



Working Paper Series

**To impute or not to impute, and how?  
A review of alternative poverty estimation  
methods in the context of unavailable  
consumption data**

Hai-Anh H. Dang

ECINEQ WP 2019 - 507

# To impute or not to impute, and how? A review of alternative poverty estimation methods in the context of unavailable consumption data<sup>\*</sup>

Hai-Anh H. Dang<sup>†</sup>  
*World Bank, USA*

## Abstract

There is an increasingly stronger demand for more frequent and accurate poverty estimates, despite the oftentimes unavailable household consumption data. We review imputation methods that have been employed to provide poverty estimates in such data-scarce contexts. These range from estimates on a non-monetary basis, estimates for specific project targeting or tracking trends at the national level, to estimates at a more disaggregated level, and estimates of poverty dynamics. We provide a concise synthesis, which serves as an introduction to the literature. We focus on intuition and practical insights that highlight existing methods' nuanced differences rather than their technical aspects.

**Keywords:** poverty, imputation, consumption, wealth index, synthetic panels, household survey.

**JEL Classification:** C15, I32, O15.

---

<sup>\*</sup>We would like to thank the editor Deborah Eade, an anonymous referee, Gero Carletto, Thesia Garner, Stephen Jenkins, Dean Jolliffe, Peter Lanjouw, and participants at the IARIW conference (Copenhagen) for helpful discussions on earlier versions. We are grateful to the UK Department of International Development for funding assistance through its Strategic Research Program (SRP) and Knowledge for Change (KCP) programs.

<sup>†</sup>Contact details: Dang is an economist in the Analytics and Tools Unit, Development Data Group, World Bank and is also affiliated with Indiana University and Vietnam's Academy of Social Sciences. [hdang@worldbank.org](mailto:hdang@worldbank.org).

## I. Introduction

Fighting poverty is a challenging and complex undertaking faced by policy makers in developing and richer countries alike. This undertaking can include different policy options ranging from crafting short-term safety-net programs that prevent vulnerable households from sliding into poverty during a time of crisis, to designing long-term plans that invest in education and skills formation to promote economic growth. But regardless of the specific poverty reduction strategies, a prerequisite for the success of the whole process is a clear understanding of poverty trends and dynamics (either at any single snapshot in time or over different time periods). Indeed, inaccurate measurement naturally results in an inefficient—and even a wasteful—use of resources if, say, short-term employment programs are employed to address a chronic poverty situation caused by inadequate infrastructure.

Perhaps the greatest challenge with poverty measurement is the fact that household consumption (or income) data—the underlying data source that provides poverty estimates—do not often meet the necessary requirements.<sup>1</sup> For example, such data may simply be unavailable, or may not be comparable from one survey round to the next. As another example, household consumption data are seldom collected on the same households (or individuals) over time, thus making it difficult—if not possible—to track the dynamics of these households' movements into or out of poverty in different periods.

We offer in this paper a review of alternative poverty estimation (or imputation) methods in contexts where consumption data are unavailable. The economic literature on poverty imputation

---

<sup>1</sup> We use the terms “income” and “consumption” interchangeably in this review. The latter is often considered to offer better measures of household welfare, especially in developing countries (see, e.g., Deaton (1997)). Consequently, our central assumption in this paper is that household consumption (or income) data provide the benchmark measure of household welfare. This practical and widely-used assumption helps operationalize and better focus the measurement of household welfare. We also focus on monetary poverty in this paper; see Alkire *et al.* (2015) for a comprehensive discussion of the alternative approach of multi-dimensional poverty.

has grown rapidly in the past 20 years, and various methods have been developed. In fact, methods that have been proposed and discussed using different frameworks and terminology may turn out to be more similar than one might think.<sup>2</sup> A typical development practitioner who does not keep regular track of the latest advances in the field may find it time-consuming, and perhaps even quite challenging, to stay abreast of this literature.

At the same time, to our knowledge, currently very few studies provide a systematic introduction to this literature. We thus aim to fill in this gap by providing a succinct but user-friendly synthesis of existing poverty imputation methods. In particular, we aim to achieve the following objectives in this review

- i) offer a classification of the various poverty imputation situations
- ii) lay out clearly the contexts where imputation is most relevant
- iii) provide an accessible description of imputation techniques, with explicit flags for common (or potential) mistakes
- iv) point out “gray” areas with current imputation methods that need more research.

Our goal is to offer a systematic discussion of imputation methods in a consistent format, which starts first with each method’s motivation, a brief description of the method, some recent application examples, and the remaining challenges. While this format may appear rigid, it offers a (somewhat) self-contained treatment of different methods. It also helps facilitate comparison and cross-reference between the various methods and highlight the nuanced differences among them. To help readers—particularly for those who are new to the topic—obtain a quick overview of the

---

<sup>2</sup> This is not to mention the (beneficial) interactions between the economic literature and the well-established, and more general statistical literature on missing data imputation. We return to more discussion in later sections. Given our focus on intuition and practical insights in this paper for researchers new to the topic, we only provide a superficial description of the imputation methods, and we leave more technical details to footnotes and the cited references. A more technical review with software instructions is offered in Dang *et al.* (2019).

literature and the “feel” behind each method, we will focus on offering the intuition rather than the technical details. Readers who are more familiar with these methods may also find some useful practical insights and suggestions for further research. While these methods can be applied to countries at different income levels, we will discuss several existing (well-cited) studies with a focus on developing countries. We also offer more technical details for these methods in Appendix A, which use similar headings as in the main text to facilitate cross-reference.

This paper consists of seven sections. We provide a simple, but new classification of poverty imputation methods in the next section, which also offers a roadmap to the other sections in the paper. This roadmap can function both as a graphical illustration that links the different methods, and as a “user’s guide” that can help readers better identify the issues of their interest. It is subsequently followed by more detailed discussion for poverty estimates on a non-monetary basis (i.e., using wealth indexes) (Section III) and monetary basis (Section IV). Section IV is further categorized according to different objectives for poverty estimation, including poverty estimates for project targeting (Section IV.1), monitoring poverty trends at the national level (Section IV.2), and poverty estimates that are disaggregated below the national level (Section IV.3). We then discuss estimates of poverty dynamics (Section V), before offering some further reflections (Section VI) and the conclusion (Section VII).

## **II. A “Roadmap” of Poverty Imputation Methods**

### **II.1. Roadmap**

The need to provide imputed poverty estimates depends on data availability, thus it varies from context to context, and nuanced differences exist between seemingly similar imputation methods. It can be useful to be explicit about the outcomes of interest, as well as the desired types of analysis and their associated challenges, before starting an imputation project. We offer in Figure 1 a

simple list of five common questions that can be asked about key aspects of the poverty imputation process, and the suggested strategy to address each question to be discussed in more detail (in the section listed in parentheses). Put differently, this figure proposes a checklist in the decision process to identify the relevant imputation method given the available data.

Figure 1 suggests that the first question a researcher can ask is whether poverty estimates are to be produced on a monetary (or money-metric) basis. If the answer to this question is no, the relevant imputation method is to construct an asset (wealth) index, and this method is discussed in more detail in Section III. If the answer is yes, the next useful question is whether the desired poverty estimate is nationally representative. If the answer to this question is no, proxy means testing is likely the relevant method (Section IV.1). A yes to this question leads to the third question of whether the poverty estimate will be used for dynamics analysis, which involves an examination of household movements into or out of poverty based on synthetic panel data (Section V).

A no to this question leads to the fourth question of whether the poverty estimate will be disaggregated below the national (or the survey's minimum aggregation) level. If yes, techniques commonly known as "poverty mapping" (or small-area estimation techniques) are most relevant. Another name for these techniques is survey-to-census imputation, where the imputation model is first built from a survey and subsequently applied to a census, since the latter offers reliable disaggregate data (Section IV.3). If no, survey-to-survey imputation methods are likely most appropriate. The final question is whether the imputation will involve surveys of the same design. Within-survey imputation and across-survey imputation should be respectively employed if imputation is done on surveys of the same design or of a different design (Section IV.2).

A related but different classification is offered in Dang, Jolliffe, & Carletto (2019), where a poverty imputation situation is classified according to the degree of missing consumption data. We present a slightly modified version of this classification in Table 1, which offers three categories in a roughly decreasing order of the severity of missing consumption data: completely missing (Category A), partially missing (Category B), and available cross-sectional data but missing panel data (Category C). In other words, while Figure 1 focuses on the functional or practical aspects of poverty imputation methods (and the associated roadmap pointing to the relevant discussion in this paper), Table 1 offers another classification that is more technical and data-oriented.

Notably, Category A can be further broken down into two data situations: (i) one where the available survey produces no consumption data, and the other (ii) where the available survey is designed for project targeting purposes. Corresponding to these data situations are the Demographic Health Surveys (DHSs) and most small-scale (or sub-national) surveys. Category B can also be further disaggregated into three different but related data situations: (i) consumption data are incomparable across survey rounds, (ii) consumption data are unavailable in the current survey but available in some other related surveys, and (iii) consumption data are unavailable at more disaggregated administrative levels than those offered in the current survey. Finally, Category C addresses the widespread situation that most surveys in developing countries do not provide (nationally representative) household panel data.<sup>3</sup> Table 1 also lists the typical poverty imputation (or corresponding data) situation, examples of surveys, and some recent studies corresponding to each missing data category, which we will discuss in more detail in later sections.

---

<sup>3</sup> Notably, the examples are for presentation purposes and can overlap. For example, depending on the discussed poverty imputation methods, a DHS can appear in either Category A (i.e., generate a wealth index when consumption data are completely missing) or Category B (i.e., implement imputation when consumption data are partially missing).

It can be useful to briefly review the country contexts that these data situations are most relevant for. While all the three main data categories are relevant for developing countries, some also apply to richer countries to varying degrees. Put differently, the decreasing order of the severity of missing consumption data in Table 1 roughly corresponds to an increasing order of country incomes. In particular, Category A and Category B (i) mostly pertain to low-income or lower-middle-income countries, where resources and technical constraints either hinder the collection of consumption data or render such process not fully effective. Category B (ii) and (iii) can apply to lower-middle-income or upper-middle-income countries. In fact, Category B (iii) also applies to a high-income country context such as the US where there is a need to provide poverty (and income) estimates at more disaggregated levels. Finally, Category C is also relevant for middle-income and high-income countries where there are no panel data, or the panel data suffer from quality issues. As an example, due to attrition, the percentage of households that remain in the panel Russia Longitudinal Monitoring Survey (RLMS) in the first 10 years after it was fielded is around 60 percent; this figure further decreases by half to 29 percent after another 10 years (Kozyreva *et al.*, 2016). We come back to more discussion on specific country examples in Sections III to V.

## II.2. Typical Estimation Equations

While we focus in this paper on the intuition behind existing studies, it can be useful to briefly discuss the commonly used empirical framework for clarity. Let  $x_j$  be a vector of characteristics that represent all the factors that determine a household's consumption, where  $j$  indicates the survey type. More generally,  $j$  can indicate either another round of the same household expenditure



survey, or a different survey (census), for  $j= 1, 2$ .<sup>4</sup> Subject to data availability,  $x_j$  can include household variables such as the household head's age, sex, education, ethnicity, religion, language (i.e., which can represent household tastes), occupation, and household assets or incomes. Occupation-related characteristics can generally include whether the household head works, the share of household members that work, the type of work that household members participate in, as well as context-specific variables such as the share of female household members that participate in the labour force, or some variables at the region level. Other community or regional variables can also be added since these can help control for different labor market conditions.

The following linear model is typically employed in empirical studies to project household consumption ( $y$ ) on household and other characteristics ( $x$ )

$$y_{ij} = \beta' x_{ij} + \mu_{ij} \quad (1)$$

for household  $i$  in survey  $j$ , for  $i= 1, \dots, N$  (see, e.g., Elbers, Lanjouw, & Lanjouw (2003), Ravallion (2016)). Equation (1) is often extended to a more general model

$$y_j = \beta_j' x_j + v_{cj} + \varepsilon_j \quad (2)$$

where the error term  $\mu_{ij}$  is now broken down into two components, one ( $v_{cj}$ ) is a cluster random effects and the other ( $\varepsilon_j$ ) the idiosyncratic error term. Hereafter we suppress the subscript that indexes households to make the notation less cluttered.<sup>5</sup> For convenience, we also refer to the survey that we are interested in imputing poverty estimates for as the target survey, and the survey that we can estimate Equation (1) on as the base survey. The former survey is usually more recent

---

<sup>4</sup> More generally,  $j$  can indicate any type of relevant surveys that collect household data sufficiently relevant for imputation purposes such as labor force surveys or demographic and health surveys.

<sup>5</sup> Conditional on household characteristics, the cluster random effects and the error terms are usually assumed uncorrelated with each other and to follow a normal distribution such that  $v_{cj}|x_j \sim N(0, \sigma_{v_j}^2)$  and  $\varepsilon_j|x_j \sim N(0, \sigma_{\varepsilon_j}^2)$ . While the normal distribution assumption results in the standard linear random effects model that is more convenient for mathematical manipulations and computation, it is not necessary for this type of model. As can be seen later, we can remove this assumption and use the empirical distribution of the error terms instead, albeit at the cost of somewhat more computing time.

(or offers more disaggregated information, as in the case of a census) and has no consumption data, while the latter is usually older and has consumption data.

We discuss the various poverty imputation methods starting from the next section and offer further technical details in Appendix A.

### **III. Non-Monetary Poverty Estimates**

#### *Motivation*

Consumption data are not always collected in a household survey for various reasons. The main reason is that since a typical consumption module (e.g., like that of a Living Standards Measurement Survey) usually consists of up to hundreds of items of food and non-food consumption, it takes time and a certain level of technical expertise to design such a module well. Furthermore, these items need to be updated from time to time to reflect changes in household consumption patterns (e.g., buying a smart phone is becoming an increasingly common purchase these days, but might not have been so just 10 years ago). Other factors such as seasonal variations in household consumption (e.g., consumption during holidays can be larger than that during regular times) could only be appropriately accounted for by fielding the consumption module in different months throughout the year. Consequently, collecting household consumption data requires considerable financial resources, time, and logistical capacity, which result in most household consumption surveys being implemented every few years rather than on an annual basis.<sup>6</sup> Given this data situation, for the years when consumption data are not collected, but other non-consumption surveys such as labor force surveys (LFSs) or small-scale surveys are available, these surveys may be “repurposed” to generate some substitute variable for consumption data. Indeed, one well-known example is the DHSs that (usually) do not collect consumption data but offer a wealth index instead.

---

<sup>6</sup> We return to more discussion on this point in Section VI.

*Method Description*

A well-established method to produce poverty estimates for surveys that have no consumption data but have data on household assets and the physical characteristics of the house is to construct a wealth index from these assets. This method typically consists of two steps. The first step is to identify the list of assets to be used in the construction of the index, and the second step is to apply some aggregation method to convert these assets into a (single-valued) index. Filmer & Pritchett (2001), who popularized the use of such indexes in the economic literature, employ the following variables from the India National Sample Survey (NSS): household ownership of consumer durables (including clock/watch, bicycle, radio, television, sewing machine, refrigerator, and car), characteristics of the household's dwelling (including indicators about toilet facilities, the source of drinking water, the rooms in the dwelling, the building materials used, and the main source of lighting and cooking), and household land ownership. They then utilize the principal component analysis (PCA) technique to create a wealth index, which offers different combinations of these assets that aim to capture as much variation in the data as possible.<sup>7</sup>

Variations exist on both of these steps. Since household questionnaires vary in different contexts, the list of assets to be employed depends on data availability. For example, Filmer & Scott (2012) examine all available asset variables in the DHS from four different countries and find these variables to vary across countries. In particular, their data sets contain between 12 (Uganda) and 29 (Nicaragua) indicators of asset ownership, and between 4 (Ghana) and 12 (Albania) and even as many as 37 (Zambia) indicators of housing characteristics. Besides the PCA

---

<sup>7</sup> See, for example, Jolliffe (2002) for a comprehensive treatment of PCA methods.

technique, the aggregation method can range from simply counting (or adding up) all the indicators of asset ownership to using other techniques such as factor analysis (Sahn & Stifel, 2000).<sup>8</sup>

#### *Examples and Remaining Challenges*

Sahn & Stifel (2000) offer a well-cited study that constructs wealth indexes for two or more periods from the DHS for 11 Sub-Saharan African countries between 1986 and 1998. They find that poverty generally declined, largely due to improvements in rural areas. A recent study by Harttgen, Klasen, & Vollmer (2013) analyses 160 DHS surveys from 33 African and 34 non-African countries and constructs wealth indexes in different ways. This study argues, however, that employing wealth indexes as a proxy for trends in household consumption is subject to various types of biases. These include changing preferences for certain assets (e.g., the increasing ownership of smart phones) and changing relative prices among different assets leading to more demand for one asset at the expense of others (e.g., the dramatically decreasing price of smart phones).

Other additional challenges exist with using wealth indexes to measure poverty. First, as suggested by the title of this section, wealth indexes offer a non-monetary or relative measure of poverty. Put differently, a wealth index can only identify a household as poor by its relative position on the population's wealth distribution. This renders it difficult, if not impossible, to compare the poverty rate among different countries without assuming that all the countries under comparison have comparable distributions of wealth. Practically speaking, the list of assets must

---

<sup>8</sup> Moser & Felton (2009) propose a variant of the PCA technique, where analysis is done on the components of each of several types of capital including physical capital, productive capital, human capital, and social capital. Another approach is to produce households' ranks in the population with the number of consumption items they own, if we make the additional assumption that households place an order of importance on their consumption items when having to reduce their consumption expenditure (Deutsch, Silber, & Wang, 2017). Notably, a common mistake when constructing wealth indexes is to convert ordinal asset variables to dummy variables and then aggregate; see Kolenikov & Angeles (2009) for a careful analysis of this issue. An alternative approach is to collect data on a reduced set of consumption items that may offer strong correlation with the total consumption aggregate (Morris *et al.*, 2000).

be either identical or comparable for all these countries, which is certainly harder to satisfy than comparing a consumption distribution denominated on a monetary basis such as the international Purchasing Power Parity (PPP) dollars. For example, an air-conditioner may be a valuable asset in countries in a tropical climate, but it may not even be found in those in a frigid climate, and vice versa with assets such as a heater.<sup>9</sup> Another inconvenience is that the poverty line (threshold) is relative and would need to be fixed for the whole distribution.<sup>10</sup> Furthermore, wealth indexes tend to provide biased estimates of poverty rates (as measured by household consumption) and may only be able to help track poverty trends over time under special conditions (Dang *et al.*, 2019).

We end this section on a cautious note that even if all the challenges discussed above are overcome, to our knowledge, several other technical challenges remain with the usage of wealth assets to measure poverty. For example, it remains unclear how many assets are sufficient in constructing a wealth index. The same challenges with cross-country comparison for wealth indexes may apply to different regions within a (large) country. Indeed, urban and rural areas likely differ in terms of consumption patterns, which can translate into differences in the assets that are used. It also remains unclear how to take into account the issue of quality versus quantity of assets (e.g., the ownership of one brand-new luxurious car can be different from that of three old and cheap cars).<sup>11</sup>

#### **IV. Poverty Estimates for Different Objectives**

---

<sup>9</sup> Asset ownership is often more “sticky” than household consumption; indeed, unless households could immediately sell assets as a coping strategy, it can be difficult to pick up the impact of shocks on household well-being through asset ownership alone.

<sup>10</sup> A similar concern applies to comparing the wealth index within a country (or countries) over time. In this case, assets for all countries from two periods must be pooled together to construct the wealth index.

<sup>11</sup> Ngo (2018) offers a method to construct a utility-based living standards index based on the values of assets that may address this issue to some extent (assuming that asset quality is correctly measured by asset values). But one practical challenge with this approach is that it requires the collection of asset price data.

We discuss in this section three popular uses of poverty imputation. The first use is to impute from a nationally representative survey into a smaller survey (project targeting), the second use is to impute from one round of the nationally representative survey into the next (or another nationally representative survey) for tracking poverty trends. The third use is to impute from a nationally representative survey into a census to obtain estimates at more disaggregated levels. These three uses represent variations of the same underlying common imputation model (Equation 2), which employ related but different assumptions.

#### **IV.1. Project Targeting**

##### *Motivation*

Identifying poor households that are beneficiaries of social transfer projects is a common task known as proxy means testing, but completing this task is not simple when there are no data on household consumption. Put differently, this case shares a similar constraint with that of the previous section: consumption data are not collected because of various costs or logistical constraints. One major difference with proxy means testing, however, is that it targets a subset of the poor population that are to benefit from (government or NGO) subsidies rather than measuring poverty for the whole population. Another important feature is that proxy means testing is (mostly) implemented when there are two sources of household survey data: one is a nationally representative survey that collects both household consumption data and the variables  $x_j$  (that are used in Equation (1)), and the other is the special and (much) lighter survey conducted to collect only the variables  $x_j$  for the purpose of proxy means testing.

##### *Method Description*

Proxy means testing involves two steps. The first step is to estimate the household consumption model using the nationally representative household (or larger) survey, as expressed in Equation (1). Once this is done, the second step is to combine the predicted coefficients  $\hat{\beta}$  obtained from the

first step with the variables  $x_j$  in the smaller survey to generate the predicted consumption (i.e., estimating the term  $\hat{\beta}'x_j$ ) for household  $i$  in the smaller survey. In other words, proxy means testing employs the predicted coefficients  $\hat{\beta}$  that come from the larger consumption survey and the variables  $x_{ij}$  in the smaller survey to generate the predicted household consumption, which is subsequently utilized to generate poverty estimates.<sup>12</sup>

### *Examples and Remaining Challenges*

A major advantage of proxy means testing over wealth indexes is that the former offers an estimate of household consumption, while the latter an estimate of household wealth. As such, proxy means testing perhaps provides a better estimate of poverty rates (that are based on household consumption data). Still, proxy means testing tends to offer biased poverty estimates. Indeed, Brown, Ravallion, & van de Walle (2018) offer a recent assessment of the performance of various proxy-means testing methods using data for nine African countries. They find that standard proxy-means testing is useful with targeting but excludes many poor people. Furthermore, even with some methodological adjustments, there is room for improvement with proxy-means tests, particularly with identifying the poorest.

The intuition behind this result is rather straightforward: the household poverty status is based on the household consumption that consists of two terms on the right hand side of Equation (1), but proxy means testing likely offers biased estimates since it offers the estimate for only one of these terms (i.e.,  $\hat{\beta}'x_j$ ).<sup>13</sup> Perhaps the largest advantages of proxy means testing is that for a smaller population group, it is rather simple to implement (relatively speaking, compared with other

<sup>12</sup> For recent reviews of proxy mean tests, see Grosh *et al.* (2008), Coady *et al.* (2014), and Brown *et al.* (2018).

<sup>13</sup> Proxy means tests offer unbiased estimates of poverty only in the special case that the poverty line is set exactly at the mean consumption level (see Dang *et al.* (2019) for a more detailed technical discussion). See also Diamond *et al.* (2016) for a careful comparison of the poverty score card—a variant of proxy means testing—and other (regression-based) poverty imputation methods.

imputation methods). Consequently, it can offer a quick and inexpensive estimate of poverty in the absence of household consumption data. This, however, may come at the expense of estimate precision.

## **IV.2. Tracking Poverty Trends at the National Level**

### *Motivation*

While proxy means testing targets a subset of the poor population that are to benefit from some social transfer programs, tracking trends in poverty rates at the national level requires poverty estimates for the whole population. The motivation for poverty imputation in this case is related to discussions regarding national poverty trends. For example, the Sustainable Development Goals (SDGs) require frequent monitoring of (national and global) poverty trends that would perhaps require much more frequent collection of household consumption data than most existing household consumption surveys allow.<sup>14</sup> Indeed, Serajuddin *et al.* (2015) find that over the period 2002- 2011, of the 155 countries for which the World Bank monitors poverty data using the World Development Indicators (WDI) database, almost one-fifth (i.e., 28) have only one poverty data point and as many as 29 countries do not have any poverty data point in the same period. A recent survey by Beegle *et al.* (2016) indicates that just slightly more than half (i.e., 27) of the 48 countries in Sub-Saharan Africa had two or more comparable household surveys for the period between 1990 and 2012. Furthermore, most developing countries, say India, do not collect household consumption data annually. As such, poverty estimates for these countries would have to be obtained using alternative methods such as poverty imputation methods.

In this case, similar to proxy means testing, we also need two sources of household survey data: one survey that collects both household consumption data and the variables  $x_j$  (that are used

---

<sup>14</sup> The first goal of the SDGs is to eliminate extreme poverty and reduce national poverty levels at least by half by 2030. For details see <https://sustainabledevelopment.un.org/sdg1>.



in Equation (1)), and the other is another non-consumption survey that offers only the variables  $x_j$ . One major difference is that both surveys should provide nationally representative data.

We can divide this case into two subcategories depending on whether the non-consumption survey is of the same design or a different design from the consumption survey. The former subcategory—hereafter referred to as within-survey imputation—includes situations where the household consumption data collected in the two surveys are not comparable over time due to changes to the consumption items (while the  $x_j$  remain similar). This situation occurs more often than one might think; for example, national statistical agencies regularly update the list of consumption items from time to time to reflect changes in household consumption patterns regarding high-technology goods such as smart phones or smart televisions. The latter subcategory—hereafter referred to as across-survey imputation—includes situations where the consumption survey was implemented a few years back, and a newer round is yet to be fielded. At the same time, there exists another more recent non-consumption survey such as an LFS that can be combined with the consumption survey to provide more recent poverty estimates. Both of these sub-categories of within-survey imputation and across-survey imputation are also commonly known as survey-to-survey imputation (which is different from survey-to-census imputation or the poverty mapping technique discussed in the next section).

We discuss next the techniques and challenges for both these subcategories.

#### *Method Description*

Similar to proxy means testing, survey-to-survey imputation methods consist of two main steps. The first step is to estimate the household consumption model using the consumption survey, as expressed in Equation (1). Once this is done, the second step is to combine the predicted coefficients  $\hat{\beta}$  and the distribution of the error term  $\mu_{ij}$  obtained from the first step with the

variables  $x_{ij}$  in the non-consumption survey to generate the predicted consumption, which is in turn converted into the (headcount) poverty rate. However, different from proxy means testing, the predicted household consumption generated using survey-to-survey imputation methods is composed of *both* the two terms on the right-hand side of Equation (1), that is  $\hat{\beta}'x_{ij}$  and  $\hat{\mu}_{ij}$ .<sup>15</sup> For consistency, the poverty line in the base survey—rather than that in the target survey—should be used together with the predicted consumption to obtain poverty estimates. In fact, since we need to use only the poverty line in the base survey, this obviates additional challenges with identifying more than one comparable poverty line over time (see, e.g., Lanjouw & Lanjouw (2001)).

Notably, the estimation framework utilized by most existing economic studies is largely based on the seminal survey-to-census imputation method offered by Elbers *et al.* (2003) (which we return to more discussion in the next section).<sup>16</sup> Most recently, building on the Elbers *et al.* (2003) method, Dang *et al.* (2017a) attempt to bring some further improvements to the survey-to-survey poverty imputation method, which include simpler variance formulas, more guidance on the selection of control variables for model building, and formulas for standardization of variables from surveys with different sampling designs. They validate estimation results against both

<sup>15</sup> Furthermore,  $\hat{\mu}_{ij}$  is usually disaggregated into a cluster random effects term ( $\hat{v}_{cj}$ ) and an idiosyncratic error term ( $\hat{\varepsilon}_j$ ) as in Equation (2). This feature, as well as the addition of  $\hat{\mu}_{ij}$  to the error term, results in more accurate estimates of household consumption and poverty estimates than proxy means testing (see Dang *et al.* (2017a) for more detailed discussion).

<sup>16</sup> Variants on this method exist. For example, Tarozzi (2007) proposes a two-step inverse probability weighting probit estimator, with the relevant weights derived in the first step from the change in the distribution of household characteristics across the two surveys. Mathiassen (2009) also employs a probit estimator, but proposes an exact expression for the standard errors and imposes a stricter parametric functional form on the error term. On the empirical front, Christiaensen *et al.* (2012) and Mathiassen (2013) apply the Elbers *et al.* (2003) framework to provide poverty estimates based on within-survey imputation for several countries including China, Kenya, the Russian Federation, Uganda, and Vietnam. Using the same technique, other studies combine the household consumption survey with other surveys such as the DHS (Stifel & Christiaensen, 2007) or the LFS (Mathiassen, 2009; Doudich *et al.*, 2016) to provide across-survey imputation. This imputation method is also related to a larger literature on multiple imputation (MI) in statistics (see, e.g., Rubin (1987) and Carpenter & Kenward (2013)). Yoshida *et al.* (2015) provide a recent application of MI method to poverty estimation. See Dang *et al.* (2019) for further discussion on these methods.

household consumption data and LFS data from Jordan before combining these two sources of data to provide more recent poverty estimates.

### *Examples and Remaining Challenges*

Poverty reduction in India has been subject to intense debates in the past, which was caused by comparability issues with different rounds of the National Sample Survey (NSS)—the country’s mainstream consumption survey data—over time (Deaton & Kozel, 2005). Similar concerns, albeit to a lesser extent, were raised about the dramatic poverty reduction between 2004 and 2012 as well. One main reason is that the questionnaire design of the consumption module in the 2011/12 (68th) round of the NSS is not comparable to that in the 2009/10 (66th) round (and 2004/05 or 61st round), which may result in inconsistently constructed and incomparable consumption data. Dang & Lanjouw (2018) apply imputation methods to provide checks on the poverty trend. They first build an imputation model using the 2004/05 round as the base survey to obtain poverty estimates for 2009/10, which are satisfactorily not different from the actual poverty rates. They subsequently employ the same model using the 2009/10 as the base survey to obtain poverty estimates for 2011/12. These estimates are close to the actual rates in this year, thus providing supportive evidence for the swift fall in poverty observed in the data.

A key assumption for survey-to-survey imputation is that the coefficients  $\beta$  estimated from the previous consumption survey can be combined with the variables in the more recent survey to obtain poverty estimates. This assumption is also commonly known as the constant parameter assumption.<sup>17</sup> While concerns exist that this assumption is likely to be valid only under normal circumstances, rather than during periods of fast (economic growth and) poverty reduction, it has been shown to hold during a period of dramatic economic growth in China and Vietnam where

---

<sup>17</sup> Notably, this assumption is also needed for consistency for the  $\hat{\beta}'x_{ij}$  part in the case of proxy means tests.

poverty incidence was cut by around half (Christiaensen *et al.*, 2012). Furthermore, a weaker version of this assumption has been proposed and validated for data from various countries such as India, Jordan, and Vietnam (Dang *et al.*, 2017a; Dang & Lanjouw, 2018; Dang *et al.*, 2019). Yet, we would like to note that the validity of this assumption can be context-specific, and it can be useful to check it using at least two previous rounds of household consumption surveys wherever such data are available. Common sense also suggests that the longer is the time interval between the base survey and the target survey, the more likely that this assumption can be violated.

Another concern with survey-to-survey imputation, or more accurately speaking, across-survey imputation methods, is that the variables used in the imputation in both the base survey and the target survey should have the same distribution. This seemingly rather innocuous condition appears often taken for granted in many studies, but if it is not satisfied, it may result in severely biased estimates (Dang *et al.*, 2017a). The intuition is rather straightforward: variables such as household size or labor force participation may be defined differently in a household consumption survey and a labor force survey, and the data can be collected accordingly in different ways. Consequently, this condition should be carefully checked with, say, t-tests for equality of distributions for the variables of interest. Furthermore, appropriate adjustments (e.g., standardizing the variables) should be done before imputation is implemented on surveys of a different design.<sup>18</sup>

### **IV.3. Disaggregated Poverty Estimates at the Sub-National Level**

#### *Motivation*

---

<sup>18</sup> The inconsistency between different surveys is well documented in studies using data from richer countries. For example, Abraham *et al.* (2013) examine the differences between employment data between the U.S. Current Population Surveys and employer-reported administrative data. See also Angrist & Krueger (1999) for a related review of comparability and other data issues with a focus on labor force surveys. Dang *et al.* (2017a) also offer a simple method to standardize the variables from the two different surveys. This study also provides more discussion on another related issue of selecting variables in estimating the consumption model as in Equation (1).

In most household consumption surveys, consumption data are rarely available at a disaggregated level (such as the state or province level) due to the typically limited sample size of household surveys. However, there exists a strong demand to produce poverty estimates at more disaggregated levels for various purposes such as social transfer targeting or budget allocation. For example, statistical agencies such as the U.S. Census Bureau routinely implement this task to identify poor communities.<sup>19</sup> This task is commonly known as “poverty mapping” in most studies on developing countries, since poverty estimates are usually graphed on a map at a lower-level administrative level (such as that of a district or a local community). In other words, this case typically involves survey-to-census imputation, since only censuses can offer more disaggregated data than those available in a household survey. Put differently, survey-to-census imputation can often be regarded as some type of geographical imputation, which differs from the (mostly) temporal imputation offered with survey-to-survey imputation.

#### *Method Description*

Similar to proxy means testing and survey-to-survey imputation methods, survey-to-census imputation methods also consist of two main steps. The first step is to estimate the household consumption model in Equation (2) using the more aggregated-level household survey. The second step is to combine the predicted coefficients  $\hat{\beta}$  and the distribution of the two error terms  $\hat{v}_{cj}$  and  $\hat{\varepsilon}_j$  obtained from the first step with the variables  $x_{ij}$  in the census to generate predicted consumption data at a more disaggregated level.

Elbers *et al.* (2003) is perhaps the first study that introduces a formal framework for survey-to-census imputation in economics. Building on this framework, Tarozzi & Deaton (2007) propose

---

<sup>19</sup> See, e.g., <https://www.census.gov/srd/csrsm/SmallArea.html>.

another condition where the conditional distribution of household consumption given  $x_j$  is the same for both the survey and the census. This assumption ensures that the estimated parameters from the smaller areas (as representative in the survey) can be imposed on the data for the larger areas (as representative in the census). Notably, another name for this topic in the statistical literature is “small-area estimation” (see, e.g., Rao & Molina (2015) for a recent textbook treatment). Recent developments have been proposed, mostly in the statistical literature, to offer extensions or alternative estimation techniques to the Elbers *et al.* (2003) method.<sup>20</sup>

#### *Examples and Remaining Challenges*

Elbers *et al.* (2007) use “poverty maps” from three low-income and middle-income countries (i.e., Cambodia, Ecuador, and Madagascar) for an ex ante evaluation of the distributional incidence of geographic targeting of public resources. They simulate the impact on poverty of transferring an exogenously given budget to geographically defined sub-groups of the population according to their relative poverty status. They find large gains from targeting smaller administrative units, such as districts or villages. They also suggest that poverty map-based geographic targeting can be combined with within-community targeting mechanisms for better estimation results. Lanjouw, Marra, & Nguyen (2017) employ small area estimation techniques to estimate the poverty indexes of Vietnam's provinces and districts in 2009. They find poverty rates to become more spatially concentrated over time, which is consistent with agglomeration-related growth processes. They offer simulation results suggesting that in both 1999 and 2009 geographic targeting for poverty

---

<sup>20</sup> See, e.g., Bilton *et al.* (2017) for a proposal to use a classification trees technique for poverty mapping, and Das & Chambers (2017) for alternative standard error formulae with the Elbers *et al.* (2003) method. See also Pratesi (2016) for a recent collection of studies discussing various technical aspects of poverty mapping. Another study by Steele *et al.* (2017) applies machine learning techniques and big data (i.e., cell phone and satellite data) to evaluate poverty mapping. Engstrom, Hersh, & Newhouse (2017) propose to employ satellite data to refine poverty estimates.

alleviation improves upon a uniform lump-sum transfer, particularly for the more spatially disaggregated target populations.

We note that survey-to-census imputation shares a similar issue as with across-survey imputation methods: the variables used in the imputation in both the (base) survey and the (target) census should have the same distribution. Perhaps we cannot overemphasize the importance of this condition, given both the theoretical results and empirical evidence (offered by Tarozzi & Deaton (2007) and Dang *et al.* (2017a)). However, inadequate attention appears to have been paid to rigorously standardize the variables in both the survey and the census. Furthermore, the constant parameter assumption is also necessary for us to impute from the survey into the census.

Another interesting and useful area that needs more research is how to produce and interpret the evolution of poverty maps over time. Multiple challenges exist with generating such dynamic poverty maps. One is that we would need survey and census data at two points in time, and both these data sources should be comparable both at each point in time (i.e., the issue discussed immediately above) and over time (i.e., the issue of the constant parameter assumption as discussed with survey-to-survey imputation methods). While some alternative techniques have been proposed,<sup>21</sup> it seems that this topic still needs more development.

## V. Dynamic Poverty Estimates

### *Motivation*

Different poverty situations are best addressed with different policy responses. In particular, transitory and chronic poverty typically require different policy instruments and no single policy may successfully address both. For example, while transitory poverty can be alleviated with safety

---

<sup>21</sup> For example, Nguyen (2011) offers an innovative study that uses panel data from household surveys to estimate the relation between expenditure in the second period and household characteristics in the first period. The estimated parameters are then applied to a census in the second period to predict expenditure and poverty measures in a future third period. This approach may address, partially but not completely, the issues raised above.

net programs, chronic poverty would need to be tackled with structural and longer-term interventions such as investment in human capital and building infrastructure.<sup>22</sup>

However, poverty estimates based on cross-sectional data provide only static snapshots of poverty rates, rather than the dynamics of poverty transitions over time. Absent a clear understanding of poverty dynamics, a seemingly unchanged poverty rate of, say, 15 percent in two periods could conceal dynamic processes ranging from zero mobility (i.e., where all households see no change in their welfare) to perfect mobility (i.e., where all poor households in the first period escaped poverty and were all replaced by households that had previously been non-poor in the first period) and any scenario between these two extremes. Dynamics analysis is crucial for the design of effective and efficient poverty reduction policies, but such analysis requires panel survey data that are usually unavailable (particularly for developing countries).

### *Method Description*

In the absence of actual panel data, Dang *et al.* (2014) and Dang & Lanjouw (2013) recently propose methods to construct synthetic panels from two (or more) rounds of repeated cross sections, which have provided encouraging results in various settings.<sup>23</sup> Their techniques share certain similarities with survey-to-survey imputation methods and include the following steps. First, estimate the household consumption model in Equation (1) using the available cross sections

---

<sup>22</sup> See, e.g., Barrett (2005) and Ravallion (2016) for more discussion on different policy approaches to reduce poverty.

<sup>23</sup> For example, recent applications and further validations of these synthetic panels methods include Ferreira *et al.* (2013), Cruces *et al.* (2015), and Vakis, Rigolini, & Lucchetti (2015) for Latin American countries, Bourguignon, Moreno, & Dang (2018) and Foster & Rothbaum (2015) for Mexico, Balcazar *et al.* (2018) for Colombia, Martinez *et al.* (2013) for the Philippines, Garbero (2014) for Vietnam, Cancho *et al.* (2015) for countries in Europe and Central Asia, Dang & Ianchovichina (2018) for countries in the Middle East and North Africa region, Dang & Dabalén (2019) and Dang, Lanjouw, & Swinkels (2017b) for Sub-Saharan African countries, and Dang & Lanjouw (2017, 2018) for India, Vietnam, and the United States. Researchers at international organizations including the UNDP and the Asian Development Bank have also applied these methods for analysis of welfare mobility (UNDP, 2016; Jha *et al.*, 2018; OECD, 2018); see also OECD (2015) for an application by the OECD to study labor transitions in richer countries. See also Gibson (2001) for a related study on how panel data on a subset of individuals can be used to infer chronic poverty for a larger sample.



to obtain the predicted coefficients  $\hat{\beta}$  and the distribution of the error term  $\mu_{ij}$  for each year (i.e., its standard deviations  $\sigma_{\mu_j}$ ). Second, estimate the correlation coefficient  $\rho$  of the error terms  $\mu_{ij}$ , using cohort-aggregated household consumption between the two survey rounds. Third, combine  $\hat{\beta}$ ,  $\hat{\rho}$ , and  $\hat{\sigma}_{\mu_j}$  to obtain estimates of poverty mobility using bivariate probability formulae.

A key difference, however, with survey-to-survey imputation methods is that the  $x_j$  in Equation (1) should consist of time-invariant characteristics alone. These include such variables as ancestry, language, place of birth, and parental education, which provide the connectors between different rounds of cross sections with the appropriate age adjustments.<sup>24</sup> The age range can be adjusted similarly if there is a different time interval between the two survey rounds. Other time-varying household characteristics can also be used if retrospective questions about the round-1 values of such characteristics are asked in the second-round survey. For example, the cohorts age 25 to 55 in the first survey round would likely be the same cohorts age 30 to 60 in the second survey round five years later, assuming the two survey rounds are comparable. The effects of the time-varying variables are thus subsumed in the correlation coefficient  $\rho$  of the error terms. This feature stands in contrast with the survey-to-survey imputation methods discussed in earlier sections that aim to capture as many relevant (time-invariant and time-varying) variables on the right-hand side in their estimation models.<sup>25</sup>

### *Examples and Remaining Challenges*

<sup>24</sup> While religion is a time-invariant variable in most contexts, it may not necessarily be so where people can convert. It is thus useful to examine the specific context. We would like to thank an anonymous reviewer for suggesting this point.

<sup>25</sup> This difference is further accentuated with some missing data imputation methods in the statistical literature where sample design variables such as sampling weights, strata and cluster identifiers are also included in the estimation model (see, e.g., Rubin (1987)).

Despite a growing collection of nationally representative panel surveys for African countries, data coverage is unfortunately available for only seven countries, and the time periods spanned by these panel surveys are mostly limited to short periods of three years or less.<sup>26</sup> To overcome this data shortage, Dang & Dabalen (2019) construct synthetic panel data for more than 20 countries accounting for two-thirds of the population in Sub-Saharan Africa; these synthetic panels span an average of six years for each country. Their analysis suggests that all these countries as a whole have had pro-poor growth, with one-third of the poor population escaping poverty, which is larger than the proportion of the population that fell into poverty in the same period. Chronic poverty, however, remains high and a considerable proportion of the population is vulnerable to falling into poverty.

Despite their increasing popularity, synthetic panels are not the perfect substitute for actual panel data. In particular, not much is known whether, and how useful synthetic panels can be utilized in applications involving studying a causal relationship or regression analysis. The analysis offered to date in terms of profiling the poverty trajectories for population groups with these synthetic panels is mostly descriptive, with little explicit attention to underlying causal mechanisms. Furthermore, in countries with high rates of immigration and emigration, one would need to be cautious in constructing synthetic panels, as the composition of the cohort may be changing over time. More research is thus needed on the application of these synthetic panels in such contexts. Put differently, it is useful to know to which extent synthetic panels can substitute for actual panel survey data.

## **VI. Further Reflections on Related Issues**

---

<sup>26</sup> More details about the LSMS-ISA panel surveys are available at <http://surveys.worldbank.org/lms/programs/integrated-surveys-agriculture-ISA>.

We are more often than not faced with contexts where we either have no consumption data, or the available consumption data have quality issues. These challenges are illustrated with the review studies by Serajuddin *et al.* (2015) and Beegle *et al.* (2016) and our examples in the preceding sections. Clearly, poverty imputation is useful, and may likely be the only choice in these cases, particularly for poorer countries (and to some degree for richer countries). But what about other contexts where we have a choice over poverty imputation and implementing a full-scale household consumption survey?

Dang *et al.* (2019) suggest that even in these contexts, there are still a couple of advantages to poverty imputation methods, particularly in the following scenarios

- i) in the immediate term (when micro survey data are not fully available for all countries)
- ii) survey costs and/or survey implementation pose a challenge, and
- iii) back-casting consumption from a more recent survey for better comparison with older surveys.

It can be useful to offer some additional commentary on these cases. Cases (i) and (ii) are closely related and are perhaps the main driving factor behind poverty imputation for poorer countries, since very few, if any, of these countries can afford the financial and logistical challenges of fielding a household consumption survey every year. Consequently, most developing countries are likely to implement the household consumption survey every few years. In such contexts, poverty imputation methods can offer a (far) less costly option to provide estimates for the intervening years between the surveys.<sup>27</sup> A special example of case (ii) that is receiving increasingly more

---

<sup>27</sup> Recent estimates suggest that the average cost of implementing a household consumption survey (in 2014 or later) ranges from approximately \$US 800,000 to US\$ 5 million depending on the context and sample sizes (Kilic *et al.*, 2017). At the same time, implementing poverty imputation may require only a fraction of this cost since its major expense is to cover analytical time. Indeed, selective pairing of international experts with national statistical agencies' staff can form small teams that provide both cost-effective analysis and local capacity building.

attention is poverty estimation for refugee populations, who are more mobile and much harder to reach than regular populations (see, e.g., Beltramo *et al.* (2019)).

Case (iii), although less common, represents the scenario where poverty imputation is the only route to provide comparable poverty estimates for surveys fielded in the past. For example, national statistical agencies usually update questions in the consumption module of the household survey from time to time to reflect latest changes in household consumption patterns. These changes may be caused by advances in technology (e.g., more individuals own a smart phone now than 10 years ago). They may also be caused by various other reasons such as survey implementation (e.g., the switch from paper-based interviews to computer-based interviews), design shifts (e.g., changes from diary to recall data collection), and seasonality (i.e., surveys in agrarian societies being collected at different points during the crop cycle). These changes result in comparability issues for the consumption aggregate in the more recent survey rounds compared with those in earlier survey rounds.<sup>28</sup> Poverty imputation methods can offer a cost-effective solution in such situations.

Another advantage with poverty imputation methods that has received little attention is the fact that such methods, particularly survey-to-survey imputation, can help us bypass the oftentimes thorny issues of obtaining the appropriate (intertemporal and intraregional) price deflators. This issue worsens for cross-country comparison, since in that case we have to employ conversion factors to convert the different currencies to the same base.

Notably, a promising direction for further development with poverty imputation methods is that they need not be restricted to the topic of poverty alone but can be utilized in other fields as

---

<sup>28</sup> Besides our earlier discussion with India, survey design issues that compromise the comparability of poverty estimates are found in various countries such as China (Gibson *et al.*, 2003), Tanzania (Beegle *et al.*, 2012), and Vietnam (World Bank, 2012).

well. For example, Fujii (2010) and Sohnesen *et al.* (2017) employ the Elbers *et al.* (2003) method to provide a map of child malnutrition respectively in Cambodia and Ethiopia; Gibson (2018) uses the same method to study the effect of deforestation on subsequent inequality in the rural Solomon Islands. As another example, a recent study by Dang & Ianchovichina (2018) constructs synthetic panels from the cross sections in the Gallup World Poll for 16 countries in the Middle East to offer analysis of dynamics of subjective well-being during the Arab Spring period. Figure 2, taken from this study, plots the percentage of the unhappy or vulnerable in the first year who move up one or two welfare categories in the second year for major population groups classified by gender, education levels, work status, migration status, and residence areas.<sup>29</sup> This figure suggests that upward mobility is weaker than downward mobility both for Arab Spring and other Arab countries, and migrants are more likely to be less upwardly mobile (and more downwardly mobile) in Arab Spring countries, but the opposite holds for non-Arab Spring countries. As such, in a similar spirit, other potentially useful applications of poverty imputation methods may include emerging policy-relevant topics such as vulnerability, multidimensional poverty, or gender equality.

Yet, we end this section on a cautious note that poverty imputation methods, like most other statistical models, rely heavily on the modelling techniques and their accompanying assumptions. If the model assumptions are satisfied by the data, poverty imputation can yield encouraging and low-cost results. However, the opposite holds where the model assumptions are invalid. It would thus be useful to offer careful checks on the relevant modelling assumptions as well as the variable selection process before providing imputation-based poverty estimation.

## VII. Conclusion

---

<sup>29</sup> Using the Gallup World Poll data, Dang & Ianchovichina (2018) define the unhappy group as those with a life evaluation scores of 4 and below, the vulnerable group as those with life evaluation scores between 5 and 7, and the happy (or satisfied) with wellbeing scores of 8 or above.

The growing demand for more frequent and accurate poverty estimates is not satisfied by the current availability of household consumption data, at least in the short run. Imputation methods offer a promising solution against this background and have been widely in use.<sup>30</sup> These methods have also received increasingly more attention. For example, the recent high-profile “Atkinson” report on monitoring global poverty (World Bank, 2017) explicitly calls for further exploration of imputation techniques for poverty measurement purposes in data-scarce contexts. Yet, there is currently a severe dearth of research that can bridge the gaps between the typical development practitioners and the latest advances in the field. We aim to help fill in this gap with this review.

---

<sup>30</sup> See, e.g., Jolliffe *et al.* (2015) for a recent review.

## References

- Abraham, K. G., Haltiwanger, J., Sandusky, K., & Spletzer, J. R. (2013). Exploring differences in employment between household and establishment data. *Journal of Labor Economics*, 31, S129-S172. <https://www.journals.uchicago.edu/doi/abs/10.1086/669062>
- Alkire, S., Foster, J., Santos, M. E., Seth, S., Roche, J. M., & Ballon, P. (2015). *Multidimensional poverty measurement and analysis*. USA: Oxford University Press.
- Angrist, J. D., & Krueger, A. B. (1999). Empirical strategies in labor economics. In *Handbook of labor economics* (Vol. 3, pp. 1277-1366). Elsevier.
- Balcazar, C. F., Dang, H.-A., Malasquez, E., Olivieri, S., & Pico, J. (2018). *Welfare dynamics in Colombia: results from synthetic panels*. Policy Research Working Paper No. 8441. Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-8441>
- Barrett, C. B. (2005). Rural poverty dynamics: development policy implications. *Agricultural Economics*, 32, 45-60. <https://doi.org/10.1111/j.0169-5150.2004.00013.x>
- Beegle, K., De Weerd, J., Friedman, J., & Gibson, J. (2012). Methods of household consumption measurement through surveys: Experimental results from Tanzania. *Journal of Development Economics*, 98(1), 3-18. <https://doi.org/10.1016/j.jdeveco.2011.11.001>
- Beegle, K., Christiaensen, L., Dabalen, A., & Gaddis, I. (2016). *Poverty in a rising Africa*. Washington, DC: The World Bank. <https://doi.org/10.1596/978-1-4648-0723-7>
- Beltramo, T., Dang, H.-A., Sarr, I., & Verme, P. (2019). Is imputing poverty efficient? An example from refugee data in Chad". *Paper to be presented at the American Economic Association's Annual Meeting* (San Diego, 2020).
- Bilton, P., Jones, G., Ganesh, S., & Haslett, S. (2017). Classification trees for poverty mapping. *Computational Statistics & Data Analysis*, 115, 53-66. <https://doi.org/10.1016/j.csda.2017.05.009>
- Bourguignon, F., Moreno, H., & Dang, H.-A. (2018). On the construction of synthetic panels. *Working paper*. Paris: Paris School of Economics.
- Brown, C., Ravallion, M., & van de Walle, D. (2018). A poor means test? Econometric targeting in Africa. *Journal of Development Economics*, 134, 109-124. <https://doi.org/10.1016/j.jdeveco.2018.05.004>
- Cancho, C., Dávalos, M. E., Demarchi, G., Meyer, M., & Páramo, C. S. (2015). *Economic mobility in Europe and Central Asia: Exploring patterns and uncovering puzzles*. World Bank Policy Research Paper No. 7173. Washington, DC: The World Bank. <https://doi.org/10.1596/1813-9450-7173>
- Carpenter, J. & Kenward, M. (2013). *Multiple imputation and its application*. Chichester: John Wiley & Sons.

- Christiaensen, L., Lanjouw, P., Luoto, J., & Stifel, D. (2012). Small area estimation-based prediction methods to track poverty: validation and applications. *The Journal of Economic Inequality*, 10, 267-297. <https://link.springer.com/article/10.1007/s10888-011-9209-9>
- Coady, D., Grosh, M., & Hoddinott, J. (2004). Targeting outcomes redux. *The World Bank Research Observer*, 19(1), 61-85. <https://doi.org/10.1093/wbro/lkh016>
- Cruces, G., Lanjouw, P., Lucchetti, L., Perova, E., Vakis, R., & Viollaz, M. (2015). Estimating poverty transitions using repeated cross-sections: a three-country validation exercise. *The Journal of Economic Inequality*, 13, 161-179. <https://link.springer.com/article/10.1007/s10888-014-9284-9>
- Cuesta, J., & Ibarra, G. L. (2017). Comparing cross-survey micro imputation and macro projection techniques: Poverty in post revolution Tunisia. *Journal of Income Distribution*, 25(1), 1-30. <https://pdfs.semanticscholar.org/678f/e7e47cb3067712ea380a6d0fcd382031c283.pdf>
- Dang, H.- A., & Dabalén, A. L. (2019). Is poverty in Africa mostly chronic or transient? Evidence from synthetic panel data. *The Journal of Development Studies*, 55, 1527-1547. <https://doi.org/10.1080/00220388.2017.1417585>
- Dang, H.-A., & Ianchovichina, E. (2018). Welfare dynamics with synthetic panels: the case of the Arab world in transition. *Review of Income and Wealth*, 64, S114-S144. <https://doi.org/10.1111/roiw.12389>
- Dang, H.-A., & Lanjouw, P. (2013). Measuring poverty dynamics with synthetic panels based on cross-sections. *World Bank Policy Research Working Paper No. 6504*. Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-6504>
- . (2016). Toward a new definition of shared prosperity: A dynamic perspective from three countries. In *Inequality and growth: Patterns and policy* (pp. 151-171). London: Palgrave Macmillan.
- . (2017). Welfare dynamics measurement: Two definitions of a vulnerability line and their empirical application. *Review of Income and Wealth*, 63, 633-660. <https://doi.org/10.1111/roiw.12237>
- . (2018). Poverty Dynamics in India between 2004 and 2012: Insights from Longitudinal Analysis Using Synthetic Panel Data. *Economic Development and Cultural Change*, 67(1), 131-170. <https://www.journals.uchicago.edu/doi/abs/10.1086/697555>
- Dang, H.-A., Lanjouw, P. F., & Serajuddin, U. (2017a). Updating poverty estimates in the absence of regular and comparable consumption data: methods and illustration with reference to a middle-income country. *Oxford Economic Papers*, 69, 939-962. <https://doi.org/10.1093/oep/gpx020>



- Dang, H.-A., Lanjouw, P. F., & Swinkels, R. (2017b). Who Remained in Poverty, Who Moved Up, and Who Fell Down?. In *Poverty Reduction in the Course of African Development*, 183-217. UK: Oxford University Press.
- Dang, H.-A., Jolliffe, D., & Carletto, C. (2019). Data gaps, data incomparability, and data imputation: A review of poverty measurement methods for data-scarce environments. *Journal of Economic Surveys*, 33, 757-797. <https://doi.org/10.1111/joes.12307>
- Dang, H.-A., Lanjouw, P., Luoto, J., & McKenzie, D. (2014). Using repeated cross-sections to explore movements into and out of poverty. *Journal of Development Economics*, 107, 112-128. <https://doi.org/10.1016/j.jdeveco.2013.10.008>
- Das, S., & Chambers, R. (2017). Robust mean-squared error estimation for poverty estimates based on the method of Elbers, Lanjouw and Lanjouw. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 180, 1137-1161. <https://doi.org/10.1111/rssa.12311>
- Deaton, A. (1997). *The analysis of household surveys: A microeconomic approach to development policy*. MD: The Johns Hopkins University Press.
- Deaton, A., & Kozel, V. (2005). *The great Indian poverty debate*. New Delhi: Macmillan.
- Deutsch, J., Silber, J., & Wan, G. (2017). Curbing one's consumption and the impoverishment process: The case of Western Asia. *Research on Economic Inequality: Poverty, Inequality and Welfare*, 137-159. <https://doi.org/10.1108/S1049-258520170000025010>
- Diamond, A., Gill, M., Dellepiane, M. R., Skoufias, E., Vinha, K., & Xu, Y. (2016). *Estimating poverty rates in target populations: An assessment of the simple poverty scorecard and alternative approaches*. The World Bank. *Policy Research Working Paper No. 7793*. Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-7793>
- Doudich, M., Ezzrari, A., Van der Weide, R., & Verme, P. (2016). Estimating quarterly poverty rates using labor force surveys: a primer. *The World Bank Economic Review*, 30, 475-500. <https://doi.org/10.1093/wber/lhv062>
- Elbers, C., Lanjouw, J. O., & Lanjouw, P. (2003). Micro-level estimation of poverty and inequality. *Econometrica*, 71(1), 355-364. <https://doi.org/10.1111/1468-0262.00399>
- Elbers, C., Fujii, T., Lanjouw, P., Özler, B., & Yin, W. (2007). Poverty alleviation through geographic targeting: How much does disaggregation help?. *Journal of Development Economics*, 83(1), 198-213. <https://doi.org/10.1016/j.jdeveco.2006.02.001>
- Engstrom, R., Hersh, J., & Newhouse, D. (2017). Poverty from space: Using high-resolution satellite imagery for estimating economic well-being. *Policy Research Working Paper No. 8284*. Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-8284>

- Ferreira, F. H., Messina, J., Rigolini, J., López-Calva, L. F., Lugo, M. A., & Vakis, R. (2012). *Economic mobility and the rise of the Latin American middle class*. The World Bank. Washington DC: World Bank. <https://doi.org/10.1596/978-0-8213-9634-6>
- Filmer, D., & Pritchett, L. H. (2001). Estimating wealth effects without expenditure data—or tears: an application to educational enrollments in states of India. *Demography*, 38(1), 115-132. <https://link.springer.com/article/10.1353/dem.2001.0003>
- Filmer, D., & Scott, K. (2012). Assessing Asset Indices. *Demography*, 359-392. [doi: 10.1007/s13524-011-0077-5](https://doi.org/10.1007/s13524-011-0077-5)
- Fujii, T. (2010). Micro-level estimation of child undernutrition indicators in Cambodia. *The World Bank Economic Review*, 24(3), 520-553. <https://doi.org/10.1093/wber/lhq016>
- Garbero, A. (2014). Estimating poverty dynamics using synthetic panels for IFAD-supported projects: a case study from Vietnam. *Journal of Development Effectiveness*, 6, 490-510. <https://doi.org/10.1080/19439342.2014.981954>
- Gibson, J. (2001). Measuring chronic poverty without a panel. *Journal of Development Economics*, 65, 243-266. [https://doi.org/10.1016/S0304-3878\(01\)00136-5](https://doi.org/10.1016/S0304-3878(01)00136-5)
- . (2018). Forest loss and economic inequality in the Solomon Islands: using small-area estimation to link environmental change to welfare outcomes. *Ecological Economics*, 148, 66-76. <https://doi.org/10.1016/j.ecolecon.2018.02.012>
- Gibson, J., Huang, J., & Rozelle, S. (2003). Improving estimates of inequality and poverty from urban China's household income and expenditure survey. *Review of Income and Wealth*, 49(1), 53-68. <https://doi.org/10.1111/1475-4991.00074>
- Grosh, M. E., Del Ninno, C., Tesliuc, E., & Ouerghi, A. (2008). *For protection and promotion: The design and implementation of effective safety nets*. Washington, DC: World Bank. <https://doi.org/10.1596/978-0-8213-7581-5>
- Harttgen, K., Klasen, S., & Vollmer, S. (2013). An African growth miracle? Or: what do asset indices tell us about trends in economic performance?. *Review of Income and Wealth*, 59, S37-S61. <https://doi.org/10.1111/roiw.12016>
- Jha, S., Martinez Jr, A., Quising, P., Ardaniel, Z., & Wang, L. (2018). Natural disasters, public spending, and creative destruction: a case study of the Philippines. *ADB Working Paper 817*. Tokyo: Asian Development Bank Institute. <https://www.adb.org/sites/default/files/publication/408351/adb-wp817.pdf>
- Jolliffe, D., Lanjouw, P., Chen, S., Kraay, A., Meyer, C., Negre, M., Prydz, E., Vakis, R., & Wethli, K. (2015). *A measured approach to ending poverty and boosting shared prosperity: Concepts, data, and the twin goals*. Washington DC: The World Bank. <https://doi.org/10.1596/978-1-4648-0361-1>

- Kilic, T., Serajuddin, U., Uematsu, H., & Yoshida, N. (2017). *Costing household surveys for monitoring progress toward ending extreme poverty and boosting shared prosperity*. World Bank Policy Research Paper no. 7951. Washington, DC: World Bank.  
<https://doi.org/10.1596/1813-9450-7951>
- Kolenikov, S. & Angeles, G. (2009). Socioeconomic status measurement with discrete proxy variables: Is principal component analysis a reliable answer? *Review of Income and Wealth*, 55(1): 128-165. <https://doi.org/10.1111/j.1475-4991.2008.00309.x>
- Kozyreva, P., Kosolapov, M., & Popkin, B. M. (2016). Data resource profile: The Russia longitudinal monitoring Survey—Higher school of economics (RLMS-HSE) phase II: Monitoring the economic and health situation in Russia, 1994–2013. *International Journal of Epidemiology*, 45, 395-401. <https://doi.org/10.1093/ije/dyv357>
- Lanjouw, J. O., & Lanjouw, P. (2001). How to compare apples and oranges: Poverty measurement based on different definitions of consumption. *Review of Income and Wealth*, 47(1), 25-42. <https://doi.org/10.1111/1475-4991.00002>
- Lanjouw, P., Marra, M., & Nguyen, C. (2017). Vietnam’s evolving poverty index map: Patterns and implications for policy. *Social Indicators Research*, 133(1), 93-118.  
<https://link.springer.com/article/10.1007/s11205-016-1355-9>
- Martinez, A. Jr., Western, M., Haynes, M., & Tomaszewski, W. (2013). Measuring income mobility using pseudo-panel data. *Philippine Statistician*, 62, 71-99.  
[http://www.psai.ph/docs/publications/tps/tps\\_2013\\_62\\_2\\_5.pdf](http://www.psai.ph/docs/publications/tps/tps_2013_62_2_5.pdf)
- Mathiassen, A. (2009). A model based approach for predicting annual poverty rates without expenditure data. *The Journal of Economic Inequality*, 7(2), 117-135.  
<https://link.springer.com/article/10.1007/s10888-007-9059-7>
- . (2013). Mathiassen, A. (2013). Testing Prediction Performance of Poverty Models: Empirical Evidence from U ganda. *Review of Income and Wealth*, 59(1), 91-112.  
<https://doi.org/10.1111/roiw.12007>
- Morris, S. S., Carletto, C., Hoddinott, J., & Christiaensen, L. J. (2000). Validity of rapid estimates of household wealth and income for health surveys in rural Africa. *Journal of Epidemiology & Community Health*, 54, 381-387. <http://dx.doi.org/10.1136/jech.54.5.381>
- Moser, C., & Felton, A. (2009). The construction of an asset index. In *Poverty dynamics: interdisciplinary perspectives*, 102-127. Oxford: Oxford University Press.
- Ngo, D. K. (2018). A theory-based living standards index for measuring poverty in developing countries. *Journal of Development Economics*, 130, 190-202.  
<https://doi.org/10.1016/j.jdeveco.2017.10.011>

- Nguyen, C. V. (2011). Poverty projection using a small area estimation method: Evidence from Vietnam. *Journal of Comparative Economics*, 39, 368-382.  
<https://doi.org/10.1016/j.jce.2011.04.004>
- OECD. (2015). *OECD employment outlook 2015*. Paris: OECD Publishing.  
[http://dx.doi.org/10.1787/empl\\_outlook-2015-en](http://dx.doi.org/10.1787/empl_outlook-2015-en)
- . (2018). *A broken social elevator? How to promote social mobility*. Paris: OECD Publishing.  
<https://doi.org/10.1787/9789264301085-en>.
- Pratesi, M. (Eds.) (2016). *Analysis of poverty data by small area estimation*. UK: John Wiley & Sons.
- Rao, J. N. K., & Molina, I. (2015). *Small area estimation*, 2nd edition, New York: Wiley.
- Ravallion, M. (2016). *The economics of poverty: History, measurement, and policy*. New York: Oxford University Press.
- Rubin, D. B. (1987). *Multiple imputation for nonresponse in surveys*. New York: Wiley.
- Sahn, D. E., & Stifel, D. C. (2000). Poverty comparisons over time and across countries in Africa. *World Development*, 28, 2123-2155. [https://doi.org/10.1016/S0305-750X\(00\)00075-9](https://doi.org/10.1016/S0305-750X(00)00075-9)
- Serajuddin, U., Uematsu, H., Wieser, C., Yoshida, N., & Dabalen, A. (2015). *Data deprivation: another deprivation to end. World Bank Policy Research Paper no. 7252*. Washington, DC: World Bank. <https://doi.org/10.1596/1813-9450-7252>
- Sohnesen, T. P., Ambel, A. A., Fisker, P., Andrews, C., & Khan, Q. (2017). Small area estimation of child undernutrition in Ethiopian woredas. *PloS one*, 12, e0175445.  
<https://doi.org/10.1371/journal.pone.0175445>
- Steele, J. E., Sundsøy, P. R., Pezzulo, C., Alegana, V. A., Bird, T. J., Blumenstock, J., Bjelland, J., Engø-Monsen, K., de Montjoye, Y., Iqbal, A. M., Hadiuzzaman, K. N., Lu, X., Wetter, E., Tatem, A. J., & Bengtsson, L. (2017). Mapping poverty using mobile phone and satellite data. *Journal of The Royal Society Interface*, 14, 4: 20160690.  
<https://doi.org/10.1098/rsif.2016.0690>.
- Stifel, D., & Christiaensen, L. (2007). Tracking poverty over time in the absence of comparable consumption data. *The World Bank Economic Review*, 21, 317-341.  
<https://doi.org/10.1093/wber/lhm010>
- Tarozzi, A. (2007). Calculating comparable statistics from incomparable surveys, with an application to poverty in India. *Journal of Business & Economic Statistics*, 25, 314-336.  
<https://doi.org/10.1198/073500106000000233>

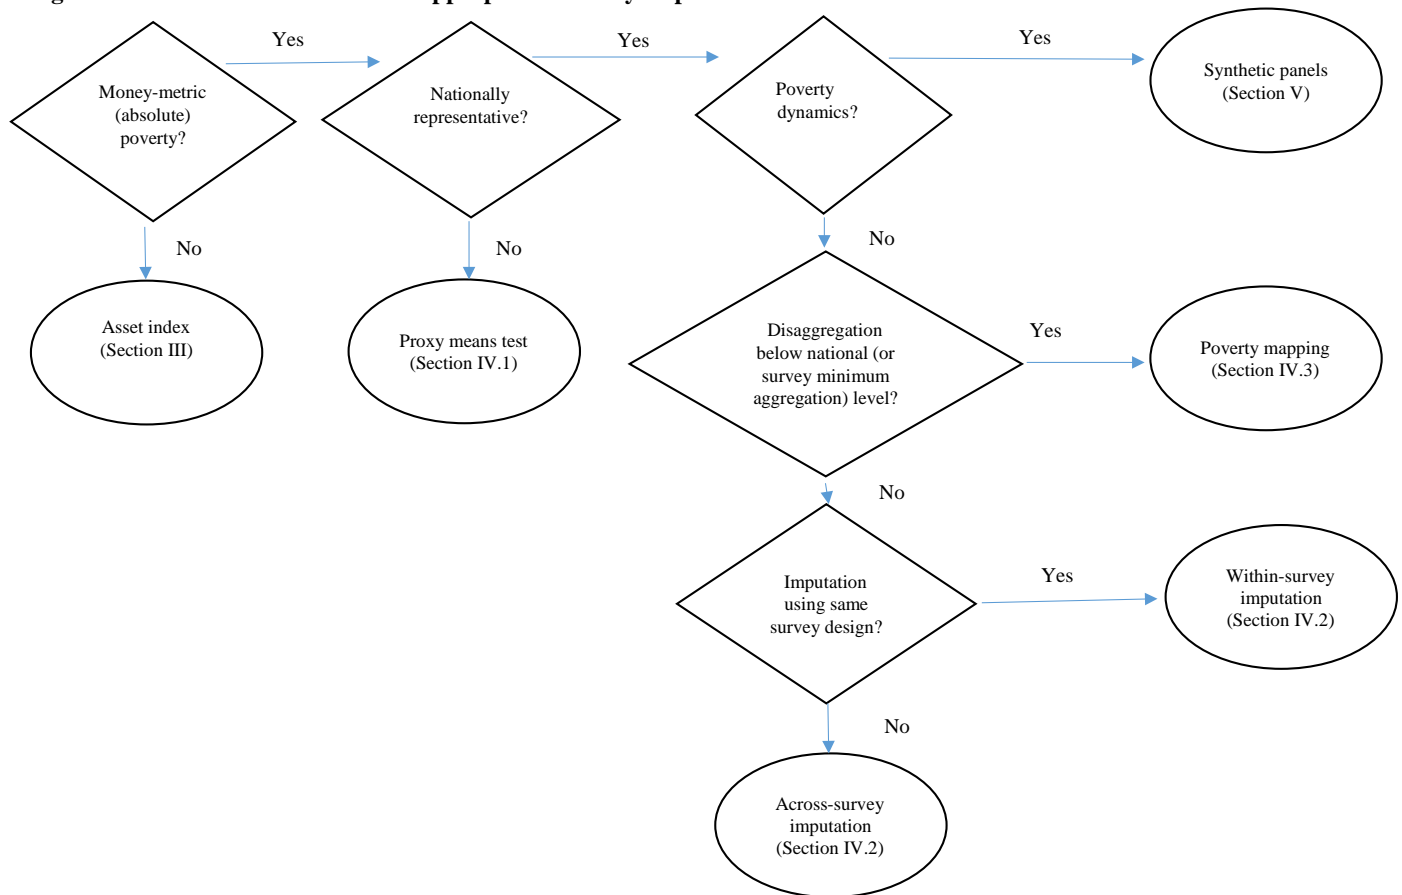
- Tarozzi, A., & Deaton, A. (2009). Using census and survey data to estimate poverty and inequality for small areas. *The review of economics and statistics*, 91, 773-792.  
<https://doi.org/10.1162/rest.91.4.773>
- United Nations Development Programme (UNDP). (2016). *Multidimensional progress: Well-being beyond income*. New York: United Nations Development Programme.  
[https://www.undp.org/content/dam/rblac/docs/Research%20and%20Publications/IDH/UNDP\\_RBLAC\\_HDR\\_ENGFinal2016.pdf](https://www.undp.org/content/dam/rblac/docs/Research%20and%20Publications/IDH/UNDP_RBLAC_HDR_ENGFinal2016.pdf)
- Vakis, R., Rigolini, J., & Lucchetti, L. (2016). *Left behind: chronic poverty in Latin America and the Caribbean*. Washington, DC: World Bank. <https://doi.org/10.1596/978-1-4648-0660-5>
- World Bank. (2017). *Monitoring global poverty : report of the commission on global poverty*. Washington, DC: The World Bank. <https://doi.org/10.1596/978-1-4648-0961-3>
- Yoshida, N., Munoz, R., Skinner, A., Lee, C.K.E., Brataj, M., Durbin, W., Sharma, D., & Wieser, C. (2015). *Swift data collection guidelines version 2*. Washington, DC: World Bank.  
<http://documents.worldbank.org/curated/en/591711545170814297/pdf/97499-WP-P149557-OUO-9-Box391480B-ACS.pdf>

**Table 1: Categories of Missing Household Consumption Data and Recent Sample Studies**

Type	Extent of Missing Consumption Data	Typical Situation	Example	Recent Sample Studies
A	Completely missing (e.g., wealth index)	i) Non-consumption surveys	Demographic and Health Surveys and most small-scale (or sub-national) surveys	Sahn & Stifel (2000); Filmer & Pritchett (2001); Filmer & Scott (2012)
		ii) Proxy means test/ project targeting	Most small-scale surveys	Grosh <i>et al.</i> (2008); Coady <i>et al.</i> (2014); Brown <i>et al.</i> (2018)
B	Partially missing (e.g., imputed consumption)	i) Consumption data not comparable across survey rounds	Some rounds of India's National Sample Surveys	Tarozzi (2007); Christiaensen <i>et al.</i> (2012); Mathiassen (2013)
		ii) Consumption data unavailable in current survey but available in another related survey	The annual LFS does not have consumption data, but the household consumption survey is implemented every few years	Mathiassen (2009); Doudich <i>et al.</i> (2016); Dang <i>et al.</i> (2017a)
		iii) Consumption data unavailable at more disaggregated administrative levels than those in current survey	Population census data are representative at lower administrative level than a household consumption survey, but do not collect consumption data.	Elbers <i>et al.</i> (2003); Elbers <i>et al.</i> (2007); Tarozzi & Deaton (2007)
C	Available cross sections, but missing panel data (e.g., synthetic panels)	Most surveys in developing countries do not offer panel data		Dang <i>et al.</i> (2014); Dang & Lanjouw (2013); Bourguignon <i>et al.</i> (2018)

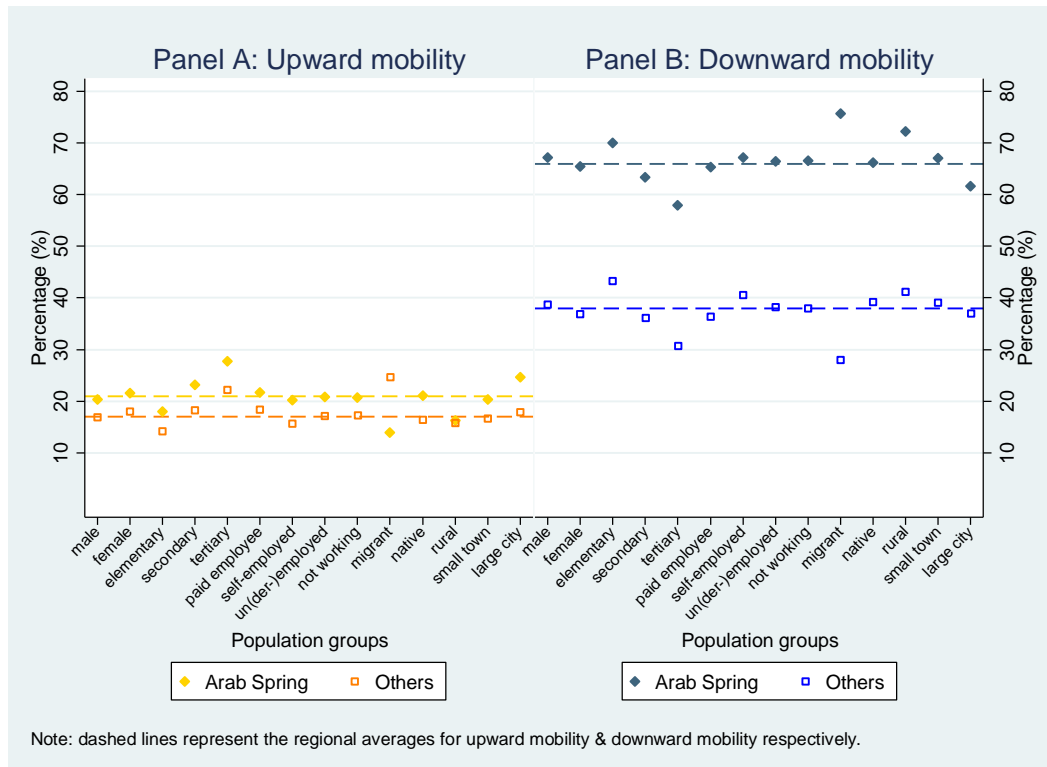
**Note:** LFS stands for Labor Force Surveys. This table is a modified and expanded version of Table 1 in Dang *et al.* (2019).

**Figure 1: Decision Process to Select Appropriate Poverty Imputation Methods**



**Note:** Rhombus represents the desired poverty estimates, and circle represent the suggested poverty imputation method.

**Figure 2: Subjective Wellbeing Dynamics in Arab Spring Countries and Other Arab Countries Based on Synthetic Panels, 2007-2012**



Source: Dang & Ianchovichina (2018).



## Appendix A: Analytical Framework

### A.1. General Setup

Household consumption is typically estimated using the following reduced-form linear model

$$y_{ij} = \beta' x_{ij} + \mu_{ij} \quad (1.1)$$

for household  $i$  in survey  $j$ , for  $i = 1, \dots, N$  (see, e.g., Elbers *et al.* (2003), Ravallion (2016)).  $x_{ij}$  can include household variables such as the household head's age, sex, education, ethnicity, religion, language (i.e., which can represent household tastes), occupation, and household assets or incomes. We will examine various versions of Equation (1.1) in the following sections.

### A.2. Non-Monetary Poverty Estimates

Consider a variant of Equation (1.1) where the left-hand side variable, household consumption  $y_{ij}$ , is now missing, but we have data on household assets  $a_{ij}$ —a subset of  $x_{ij}$ . We can generate a wealth index  $w_{ij}$  which offers the best combination of (the elements of the different) household assets  $a_{ij}$  as follows

$$\gamma' a_{ij} = w_{ij} \quad (1.2)$$

Note that household assets are just a component of the characteristics  $x_{ij}$  in Equation (1.1).

There are two common approaches to obtaining these weights: the first is to simply count the number of assets a household possesses, and the second is to obtain a (linear) combination of weights that captures the most variation of the  $a_{ij}$  by statistical techniques such as the principal component analysis (PCA). The second approach is popularized by Filmer & Pritchett (2001).

### A.3. Project Targeting

Most of the estimates based on proxy means testing are usually estimated as

$$y_{ij}^p = \beta_j^{p'} x_{ij,p} \quad (1.3)$$

where the vector of coefficients  $\beta_j^p$  is often obtained from the regression using another data set (see, e.g., Grosh *et al.* (2008), Coady *et al.* (2014), Brown *et al.* (2018)). For example,  $\beta_j^p$  can be obtained from a regression using the household consumption data, and then be imposed on the data from another—often smaller—survey that aims at targeting poor households for a social protection program.

### A.4. Tracking Poverty Trends at the National Level

We extend Equation (1.1) to a more general model

$$y_j = \beta_j' x_j + v_{cj} + \varepsilon_j \quad (1.4)$$

where the error term  $\mu_{ij}$  is now broken down into two components, one ( $v_{cj}$ ) a cluster random effects and the other ( $\varepsilon_j$ ) the idiosyncratic error term. We suppress the subscript that indexes households to make the notation less cluttered in this sub-section. Conditional on household characteristics, the cluster random effects and the error terms are usually assumed uncorrelated with each other and to follow a normal distribution such that  $v_{cj} | x_j \sim N(0, \sigma_{v_j}^2)$  and  $\varepsilon_j | x_j \sim N(0, \sigma_{\varepsilon_j}^2)$ . While the normal distribution assumption results in the standard linear random effects model that is more convenient for mathematical manipulations and computation, it is not necessary for this type of model. We can remove this assumption and use the empirical distribution of the error terms.

Assume that the explanatory variables  $x_j$  are comparable for both surveys (Assumption 1), and that the changes in  $x_j$  between the two periods can capture the change in poverty rate in the next period (Assumption 2). Suppressing the subscript for households in the following equations, Dang *et al.* (2017a) define the imputed consumption  $y_2^1$  as

$$y_2^1 = \beta_1' x_2 + v_1 + \varepsilon_1 \tag{1.5}$$

and estimate it as

$$\hat{y}_{2,s}^1 = \hat{\beta}_1' x_2 + \tilde{v}_{1,s} + \tilde{\varepsilon}_{1,s} \tag{1.6}$$

where the parameters  $\beta_1'$  are estimated using Equation (1.4), and  $\tilde{v}_{1,s}$  and  $\tilde{\varepsilon}_{1,s}$  represent the  $s^{th}$  random draw (simulation) from their estimated distributions, for  $s= 1, \dots, S$ . We suggest using 1,000 simulations or more. The poverty rate in period 2 and its variance can then be estimated as

$$\text{i) } \hat{P}_2 = \frac{1}{S} \sum_{s=1}^S P(\hat{y}_{2,s}^1 \leq z_1) \tag{1.7}$$

$$\text{ii) } V(\hat{P}_2) = \frac{1}{S} \sum_{s=1}^S V(\hat{P}_{2,s} | x_2) + V\left(\frac{1}{S} \sum_{s=1}^S \hat{P}_{2,s} | x_2\right) \tag{1.8}$$

Dang *et al.* (2017a) offer a test for Assumption 2 if consumption data exist in both survey rounds, which break down the change in poverty between the survey rounds into two components, one due to the changes in the estimated coefficients (the first term in square brackets in Equation (9) below) and the other the changes in the  $x$  characteristics (the second term in square brackets in Equation (9) below). Assumption 2 would be satisfied if the poverty change is mostly explained by the latter component. This can be expressed as

$$\begin{aligned} P(y_2) - P(y_1) &= [P(y_2) - P(y_2^1)] + [P(y_2^1) - P(y_1)] \\ &= [P(\beta_2' x_2 + \eta_2) - P(\beta_1' x_2 + \eta_1)] + [P(\beta_1' x_2 + \eta_1) - P(\beta_1' x_1 + \eta_1)] \end{aligned} \tag{1.9}$$

where  $\eta_j$  is defined as  $v_j + \varepsilon_j$ ,  $j= 1, 2$ , for less cluttered notation.

However, a practical use of Assumption 2 is model selection; in particular, we can test for this assumption in contexts where consumption data exist for earlier survey rounds, and select the imputation model that offers estimation results that are most satisfactory for this assumption.

### A.5. Disaggregated Poverty Estimates at the Sub-National Level

The most commonly used framework developed by Elbers, Lanjouw, and Lanjouw (2003) proposes that we impose the estimated parameters from Equation (1.4) on the same variables in a population census (instead of another household consumption survey as with Equation (1.5)). But note that Elbers *et al.* (2003) suggest a Bayesian estimation method where the distribution of  $\hat{\beta}_1'$ , rather than the point estimates of  $\beta_1'$ , is employed for the imputation.

### A.6. Missing Panel Consumption Data

Let  $x_{ij}$  be a vector of time-invariant household characteristics observed in survey round  $j$  ( $j= 1$  or  $2$ ) that are also observed in the other survey round for household  $i$ ,  $i= 1, \dots, N$ . These household characteristics can include such variables as ancestry, religion, language, place of birth, parental education, and other time-varying household characteristics if retrospective questions about the round-1 values of such characteristics are asked in the second-round survey. To reduce spurious changes due to changes in household composition over time, we usually restrict the estimation samples to household heads age, say 25 to 55 in the first cross section and adjust this age range accordingly in the second cross section. This age range is usually used in traditional pseudo-panel analysis but can vary depending on the cultural and economic factors in each specific setting.

Let  $y_{ij}$  represent household consumption or income in survey round  $j$ ,  $j= 1$  or  $2$ . The linear projection of household consumption (or income) on household characteristics for each survey round is given by

$$y_{ij} = \beta_j'x_{ij} + \mu_{ij} \tag{1.10}$$

Let  $z_j$  be the poverty line in period  $j$ . We are interested in knowing the unconditional measures of poverty mobility such as

$$P(y_{i1} < z_1 \text{ and } y_{i2} > z_2) \tag{1.11}$$

which represents the percentage of households that are poor in the first survey round (year) but nonpoor in the second survey round, or the conditional measures such as

$$P(y_{i2} > z_2 | y_{i1} < z_1) \tag{1.12}$$

which represents the percentage of poor households in the first round that escape poverty in the second round.

If panel data are available, we can estimate the quantities in (1.11) and (1.12); but in the absence of such data, we can use synthetic panels to study mobility. To operationalize the framework, we make two standard assumptions. First, we assume that the underlying populations being sampled in survey rounds 1 and 2 are identical such that their time-invariant characteristics remain the same over time (Assumption 3). More specifically, coupled with Equation (1.10), this implies the conditional distribution of expenditure in a given period is identical whether it is conditional on the given household characteristics in period 1 or period 2 (i.e.,  $x_{i1} \equiv x_{i2}$  implies  $y_{i1}|x_{i1}$  and  $y_{i1}|x_{i2}$  have identical distributions). Second, we assume that  $\mu_{i1}$  and  $\mu_{i2}$  have a bivariate normal distribution with positive correlation coefficient  $\rho$  and standard deviations  $\sigma_{\mu_1}$  and  $\sigma_{\mu_2}$  respectively (Assumption 4). (Note that we refer to these assumptions as Assumptions 3 and 4 for presentation purposes only, they are not related to Assumptions 1 and 2 discussed earlier). Quantity (1.11) can be estimated by

$$P(y_{i1} < z_1 \text{ and } y_{i2} > z_2) = \Phi_2\left(\frac{z_1 - \beta_1'x_{i2}}{\sigma_{\mu_1}}, -\frac{z_2 - \beta_2'x_{i2}}{\sigma_{\mu_2}}, -\rho\right) \tag{1.13}$$

where  $\Phi_2(\cdot)$  stands for the bivariate normal cumulative distribution function (cdf) (and  $\phi_2(\cdot)$  stands for the bivariate normal probability density function (pdf)). Note that in Equation (1.13), the estimated parameters obtained from data in both survey rounds are applied to data from the second survey round ( $x_2$ ) (or the base year) for prediction, but we can use data from the first survey round as the base year as well. It is then straightforward to estimate quantity (1.12) by dividing quantity (1.11) by  $\Phi\left(\frac{z_1 - \beta_1'x_{i2}}{\sigma_{\mu_1}}\right)$ , where  $\Phi(\cdot)$  stands for the univariate normal cumulative distribution function (cdf).

In Equation (1.13), the parameters  $\beta_j$  and  $\sigma_{\mu_j}$  are estimated from Equation (1.1), and  $\rho$  can be estimated using an approximation of the correlation of the cohort-aggregated household consumption between the two surveys, which yields an estimate for  $\rho_{y_{i1}y_{i2}}$ . The partial correlation coefficient  $\rho$  can then be estimated by

$$\rho = \frac{\rho_{y_{i1}y_{i2}}\sqrt{\text{var}(y_{i1})\text{var}(y_{i2}) - \beta_1'\text{var}(x_i)\beta_2}}{\sigma_{\mu_1}\sigma_{\mu_2}} \tag{1.14}$$

Note that the standard errors of estimates based on the synthetic panels can in fact be even smaller than that of the true (or design-based) rate if there is a good model fit (or the sample size in the target survey is significantly larger than that in the base survey; see Dang & Lanjouw (2013) for more discussion).