that the rapid consumption methodology can reliably estimate consumption and poverty. The experience in Mogadishu also shows that the rapid consumption methodology can be implemented in extremely high-risk areas, due to its success in limiting face-to-face interview time to less than one hour. While these results are encouraging, the rapid consumption methodology has some limitations.

The rapid consumption questionnaire varies in comprehensiveness and the order of items in the consumption module between households.

⁸The reported numbers are corrected against correlation with household characteristics included in the welfare model. As the welfare model for the prediction of consumption includes household size, we have run a robustness check excluding household size from the welfare model used for prediction. The correlation between consumption per capita and household size is still significant (coefficient: -0.03 , t-statistic: -2.17 , p -value: 0.03).

Table 2 Number of items and consumption shares captured per module

	Food consumption				Non-food consumption			
	Number of items	Share Hargeisa (%)	Share Mogadishu (%)	Share Mogadishu imputed (%)	Number of items	Share Hargeisa (%)	Share Mogadishu (%)	Share Mogadishu imputed (%)
Core	33	91	64	54	26	76	62	52
Module 1	19	3	9	16	15	7	9	12
Module 2	20	2	14	14	15	5	9	12
Module 3	15	2	5	6	15	6	8	9
Module 4	15	2	8	9	15	6	11	15

Note simulated for Hargeisa, estimated for Mogadishu before imputation of non-assignment modules (normalized to 100%), and after imputing full consumption




Fig. 4 Cumulative consumption distribution (in USD) per day and per capita (Color figure online) (Note For core module (dark blue), core and assigned optional modules (medium blue), and imputed consumption (light blue). The presented consumption aggregate does not include consumption from durable goods

The effect of a response bias due to this can neither be estimated from the simulations nor from the data collected in Mogadishu. However, an enhanced design with different optional modules varying in their comprehensiveness can shed light on this bias. Comparison between responses for the same item in a comprehensive and an incomprehensive list would indicate a lower bound for response bias. Assuming that a comprehensive list results in a better estimate, the response bias could be corrected.

The rapid consumption methodology can increase the gap between capacity at enumerator level and the complexity of the survey instrument. Capacity at the enumerator level is often low in developing countries, especially in a fragile context. The rapid consumption methodology increases the complexity of the questionnaire, which can further increase the gap between existing and required enumerator capacity. However, CAPI technology can seal off complexity from enumerators, as software can automatically create the consumption module

based on core and optional modules for each household without showing the partition to the enumerator. In Mogadishu, advanced CAPI technology was used to automatically generate the questionnaire based on the assignment of the household to an optional module. While enumerators were made aware that different households would be asked about different items, administering the rapid consumption questionnaire did not require any additional training of enumerators beyond that needed for a standard consumption questionnaire.

Analysis of rapid consumption data requires high capacity. Analysis capacity is usually limited in developing countries, and especially in fragile contexts. While the general idea of optional consumption modules being assigned to households is digestible by local counterparts, poverty analysis based on a bootstrapped sample of consumption distribution is likely to overwhelm local capacity. However, even standard poverty analysis is often beyond the limits of local capacity in fragile countries. Therefore, capacity building usually focuses on data collection skills with a longer-term perspective on increasing data analysis capacity. In addition, the rapid consumption methodology might be the only way of creating poverty estimates in certain areas, for example, in Mogadishu.

The results of the *ex post* simulation and the application of the methodology in Mogadishu suggest that the rapid consumption methodology is a promising approach to estimating consumption and poverty in a cost-efficient and fast manner, even in fragile areas.⁹ A similar *ex post* simulation for South Sudan and Kenya (data not shown) indicates that the rapid consumption methodology can also be applied at the country-level, with large intra-country consumption variation.¹⁰ The rapid consumption methodology has been implemented in Somalia, South Sudan, and Kenya, with additional countries in the pipeline.

⁹Costs for implementing a rapid consumption survey are lower than conducting a full consumption survey due to the reduced face-to-face time needed, allowing enumerators to conduct more interviews per day.

¹⁰Ongoing fieldwork is currently employing the rapid consumption methodology in South Sudan to update poverty numbers.

Annex

Consumption of non-assigned optional modules can be estimated by different techniques. Three classes, each with two techniques, are presented here, differing in their complexity and theoretical underpinnings. The first class of techniques uses summary statistics such as the average, to impute missing data. The second class is based on multiple univariate regression models. The third class uses multiple imputation techniques, taking into account the variation absorbed by the residual term.

Summary Statistics (Mean and Median)

This class of techniques applies a summary statistic on the module-specific consumption data collected and applies the result to the missing modules. Each household is assigned the same consumption per missing module. Here, the mean and the median are used as summary statistics. The median has the advantage of being more robust against outliers but cannot capture small module-specific consumption if more than half of the households have zero consumption for the module.

Module-Wise Regression (Ols and Tobit Regression)

Module-wise estimation applies a separate regression model for each module. This allows for differences in core consumption to be captured, as well as other household characteristics. Coefficients are estimated based only on the subsample assigned to the module under consideration. In general, a bootstrapping approach using the residual distribution could mimic multiple imputations, but this is not applied here. Given the impossibility of negative consumption, a Tobit regression with a lower bound of zero is used in addition to a standard OLS regression approach. For the OLS regression, negative imputed values are set to zero.

Multiple Imputation Chained Equations (Mice)

Multiple Imputation Chained Equations (MICE) uses a regression model for each variable and allow missing values in the dependent and independent variables. As missing values are allowed in the independent variables, the consumption of all optional modules can be used as explanatory variables. As a first step, missing values in the explanatory variables are drawn randomly. These values are substituted iteratively with imputed values drawn from the posterior distribution estimated from the regression. While the technique of chained equations cannot be theoretically shown to converge in distribution, the results in practice are encouraging, and the method is widely used.

Multivariate Normal Regression (MImvn)

Multiple Imputation Multivariate Normal Regression uses an expectation-maximization (EM)-like algorithm to iteratively estimate model parameters and missing data. In contrast to chained equations, this technique is guaranteed to converge in distribution with the optimal values. An EM algorithm draws missing data from a prior (often non-informative) distribution and runs an OLS to estimate the coefficients. The coefficients are iteratively updated based on reestimation using imputed values for missing data drawn from the posterior distribution of the model. MImvn employs a data-augmentation (DA) algorithm, which is similar to an EM algorithm, but updates parameters in a non-deterministic fashion, unlike the EM algorithm. Thus, coefficients are drawn from the parameter posterior distribution rather than chosen by likelihood maximization. Hence, the iterative process is a Markov chain Monte Carlo (MCMC) method in the parameter space, with convergence with the stationary distribution that averages the missing data. The distribution for the missing data stabilizes at the exact distribution to be drawn from, to retrieve model estimates averaging over the missing value distribution. The DA algorithm usually converges considerably faster than using standard EM algorithms.

Estimation Performance

The performance of the different estimation techniques is compared based on the relative bias (mean of the error distribution) and the relative standard error. We define the relative error as the percentage difference between the estimated consumption and the reference consumption (based on the full consumption module). The relative bias is the average of the relative error. The relative standard error is the standard deviation of the relative error. For estimations based on multiple imputations, the error is averaged over all imputations.

Each proposed estimation procedure is run on the random assignments of households to the optional modules. A constraint ensures that each optional module is assigned equally often to a household per enumeration. The relative bias and the relative standard error are reported across all simulations.

The performance measures can be calculated at different levels. At the household level, relative error is the relative difference in household consumption. At the cluster level, relative error is defined as the relative difference of the average reference household consumption and the average estimated household consumption across the households in the cluster. Similarly, the simulation level compares total average consumption for all households.

References

- Beegle, Kathleen, Joachim De Weerd, Jed Friedman, and John Gibson. (2012). "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. (2011). "Small Area Estimation-Based Prediction Methods to Track Poverty: Validation and Applications." *Journal of Economic Inequality* 10 (2): 267–297.
- Doudich, M., A. Ezzrari, R. van der Weide, and P. Verme. (2013). "Estimating Quarterly Poverty Rates Using Labor Force Surveys." World Bank Policy Research Working Paper no. 6466.
- The World Bank. (2018). *Poverty and Shared Prosperity 2018: Piecing Together the Poverty Puzzle*. Washington, DC: World Bank.

The opinions expressed in this chapter are those of the author(s) and do not necessarily reflect the views of the International Bank for Reconstruction and Development/The World Bank, its Board of Directors, or the countries they represent.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 3.0 IGO license (<https://creativecommons.org/licenses/by/3.0/igo/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the International Bank for Reconstruction and Development/The World Bank, provide a link to the Creative Commons license and indicate if changes were made.

Any dispute related to the use of the works of the International Bank for Reconstruction and Development/The World Bank that cannot be settled amicably shall be submitted to arbitration pursuant to the UNCITRAL rules. The use of the International Bank for Reconstruction and Development/The World Bank's name for any purpose other than for attribution, and the use of the International Bank for Reconstruction and Development/The World Bank's logo, shall be subject to a separate written license agreement between the International Bank for Reconstruction and Development/The World Bank and the user and is not authorized as part of this CC-IGO license. Note that the link provided above includes additional terms and conditions of the license.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

