

World Bank Center for Development Data (C4D2)

Partnership for Capacity Development in Household Surveys for Welfare Analysis

Measuring Consumption

Instructor Manual

This draft: April 2020

This draft was prepared by Giulia Mancini and Giovanni Vecchi (University of Rome Tor Vergata). We would like to acknowledge the contribution provided by Nicola Amendola and Sédi-Anne Boukaka (University of Rome Tor Vergata) in preparing the material underlying the course. We are grateful to Gero Carletto, Michelle Jouvenal, Shelton Kanyanda and Alberto Zezza (World Bank) for helpful suggestions.

Introduction: how to use this Manual

This Manual is addressed to the Instructor of the course “Measuring Consumption”, prepared by the World Bank’s Center for Development Data (C4D2). It is part of a **Teaching Package** that also contains the following materials:

1. Course syllabus
2. Lecture slides (classroom version, PowerPoint, and handout version, pdf)
3. Reading package
4. Final exam and grading sheet
5. Course evaluation form

The Manual explains how best to make use of each of these items, and provides practical guidance to support the Instructor in the preparation and delivery of the course.

The rest of the Manual is organized as follows. First, the section titled **The Course** gives a general outline of the course, describes its overall learning objectives, explains how they will be reached, and lays out the Syllabus. This section also contains some important logistical information (required facilities, teaching materials, etc.), that should be assessed in preparation for the course.

The remaining sections provide information for each of the **15 lectures** included in the course. For each lecture, you will find *(i)* learning objectives; *(ii)* a description of any preparatory work that the Instructor is advised to do before giving the lecture; *(iii)* a time allocation plan, which advises Instructors on how best to allocate the available classroom time; *(iv)* a writeup of the topics covered in the lecture, designed to support the Instructor’s understanding of the slides; and *(v)* keys to any exercises associated with the lecture.

The Course

Target audience and pre-requisites

The target audience is Master students enrolled in Graduate Programs in Statistics, Official Statistics, and related fields, offered by the Regional Training Centers.

This course is designed to be self-contained, therefore it has no compulsory pre-requisites. However, it is an advantage if students have a basic command of Statistics, Sampling Design, and Microeconomics.

Learning objectives of the course

By training students who will pursue careers within Statistical Offices and Institutions, this course aims at improving the **quality** and **comparability** (both over time and across countries) of household survey data for living standards measurement.

To do so, the course focuses on two main topics: first, it lays out the **conceptual framework** underlying the measurement of living standards; second, it offers practical **guidelines for survey design and data collection**, in the specific context of household consumption and expenditure modules.

The **emphasis** of the course **is on data collection, rather than data analysis**. At the end of the course, participants will have learned how to best approach the design of consumption modules in household budget surveys. The computation of poverty and inequality indicators is covered only briefly in this course, as a way of informing data producers about the uses of household consumption data.

Course syllabus

The course contains **15 lectures**. Teaching materials are optimized for each lecture to last two hours, including a short break of about 15 minutes.

A number of **readings** are associated to each lecture: all of them are provided in the form of a Reading Package, to be shared with students. Students should be encouraged to complete readings before each lecture. Then, classroom presentations provide guidance for the sometimes technical and challenging published material. A few selected readings are to be considered essential to a student's learning, and are marked **required**. All other readings are optional, but useful for those interested in knowing more.

Lectures also contain a number of **exercises**, to be assigned as homework. This Manual includes keys and evaluation guidelines for the exercises.

Plan of the lectures:

- 1) **Measuring living standards: a conceptual framework**
From theory to practice: Income vs. Expenditure/Consumption vs. Wealth
- 2) **The Consumption Aggregate**
Four building blocks: (i) food, (ii) non-durable non-food items, (iii) housing, and (iv) durable goods. Adjustments: for household size and needs, for within-year inflation.
- 3) **Understanding household surveys**
Types of surveys; defining features of household surveys that measure consumption, as opposed to other statistical instruments.
- 4) **Principles of questionnaire design**
Overview of the survey process; general principles of questionnaire design, as they apply to the measurement of consumption expenditures.
- 5) **Measuring food consumption: the foundations**
Concepts (*e.g.* acquisition vs. consumption). Questionnaire design (recall vs. diary)
- 6) **Measuring food consumption: questionnaire design**
Questionnaire design (list of food items, meal participation, seasonality)
- 7) **Food away from home and the use of non-standard units for measuring food consumption**
Questionnaire design (non-standard measurement units, food away from home)
- 8) **Measuring consumption of non-durable non-food goods**
Analytical needs; questionnaire design.
- 9) **Durable goods**
Analytical needs and questionnaire design.
- 10) **Housing**
Analytical needs; questionnaire design.
- 11) **Data validation and diagnostics**
Key principles of statistical data editing.
- 12) **Outlier detection and treatment**
Key principles of outlier detection and treatment.
- 13) **Measuring inequality**
Lorenz curve and Gini index; other selected inequality measures.
- 14) **Measuring poverty**
Poverty lines and measures.
- 15) **Describing data**
Preparation of a “tabulation report”.

Course evaluation

The Teaching Package contains a **course evaluation form** designed to collect feedback from students. The evaluation form is anonymous, and should be handed out after the last lecture. Instructors are encouraged to share results with the C4D2 team, to help improve future iterations of the course.

Final exam

The final exam is a take-home assignment. The text of the final exam is provided as part of the Teaching Package, to serve as a blueprint for the Instructor's convenience.

Logistics

The total amount of **classroom time** required for the course is 30 hours (there are 15 lectures, each lasting two hours). Training Centers are advised to schedule one or, at most, two of these lectures in one day. Because the final exam is a take-home assignment, no further classroom time is required after the end of the course.

The course calls for the use of **PowerPoint** presentations to support the Instructor's delivery: therefore, a room outfitted with a **projector** is required. The use of a **microphone** by the Instructor is strongly recommended: it greatly facilitates the engagement of participants throughout the presentations. One or more microphones should be available for participants as well, to facilitate class discussions.

Instructors can use the following **checklist** to organize the logistics of the course:

Before registration:

- Share the syllabus of the course, clarify learning objectives: make sure participants' expectations are aligned with the contents of the course
- Confirm availability of required facilities (projector, microphones)

Before the first lecture:

- The Instructor gets familiar with the readings associated to each lecture
- Edit Lecture 1's "Contacts" slide, adding the Instructor's email address, office hours, etc.
- Print out slides (handout version) for each student
- Prepare to share digital copy of reading package with students

Before the final lecture:

- Print out evaluation forms for each student
- Print out (or prepare to share digital copy of) final exam

Please note: printing out the **handout version** of the course slides, so that each student can use the printed slides for notes and exercises, is strongly encouraged.



Lecture 1

Measuring living standards: a conceptual framework

Learning objectives

The goal of this lecture is to provide the theoretical basis for the choice of a measure of living standards, which analysts need in order to estimate both poverty and inequality. The material covered in the lecture justifies the focus on a consumption-based measure of living standards.

Suggested preparation

The Instructor is assumed to be familiar with the material contained in section 2.2 (“Money-metric utility”) of Deaton and Zaidi (2002), as well as with references therein listed.

Time allocation

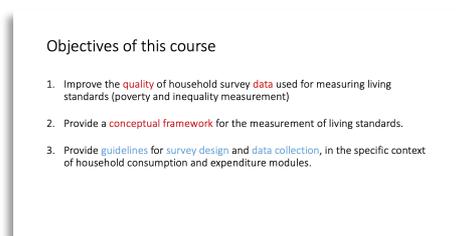
Students’ introductions (name, background)	15 min
Introduction	15 min
What is the standard of living?	30 min
Break	15 min
Choosing a measure of living standards	45 min

Annotated lecture

As the first lecture begins, it is a good idea to break the ice by getting to know each student, and setting their **expectations** for the course. The Instructor can do a round of introductions and ask each student what they expect to learn. The students should be informed in advance of the objectives and contents of the course; the description of the course and the syllabus should be circulated well in advance of registration: please check the Logistics section of this manual.

If that has been done, students' expectations should not differ too much from what the Instructor is about to deliver. The course is *not* about data analysis: it is about data collection. The course is *not* about hands-on practical sessions: it is about concepts. This moment gives the Instructor the opportunity to clarify this once more.

The first slide helps to do just that, giving students a bird's-eye view of the whole course. This task does not require much time, but clarifying what the course is about, and why it is important – and hopefully interesting, too – is key. The aim is to convey **motivation** (students should feel they are about to learn important concepts and tools). This is why the first bunch of slides gives an overview of the topics that will be covered in the following classes, and, even more importantly, explains why students should take the time to study them.

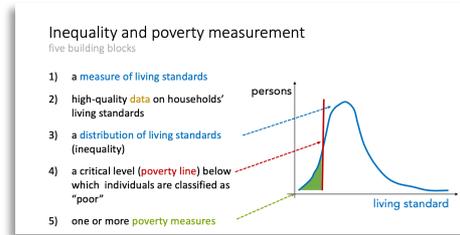


The broad, long-term objective of the course, that justifies an interest in the topics covered, is that of improving the quality of household survey data used for measuring living standards (which in turn are the basis for poverty and inequality measurement).

What do we mean by **data quality**, exactly? Before going into details, we can take this opportunity to make a general point: learning to use words and terms as precisely and technically as possible is part of the aims of this course. The attitude that we encourage in students is to ask questions such as “what do we mean by ‘something’, exactly?” where emphasis is on ‘**exactly**’. Answers are in the **literature**, which students will be encouraged to consult throughout the course, instead of relying entirely on ‘common sense’. Here, we can consider engaging with students, and asking *them* “how would you define and assess ‘data quality’, exactly?”. We allocate only a couple of minutes to discussing answers, then move on.

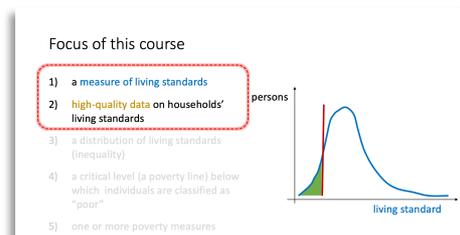
Back to our question ‘What do we mean by **data quality**, exactly?’. In their manual on survey quality, Biemer and Lyberg (2003: 13-18) review different ways in which data quality has been defined by various statistical institutions: some common criteria emerge. One of the most used criteria is that, in order to be of good quality, data should be **relevant**

for the needs of the users. Therefore, it is crucial for data producers to understand and clearly define the research objective that justifies their data collection efforts (what will the data be used for?). This is why we have the next two bullets: the course will provide a conceptual framework, to clarify the goals and methods of welfare analysis, and in light of those, it will offer specific guidelines for survey design and data collection.



Because the course is focused on the collection of data for the purpose of inequality and poverty measurement, this slide explains the 5 building blocks that are needed to construct estimates of inequality and poverty:

- 1) a **measure of living standards**, that is, a quantifiable concept of what living standards are and how they can be defined in practice;
- 2) high-quality **data** on households' living standards, typically collected through surveys;
- 3) a distribution of living standards, that is, the actual data that come from the field, and that describe the variation of living standards across the population of interest. At this stage, it is possible to produce **inequality** estimates;
- 4) a critical level, a **poverty line**, below which individuals are classified as "poor", that is, a value, a fixed amount of whatever quantity indicates individual living standards, that gives us a threshold below which individuals are classified as poor;
- 5) one or more **poverty measures**, such as a simple headcount of how many poor individuals there are, or any other more complex measure of our choice, to summarize the results of the analysis.



The focus of this course will be on the first 2 building blocks: a measure of living standards (again, this is the conceptual foundation of the analysis, which must be embedded in the design of the survey); and high-quality data that describe households' living standards. This is a good time to comment on the importance of **'theory'**: the course will combine 'theory' (first building block) with 'practice' (second building block). There is no such a thing as good data collection in the absence of a theoretical framework. Building blocks 3) to 5)

venture into analytical territory, and are complex enough to be the subject of a different course.

Course overview	
fifteen two-hour lectures	
1. Measuring living standards: a conceptual framework	8. Measuring consumption of non-durable non-food items
2. The consumption aggregate	9. Durable goods
3. Understanding household surveys	10. Housing
4. Principles of questionnaire design	11. Data validation and diagnostics
5. Measuring Food Consumption – I	12. Outlier detection and treatment
6. Measuring Food Consumption – II	13. Measuring inequality
7. Measuring Food Consumption – III	14. Measuring poverty
	15. Describing data

This slide gives an overview of the topics to be covered in each of the 15 lectures of the course. We should *not* read and comment item by item – this would take too long, and probably bore students to death. We simply remind students that there is a Syllabus, and encourage them to check it and use it as a roadmap throughout the duration of the course.

<p>Practical instructions</p> <ol style="list-style-type: none">Breaks Expect a 15-minute break for each lectureReadings Some compulsory, some optional (reading package available)Homework No stars is basic, one star (*) is difficult, two stars (**) is very difficultFinal exam Take-home assignment	<p>Contacts</p> <p>Please add the instructor's email address, office hours, etc.</p>
--	--

Next, some practical instructions that are of interest for students.

Do not forget: the **Instructor's contacts** should be personalized.

Measuring living standards: a conceptual framework
LECTURE 1

After the introduction, we enter into the actual contents of lecture 1. This lecture is about **concepts**: it provides the theoretical foundation for the need of data on household consumption. It uses a bit of economic theory and technical jargon in order to give students some common ground with data analysts. The importance of a theoretical foundation for the way we measure living standards is one of the main takeaways of the lecture: to make the point effectively, the Instructor should be familiar with the literature cited throughout.

A foundational question

- What is the “standard of living”?
- It is a profound question, that (apparently) defies simplification
- Our aim is to provide a **quantifiable** answer

The first part of the lecture answers the question: “What is the standard of living?”. Here, the Instructor can involve students, and encourage them to give their own answers to the question. The point that should emerge from this short discussion is that the answer is highly complex, highly subjective, and in no way easy to agree on. We emphasize that is good to keep this **conceptual complexity** in mind, although statisticians, economists, data analysts are interested in providing a quantifiable answer – one that allows for an estimate of how many poor people live in a country, for instance.

Amartya Sen
(1933 -)



- 1998 Nobel Prize in Economics
- Why?
- “(...) for his contributions to welfare economics”

Amartya Sen
Commodities and Capabilities, 1987, p. 1.



«There are many different approaches (...) to judging whether the person is doing well (...): is he well off? Is she **happy**? Does he feel **fulfilled**? Does she have much **freedom**? Can he get what he **wants**? Can she do what she would like to do? Is **society** being good to him? Is she having a good life? These distinct questions have their own peculiar relevance in particular contexts and each has an importance of its own.»

The next few slides introduce the views of a notable economist and philosopher, Amartya Sen. Sen’s position is very critical of the simplification that mainstream economics adopts in order to arrive at a measurement of living standards – which will be covered in the rest of the lecture. However, the quote from Sen’s book ‘Commodities and Capabilities’ is a good summary of the complexity of the concept of well-being, and conveys the need of delimiting the issue, if one wants to make this concept measurable.

How to narrow Sen’s list down?

- Happiness
- Fulfillment
- Money
- Health
- Freedom
- ...

• Q. How did mainstream economics eschew this complexity?
• A. They introduced one more concept: ‘**utility**’.

This slide brings the audience back to the fundamental question of defining the standard of living, after having established that such a definition is a complex matter, not just intuitively (see the short discussion that the class just had about well-being), but also in the literature (see Sen’s thought). However, mainstream economics provides a way to **reduce** this **complexity**. It does so by introducing the concept of “utility”.

The next slides explain the basics of **utility theory**. If more details are needed, students can be encouraged to consult a manual of microeconomics (chapters on consumer choice). One of the most widely used is Varian (2010), while Deaton and Muellbauer (1980) or Varian (1992) are at a more advanced level.

Mainstream economists

- Economists assume that the standard of living derives from the consumption of goods and services.
- Any given basket of goods and services gives a certain 'utility' to the consumer.
- They assume that utility depends on q .
- The simplest case is with one good: this is how the utility $u(q)$ varies with q .
- 'Utility' is clearly not observable, but provides the foundation for much of the conceptual framework that poverty analysts use.

The answer that mainstream economists give to the complexities of defining well-being is based on the concept of utility, introduced in this slide. “Utility” is an abstract term, a catch-all, standing for whatever a person pursues. It is a useful tool for standard economic theory, which is concerned with individual choice.

Angus Deaton
(born in Edinburgh, Scotland, 1945)



- 2015 Nobel Prize in Economics
- Why?
- “(…) for his analysis of consumption, poverty, and welfare”

Deaton and Zaidi (2002)



- The *Guidelines* have been downloaded 2,716 times in the last 5 years alone; Doemeland and Trevino (2014) find that only 2% of World Bank “knowledge products” surpass 1,000 downloads over a 5-year period.
- Must read
- A review of the *Guidelines* in light of recent literature is forthcoming: Mancini and Vecchi (2020)

These slides introduce the work of another notable economist, Angus Deaton, who will be cited extensively in this lecture: he will help us connect utility theory with the measurement of living standards. Students should become familiar with Deaton’s work, especially with Deaton and Zaidi (2002), a paper entitled *Guidelines for Constructing Consumption Aggregates for Welfare Analysis*. In this lecture, we will cover the theoretical section of the *Guidelines* (in turn based on the textbook by Deaton and Muellbauer).

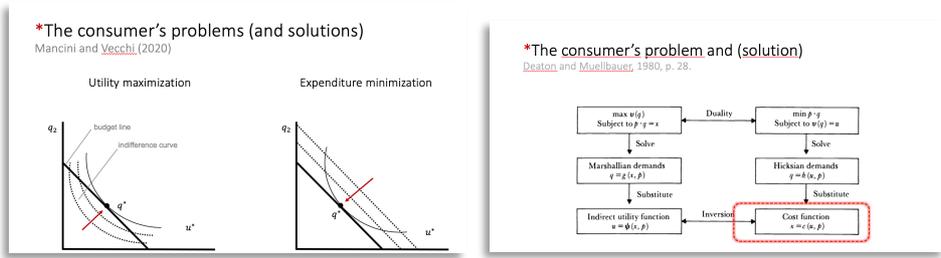
The link between ‘utility’ and the standard of living

- Consumers are assumed to maximize utility
- To do so, the consumer chooses an optimal bundle of goods and services: ‘optimal’ depends on a) her tastes (preferences), b) the prices she faces on the market, and c) her budget.
- Let us denote the optimal bundle with q^* .
- Maximum utility is then $u(q^*)$.

Individuals (consumers) are assumed to derive utility from the consumption of goods and services, which means that utility depends on q (a vector of goods and services). In this theoretical setting, consumers are assumed to make their choices based on utility. How? They choose the particular combination of goods and services (also called a “bundle”) that gives them the highest possible utility, given the circumstances. These circumstances include the consumers’ individual tastes (also called preferences), the prices that they face on the market, and their budget (the resources they are able to spend on consumption). In other words, consumers are assumed to maximize utility. Consider a single consumer, and let us introduce a bit of notation: the optimal bundle (the one that the consumer ends up picking) is indicated by q^* , and therefore the consumer’s maximum utility is $u(q^*)$.

- 'Utility' in concrete terms
- Maximum utility is $u(q^*)$.
 - Q. How much is $u(q^*)$?
 - A. We cannot tell – utility is not observable.
 - Q. Alternatively, we can ask: how much does q^* cost?
 - A. $x = p \times q^*$
this is the cost of the optimal bundle, the one that gives the maximum utility to the consumer
 - Welfare analysts follow Deaton and Zaidi (2002), a paper that shows how to calculate the value of utility $u(q^*)$, given the cost of the bundle $x = p \times q^*$

We would like to know how much is $u(q^*)$, because that would be a measure of the consumer's living standard. However, we cannot observe utility: it is just an abstract concept. What we can observe is the cost of q^* , which can be indicated by $x = q^* \times p$ (where p are the prices paid by the consumer– remember that both q^* and p are vectors, so in our notation, $q^* \times p = \sum_k p_k q_k^*$, where k indicates each of the goods consumed). Welfare analysts follow Deaton and Zaidi (2002), a paper that shows how to calculate the cost of utility given the cost of the bundle.



Optional slides are denoted by a 'star' (*). Depending on the circumstance, we suggest that, especially the first times the course is taught, these slides are omitted from the presentation, unless we feel confident that it will not take to long to convey the message, and the audience is well equipped to absorb its content. These optional slides go into more detail on one key point: the fact that the cost of the optimal bundle indicates the utility achieved by the consumer when she chooses a given bundle. The justification for this is the equivalence between the utility maximization and expenditure minimization problems, known as *duality*. A detailed discussion is in Deaton and Muellbauer (1980), p. 28; also refer to Mancini and Vecchi (2020) for a non-technical version. In what follows, we assume the Instructor has *not* shown to the class and commented on the primal and dual consumer problems in the starred slide above.

- A utility-consistent definition of standard of living
Deaton and Zaidi 2002, p. 9, eq. (2.6)
- D&Z show that the value of the utility associated to the optimal bundle can be calculated as household expenditure (x) adjusted for purchasing power. Either:
 - $u = x/P$ (P is a Paasche price index)
 - or
 - $u = x/L$ (L is a Laspeyres price index)
 - Economists refer to x/P as to money metric utility (MMU) function.
 - The ratio x/L is called welfare ratio (WR).
 - D&Z argue that for poverty measurement the best choice is x/P (eq. 2.6).

Deaton and Zaidi (2002) show that the utility associated to the optimal bundle can be approximated by (calculated as) the cost of the optimal bundle, divided by a price index, which transforms the expenditure into real terms, correcting for differences in purchasing power that arise across different consumers. Price indices can be computed according to

the Paasche or Laspeyres formula: these details are not important for now, and they will be covered in the following lecture. Suffice it to say that x/P is called money-metric utility (MMU), and it is the quantifiable concept of living standard that modern economic theory offers. In practice, this concept corresponds to total consumption expenditure adjusted with a Paasche price index (the index will be explained in detail in lecture 2 – so the suggestion is not to waste time going into details). This is what equation (2.6) in Deaton and Zaidi (2002), possibly the single most important result underlying the way we measure consumption in poverty and inequality analysis, establishes. We do not expect students to get down to the details and the math required to derive this equation, but it is important to explain that this recommendation is not one of many available, but the only one consistent with modern economic theory. Other options, no matter how attractive, do not have this advantage.

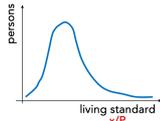
Deaton and Zaidi's recommendation # 1

Box 1. Summary of Theoretical Issues and Recommendations	
Issue	Recommendation
<p>Money Metric Utility (MMU) vs. Welfare Ratio(WR)</p> <p>MMU is the amount required to sustain a level of living and requires that consumption be adjusted by a Paasche price index that reflects the prices the household faces and whose weights are different for each household.</p> <p>WR is an indication of how much better or worse off a household is than a reference household (usually at the poverty line) and requires consumption to be adjusted by a Laspeyres price index that reflects the prices faced by the reference household but whose weights are the same for all households.</p> <p>The use of MMU can cause difficulties in analyzing the impact of redistributive policy but, on the other hand, WR does not necessarily represent welfare correctly. The latter is the more serious drawback in practice.</p>	<p>Attempt should be made to use Money Metric Utility and to calculate the Paasche price indices with individual household weights.</p>

This slide is meant as a final reminder for students that the topics discussed so far, although technical, are essential. The use of money-metric utility (MMU) as a welfare indicator is no less than **number 1** in the list of recommendations made by Deaton and Zaidi in their *Guidelines* for the welfare analyst. MMU is the pillar on which everything rests! *If this concept is unclear, the importance of measuring consumption will be lost on trainees.*

The living standard in practice

- Unlike the utility $u(q)$, the MMU $u = x/P$ is observable and can be calculated based on household budget and price data.
- This is what underlies a key decision: "deriving total consumption expenditure and dividing it by a price index is our basic strategy to measure welfare" (D&Z, p. 10)
- Standards of living can be proxied by x/P , that is using total consumption expenditure adjusted with a Paasche price index.



In practice, Deaton and Zaidi's recommendation – to use total consumption expenditure divided by a Paasche price index to proxy living standards – answers the question posed at the beginning of the lecture: **what is the standard of living?** x/P is no longer abstract, or unobservable, like the concept of utility. Instead, it is a concrete and achievable measure of well-being, consistent with economic theory.

*More on the theoretical foundations of standard of living measurement

Welfare ratios: definition, pros and cons
Blackaby and Donaldson (1987)

WELFARE RATIOS AND INTERPERSONALLY SENSITIVE CONSUMPTION STANDARDS

Blackaby, P. and Donaldson, I. (1987) Welfare Ratios and Interpersonally Sensitive Consumption Standards. *Journal of Economic Surveys*, 1(1), 1-14.

Why is MMU preferable to welfare ratios?
(Deaton 2003)

Household Survey, Consumption, and the Measurement of Poverty

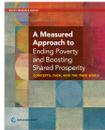
Deaton, A. (2003) Why is MMU preferable to welfare ratios? *Journal of Economic Surveys*, 17(1), 1-14.

This optional slide suggests further readings on the foundations of living standards measurement to interested students.

Expenditure, consumption, and consumption expenditure
 Browning, Crossley and Winter (2014, 477)

- We define household **expenditure** as the nominal money outlay of the household.
- Household **consumption** is the **quantity** of goods and services that the household enjoys in a given period.
- Household **consumption expenditure** is those money outlays directed toward consumption (e.g., it excludes outlays for investment purposes).

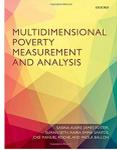
Because we are now dealing more and more with the *practice* of measuring living standards, some clarifications are in order. This slide makes some important distinctions: “expenditure”, “consumption”, and “consumption expenditure” are different concepts. “**Expenditure**” is the nominal money outlay of the household: simply put, it is the total amount of money spent by the household over a given reference period. “**Consumption**” refers to the quantity of goods and services that the household enjoys, or uses, during the reference period. “**Consumption expenditure**” refers only to the expenditure that is directed toward consumption: in other words, the amount of money spent to buy goods and services that are actually used (as opposed to being stored, which qualifies as an investment) over the reference period. An example may help: suppose that over a given reference period, say one month, a household spends \$30 to buy 30 kg of rice. Of these, 20 kg are actually eaten by the members of the family, while the remaining 10 kg are stored, and remain available for future needs (maybe they will be consumed, maybe they will be given away, or wasted: for now, we do not know). In this example, \$30 is household expenditure; 20 kg is household consumption of rice; \$20 is household consumption expenditure.

<p>Recap</p> <ul style="list-style-type: none"> ▪ There are a number of approaches to measuring living standards ▪ Economists seek to measure utility, which they approximate by consumption expenditure, adjusted for purchasing power: x/P ▪ This consumption-based measure is a simple (remember Sen's critique) but strong candidate to proxy the concept of living standard. 	<p>This explains why ... <small>World Bank, 2015, p. 31</small></p>  <p>“Consumption per capita is the preferred welfare indicator for the World Bank's analysis of global poverty.”</p>
--	---

The next slides recap the take-home points from the first part of the lecture. There are a number of different approaches to measuring living standards, which were not all covered here; so far, the focus has been on one approach, which is put forward by mainstream economics. Economists seek to measure utility, which they approximate by consumption expenditure, adjusted for purchasing power. Granted, expenditure excludes potentially important contributors to utility, such as publicly provided goods, or leisure – anything that does not have a measurable price. It also measures a “simplified” concept of well-being which Amartya Sen critiques, as seen in the opening slides. However, despite these limitations, this approach is advantageous enough to be the preferred one for the World Bank’s analysis of global poverty.

***Multidimensional poverty measurement**
 Alkire et al (2015)

- The multi faceted nature of standard of living has prompted a second approach, **multidimensional poverty measurement**.
- This builds on Amartya Sen's **functionings and capabilities** approach.
- **Functionings** are the beings and doings that people value, e.g. being well fed.
- **Capabilities** are the freedoms to achieve valuable functionings.
- Alkire et al. have developed a method to deal with a multi-dimensional definition of standard of living



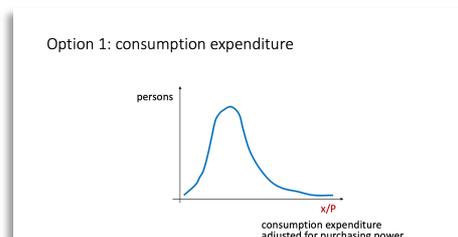
This optional slide mentions an important exception to the points just made: **multidimensional poverty measurement** is increasingly popular around the world. The cited reference is suggested to interested students.



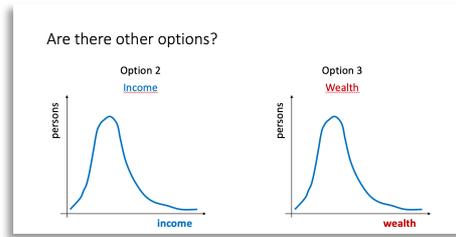
Now is a good time for a break.

2. Choosing a measure of living standards

After having laid the theoretical foundations for the choice of consumption expenditure as a measure of living standards, we move to considering other potential candidates, some of which are very common in practice: the final part of the lecture deals with the pros and cons of these competing measures.



The choice of **consumption expenditure** adjusted for purchasing power follows directly from the theoretical framework of money-metric utility, which was just discussed.



Other potential candidates are: household **income**, that is, receipts, whether monetary or in kind (goods and services) that accrue to the household or to individual members of the household over a given period of time (labor income, capital income, public and private transfers received); and household **wealth** or net worth, that is, the value of savings, investments, real estate and cash, less any debts, possessed by the household or by its members at one point in time.

Option 3: **Wealth**

- Wealth contributes to the standard of living. It does so indirectly, but it certainly does it.
- Economic theory says that wealth is a **stock** of resources. It is accumulated via **past** choices, and it may or may not be used to generate consumption in the **present**, which is what we care about.
- Conclusion: we put wealth aside.

Should these other options be considered viable candidates for the purpose of measuring living standards? **Wealth** is easy to rule out, therefore it is examined first. While it seems logical to believe that wealth contributes to the standard of living, there are both conceptual and practical arguments against it. Conceptually, wealth can be defined as a stock of resources. While consumption expenditure is a flow, which means that it is defined over an interval of time (what was spent over a day, or a month, or a year), wealth is measured at one specific point in time (say, today at midnight), and may have accumulated in the past. Individuals may or may not choose to use their stock of wealth to generate consumption (and thus utility) in the present, which is what we care about. There are also practical reasons that discourage the use of wealth as an indicator of living standards: it is difficult to measure, and data on wealth are only rarely available.

Option 2: **Income**

- «Among economic measures of living standards, the **main competitor** to a consumption based measure is a measure based on income» (Deaton and Zaidi 2002: 13)
- «In some countries, notably in Latin America, income is the only available indicator of economic welfare.» (World Bank 2015: 32)

Income, on the other hand, seems like a viable alternative: in fact, many countries use it.

Two identical households: A and B

Example

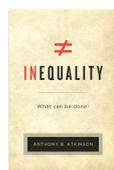
- Household A has a monthly **income** of \$1,000. This month, members of the household have consumed goods and services for a total value of \$900. The leftover \$100 is saved.
- Household B runs a family business, which did not do too well this month: income has been \$0. However, members of the household have financed their needs by using past savings, so they also have consumed goods for a total value of \$900.
- If we used **consumption** for measuring living standard, A and B would be equally well-off.
- If we used **income**, A would be better off than B.
- Which of these conclusions is correct?

An answer

- The use of **consumption** is justified by the concept of standard of living that was covered earlier: it captures the value of **use** of commodities (money-metric utility function).
- The use of **income** fits a slightly different concept of standard of living, where the emphasis is on **potential** rather than **actual** consumption
- We saw that when it comes to measuring **poverty**, microeconomic theory suggests to use (price-adjusted) **consumption** expenditure.
- What if interest were on **inequality**?

An example of the difference between income and consumption expenditure clarifies that we are talking about two very different objects. The answer to the final question, “which of these two conclusions is correct?”, is, in a way, ambiguous. If one accepts the conceptualization of standard of living that was covered in the first part of the lecture, then consumption expenditure is the correct measure, because it captures the value of use of commodities (and therefore, money-metric utility). However, the choice of income can be justified by a slightly different concept of standard of living, where the emphasis is on potential rather than actual consumption. This concept turns out to be particularly fitting when the main focus is inequality, rather than poverty. This is a complex point, which will be the subject of a future course prepared by the World Bank’s Center for Development Data (C4D2, focused on *Measuring Income and Wealth*).

*Sir Anthony B. Atkinson
2015, p. 37



- “I continue to focus on income as an indicator of potential control over resources. The use of **income** is indeed recognition that the use of resources **goes beyond consumption**.”
- “When measuring **inequality**, we are concerned not only with the consumption but also with the **power** that wealth can convey.”

This optional slide elaborates on the point of income vs. consumption. This slide shows quotes from a book by Anthony B. Atkinson, a well-known scholar and expert in the field of inequality, where this alternative conceptual framework is briefly explained. This is just a quick hint, as the main frame of reference for poverty measurement, and for this course, remains money-metric utility; interested students can go to Atkinson’s book for more information.

Income vs. consumption: which one to choose?

- The choice of the measure depends on:
 - 1) the **question** one is addressing
 - 2) a number of **practical considerations**
- The next few slides summarize some **advantages** (▲) and **disadvantages** (▼) of each measure.

The next slides list some of the **advantages and disadvantages** of income and consumption expenditure as measures of living standards. We suggest to comment on each bullet rapidly, inviting students to think about how convincing each argument is given the

context – we emphasize that context matters, that there is an on-going discussion among experts on the relevance of these arguments, and no easy way exist to reach an agreement. What is important is awareness of these issues, and ability to assess each of them in a specific context.

Income
Advantages

- ▲ Limited number of **sources** of income (less than items for consumption); in principle, easier to collect the information
- ▲ It is often possible to **assign** certain **sources** of income to **specific members** of the households
- ▲ Measures the **potential command over resources** (an advantage if this is the concept of interest)

The **advantages of income** are the following: *(i)* in practice, household-level totals of both consumption expenditure and income are often computed as a sum of sub-components (expenditures in various items, various sources of income), and sources of income are fewer in number than consumption items: this makes it easier to gather information on income, at least in this regard; *(ii)* some sources of income are directly connected to the individuals earning them, which allows for analyses of individual living standards in a way that consumption expenditure does not; *(iii)* the final point reiterates what was covered earlier: income is consistent with a concept of well-being as potential command over resources, and this makes income a good candidate when analysts are interested in such an approach.

Income
Disadvantages

- ▼ May be affected by **short-term fluctuations** (e.g. seasonal fluctuations in rural areas)
- ▼ **Under-reporting** (forgetting, reluctance to disclose, difficult to measure, etc.)
- ▼ Some income components are **difficult to observe** (e.g., income from informal labor activity, from home agricultural production)

The **disadvantages of income** are the following: *(i)* income is frequently affected by short-term fluctuations, especially in rural areas: we are usually interested in a measure of living standards that is representative of a longer reference period, say a year (more on this later), and it is challenging to adjust for these fluctuations at the data collection stage, and at the analysis stage; *(ii)* under-reporting is a serious issue when it comes to collecting data on income, and exposes the measure to significant measurement error; *(iii)* some income components are difficult to observe: two cases that are especially relevant in poorer countries are that of revenues from informal labor, and that of valuing incomes from agricultural home production (which are in-kind).

Consumption expenditure
Advantages

- ▲ Sound **theoretical foundations** (utility theory)
- ▲ Shows **long-term average well-being**, taking both consumption smoothing and insurance opportunities into account
- ▲ Measures the **use of resources** (an advantage if this is the concept of interest)

The **advantages of consumption expenditure** are the following: *(i)* the choice of this indicator is grounded in utility theory (the framework that was developed at the beginning of the lecture), and this is a desirable feature for empirical analysis; *(ii)* consumption is typically easier to recall than income, which decreases measurement error; *(iii)* again, the final point reiterates the consistency of consumption expenditure with a concept of well-being that coincides with the actual fruition of resources, which makes it a good candidate if that is the target of analysis.

Consumption expenditure

Disadvantages

- ▼ Households may **not be able to smooth consumption** (e.g. borrowing, insurance, social networks)
- ▼ Some expenses are **not made regularly**, which adds **noise** to the data
- ▼ Some components (durable goods and housing rents) are **difficult to capture**

The **disadvantages of consumption expenditure** are the following: *(i)* certain expenditures are directed toward items that do not enhance well-being (some economists call these items “bads”, in contrast with the term “goods” that is used generally): fines are an example; *(ii)* families may make extraordinary expenditures, that are not representative of their usual living standards: weddings and other social functions are an example; *(iii)* some components of consumption expenditure, like the use-value of durable goods and housing, are difficult to measure and require estimation procedures (more on this in a coming lecture).

What is the balance of the pros and cons?

- There is no hard-and-fast rule: the choice ultimately depends on the **type of analysis**, and the **context**.
- If the focus is on poverty measurement, analysts consider one last and important dimension, **time**.
- We define as **reference period** the period over which we want to measure welfare.

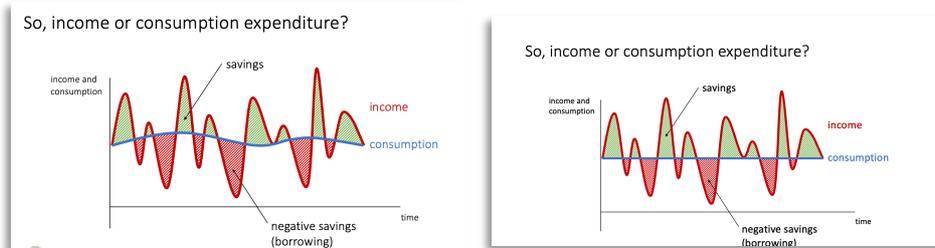
So, what is the balance of the pros and cons that have just been reviewed? As previously mentioned, the choice depends on the goals of the analysis and on the context. However, a final argument can be made regarding pros and cons of the two competing measures, and that concerns the reference period, which measured living standards are supposed to be representative of.

Time matters

- In the simplest models of **textbook economics**, time does not exist. Individuals spend all money, and they consume all goods. Hence, the choice of the indicator does not matter:
income = consumption expenditure.
- In **real life**, time exists and matters (as in other models in textbook economics). If we assume a reference period equal to, say, a year, then income and consumption expenditure can differ:
income = consumption expenditure + savings.

According to the simplest models of consumer choice presented in textbooks of economics, **time** does not exist. Individuals make their consumption choices just once, selecting the

optimal bundle, and spending all of their available budget. In this extreme case, income and consumption expenditure are the same. However, in real life, time matters. In general, income and consumption expenditure can differ: $\text{income} = \text{consumption expenditure} + \text{savings}$ (where savings may be negative, which is called dissaving, and occurs when there is borrowing). We are interested in the living standards of households over a given period of time, typically **a year**, and we want our estimates to be representative of that period.



The implication of this discussion for our choice of income versus consumption expenditure is explained by this figure. If the reference period we are interested in is a year, then we must evaluate income and consumption expenditure (or even wealth, or a combination of all three) over that period of time. Empirically, consumption expenditure fluctuates less (in other words, it is **smoother** over time) in the short-run than income does – refer to Deaton and Zaidi (2002: 14) for a discussion of this topic. The graph helps intuitively grasp what “smoothing” is.

In conclusion

- Smoothing gives **consumption** a **practical advantage** over income in the measurement of living standards.
- Observing consumption over a relatively short period – even a week or two – tells us a great deal more about annual (or even longer period) living standards than income can tell.

This slide wraps up the discussion on **income vs. consumption**: among the advantages of using consumption expenditure as an indicator of living standards, we must count its (comparative) stationarity over time: because of smoothing, observing consumption over a relatively short period – even a week or two – tells us a great deal more about annual (or even longer period) living standards than income can.

The international practice

- Where do countries around the world fall when choosing between income and consumption expenditure as the preferred indicator of living standards?

Sub-Saharan Africa



Surveys	Consumption vs Income
Côte d'Ivoire 2015	Consumption
Kenya 2015	Consumption
Malawi 2016	Consumption
Mali 2018	Consumption
Mozambique 2014	Consumption
South Africa 2014	Consumption
Tanzania 2018	Consumption
Uganda 2016	Consumption
Zambia 2015	Consumption
Zimbabwe 2017	Consumption

At this point of the presentation the audience is most likely tired... so the way we conclude is by showing them something concrete, and hopefully interesting. This bunch of slides provides some examples of the choices made by countries around the world when it comes

to measuring poverty and inequality, to provide a sense of what the common practice is, and how it varies. This overview is meant to give students some practical examples, and to ground the concepts covered in the lecture into practice.

Lessons learned

- The concepts and vocabulary introduced in this lecture are of paramount importance for data providers, not just for analysts, because data quality depends on **relevance** for a specific research objective.
- Poverty analysts need a proxy variable for the standard of living. Economic theory combined with practical arguments suggest to use **consumption expenditure adjusted for purchasing power**.
- Economists call it **money-metric utility function**, defined as x/P , where x is consumption expenditure and P is a Paasche price index.
- **Income** remains a strong candidate.

The final slide is a recap of the whole lecture: it is a good idea to spend a few minutes on it, to make sure that the main bring-home conclusions are clear. The first point is especially important, and justifies the time and effort spent on the theoretical issues that were the focus of the lecture: the concepts and vocabulary just introduced are of paramount importance for data providers, not just for analysts, because data quality depends on **relevance** for a specific research objective. If data providers and analysts do not share the same conceptual tools, data will never match user needs and expectations, no matter the amount of money and effort spent in collecting them. The remaining points summarize the key messages of the lecture.

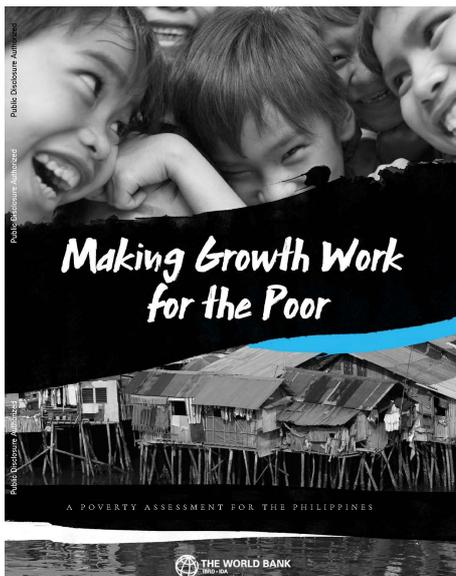
Homework

Exercise 1 – Engaging with the literature

This exercise is conceived as an invitation to get involved with relevant pieces of literature that are related to the topics covered in the lecture. The emphasis here is not on the essay itself, which should be brief and synthetic, but on the student’s understanding of the key conclusions of the papers, and how they relate to the lecture. Obviously, there is not a specific and unique solution for this exercise: the Instructor should evaluate each short essay by keeping in mind the intent that was just described. The same holds for all “Engaging with the literature” exercises in the course.

Exercise 2 – Income or consumption?

This exercise gives students an opportunity to see how the theoretical principles covered in the lecture operate in practice, by exploring actual Poverty Assessment reports and the analytical choices they make. Students may find the cross-country variation in the choice of income or consumption interesting, and may also face instances in which the indicator of choice is not entirely clear – which is a valuable lesson in transparency and accountability. An example of what students may find is from the latest for the Philippines (2018), “Making Growth Work for the Poor: A Poverty Assessment for the Philippines”:



Box 1.1. Poverty estimates using national and international poverty lines

National official poverty estimates in the Philippines are produced by the Philippine Statistics Authority (PSA). These are derived using income welfare aggregates evaluated against per capita poverty lines that are set broadly following the cost-of-basic-needs (CBN) approach. Using a national reference food bundle based on expert opinion of what constitutes a nutritionally adequate bundle, province-specific bundles are set separately for urban and rural areas and reflect locally consumed commodities. These locally

Exercise 3 – Multidimensionality of well-being

The goal of this exercise is to show that drawing conclusions on a country's living standards based on multiple indicators of well-being, while attractive (and useful in some contexts), is challenging, as different indicators do not normally agree with each other. Once again, the emphasis here is not on the essay itself, but on students' ability to engage critically with some of the concepts covered in the lecture.

Lecture 2

The consumption aggregate

Learning objectives

The goal of this lecture is to complete the conceptual framework underlying the construction of a measure of living standards. Building on the conclusion reached by Lecture 1 – preference should be given to a consumption-based measure of living standards – this lecture provides a working definition of consumption expenditure, and discusses the adjustments that analysts should apply to it, in order to construct a *living standard indicator*.

Suggested preparation

The Instructor is assumed to be familiar with the material contained in sections 4 and 5 of Deaton and Zaidi (2002).

Time allocation

Roadmap and ‘which expenditures?’	15 min
Adjusting for household size and composition	30 min
Adjusting for purchasing power	
Terminology	10 min
Break	15 min
Inflation example, price indices	45 min
Lessons learned	5 min

Annotated lecture

The lecture opens with the agenda for the day.

It takes four to construct a living standard indicator

$$\text{Living standard indicator} = \frac{\text{nominal household consumption expenditure}^1}{\text{household size}^2 \times \text{temporal CPI}^3 \times \text{spatial CPI}^4}$$

Numerator

- Which expenditures, exactly?

Denominator

- Household size is not self-defining (boarders? guests? servants?..)
- How to account for inflation (temporal CPI) and differences in price levels across the national territory (spatial CPI)?

This slide provides the roadmap for the whole lecture. Our objective, constructing a living standard indicator on the basis of survey data, consistently with the theoretical framework covered in Lecture 1, is going to be reached by combining four elements, as shown here. This lecture will be discussing each of the four elements.

1. Which expenditures, exactly?

The nominal consumption aggregate

Nominal consumption aggregate =

- monetary expenditures on food and non-food non-durable goods and services consumed
- + value of in-kind consumption
- + value of use (not purchase) of durables
- + value of use of owner-occupied housing.

No allowance for the value of time and leisure and no allowance for public goods.

The next few slides discuss the numerator of the living standard indicator: nominal household consumption expenditure. The takeaway for this section is that this step does not simply imply adding up all household expenditures, precisely because what matters is *consumption*, rather than *expenditure* per se. The discussion of this topic is intentionally kept short, as it relates more to analysis than to data collection. More space will be given to adjustments (the denominator of the living standard indicator), which have more bearing on survey design choices.

*Definition: public good

- In economics, a public good is a good that is both non-excludable and non-rivalrous
 - Non-excludable – individuals cannot be excluded from use (they can benefit from it even if they are not paying for it)
 - Non-rivalrous – use by one individual does not reduce availability to others (the good can be used simultaneously by more than one person)
- Many public goods are provided by the government. For example, national defense – but also streets, street lighting, and much more.
- We can think of ‘public goods’ within a household, too.
- Housing, heating, transportation, ... are private goods, by all means, but non-excludable and non-rivalrous within the household – they can be shared by household members.

A quick note on the slide above: slides with the “encyclopedia” icon provide **definitions** for key concepts. They may be technical concepts, or, as in this case, ideas that are quite well-known, but for which it is useful to give a *precise* definition, so that everybody in the room is on the same page. Here we use the concept of ‘public good’ in a general fashion, to indicate all those goods that are owned by the household, but are shared between household members (*i.e.* they are on-exclusive and non-rivalrous within the household). For a precise definition see Stiglitz (2015), *Economics of the Public Sector*, ch. 5.

Living standard indicator = $\frac{\text{nominal household consumption expenditure} \quad \checkmark}{\text{household size} \times \text{temporal CPI} \times \text{spatial CPI} \quad \square}$

This slide signals that the numerator of the living standard indicator has been dealt with: it is time to move on to the next component, that is, the adjustment for household size and composition.

2. Adjusting for household size and composition

Adjusting for household size and economies of scale

- Larger households consume more, because there are "more mouths to feed"
- One possibility is to consider **per capita** consumption
- There is a subtler issue: housing, heating, transportation etc. are **shared** between members: an analogy can be drawn with the idea of **public goods**
- Example: **housing**. Consumption by one member of the household does not necessarily reduce the amount available for consumption by another person within the same household. Economists say that there are significant **economies of scale** for housing.
- By **failing to adjust** for economies of scale, one might underestimate the wellbeing of large households (and overestimate that of small households).

Some introductory slides explain why we should concern ourselves with differences in household size (including the possibility of accounting for economies of scale) and household composition.

Economies of scale: adjustment

- A popular strategy is to rescale household consumption expenditure as follows:

$$\bar{x}_i = \frac{x_i}{(n_h)^\alpha} \quad \alpha \in [0,1]$$
- $\alpha = 1$ means we assume that **no** goods consumed are **public** in the household, in which case consumption is equally divided among household members. No adjustment for economies of scale is made.
- $\alpha = 0$ means we assume that all goods consumed are **public** in the household. This is a purely hypothetical situation in which each individual is assumed to consume the total consumption of the household.
- In practice, α assumes conventional values. E.g. $\alpha = 0.5$ implies that a household of **four** persons needs twice as much as a single-person household.

Economies of scale: to adjust or not to adjust?

- When a high percentage of budget is devoted to public goods (that is, if price and quantity of housing, utilities, and durable goods are high), **economies of scale** are likely to be significant.
- Analysts look at the **shares in the data**, and decide whether to adjust.
- Rule of thumb:
large share = adjust
small share = do not adjust

The following slides explain the most common approach used in practice to adjust for economies of scale (using the coefficient 'alpha' to modify household size), and shares a general principle to help decide whether or not such an adjustment is necessary.

Adjusting for household composition

- It is usually assumed that **children** and the **elderly** need less than working-age **adults**.
- Similarly, it is thought that **women** need less consumption than **men**.
- If that is the case, our standard of living indicator should account for differences in household composition.



Equivalence scales: adjustment

- If adjustments are to be made, we use **equivalence scales**.
- An equivalence scale calculates the **number of equivalent adults** in the household.
- For **example**, an equivalence scale may look like this:

$$ES = (n_{\text{males } 15+} \times 1) + (n_{\text{females } 15+} \times 0.8) + (n_{\text{kids } 0-14} \times 0.5)$$
 where ES denote the **equivalent household size**, that is, the number of equivalent adults.
- Different categories have different **'weights'**: adult males may count for 1, adult females for 0.8, ...
- Once the equivalent household size has been calculated, we use it to rescale household consumption expenditure as follows:

$$\bar{x}_i = \frac{x_i}{ES}$$

Next, we introduce the concept of adjusting for differences in household composition. This is relatively intuitive: a household made up by one adult and two kids may need less consumption than a household made up of three adults, therefore the two households may enjoy a different level of welfare, even though they may spend the same for consumption. Differences in average human energy requirements (minimum calories needed) are an obvious demonstration of the fact that needs vary with demographic characteristics.

To account for these differences in needs, analysts use equivalence scales. An example illustrates: the basic idea is that to each relevant demographic category, we associate a weight, that describes how much lower the needs of that category are with respect to a reference (usually adult males).

Equivalence scales: OECD and Eurostat

- The **OECD equivalence scale (OECD-I)** is defined as follows:
 $ES_{OECD-I} = 0.3 + 0.7 \times A + 0.5 \times K$
- where ES_{OECD-I} is the number of 'adult equivalents', that is, the household equivalent size, as measured by the OECD type-I scale.
- The first adult (A) is given a weight of 1. Other adults are given a weight of 0.7, to reflect economies of scale. Children (K) are given a weight of 0.5 to reflect their lower needs.
- In the late 1990s Eurostat adopted the so-called **OECD-modified scale (OECD-II)**:
 $ES_{OECD-II} = 0.5 + 0.5 \times A + 0.3 \times K$

Equivalence scales: USA

- Originally suggested by Cutler and Katz (1992), and subsequently adopted by the **US National Research Council (1995)**:
 $ES_{NRS} = (A + \alpha \times K)^\theta$
- where ES_{NRS} denotes the number of adult equivalents that is, the household equivalent size, as measured by the NRS scale;
- As before, A = number of adults, K = number of children;
- α in $[0,1]$ is cost of a child relative to that of an adult;
- θ in $[0,1]$, $(1 - \theta)$ measures the extent of **economies of scale**.

The choice of a particular equivalence scale amounts to the choice of the weights. The most commonly used scales are listed here: the slides show the general formulas for the OECD-I and OECD-II scales, and for the so-called NRC scale. The latter accounts for both economies of scale (note the theta coefficient) and household composition, in a single formula.

Equivalence scales in practice

Household composition	Equivalent household size				
	Per capita	OECD-I scale	OECD-II scale	Square root scale ($\alpha = 1, \theta = 1/2$)	Per household
1 adult	1	1	1	1	1
2 adults	2	1.7	1.5	1.4	1
2 adults, 1 child	3	2.2	1.8	1.7	1
2 adults, 2 children	4	2.7	2.1	2.0	1
2 adults, 3 children	5	3.7	2.4	2.2	1

The table compares common adjustments for household size. The first column, per capita, is the baseline: it simply shows the size of each of the household types listed. The other categories show the equivalent size that results from using OECD-I and II, and the 'square-root scale' (a simple adjustment for economies of scale, that does not account for differences in composition). Finally, the last column shows what happens if we do not adjust at all, and simply use total household consumption.

Equivalence scales: to adjust or not to adjust?

- If **children/elderly** are as "expensive" as **adults** despite their lower nutritional requirement (e.g. because of very high costs for education or health), **less need for adjustment**.
- Rule of thumb:
large differences in the "cost" of different household members = **adjust**
small differences = **do not adjust**.

To wrap up, we share some advice on how to decide whether or not we should adjust for differences in household composition.

The international practice

Where do countries around the world fall when adjusting for household size and composition?

Sub-Saharan Africa



Surveys	Consumption vs Income	Household Size
Côte d'Ivoire 2015	Consumption	Per Capita
Kenya 2015	Consumption	Per Adult Equivalent
Malawi 2016	Consumption	Per Capita
Mali 2018	Consumption	Per Capita
Mozambique 2014	Consumption	Per Capita
Nigeria 2010	Consumption	Per Adult Equivalent
South Africa 2014	Consumption	Per Capita
Tanzania 2018	Consumption	Per Adult Equivalent
Uganda 2016	Consumption	Per Adult Equivalent
Zambia 2015	Consumption	Per Adult Equivalent
Zimbabwe 2017	Consumption	Per Capita

Then, some evidence on what countries around the world actually implement is presented. Different regions of the world adopt different strategies, but it is shown that expressing the living standard indicator *at least* on a per capita basis is routine.

$$\text{Living standard indicator} = \frac{\text{nominal household consumption expenditure}}{\text{household size} \times \text{temporal CPI} \times \text{spatial CPI}}$$

✓ 1 4

Again, this slide signal that the lecture is now moving on to the last topic: the adjustment for **differences in purchasing power**, which encompasses elements 3 and 4.

3. & 4. Adjusting for purchasing power

...

Temporal and spatial deflators

- 1) Inflation (time)
Temporal (monthly, yearly) price index
- 2) Cost-of-living differences across the national territory (space)
Spatial price index

Price indices are typically expressed as a proportion of some reference price level:

- Price index = 1 (or 100) → current price level is the **same** as the reference level
- Price index > 1 (or 100) → current price level is **higher** than the reference level
- Price index < 1 (or 100) → current price level is **lower** than the reference level

These slides provide some basic terminology: nominal vs. real, consumer price indices (CPI), temporal and spatial deflation. The way CPIs should be interpreted is also clarified.



Now is a good time for a break.

Example: inflation matters

- Assume all households in the country are identical (same size, composition, etc.)
- Assume consumption expenditure $x = \$1,000$ for all households
- Assume inflation = 5% per month during the survey year
- Note: this is a high inflation rate... what would that be on a yearly basis?
- Assume that each month 1/12 of the households are interviewed
- Assume that the poverty line equals \$ 750

...

Kenya 2015
KIHBS 2015/16,
Paasche Index

These slides present a hypothetical example, that demonstrates why adjusting for differences in purchasing power is important. The example focuses on temporal price differences, i.e. inflation, but the same reasoning holds for spatial price differences. The upshot is that using nominal or real expenditure produces different poverty estimates, and that the second measure best represents actual living standards. A real example from Kenya concludes the discussion, showing the relevance of this issue in practice.

What are price indices (or deflators) exactly?

- Many indices exist:
 - Laspeyres
 - Paasche
 - Fisher
 - Törnqvist
 - ...
- We focus on Laspeyres and Paasche.
- Q. Why?
- A. In lecture 1 we concluded that our best strategy for proxying living standard is either x/P or x/L , with a preference for the former.

The final batch of slides discusses different price indices, namely Laspeyres and Paasche. Section 4 of Deaton and Zaidi (2002) is the main reference for this topic. The goal is for students to understand that there are several ways to adjust for price differences, and that they are not equivalent.

The Laspeyres index

- The single **most popular** index among both economists and international statistical agencies.
- According to the ILO Bureau of Statistics, **114 out of 187 countries** use the Laspeyres formula.



The Laspeyres index: definition

▪ The Laspeyres index answers the question:

“what is the **cost** of a **fixed** basket of commodities purchased in the **base period** relative to its cost at the **base period** market prices?”

$$L^t = \frac{p^t \cdot q^0}{p^0 \cdot q^0}$$

- q^0 reference vector of quantities;
- p^t vector of prices faced in period t ;
- p^0 reference set of prices.

The Laspeyres index: interpretation

- When applied to bundles consumed by individual households, a **Laspeyres index that equals 1** (or 100) implies that, if the household could afford to buy the reference consumption bundle in the **base period**, then she can also afford it in the **current period**.

The presentation of both indices follows a simple structure: first, we introduce the index using its formula. Then, we give the interpretation of the index ‘in words’. Finally, we comment on its advantages and disadvantages.

Regarding the formula, the only real difficulty here is understanding the notation: keeping in mind the example of temporal differences in prices (but everything applies to spatial differences, too) we can think of the apex ‘0’ as the reference or base period, and of the apex ‘t’ as the current period. The letter q indicates a set of quantities of goods consumed by households; p indicates the corresponding prices. Note that the ‘dot-product’ here indicates the sum of each quantity times its price (that is, the total cost of the consumed bundle).

Once this is clear, the expression for the Laspeyres index has no secrets: it is the ratio of the costs of a reference bundle of goods, evaluated at the base period and at the current period.

The Laspeyres index: comment

- A key feature of the Laspeyres formula is that it tends to **overstate** the rise in the cost of living by **not allowing any substitution** between goods to occur (Diewert, 2001).
- To the extent to which price and demanded quantity are negatively correlated, **the Laspeyres index provides an upper bound** to the “true cost of living” faced by a household.

To comment on the Laspeyres index, we share that one of its key features is that it tends to *overstate* cost of living differences, by not allowing any substitution between goods to occur when prices vary. This is a consequence of holding quantities q^0 fixed at the base (reference) level: in reality, when faced with a price change, households tend to reallocate their resources away from relatively pricier and towards relatively cheaper goods. Consequently, the Laspeyres index provides an upper bound to the cost of living faced by a household.

The Paasche index

- The Paasche index is the one that most **welfare analysts** opt for.
- **Deaton and Zaidi (2002)** explain why.



The Paasche index: definition

- The Paasche index:
“what is the cost of a fixed basket of commodities purchased in period t relative to its cost at the base period market prices?”

$$p^t = \frac{p^t \cdot q^t}{p^0 \cdot q^t}$$

- q^t : vector of quantity purchased in period t ;
- p^t : vector of prices faced in period t ;
- p^0 : reference set of prices

The Paasche index: interpretation

- When applied to bundles consumed by individual households, a **Paasche index that equals 1** (or 100) states that, in the **base period**, a household could have consumed the same bundle as she is consuming in the **current period**.

Using the same logic, we define and interpret Paasche. The key difference between the two indices lies in how prices are *weighted* by the quantities consumed: while both indices account for *price relatives*, that is, the prices in the current period relative to the reference prices (p^t/p^0), the Paasche index also accounts for the specific consumption pattern of each period t (q^t), something that is not true of a Laspeyres index, which instead uses a set of reference quantities, q^0 , that is the same when the index is computed in all subsequent periods. The importance of this distinction is clearer in the spatial context: Laspeyres uses weights (quantities) that are common to all households, while Paasche uses household-specific weights.

The Paasche index: comment

- The Paasche formula **does not allow for the substitution** of products or services at the **base period prices**.
- To the extent to which price and demanded quantity are negatively correlated, it provides a **lower bound** to the “true cost of living” faced by the household.

The Paasche formula suffers from the opposite drawback as the Laspeyres formula, because the weights are set at the level that is optimal at prices – in a way, after consumers have completed all substitutions of goods and services in response to a change in prices. To the extent that price and demanded quantity are negatively correlated, the Paasche index provides a lower bound to the cost of living faced by the household.

Paasche vs. Laspeyres

- While calculating the **Laspeyres index** for a new period requires only new price data, calculation of the **Paasche index** for a new period requires new price data and new quantity data (or alternatively new price data and new expenditure data) for each new period.
- The **Paasche index** is rarely calculated by statistical agencies because it is **data demanding**.
- Given the poverty analyst's preference for Paasche, it is common to estimate it based on household budget surveys.
- This is not as straightforward as it might seem.

In our context, however, other considerations are more important. When the focus is welfare measurement, the preferred price index is Paasche, for the reasons explained in Lecture 1. The challenge for NSOs is to collect the data necessary to its estimation.

*The Fisher index

- When **both indices** can be obtained, a recommended solution is to calculate an average of the two.
- The **Fisher index** is defined as the **geometric average** of the Laspeyres and Paasche indices:

$$P^t = \sqrt{L^t \times P^t}$$

This optional slide touches on another index that is widely used in practice, the Fisher index.

Measuring prices

- A **beautiful** and **useful** paper
- It reviews the academic **literature** about **prices** for **poverty** measurement in **Africa**

Market prices are not unit values (and vice-versa)

- Unit values are defined as the ratio of expenditure to quantity.
- **Unit value** for household h , good j : $uv_j^h = \frac{x_j^h}{q_j}$
- Unit values suffer from **quality bias**: richer households tend to buy higher quality foodstuffs, for instance.

Empirical distributions of unit values for selected food items
Maldives (2016)

Unit values cannot be treated as if they were market prices

- Analysts are increasingly dissatisfied with unit values
- **Gibson and Kim (2019)** is the last of a string of papers advocating for better data on market prices
- We are still 'in transition'

A common source of information on prices comes from household surveys themselves: *unit values*, that is, the ratio of expenditures to quantities, are often easy to compute based on the food modules of most expenditure surveys. Historically, they have been used as

proxies for market prices. However, unit values have become increasingly embattled in the recent literature on price indices, because of the many empirical problems they have. Goods of very different qualities may be grouped under the same category in household surveys, and the resulting unit values may be averages of a miscellaneous bunch. The graphs from Maldives illustrate this point: the top-left distribution looks reassuring, because it indicates that reported unit values for rice are concentrated around a single value, but the other examples suggest the presence of underlying quality differences. However, the literature on this topic has not yet converged on a concrete alternative: the transition away from unit values is not yet complete.



Again, the discussion wraps up by showing a real example of spatial differences in prices.

Lessons learned

- Data providers should be mindful of the definition of **household membership**, because of the importance of adjusting for household size.
- **Household characteristics** (gender, age, etc.) are key for the computation of equivalence scales.
- Adjustment for cost-of-living differences:
 - **temporal CPI** is needed to adjust for within-survey inflation;
 - **spatial CPI** is typically computed from household surveys, which must allow for it.
- Market prices (collected through price surveys) are different from unit values (calculated on the basis of household budget surveys).

The final slides summarize the topics discussed in the lecture, with a focus on their implications for survey data collection.

Homework

Exercises 1 and 2 – Engaging with the literature

See exercise 1, Lecture 1.

Lecture 3

Understanding household surveys

Learning objectives

The goal of this lecture is to provide an overview of the many different survey instruments developed and implemented by statistical institutions around the world. The material covered in the lecture puts household consumption and expenditure data – the main focus of the course – into a broader context.

Suggested preparation

The paper by Grosh and Glewwe (1998) provides background information on the Living Standard Measurement Study and the surveys developed under its umbrella. Familiarity with it is assumed.

Time allocation

Overview of statistical instruments	20 min
Overview of household surveys	
Introduction and terminology	30 min
Break	15 min
Taxonomy: common surveys, quick HCES, large HCES	50 min
Lessons learned	5 min

Annotated lecture

The set opens with a roadmap: after concepts, which were covered in the previous lectures, it is now time to move to data. There are many different ways for statistical institutions to collect data, and this lecture provides broad classifications and terminology for different data collection efforts.

1. An overview of statistical instruments

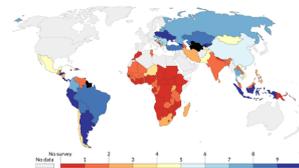
The first part of the lecture provides an overview of statistical instruments *in general*.

Demand for data is increasing

- **Perception vs. evidence**
Data is always in demand for policymakers, donors, academics, researchers, civil society, citizens, and anyone who wants evidence to support their ideas.
- **Evidence-based policy making**
Is the public sector delivering good services? Are they properly targeted? Are government policies, donor-funded development plans reducing poverty?
- **Monitoring social indicators**
Sustainable Development Goals (SDGs) comprise more than 200 indicators: need for understanding progress (or lack thereof) in each indicator

Global overview

Number of poverty surveys available via the World Bank in the last decade (2005-2014)

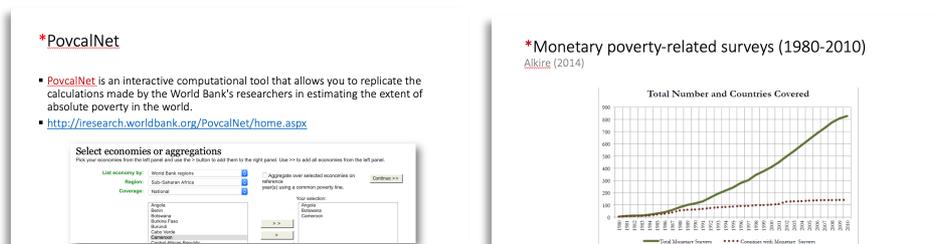


By way of motivation, the first few slides make the point that the demand for data (by policymakers, researchers, society in general) is increasing; the **supply of data** is increasing too, but there is still much work to do. The maps presented here come from the website Our World in Data (<https://ourworldindata.org/extreme-poverty>), and represent the number of recent household surveys that are available around the world as a basis for the World Bank's estimates of global poverty.

The countries in grey are those where not a single survey is available to the World Bank in the last decade. Many of these countries are rich countries, in which extreme poverty is very low. But there is also missing data for some poorer countries, in which surely a considerable share of the population is living in extreme poverty.

The countries in red and orange are those with very few available surveys. This is the case for many African countries, where there is only one survey available in the last decade.

The last map gives a general overview of 'data deprivation' in the field of poverty measurement. The key point is that efforts toward better data collection are important.



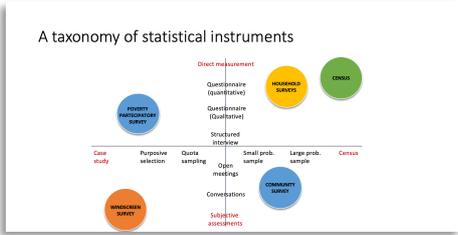
These two optional slides provide two more examples of data availability in the field of poverty measurement. **PovcalNet**, in particular, can be an interesting tool for students to explore: it allows to replicate the World Bank's official poverty estimates, and can be used to compute poverty measures on imported data as well.¹

The evidence of an increasing demand for data is used to justify the multitude of different statistical instruments currently in use.

¹ For an introduction to PovcalNet, see <http://iresearch.worldbank.org/PovcalNet/introduction.aspx>.

A useful classification

- How to make sense of the multitude of data collection efforts that are in use?
- A convenient starting point is to consider two criteria:
 - Representativeness
 - Objectivity



A classification based on two dimensions – **representativeness and objectivity** – can be a useful tool to evaluate these different instruments. Notable examples of information gathering are placed on a diagram defined by these two dimensions. Note that a *windscreen survey* is not a survey at all, but rather an informal, subjective assessment of one’s surroundings, done from one’s seat in a car; and a *poverty participatory survey* is similar to a focus group, ran with the aim of understanding poverty determinants by incorporating the perspectives of the poor themselves.

Restricting the focus

- Household surveys** are in upper-right corner, that is, they score well in terms of both **representativeness** and **objectivity**
- Once agreed on this, we can narrow them down further, based on which types of surveys are **relevant** for our purposes

Representativeness and objectivity are desirable properties, and they help restrict our focus to statistical instruments that possess these qualities: **household surveys**.

2. An overview of household surveys

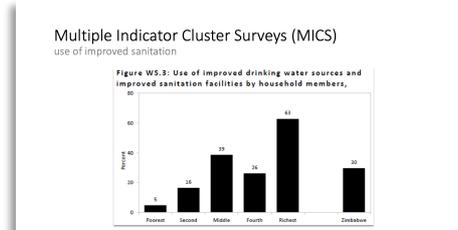
Household Surveys		
Common surveys	Household Consumption and Expenditure Surveys (HCES)	
Labor Force Surveys (LFS)	"Quick Survey" Monitoring of various socio economic indicators	"Large survey" Comprehensive income & consumption info
Demographic and Health Surveys (DHS)	Priority Surveys (PS) Core Welfare Indicator Questionnaires (CWIQ)	Household Budget Surveys (HBS) Household Income and Expenditure Surveys (HIES)
Multiple Indicator Cluster Surveys (MICS)	Comprehensive Food Security and Vulnerability Analysis (CFSVA) Welfare Monitoring Surveys (WMS) Survey of Well-being via Instant and Frequent Tracking (SWIFT)	Living Standards Measurement Studies (LSMS) Integrated Household Surveys (IHS)

The second part of the lecture focuses precisely on household surveys. The topic is introduced with some historical background and some terminology (for ‘**household**’ – the case of polygamy, which is common in some countries, is used to engage students and give an example of a local context where defining a household may be challenging – and ‘**probabilistic sample**’). Once the basics are covered, a slide introduces a classification of the main types of household surveys. On the one hand, we have ‘**common**’ surveys – common in the sense that they are virtually ubiquitous, frequently mentioned, essentially a staple for most statistical offices around the world. On the other hand, we have Household Consumption and Expenditure Surveys (**HCES**): not necessarily less ‘common’ than other surveys, what characterizes them is their focus on measuring consumption and expenditure; they can be divided into ‘**quick**’ and ‘**large**’ surveys. The following slides go into each of these groupings.



Now is a good time for a break.

Household Surveys	
Common surveys	Household Consumption and Expenditure Surveys (HCES)
Labor Force Surveys (LFS)	"Quick Survey" Monitoring of various socio economic indicators
	"Large survey" Comprehensive income & consumption info
Demographic and Health Surveys (DHS)	Priority Surveys (PS) Core Welfare Indicator Questionnaires (CWIQ) Comprehensive Food Security and Vulnerability Analysis (CFSVA) Welfare Monitoring Surveys (WMS) Survey of Well-being via Instant and Frequent Tracking (SWIFT)
Multiple Indicator Cluster Surveys (MICS)	Household Budget Surveys (HBS) Household Income and Expenditure Surveys (HIES) Living Standards Measurement Studies (LSMS) Integrated Household Surveys (IHS)



The group of slides on **common surveys** is structured as follows: each survey type (LFS, DHS, MICS) is introduced with a general description; a recent example from a country is presented, to show what the survey might look like in practice; a final slide shows an indicator or statistic that is usually computed from that type of survey. This section can be covered rather quickly, as students may already be familiar with some of the material.

Household Surveys	
Common surveys	Household Consumption and Expenditure Surveys (HCES)
Labor Force Surveys (LFS)	"Quick Survey" Monitoring of various socio economic indicators
	"Large survey" Comprehensive income & consumption info
Demographic and Health Surveys (DHS)	Priority Surveys (PS) Core Welfare Indicator Questionnaires (CWIQ) Comprehensive Food Security and Vulnerability Analysis (CFSVA) Welfare Monitoring Surveys (WMS) Survey of Well-being via Instant and Frequent Tracking (SWIFT)
Multiple Indicator Cluster Surveys (MICS)	Household Budget Surveys (HBS) Household Income and Expenditure Surveys (HIES) Living Standards Measurement Studies (LSMS) Integrated Household Surveys (IHS)

"Large" and "Quick" Surveys

- Different surveys for different aims
- Quick surveys usually enable regular computation of welfare estimates contingent upon the existence of an underlying comprehensive large survey
- The administration of Quick and Large surveys can be combined for better coordination and value addition – "SWIFT estimations are only as good as the underlying models which require access to recent large scale data sets collected by governments and multilateral agencies." (p. 3)

The discussion then moves to **HCES**. 'Quick' surveys are described synthetically, as they are complementary to 'large' surveys. Quick surveys are especially relevant in specific context, like countries impacted by conflict and violence, where efficiency and cost-effectiveness are paramount. A full discussion of this rather specialized instrument goes beyond the scope of the course: interested students are directed to the suggested references for more details.

Optional (starred) slides discuss an example from Indonesia, and the Nigeria CWIQ 2006.

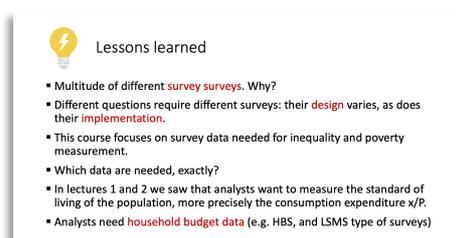
Household Surveys	
Common survey forms	Household Consumption and Expenditure Surveys (HCES)
Labor Force Surveys (LFS)	"Quick Survey" Monitoring of various socio economic indicators
	"Large survey" Comprehensive income & consumption info
Demographic and Health Surveys (DHS)	Priority Surveys (PS) Core Welfare Indicator Questionnaires (CWIQ) Comprehensive Food Security and Vulnerability Analysis (CFSVA) Welfare Monitoring Surveys (WMS) Survey of Well-being via Instant and Frequent Tracking (SWIFT)
Multiple Indicator Cluster Surveys (MICS)	Household Budget Surveys (HBS) Household Income and Expenditure Surveys (HIES) Living Standards Measurement Studies (LSMS) Integrated Household Surveys (IHS)

HBS vs LSMS

- Both are multi-topic, integrated, complex and nationally representative surveys
- Note that country-specific surveys names do not always mention 'HBS' or 'LSMS' explicitly, but that does not mean they are not in these categories
- HBS/HIES: Originally designed to provide input into the CPis and National Accounts (improve macro economic statistics)
- LSMS/IHS: Designed primarily for detailed welfare and poverty measurement and monitoring (improve development statistics)
- LSMS has typically a smaller sample size w.r.t. HBS to minimize non-sampling error
- The period of data collection/fieldwork is (usually) around 12 months for both in order to account for seasonality

Finally, the lecture zooms into 'large' HCES. This part should be emphasized over the others, as it is more important for what will follow in the rest of the course. In particular,

HBS and **LSMS** are two main approaches to conducting HCES, they are both routinely used for measuring living standards, poverty and inequality, and the differences and similarities among them should be adequately discussed.



Lessons learned

- Multitude of different **survey surveys**. Why?
- Different questions require different surveys: their **design** varies, as does their **implementation**.
- This course focuses on survey data needed for inequality and poverty measurement.
- Which data are needed, exactly?
- In lectures 1 and 2 we saw that analysts want to measure the standard of living of the population, more precisely the consumption expenditure x/P .
- Analysts need **household budget data** (e.g. HBS, and LSMS type of surveys)

The bring-home conclusion is that welfare analysts need household budget data (e.g. HBS, and LSMS type of surveys). By the end of the lecture, the characteristics of this type of surveys, and the general context in which they are situated, should be clear to students.

Homework

Exercise 1 – Engaging with the literature

See exercise 1, Lecture 1.

Exercise 2 – Household?

This exercise gives students an opportunity to see how the theoretical principles covered in the lecture operate in practice, by exploring actual survey documentation material and the definitions underlying the design of questionnaires. An example of what students may find is from the enumerator manual of Malawi's Integrated Household Panel Survey 2010-2013:

“A household may be either a person living alone or a group of people, either **related or unrelated**, who **live together** as a single unit in the sense that they **have common housekeeping arrangements** (that is, share or are supported by a common budget). A standard definition of a household is: “**a group of people who live together, pool their money, and eat at least one meal together each day**”. It is possible that individuals who are not members of the household may be residing with the household at the time of the survey. In most cases, but not all, someone who does not live with the household during the survey period is not a current member of the household.”

The students' findings can be used as a basis for discussion of the differences, or similarities, of such a fundamental definition across countries.

Exercise 3 - Household surveys

This exercise encourages students to identify the main features that characterize a survey, and pushes them to read and get familiar with technical documentation. Again, there is no single right answer for this exercise: the students' findings can be used to start a discussion on the topic of harmonization. A good solution might look like this:

country	survey	year	Sample size	Sampling method
Tanzania	DHS	2015-2016	13,376 hh	Stratified sample selected in two stages. Probability proportional-to-size selection at the first stage of sampling
Tanzania	NPS	2014-2015	3,265 hh	multi-stage cluster sample design.
Ethiopia	DHS	2016	18,008 hh	Stratified sample selected in two stages. Probability proportional-to-size selection at the first stage of sampling
Ethiopia	ESS	2015-2016	5,469 hh	Two stage sampling. Probability proportional-to-size selection (of EAs) at the first stage of sampling
Zambia	LCMS	2015	12,260 hh	two-stage stratified cluster sample design. Probability Proportional to Estimated Size (PPES) at the first stage of sampling
Zambia	LFS	2014	11,520 hh	disproportionate allocation to strata

Exercise 4 - Sherlock Holmes

The goal of this exercise is to strengthen the students' awareness that any aggregated indicator originates from survey data, and that different surveys are used for different purposes (typically, labor force and poverty statistics are computed on the basis of Labor Force Surveys and Household Budget Surveys, respectively).

Exercise 5 - The Gemini Project

The goal of this exercise is to acquaint students with the Gemini Project, an important initiative launched to help redesign the US Consumer Expenditure Survey, addressing issues of measurement error and respondent burden. The Project has spurred considerable research effort, and many of the papers published under this umbrella contain findings that are relevant for other countries.

Exercise 6 – Historical Household Budget Surveys

This exercise is similar to exercise 3, but this time the focus is on considering the issue of harmonization of survey methods over time.

Lecture 4

Principles of questionnaire design

Learning objectives

The goal of the lecture is to review some general principles of questionnaire design, as they apply to the design of surveys to measure consumption. Principles include both general planning recommendations, and specific tips on formulating questions, applicable to any module of the survey.

Suggested preparation

Chapters 2, 3 and 5 of Grosh and Glewwe (2000) are a classic reference for the topics reviewed in this lecture, and are especially useful given their focus on LSMS surveys. The more recent Glewwe (2005) revisits some of the same material, in a more synthetic fashion. Familiarity with both works is assumed. Reviewing chapter 2 of Biemer and Lyberg (2003), and chapter 3 of Iarossi (2006) is also advised.

Time allocation

The survey process	20 min
Principles of questionnaire design	
Choice of topics and respondents	40 min
Break	15 min
Formulation of questions, order of questions, field testing	40 min
Lessons learned	5 min

Annotated lecture

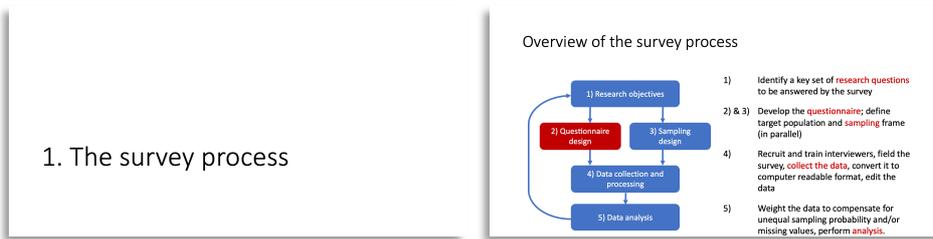
The annotated lecture consists of two slides. The first slide, titled 'Today's agenda', lists the following points:

- Overview of the **survey process**: where does questionnaire design fit?
- General principles of **questionnaire design**, with focus on the measurement of consumption
- Note: today is about **general principles**, whereas **specific guidelines** for each type of consumption expenditure will be provided in lectures 5-10.

The second slide, titled 'Useful readings – II', shows two book covers. The first is 'The Power of Survey Design' by Iarossi (2006), with a focus on Chapter 3: 'How Easy It is to Ask the Wrong Question'. The second is 'Introduction to Survey Quality' by Biemer and Lyberg (2003), with a focus on Chapter 2: 'The Survey Process and Data Quality'.

The set opens with the agenda for the day. We remind students that the lecture will focus on **general principles** underlying the questionnaire design stage, and that guidelines that apply to the design of individual modules will be discussed in upcoming lectures. We also

invite students to engage with the recommended readings, which provide the material covered in this lecture.



First, it is useful to put questionnaire design into **context**, among the different steps that make up the survey process. The design of the questionnaire is informed by research objectives, which in turn are the goal of data analysis: the idea that the last phase of the whole process – the actual use of the data – defines priorities at the very beginning of it is represented by the arrow that circles back from step 5 to step 1 in the figure.

The importance of questionnaire design
Joliffe (2001)

Measuring absolute and relative poverty: The sensitivity of estimated household consumption to survey design!

Diana Joliffe
Economic Research Service, US Department of Agriculture, 1800 M Street NE, Washington, DC 20036-5052
and William Zwinckel, Indiana University, 1300 University Blvd, Bloomington, IN 47405-1324, USA
d.joliffe@ers.usda.gov

The paper illustrates that questionnaire design significantly affects estimates of household consumption and absolute poverty. In a between-group designed experiment in 22 Indian states, seven detailed questionnaires were compared to an existing design. Household consumption for a 10 person group that ate chicken, ground rice, and vegetable soup was measured. The distribution of household consumption for the two groups was not significantly different. However, the distribution of the last questionnaire was significantly different from the first questionnaire. This difference in estimated consumption results in a system of absolute, relative poverty that the short questionnaire that a 40 person group from the village derived from the long questionnaire. In contrast, the level of relative poverty is unaffected by the change in questionnaire design. An implication of the paper is that modifications over time to questionnaires will result in spurious estimates of change in consumption and absolute poverty levels.

Impact on poverty indicators

Table 8
Poverty Indicators by Type of Questionnaire Using the National Poverty Line (Ngari)

	7-day recall		10-day recall		7-day diary		Test
	Mean	SE	Mean	SE	Mean	SE	χ^2 Level
Poverty	0.423	0.052	0.465	0.050	0.319	0.051	3.91 0.02
Poverty Gap	0.150	0.024	0.136	0.028	0.199	0.028	8.07 0.00
Squared poverty	0.079	0.014	0.053	0.010	0.100	0.010	11.46 0.00

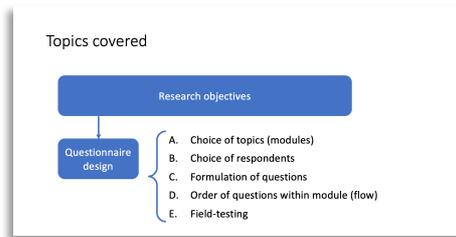
- Differences in questionnaire design are responsible for different poverty estimates
- A similar result applies to inequality estimates
- Questionnaire design matters

The papers cited in these slides, Joliffe (2001) and Backiny-Yetna et al. (2017), are here as an example of the fact that questionnaire design makes a great deal of difference to final estimates. Both papers will be discussed further in the next lecture, so there is no need to spend much time on this part. Suffice it to say that these are two examples of research comparing questionnaires that differ only in seemingly small and irrelevant features, and conclude that these features are responsible for large differences in the estimates of interest. The takeaway is: **details matter**.

So, is there a 'best design'?

- The short answer is negative: no single 'ideal' template has been designed and agreed upon.
- Many different survey design features are being tested by researchers, slowly building up **evidence** on 'what works'.
- The recent literature offers a number of practical **guidelines** (coming up in the next lectures).

At this point, seeing that there is no definite wrap-up of the findings of this literature, students may ask the question: which one among those tested is the 'best' questionnaire design? This slide gives the Instructor a chance to reiterate that the idea of a single, definitive 'best' design is a fallacy. Empirical evidence helps to define best practices for deciding among alternatives. Many of those best practices and recommendations will be reviewed in the next lectures.



This diagram shows the steps that a questionnaire design team follows when constructing a questionnaire. We will discuss them one by one.

Definition of research objectives
principles

- What research questions should the data answer?
- The answer influences the whole survey process
- Two priorities:
 - Organizing a survey design team
 - Formulating objectives as questions

Formulating objectives
Grosh (2005: 36-40)

- Useful to formulate objectives as questions, for instance:
 - "What proportion of the population is poor?"
 - "Has poverty increased or decreased over time?"
- Objectives are likely to multiply with consultations.
- Balance with constraints:
 - budget
 - capacity of the organization (experience, know-how)
 - respondents' willingness and ability to cooperate

First, a prerequisite for the whole process: the **definition of research objectives**. The Instructor should review the references cited in grey for more details.

Choice of topics
principles

- What pieces of information are needed to attain research objectives?
- Unfortunately, few general guidelines exist
- In practice, a popular solution is the **module approach** (typical of LSMS-type surveys): choosing modules (that is, topics), then moving to drafting each module.
- The **order** of modules matters: group together modules answered by same household member, and put **sensitive modules last**.

Choice of topics – an example
(Grosh, and Glewwe, 2000: 30)

<p>LSMS "core" modules</p> <ol style="list-style-type: none"> Household roster Housing Education Consumption Health Employment Transfers and other nonlabor income Metadata Prices Credit Agriculture 	<p> } Daily expenditures Food and fuel Non-food consumption Expenditures on private interhousehold transfers Durable goods </p>
--	--

After having defined the research objectives addressed by the survey, it is good practice to proceed by choosing broad **topics** to be included in the questionnaire. These two slides deliver general advice – adopting a **module approach** – and cite the LSMS survey template as an example.

Choice of respondents – I
principles

- Who should answer the questions?
- Answer: "the most knowledgeable person"
- Individual questions (e.g. employment)
 - Individuals themselves should answer
 - When forced to use "proxy respondent" (one person responds for another) consider recording who is answering on behalf of who
- Household questions
 - Household is asked to determine the "most informed individual", who will respond to questions (may change for different sub-modules)

Choice of respondents – II
in practice (Deaton and Grosh, 2000: 118)

- For expenditures, "most informed person" is a good approach, because it does not pre-judge division of labor in household (who does the shopping? who manages budget?)
- Has worked well when food is large share of budget, and when most of the household resources are pooled
- Problem: "There are expenditures on which no single person may have an accurate picture. (...) There may be items, such as clothing, that individuals purchase without any other household member knowing how much was spent."
- No easy fix: interviewing each household member individually on own expenditures is very expensive

Next, who in the household should be addressed as a **respondent**? It is recommended to select whomever is **most knowledgeable** on the question or module at hand. Deaton and Grosh (2000) point out that for some consumption items, it is unrealistic to expect a single person to have an accurate picture of each household member's own expenditures. In some cases, individual modules are indeed preferred (some examples will be cited in the next

lectures: one is food away from home) but the benefits of implementing such an approach should be weighted against the increased costs with respect to a ‘proxy respondent’ module.



Now is a good time for a break.

Formulation of questions
principles

- What to ask, exactly – and how?
- When developing a question, the designer should first of all put himself in the position of the typical, or rather the **least educated, respondent**.
- A good rule to remember in designing questions is that the respondent has probably not thought about these questions at the level of **detail required** by the survey.
- Details will be the subject of the next lectures. For now, a few **general principles** on two **specific issues**:
 1. Question wording
 2. Question type

What to ask, and how? The next lectures will address the ‘what’; for now, a few words of advice on the ‘how’ (formulation of questions).

Question wording
Iarossi (2006: p. 30-43)

- A number of studies have shown that changing even a single word in a question can significantly alter response distribution and accuracy.
- Useful checklist: the **“BOSS” principle**
- Four criteria should be followed when wording any question: **brief, objective, simple, and specific**

...

Question wording – II/II
explanation

- “During the past seven days, were you employed for wages or other remuneration, or were you self-employed in a household enterprise, were you engaged in both types of activities simultaneously, or were you engaged in neither activity?”
- **Long, unclear, and contains technical jargon**
- Possible fix: revising it as two separate questions that are **brief, simple, specific**
- 1. During the past seven days, did you work for pay for someone who is not a member of this household?
- 2. During the past seven days, did you work on your own account, for example, as a farmer or a seller of goods or services?

In terms of **question wording**, a useful rule to remember is the so-called ‘BOSS principle’: questions should be brief, objective, simple, and specific. A few slides explain each criterion with an example: students can be encouraged to try and give their own answers before we discuss ‘possible fixes’.

Question type
Grosh (2005: 45-46)

- A key decision is whether to make use of **open questions** (permitting respondents to answer in their own words) or **closed questions** (requiring respondents to select an answer from a set of choices).
- The use of **“closed questions”**, that is questions with **pre-coded answers**, is recommended
- Codes should be **mutually exclusive and collectively exhaustive**
- Coding schemes should be consistent across questions, e.g. if one question uses 1 for yes, 2 for no, then this should be maintained throughout the questionnaire, and should be clearly available for interviewer to consult

What’s wrong with my question?
Burgess (2001: 9)

What is your most usual means of travelling to college?

(Please tick one only)

Bus

Car

Bike

- This is an example of a **closed question**.
- Assume that the computer codes Bus by 1, Car by 2 and Bike by 3. If the respondent omitted to answer then this could be coded as 0 or some other missing value.

The choice between an open or closed question is about **question type**. Closed questions with pre-coded answers are most common in expenditure surveys: there are a few basic recommendations that help to design them correctly. The example reproduced in the ‘bus, car, bike’ question highlights the importance of well-designed response codes. The problem with this questions is that, although codes are mutually exclusive, they are not

exhaustive: there is no option for ‘none’, which leads to confusion between ‘none’ and missing.

Question type
Grosch (2005: 45-46)

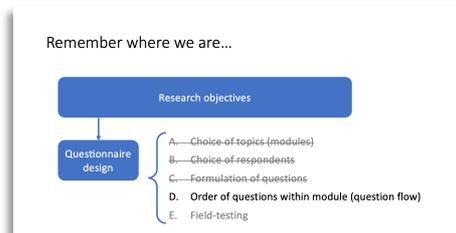
- Allow respondents to **answer on their own terms** as much as possible
- **Example:** “How much do you pay per month to rent your dwelling?”
- **Problem:** reference period is forced as 1 month. If giving amount per week or per year comes more natural to respondent, he/she is forced to convert, making room for mistakes
- **Possible fix:** “How much do you pay in rent for your dwelling?”, with option to associate pre-coded reference period, such as week, month, year, to declared amount.

Another recommendation on pre-coded question is to allow enough flexibility to accommodate the heterogeneity of respondents and reduce the cognitive burden of answering an unfamiliar question.

What happens when people are asked a question about which they have no relevant knowledge?

- In theory, respondents will say that they do not know the answer
- In practice, they may wish not to appear uninformed and may therefore give an answer to satisfy the interviewer.
- In order to reduce the likelihood of such behavior, some researchers have recommended that **don't know (DK) options** (or filters) routinely be included in questions.
- Do **DK filters** work? Evidence is mixed. Krosnick and Presser (2010: 282) argue that DK filters do *not* improve measurement.

This slide touches on ‘DK filters’. The Instructor is directed to the reference provided for more details.



We remind students how much ground was covered, and where we are in the survey process.

Question flow

- In **what order should questions be asked?**
- Early questions should be easy and pleasant to answer, and should build rapport between the respondent and the interviewer
- Flow should be tuned to logical reasoning **of the respondent**
- Related questions grouped together, minimize abrupt changes of topic
- “Filter” questions (**skips**) are important to minimize irrelevant questions
- **Sensitive questions last**

• • •

The Bank of Italy
SHIW

Q1. This is a list of different forms of saving and investment. Did the household have ... (form of saving or investment) on 31-12-2016?
(1=Yes or 2=No)

Q2. (SHOW CARD C25)
(For each form of saving or investment held on 31-12-2016)
What was the value on 31-12-2016? Answer using one of the ranges on this card.

The order in which questions are asked (the question ‘flow’) is also important. One of the main topics to have been discussed by the literature is that of sensitive questions, and how to ask them: the paper by Barton (1958) is a funny way to convey the issue. A common solution in modern surveys is the **unfolding brackets** approach, which is often used for questions on income.

Field-testing the questionnaire

- Pre-testing is the word
- Some evaluation methods require administration of the questionnaire to respondents, whereas others do not.
- The least structured evaluation method is [expert review](#), in which one or more experts critiques the questionnaire
- The most common form of pretest data collection — conventional pretesting — involves administering a questionnaire to a small sample of the relevant population under conditions close to, or identical to, those of the main survey.

Accuracy vs. Comparability

- Trade-off between following best practices and [improving the questionnaire](#), vs. ensuring [comparability](#) with previous data
- No easy solution. Incremental progress, when benefits from accuracy outweigh disadvantages of non-comparability

Finally, a few tips about pre-testing the questionnaire. The fact that no questionnaire is ever perfect, and that arriving to a good design is likely to be an incremental process, leads to some considerations about the benefits and pitfalls of accuracy (continually improving the design of the questionnaire) and those of comparability (adopting the same design as previous waves. These considerations are echoed in the final slide (lessons learned).

Homework

Exercise 1 – Engaging with the literature

See exercise 1, Lecture 1.

Exercise 2 – Question wording

Questions 1 and 3 are too subjective: they will not necessarily yield the present value of the house (according to the market), but how much it is worth according to the owner. Question 2 is worded in a much too technical way, with excessive jargon. Question 4 is just right: it is easy to understand, but at the same time grasps the right concept (what people would be *willing* to pay for the house). In fact, this question is often found in the housing section of many HCES. Question 5 is inaccurate: it does not ask for the net present value of the house, but for its original purchase price.

Lecture 5

Measuring food consumption: the foundations

Learning objectives

This lecture, together with lectures 6 and 7, has the goal of providing detailed recommendations on the collection of food consumption and expenditure data through household surveys. Lecture 5 focuses on the concepts of *acquisition* and *consumption*, and on determining the mode of data collection (*recall* vs. *diary*).

Suggested preparation

Much of the material covered in lectures 5, 6, and 7 is adapted from two references: FAO and The World Bank (2018), and Smith, Dupriez and Troubat (2014). The instructor is assumed to be familiar with both. This lecture also summarizes the evidence from a few important studies (Beegle et al. 2012, de Weerd et al. 2016, Backiny-Yetna et al. 2017), which the instructor is encouraged to review.

Time allocation

Introductory slides	5 min
Acquisition vs. consumption	
Concepts	20 min
Examples from questionnaires	15 min
Recommendations	10 min
Break	15 min
Recall vs. diary and length of reference period	
Concepts	20 min
Evidence from Tanzania and Niger	20 min
Recommendations	10 min
Lessons learned	5 min

Annotated lecture

The set opens with a brief introduction, which has the goal of re-acquainting the students to the main topic of the course – measuring consumption – after the brief detour taken during Lecture 4 on general principles of questionnaire design.

Where we stand

- What justifies our interest in collecting **data on food** consumption?
- **Food consumption expenditure** is a key component of any measure of **living standards** (lecture 1), **poverty** (lecture 14), and much more
- There are **additional research objectives**, which are useful to keep in mind when designing the food module of the questionnaire:
 - nutrition and food security
 - consumer price indices
 - Informing National Accounts
 - ...

We begin by saying that today's lecture is the first of a group of three lectures (5, 6, and 7) focused on **measuring food consumption**. What justifies such a strong focus on this topic? First and foremost, food consumption expenditure is a key component of any measure of living standards (lectures 1 and 2), and therefore, it is needed for poverty measurement. But there are several other important research topics that justify an interest in measuring food consumption with household surveys. We may 'wake up' the audience by soliciting answers (2 mins discussion might help). Either way, it should be clear that accurately measuring food consumption is important, for different types of data users and researchers.

Main references for this lecture
most useful also for the next two lectures



This and the following lectures go into the specifics of questionnaire design, and in particular, of the design of the food module. The main references for this group of lectures are the FAO/World Bank Guidelines on food data collection (left), and a recent paper by Smith, Dupriez and Troubat (right). The former illustrates best practices, while the latter is more focused on what countries actually do.

Questionnaire design challenges for food module

1. Acquisition vs. consumption
2. Recall vs. diary and length of reference period **today**
3. List of food items
4. Meal participation
5. Timing of visits
6. Food away from home
7. Non-standard measurement units

This slide is an outline of what is to come: the topics covered are the main 'challenges' that a questionnaire design team will have to face when implementing the food module of a household consumption and expenditure survey. The red rectangle indicates that Lecture 5

will deal with the first two topics (acquisition vs. consumption, and recall vs. diary), while the remaining issues will be tackled in Lectures 6 and 7.

1. Acquisition vs. consumption

Acquisition vs. consumption is our first topic.

<p>Definitions</p> <ul style="list-style-type: none"> ▪ Acquisition coming into possession, taking control of goods ▪ Consumption utilizing goods (i.e. eating, in the case of food) ▪ Mode of acquisition: <ul style="list-style-type: none"> ▪ purchase ▪ own-production ▪ in-kind receipt 	<p>Acquisition vs. consumption what to do with the chicken?</p> <ul style="list-style-type: none"> ▪ All goods that are consumed have been acquired in some way ▪ However, acquisition and consumption do not necessarily take place during the same reference period ▪ During a given period, say previous week, three possibilities: <ul style="list-style-type: none"> ▪ a chicken is acquired and eaten ($A = C$) ▪ a chicken is acquired, but not eaten ($A > C$) ▪ a chicken is eaten, but has been acquired earlier ($A < C$)
---	---

The distinction between acquisition and consumption is a conceptual one, regarding *what* exactly is captured by the food module. The first of the slides pictured above gives rigorous *definitions* of acquisition and consumption (and adds some terminology on *modes* of acquisition); the second slide, “What to do with the chicken?”, clarifies these concepts by means of an *example*. For more details, you can check Section 2 (p. 44) of Conforti, Grünberger and Troubat (2017), which elaborates further.

Why definitions matter

- Acquisition and consumption are measured for **different purposes**:
 - 1) Interest in consumption is justified by interest in estimating a number of things: standard of living, calorie intake, etc.
 - 2) Interest in acquisition is justified by interest in food security (availability)
 - 3) Interest in acquisition from purchases (i.e. food expenditure) is justified by CPI weighting, and informing national accounts
- Based on survey objectives, **concept(s)** of interest must be clear, and the **questionnaire** must be unambiguous

Does the distinction between acquisition and consumption really matter? Acquisition and consumption, as we have just seen, are different. Neither is intrinsically superior: each may be useful for different analytical purposes (the slide lists only a few examples). What is important is that the choice of what to measure be consistent with these objectives, and that the measure of interest be clear and unambiguous for both interviewers and respondents.

<p>Current practice Smith et al. (2014)</p> <ul style="list-style-type: none"> ▪ Smith et al. (2014) review 100 surveys from developing countries ▪ They find that both consumption and acquisition are commonly collected, but poor questionnaire design is common ▪ About 25% of surveys were found to include poorly worded questions, ambiguity 	<p>Approaches to data collection Conforti et al. (2017: 44)</p> <p>Typically, data on food are collected in one of three ways:</p> <ol style="list-style-type: none"> Acquisition Households report on food they acquired through purchases, own production and in-kind transfers. Actual consumption of the same food is not reported. Combination of acquisition and consumption Households report on food they acquired through purchases, without specifying the amount of food consumed. Food consumption derived from own-production or received from transfers is reported. Consumption Households report on food actually consumed, and on whether that same food was purchased, own-produced or received as a transfer.
--	---

We now transition **from concepts to practice**. What do household surveys around the world record: acquisition or consumption? The assessment by Smith et al. (2014) shows that both concepts are routinely measured in practice. However, poor questionnaire design is common: it will soon be clear what this means, specifically (slides on ‘questionnaire design issues’ are coming up). Conforti, Grünberger and Troubat (2017) summarize the main approaches to data collection found in practice (neither of which, again, is intrinsically superior): check page 44 of their paper for more details.

Common questionnaire design issues

Consider the following **examples**. Comment on each of them by answering these questions:

1. From collected data, could we estimate food consumption? Acquisition? Purchase? All of the above?
2. What about unit values?
3. Can you see any flaws in questionnaire design?

The idea of flaws in questionnaire design is introduced by some examples that students are encouraged to interact with. We suggest to ask students to comment on a few questionnaires, by answering three questions. Questions 1 and 2 (see slide above) are purely descriptive: would the data collected through each questionnaire be suitable to measure food consumption, acquisition, purchases, or unit values? These questions check whether the concepts covered in the previous slides are clear to students. Question 3 encourages students to think about possible flaws in the way the questionnaires are designed. More of these examples are used as homework for this lecture.

Rule out (or ‘filter’) question
Burundi (2006)

Food expenditures in the last 15 days

11. Did you consume (product)? Yes No → next product	12. Did you buy any? Yes No → q15	13. How much did you pay? (value)	14. How much did you buy? (quantity)	15. Did you harvest or take from stocks? Yes No	16. How much? (quantity)
FOOD PRODUCTS					
1. Haricot	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
2. Patate douce	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
3. Epinard a chair	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

This page is from **Burundi QUIBB 2006** (translated from French). With this questionnaire, it is impossible to fully capture either consumption or acquisition. If the objective is to measure *consumption*, after the filter question (question 11) data are collected on food acquired (purchased, harvested) rather than food consumed. When this occurs, quantities and expenditures on food acquired include those entering into the households’ stocks – not the household pantry for immediate consumption – and are systematic overestimates of food consumed from home production. If the objective is to measure *acquisition*, then question in M11 rules out any following question for food that were acquired but not consumed in the reference period. These are flaws in questionnaire design, because they lead to incomplete enumeration of both acquisition and consumption.

On the other hand, *unit values* can be computed based on this module: it would be sufficient to divide question 13 by question 14 (provided local units can be converted into standard units).

Usual month?
Madagascar (2005)

MONTHLY CHARACTERISTICS		GIFTS	
6-	7-	8-	9-
In the past 12 months, how many articles have you had expenditures on this... [ARTICLE] YES or NO	How much have you spent on average per month on this... [ARTICLE] [ARTICLE] [ARTICLE] [ARTICLE]	Quantity purchased of this... [ARTICLE], on average per month, in the unit... [ARTICLE] [ARTICLE] [ARTICLE]	What is the total quantity per month of the gifts you have received as gifts over the past 12 months? Same unit as in Q6
MONTH	AVG (average per mo)	QUANTITY	UNIT
		QUANTITY	NUMBER

Average per month

This example is from the **Madagascar** Enquete Permanent Menages (EPM) 2005. The questionnaire captures *acquisition*, rather than consumption, although in this page we cannot see any question on own-produced foods, which would therefore be missing from the total amount and value of acquisitions.

Unit values can be obtained as question 6 divided by question 7.

A problem with the design of this questionnaire, besides the apparent lack of questions on own-production, is that it is not clear if the “average per month” in questions 6 and 7 refers to the average of the past 12 months, or the average over those months in which consumption actually took place. A better approach would have been to clarify the ambiguity, or even to get rid of the “usual month” approach (more on this later).

Only acquisition
Iraq (2012)

SECTION 12: DIARY OF FOOD AND RECURRING NON-FOOD COMMODITIES

DATE: _____ NUMBER OF LINES COMPLETED: _____ DAY: _____

ITEM DESCRIPTION	Q1201	Q1202	Q1203	Q1204	Q1205	Q1206	Q1207	Q1208	Q1209	Q1210	Q1211	Q1212	Q1213	Q1214	Q1215	Q1216	Q1217	Q1218	Q1219	Q1220	Q1221	Q1222	Q1223	Q1224	Q1225	Q1226	Q1227	Q1228	Q1229	Q1230	Q1231	Q1232	Q1233	Q1234	Q1235	Q1236	Q1237	Q1238	Q1239	Q1240	Q1241	Q1242	Q1243	Q1244	Q1245	Q1246	Q1247	Q1248	Q1249	Q1250	Q1251	Q1252	Q1253	Q1254	Q1255	Q1256	Q1257	Q1258	Q1259	Q1260	Q1261	Q1262	Q1263	Q1264	Q1265	Q1266	Q1267	Q1268	Q1269	Q1270	Q1271	Q1272	Q1273	Q1274	Q1275	Q1276	Q1277	Q1278	Q1279	Q1280	Q1281	Q1282	Q1283	Q1284	Q1285	Q1286	Q1287	Q1288	Q1289	Q1290	Q1291	Q1292	Q1293	Q1294	Q1295	Q1296	Q1297	Q1298	Q1299	Q1300	Q1301	Q1302	Q1303	Q1304	Q1305	Q1306	Q1307	Q1308	Q1309	Q1310	Q1311	Q1312	Q1313	Q1314	Q1315	Q1316	Q1317	Q1318	Q1319	Q1320	Q1321	Q1322	Q1323	Q1324	Q1325	Q1326	Q1327	Q1328	Q1329	Q1330	Q1331	Q1332	Q1333	Q1334	Q1335	Q1336	Q1337	Q1338	Q1339	Q1340	Q1341	Q1342	Q1343	Q1344	Q1345	Q1346	Q1347	Q1348	Q1349	Q1350	Q1351	Q1352	Q1353	Q1354	Q1355	Q1356	Q1357	Q1358	Q1359	Q1360	Q1361	Q1362	Q1363	Q1364	Q1365	Q1366	Q1367	Q1368	Q1369	Q1370	Q1371	Q1372	Q1373	Q1374	Q1375	Q1376	Q1377	Q1378	Q1379	Q1380	Q1381	Q1382	Q1383	Q1384	Q1385	Q1386	Q1387	Q1388	Q1389	Q1390	Q1391	Q1392	Q1393	Q1394	Q1395	Q1396	Q1397	Q1398	Q1399	Q1400	Q1401	Q1402	Q1403	Q1404	Q1405	Q1406	Q1407	Q1408	Q1409	Q1410	Q1411	Q1412	Q1413	Q1414	Q1415	Q1416	Q1417	Q1418	Q1419	Q1420	Q1421	Q1422	Q1423	Q1424	Q1425	Q1426	Q1427	Q1428	Q1429	Q1430	Q1431	Q1432	Q1433	Q1434	Q1435	Q1436	Q1437	Q1438	Q1439	Q1440	Q1441	Q1442	Q1443	Q1444	Q1445	Q1446	Q1447	Q1448	Q1449	Q1450	Q1451	Q1452	Q1453	Q1454	Q1455	Q1456	Q1457	Q1458	Q1459	Q1460	Q1461	Q1462	Q1463	Q1464	Q1465	Q1466	Q1467	Q1468	Q1469	Q1470	Q1471	Q1472	Q1473	Q1474	Q1475	Q1476	Q1477	Q1478	Q1479	Q1480	Q1481	Q1482	Q1483	Q1484	Q1485	Q1486	Q1487	Q1488	Q1489	Q1490	Q1491	Q1492	Q1493	Q1494	Q1495	Q1496	Q1497	Q1498	Q1499	Q1500

Quantity acquired

Value

Mode of acquisition

This page is from the **Iraq** Household Socio Economic Survey 2012. The questionnaire only records food *acquisition* and not consumption. However, all modes of acquisition are recorded, and there are no apparent flaws in the design of the questionnaire. *Unit values* could be computed as question 1203 divided by 1202 (provided that non-standard units are convertible into kg).

Both acquisition and consumption
Lebanon (2012)

Did you purchase...		Did you consume...		Mode of acquisition of consumed food	
Q1401	Q1402	Q1403	Q1404	Q1405	Q1406
quantity	value	quantity	value	quantity	value
Q1401	Q1402	Q1403	Q1404	Q1405	Q1406

This page is from the **Lebanon** Household Budget Survey. This questionnaire partially records food *acquisition*: it focuses on *purchases* only, and does not record food acquired through other sources. On the other hand, the questionnaire captures food *consumption* in full: question 7 allows to record all sources of food consumed. There are no flagrant mistakes in questionnaire design, assuming that the goals of the survey did not include

measurement of total food acquisition. Again, *unit values* may be computed both from purchases and consumption, by computing the ratio between reported value and quantity.

Common questionnaire design mistakes
Evidence from Smith et al. (2014: 14-15)

1. Acquisition surveys: **filter question** on something else (18%)
2. Routine month surveys: ambiguity about whether respondents should report on the **routine month** in the recall period, or only those months in which the food item in question is consumed (13%)
3. **Ambiguity** on whether to report on acquisition or consumption (7%)
4. Data collected on **food harvested** rather than **food consumed** from home production (3%)

This slide summarizes the most common issues found in household survey questionnaires in developing countries, as listed by Smith et al. (2014). Some of these were exemplified by the questionnaires discussed with students (point 1 can be seen in the Burundi example, point 2 in Madagascar, point 4 in Burundi again). These should be seen as a catalogue of the most common mistakes to be avoided when designing the food module.

Should we collect data on acquisition or consumption?

- It depends on the purpose of the survey (lecture 4)
- **Welfare analysts** would want **consumption** (lectures 1-2)
- **Statisticians** (and others) are also interested in **acquisition** to construct weights for their CPIs
- Conforti et al. (2017) evidence from 81 recent surveys says that the difference in estimated mean acquisition and mean consumption is small, but acquisition is much more variable

Provided that mistakes and incomplete enumeration are avoided, one question remains: **is it best to record acquisition, or consumption?** The answer depends on the goals of the survey. Based on the conceptual framework covered in Lectures 1 and 2, we could argue that welfare analysts prefer food *consumption*. However, recording food *acquisition* is important for other purposes, the most common of which is probably the computation of consumer price index (CPI) weights. The evidence presented by Conforti et al. (2017) is reassuring: the two measures are not too distant, so that one may be used in place of the other, in case of necessity. A common approach in practice, as we have seen, is to record both.

Recommendations
FAO and WB (2018: 53-55)

1. Always collect data on all **modes of acquisition** (**purchase, own-production, in-kind receipts**), irrespective of whether focus is on amount consumed or acquired.
 - If questionnaire only records food obtained through some sources (such as purchases) there will be **underestimation** of both consumption and acquisition
 - Pay special care to **in-kind receipts** that are likely to be missed, such as payments for labor and social programs,
 - Be careful not to **duplicate information** captured in other modules (e.g., employment or social assistance)



Recommendations
FAO and WB (2018: 53-55)

3. Avoid sources of **incomplete or ambiguous enumeration**
 - Do not use **filter questions** on consumption to rule out acquisition (and vice versa)
 - Avoid filter questions that focus on food purchases
 - For own-production, the question must be worded to clearly indicate food consumed from **own-production** rather than **food harvested**. If not, values reported may include food entering the household's production stocks (that is, not for immediate consumption).

A set of recommendations closes this topic. Recommendations are adapted from section 3.3 of the FAO/World Bank Guidelines (2018): further details are found there.



Now is a good time for a break.

2. Recall vs. diary and length of reference period

Recall vs. diary, and the determination of the length of the reference period is our next topic for the lecture.

Definitions: recall and diary

Data on household food consumption (or acquisition) commonly collected in one of two ways:

1. Respondents are interviewed and asked to **recall** consumption during a specified period (past week, past month...).
2. Households are asked to keep a **diary** over a reference period (days, weeks...) and record consumption at the moment it takes place.

...

Example of diary
Kenya integrated household budget survey 2015, seven-day reference period

| DATE |
|------------|------------|------------|------------|------------|------------|------------|------------|
| 2015-01-01 | 2015-01-02 | 2015-01-03 | 2015-01-04 | 2015-01-05 | 2015-01-06 | 2015-01-07 | 2015-01-08 |
| 2015-01-09 | 2015-01-10 | 2015-01-11 | 2015-01-12 | 2015-01-13 | 2015-01-14 | 2015-01-15 | 2015-01-16 |
| 2015-01-17 | 2015-01-18 | 2015-01-19 | 2015-01-20 | 2015-01-21 | 2015-01-22 | 2015-01-23 | 2015-01-24 |
| 2015-01-25 | 2015-01-26 | 2015-01-27 | 2015-01-28 | 2015-01-29 | 2015-01-30 | 2015-01-31 | 2015-02-01 |

First, the slides give definitions for diary and recall, recall period and reference period. Two examples are used to substantiate the difference between diary and recall.

Diary or recall?

- Which approach is better, in terms of the quality of collected data?
- Both methods have **pros and cons**
- In particular, they both have the potential to generate **measurement error**, for different reasons
- Let us see how to evaluate the risks associated to either choice

Which of the two approaches guarantees the collection of quality data? The answer is not easy: both diary and recall surveys have the potential to generate **measurement error**. Empirical evidence helps to evaluate the risks implied by each of the two methods. The next slides illustrate.

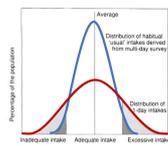
Problems with recall

- Memory can fail: **biases** related to length of recall period
- **Long recall period**
 - Tendency to forget, or **memory decay**
 - More likely if expenditure is perceived as ordinary, **not salient**
 - Leads to **under-reporting** of consumption
- **Short recall period**
 - **Telescoping**: tendency to mistakenly report consumption that has actually taken place outside the recall period
 - More likely if expenditure is perceived as extraordinary, **salient**
 - Leads to **over-reporting** of consumption

The most problematic aspect of using **recall** methods is related to memory and its failures: respondents' inability to remember correctly leads to biases in data collected through recall surveys. Depending on the length of the recall period, different types of bias may arise. *Memory decay* and *telescoping* are related to long and short recall periods, respectively. For a detailed explanation of these biases, see the discussion at pages 110 and 111 of Grosh and Glewwe (2000) (the 'Consumption' chapter by Deaton and Grosh).

Length of recall period and shape of the distribution

Rosalind Gibson (2005: 139)



- Suppose we interview individuals on food intake using a **multi-day recall**. The distribution of food intake is the **blue curve**.
- Now suppose we interview the same individuals using a **1-day recall**. We obtain the distribution of intakes represented by the **red curve**.
- Due to the short recall, total variance will be higher, and the distribution will be "wider and fatter".
- Conclusion: a **short recall overestimates variance**, that is exaggerates the incidence of both inadequate and excess intake.

Another issue that emerges when using recall surveys, and that is related to the **length of the recall period**, is that of the variability of reported consumption. Shorter recall periods record 'episodic' consumption (consumption that takes place on a given day or week) which may be randomly higher or lower than usual. These anomalies are smoothed out by longer recall periods. As a result, data recorded with short recall periods are *more variable* than those recorded with longer ones. This slide illustrates the point by displaying the distribution of reported food intake from a short recall (red curve) and a long recall (blue curve). The two curves have the same mean (on average, both estimates are correct) but the red curve has higher variance. This is a problem if data users are interested in the 'tails' of the distribution of consumption – households with very low calorie intakes, for instance – because the likelihood of extremely high and extremely low will be overestimated by surveys using short recall periods. Crossley, Fiedler and Mwangi (2016) discuss this issue in rather technical terms on page 19 of their IFPRI Discussion Paper.

Problems with diary

- In principle, **diary avoids memory fails**, as it is compiled close to the moment in which event (consumption or purchase) occurs
- In practice, diary keeping introduces other problems:
 - **Respondent burden and fatigue**, particularly when length of diary increases: evidence of "diary exhaustion" (Brzozowski, Crossley and Winter 2017; Gibson 2013)
 - To reduce these issues, high levels of supervision are needed, which imply **high implementation costs** (FAO study of Bangladesh 2010 HIES showed good results with enumerator visits every two or three days)

Using **diary** surveys to collect data on food consumption may seem like a solution to the memory biases that characterize recall surveys: with diaries, respondents note down their consumption as soon as it happens. In practice, however, the diary approach introduces

other problems, mainly linked to respondent burden and fatigue, and to the high implementation costs needed to minimize them.

Alternative methods are unsatisfactory – I/II
Usual month approach

- Respondents are asked to report consumption for the “usual month” during the previous year
- Advocated by Deaton and Grosh (2000) to capture typical consumption
- At best, it is not more effective than simple recall; at worst, it introduces errors related to education of respondents, due to cognitive burden (Fiedler and Mwangi 2017: 25; Friedman et al. 2017)

Alternative methods are unsatisfactory – I/II
Usual month approach

- Respondents are asked to report consumption for the “usual month” during the previous year
- Advocated by Deaton and Grosh (2000) to capture typical consumption
- At best, it is not more effective than simple recall; at worst, it introduces errors related to education of respondents, due to cognitive burden (Fiedler and Mwangi 2017: 25; Friedman et al. 2017)

Some authors have suggested **alternative methods**, a ‘third way’ besides diary and recall. However, empirical evidence has failed to support these alternatives.

Do these “details” matter?

- Large body of evidence finds that the choice between **diary and recall**, and of the **length of recording periods**, can **significantly affect results**
- Important papers that studied the impact of survey methodology on consumption and poverty statistics:
 - SHWALITA study in Tanzania (Beegle et al. 2012, Gibson et al. 2015, de Weerd et al. 2016)
 - Niger (Bakiny-Yetna et al. 2017)

Now that the trade-offs of the diary and recall approaches have been covered, one may ask: **do these details really matter?** The short answer is yes. The following slides summarize results from a few notable experimental studies that tested the effects of different questionnaire designs on estimated food expenditure and other outcomes.



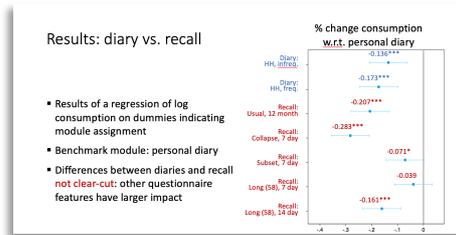
Table 1
Survey experiment consumption modes

Mode	Consumption measurement	Food source	Households
1	Long list (58 food items) 14 day	Quantity from purchases, own production, and gifts/other sources; Tabling value of consumption from purchases	504
2	Long list (58 food items) 7 day	Quantity from purchases, own production, and gifts/other sources; Tabling value of consumption from purchases	504
3	Short list (17 food items; subset of 58 food items) 7 day	Quantity from purchases, own production, and gifts/other sources; Tabling value of consumption from purchases	504
4	Call-based list (11 food items covering 90% of food consumption) 7 day	Quantity from purchases, own production, and gifts/other sources; Tabling value of consumption from purchases	504
5	Long list (58 food items) Usual 12 months	Consumption from purchases; number of months consumed; quantity per month; Tabling value per month; Consumption from own production; number of months consumed; quantity per month; Tabling value per month; Consumption from gifts/other sources; total estimated value for last 12 months	504
6	Household diary, frequent visits 14-day diary		503
7	Household diary, infrequent visits 14-day diary		503
8	Personal diary, frequent visits 14-day diary		503

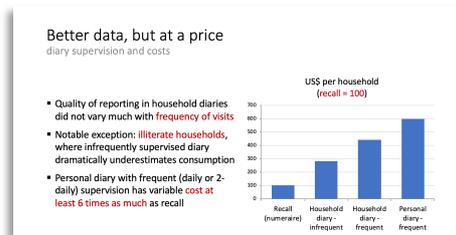
Notes: Frequent visits resulted data visits by the food amount and visits every other day by the survey enumerator for the duration of the 2-week diary. Infrequent visits entail 1 visit to deliver the diary (day 1), to pick up week 1 diary and drop off week 2 diary (day 8), and to pick up week 2 diary (day 15). Households assigned to the infrequent diary but who had no longer responses (about 10% of the 503 households) were visited every other day by the food amount and the enumerator.

Food list items are divided into three groups based on frequency of purchase. Frequently purchased items: flour, rice, beans, lentils, pulses, peas, soybeans, sugar, cooking oil, soap, kerosene, salt, tea, coffee, tobacco, and other. Infrequently purchased items: maize, sorghum, millet, cowpeas, green beans, chickpeas, lentils, pulses, peas, soybeans, and other. Non-purchased items: alcohol, tobacco, and other.

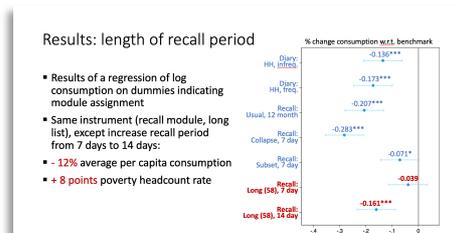
First is **Beegle et al. (2012)** with an experimental study conducted in Tanzania. The instructor is encouraged to read the paper: the slides offer a quick summary of the main findings, as they relate to the topics of Lecture 5. The study implemented eight different questionnaire designs, assigned them randomly to 4,000 households, and then compared average food consumption expenditure resulting from each of the questionnaires. One design – a personal diary with frequent visits – is considered the most accurate way of collecting data on food consumption, and is therefore used as a benchmark, against which all other questionnaire designs are stacked up.



This slide shows the **comparison of diary and recall** questionnaires, by using a type of chart that is going to appear several more times in the lecture. *The chart requires some explanation in order to make its interpretation clear, but it is an important one, and should be given as much time as needed.* It makes a crucial point: we can find answers to questions on ‘best’ questionnaire designs, if we look to the evidence. Allowing students to become comfortable with this literature, and its findings, is the real added value of the lecture. Each design is compared to the benchmark (personal diary, frequent visits). The chart plots the coefficients of a regression of log consumption on dummies indicating assignment to questionnaire designs. Asterisks indicate statistical significance: *** significant at 1%, ** at 5% level. No asterisks indicate no statistical significance (at traditional significance levels). Interpretation of coefficients is as follows (take the first line as an example): average food consumption expenditure obtained with a household diary with infrequent visits is 13.6% lower than the value obtained with a personal diary with frequent visits (the benchmark design). Now, if we compare the coefficients for the two different *diary* designs (in blue), to those for the five different *recall* designs (in red), we note that results are mixed. All numbers are negative, meaning that all alternative designs return a lower average food expenditure than the benchmark; but *how much* lower seems to depend on characteristics of the design *other than* the choice of diary vs. recall.



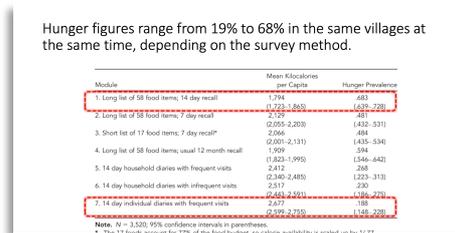
Other selected results of the study concern implementation costs for different designs: diaries give better results when closely supervised, at least for less-educated households; however, they are also much more expensive to implement than recall modules, as shown by the bar chart).



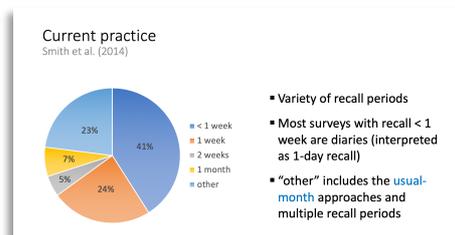
Another important finding concerns the length of the recall period: comparing the final two designs in the chart (in red), that only differ in recall length, shows that the 7-day recall comes much closer to the benchmark than the 14-day recall. In other words, if we increase the recall period from 7 to 14 days, all else being equal, the average reported consumption falls by 12% (we obtain this as the difference between the two coefficients corresponding to the 7 and 14-day recall, 0.161-0.039). Elsewhere in the paper (page 14) the author reports that this also corresponds to an increase of 8 percentage points of the poverty headcount rate.



Backiny-Yetna et al. (2017) is another experimental study, conducted in Niger. In this case, the comparison between diary and recall gives somewhat unexpected results: a 7-day recall module returns higher estimated food expenditure than a 7-day diary. This is reflected by poverty estimates (the poverty headcount is higher according to the diary module). Respondent burden is likely to be responsible for the low results obtained with the diary, which suggests that supervision matters.



Finally, De Weerd et al. (2016) use the same data as Beegle et al. (2012), but compare results on nutrition. In this case, recall questionnaires yield lower estimated calorie intakes than diary designs.



To close the discussion of empirical evidence, we once again turn to the review of the current practice by Smith et al. (2014), to see which approaches are actually implemented in practice. In this graph, diaries are interpreted as surveys using a recall period of 1 day. Note that using different recall periods for different foods ('other' in the graph) is common.

Recommendations

FAO and WB (2018: 50-53)

1. While a **diary approach** may be the “gold standard” with close supervision and careful implementation, it is not suitable for resource-constrained statistical offices in low- and middle-income countries
2. Low-income countries are advised to adopt **recall interviews and a 7-day recall period**, as this method provides a good balance between accuracy and cost-effectiveness
3. Any survey using diary methods must be closely supervised to ensure compliance. The reference period should not exceed **2 weeks**.

Recommendations

FAO and WB (2018: 50-53)

3. The “**usual month**” approach should not be used.
4. Any change in recall period or data collection method (diary vs. recall) should be accompanied by an **experimental component** aimed at assessing the change in survey estimates.
5. The evidence in Beegle et al. (2012), De Weerd et al (2016), and Backiny-Yetna et al. (2017) will hopefully serve as a useful reminder.

Recommendations are adapted from Section 3.1 of FAO and World Bank (2018).



Lessons learned

- Quality data on **food consumption** are **crucial** for several research objectives, living standards measurement being one of them
- **Questionnaire design matters**: large impact on final results
- This lecture has explored some **foundational choices** in the design of the food module:
 - Should we measure **consumption or acquisition**?
 - Should we use **diary or recall**? How should the reference period be set?
- **Experimental evidence** provides guidance (recommendations 1-5 in the previous two slides).

The final slide summarizes the main takeaways of Lecture 5: it is important to take some time for this final recap, as it provides an opportunity to cement the students’ understanding of topics covered, and take any final questions.

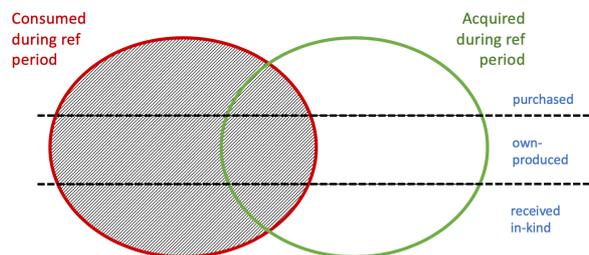
Homework

Exercise 1 – Engaging with the literature

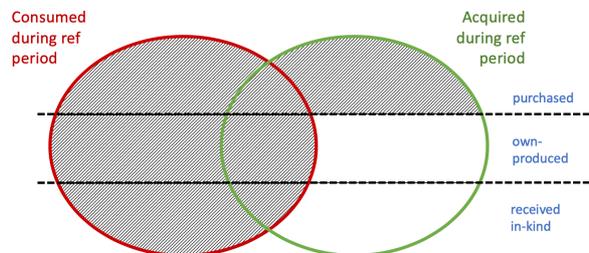
See exercise 1, Lecture 1.

Exercise 2 – Acquisition vs. consumption

Example 1 allows to estimate the total value of consumption. All three modes of acquisition are accounted for. The corresponding diagram is pictured below.



Example 2 also allows to estimate the total value of consumption. However, question C02 asks for the *total* value of purchases during the last 7 days (not just those going towards consumption): this makes it possible to estimate at least part of the value of acquisitions, precisely the portion acquired through purchases. The corresponding diagram is pictured below.



Lecture 6

Measuring food consumption: questionnaire design

Learning objectives

This lecture, together with lectures 5 and 7, has the goal of providing detailed recommendations on the collection of food consumption and expenditure data through household surveys. Lecture 6 focuses on the definition of the list of food items for which information is collected, on why and how to record meal participation, and on the timing of visits to interviewed households.

Suggested preparation

Much of the material covered in lectures 5, 6, and 7 is adapted from two references: FAO and The World Bank (2018), and Smith, Dupriez and Troubat (2014). The Instructor is assumed to be familiar with both. This lecture also summarizes the evidence from several studies (Jolliffe 2001, Beegle et al. 2012, Pradhan 2009, Gibson and Rozelle 2002, Troubat and Grunberger 2017), which the Instructor is encouraged to review.

Time allocation

Introductory slides	5 min
List of food items	40 min
Meal participation	
Concepts: what partaking is and why it matters	10 min
Break	15 min
Evidence and examples, recommendations	20 min
Timing of visits	25 min
Lessons learned	5 min

Annotated lecture

After a brief introduction, aimed at reminding students the general framework in which the current lecture is inscribed and the outline for the day, we dive into the first topic: determining the list of items for the food module.

3. List of food items

Length of the food list

- Designing the food list in all of its details is a daunting task:
 - How many different foods should be included?
 - Should some items be grouped together? Which ones?
 - Should all foods be listed, including "difficult" items like prepared foods?
 - ...
- Answers to these questions determine the **length** of the food list, which in turn **influences final results** (evidence in next slides)

The goal of these first slides is to define what is meant by ‘**food list**’, and to convey the fact that defining it from scratch is a difficult, non-trivial, and important task. The decisions that go into the definition of the food list are sometimes just summarized as “determining the length of the list”.

Comprehensiveness vs. specificity

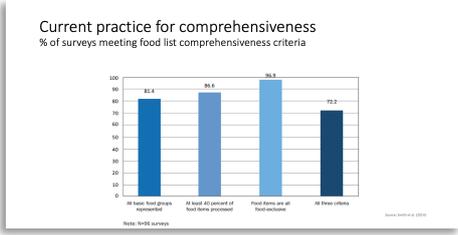
The length of the food list is actually the result of two distinct design choices:

1. **Comprehensiveness**
Whether or not all types of foods and beverages that make up the diet of the surveyed population are represented in the food list
2. **Specificity**
The degree of detail and disaggregation of the food list

It helps to think about the ‘length’ of the food list as resulting from two separate choices: the level of **comprehensiveness** of the list, and its **specificity**.

Comprehensiveness

- It is considered good practice that the food list be **as comprehensive as possible**
- By definition, excluding entire categories of foods leads to **under-estimation** of consumption
- How to **evaluate** comprehensiveness?



Comprehensiveness is, in a way, uncontroversial: the more comprehensive the list, the better. For this reason, the point is discussed quickly. Smith et al. (2014) set useful **criteria** that help to evaluate the level of comprehensiveness of the food list, and their assessment shows that most questionnaires perform well in this dimension.

Specificity

- A **detailed food list** should help respondents remember consumption more completely and accurately, a certain level of detail is also required to obtain accurate nutritional data (difficult to estimate calorie intakes from heterogeneous food aggregates)
- But the food list can be **too detailed**, and risk increasing respondent and enumerator fatigue.
- Take the case of **bananas** in Uganda...

Bananas in Uganda

- A total of 95 banana varieties are currently grown in Uganda (FPRI, 2006).
- Each variety has a local name.
- Specificity involves **trade-offs**: it is not always true that the more specific the food list, the higher the quality of the data.

Unlike comprehensiveness, **specificity** – the level of detail, or disaggregation, of the food list – has pros and cons. According to the literature, both too specific and too aggregated food lists can lead to measurement error. The case of Uganda, and its many varieties of bananas, is cited as an example: would it be a good idea to ask about consumption of all 95

varieties, specifically, in the food module? The Instructor can engage students for a couple of minutes on this, just to liven up the discussion.

Recap

- **Comprehensiveness**
It is required for the production of reliable data. The literature offers some criteria to check that food categories are adequately represented in the food list
- **Specificity**
There is a widely acknowledged trade-off involving the level of detail of the food list, but the optimal balance depends on the local context. The literature offers some general rules that guide the compromise

• • •

Empirical evidence on specificity

- **Pros**
highly aggregated food lists are linked to underreporting of consumption
 - Jolliffe (2001), Beegle et al. (2012), Pradhan (2009), Statistical Institute and Planning Institute of Jamaica (1996)
- **Cons**
longer food lists push enumerators and respondents to reduce compliance
 - Deaton and Gosh (2000), Finn and Ranchhod (2015), Statistics Indonesia and World Bank (2014)

At this point, we once again bring in the empirical evidence. The studies cited make the point that specificity matters, in that it greatly affects final estimates. The example from El Salvador (Jolliffe 2001) shows that there was a 27% increase in median reported consumption when the specificity of the food list was increased. The SHWALITA study of Tanzania is cited again, as specificity was among the many design features tested. The chart with the blue and red coefficients has been introduced and explained in lecture 5, and can be interpreted in the same way (the Instructor is encouraged to revisit it). In this case, a more specific food list caused a 23% increase in average reported consumption (this is obtained as 0.283-0.039).

Recommendations – 1/4
Comprehensiveness

1. All major food groups should be represented
2. There should be adequate representation of processed foods (including prepared meals), when these are part of the population's diet
3. List should be kept up to date, to take into account changing dietary habits

• • •

Recommendations – 4/4
Specificity

9. After a reasonable number of items to be listed for each food group has been selected, a residual category (e.g. "other fruit", "other vegetables") may be added if relevant; it is best if such categories remain marginal, as they do not allow the collection of data on quantities or the computation of nutrient intakes
10. Adoption of a food classification system can help in meeting all previous criteria. For many of the basic purposes of household consumption and expenditure surveys, the recommended standard of classification is COICOP.

The final **recommendations** give general guidelines to construct the food list. Section 3.6 of the FAO/World Bank Guidelines can be checked for further details.

4. Meal participation

• • •

Why it matters

- "The adequacy of the consumption of the household's food can be divided into two issues: **how much** food is being consumed and **who** is consuming it." (Fiedler and Mwangi, 2016: 47)
- **Per capita** measures of food consumption should be based on the **number of people sharing meals**

$$\text{per capita consumption} = \frac{\text{household total consumption}}{\text{household size} - \text{absent members} + \text{additional partakers}}$$

Next, the lecture tackles the issue of adjusting for **meal partakers**. The goal of the first slides is to give a definition of partakers, and to motivate our focus on the issue, by conveying the message that knowing the exact *denominator* of per capita consumption (*how many people are sharing the total?*) is fundamental to obtaining an accurate estimate of individual welfare (be it consumption, expenditure, or nutrition).



Now is a good time for a break.

Evidence on the impact of partakers

- Accounting for partakers **reduces inequality of consumption**
- Bouis, Haddad, and Kennedy (1992) and Bouis (1994) show that the difference between mean calorie intakes of the poorest and richest quartiles is much lower when partakers are accounted for (Kenya and the Philippines)
- Gibson and Rozelle (2002) finds similar evidence (Papua New Guinea)

% difference of per capita kcal intake after accounting for partakers, by expenditure quartile



Empirical evidence is brought in to show that **adjusting for partakers matters** in practice – that is, dividing total household consumption, or total household calorie intake, by the number of meal partakers instead of simply the number of household members, significantly changes the distribution of per capita intakes. With the adjustment, estimated per capita consumption decreases among the rich, and increases among the poor. The graph taken from Gibson and Rozelle (2002) shows the percentage difference between unadjusted (i.e. per household member) and adjusted (i.e. per partaker) per capita calorie intake, for each quartile of per capita expenditure: the difference is positive for the poorest households (the adjustment has increased their per capita intake) and negative for the richest (the adjustment has decreased their intake). This is a key reason why the adjustment is needed, given that estimating the per capita *distribution* of consumption (not just the average) is one of the main uses for household consumption and expenditure survey.

Current practice

- Assessment of 81 recent surveys by Fiedler and Mwangi (2016)
- Most commonly, surveys **do not** collect information on meal partakers
- When they do, approaches are **heterogeneous**
- Lack of research to tell us what works

In contrast to the importance of collecting information on partakers, Fiedler and Mwangi (2018) show that most countries fail to collect data on this aspect, and those that do have very heterogeneous approaches.

Some examples

Afghanistan 2007 **Mongolia 2007/08**

SECTION 15: FOOD CONSUMPTION IN LAST 7 DAYS

15.1 How many household members were outside and/or at work during the household during the last 7 days?
 people

15.2 How many meals were eaten by guests from the household cooking pot during the last 7 days? (How many guests at the house on the last 7 days?)
 person-meals

15.3 How many visitors stayed with your household for the last 7 days?
 (Specify number of visitors) Number
 How many days? person-days
 (Number of visitors multiple by days stayed here)

On average, how many people were present in the last 7 days? In this section children are defined as less than 18 years.
 Household Members

Male adults	Female adults	Male children	Female children
-------------	---------------	---------------	-----------------

Iraq 2012 non-members

Table 2: Number of participants other than the household members in meals within the household while the number of people who shared every meal (only include household members)

DAY NUMBER	DAY 1		DAY 2		DAY 3		DAY 4		DAY 5		DAY 6		DAY 7	
	B	L	B	L	B	L	B	L	B	L	B	L	B	L
TYPE OF MEAL	BREAKFAST	LUNCH												
Number of individuals other than household members sharing meals within the household expenditure														

These slides show some miscellaneous **examples**, in order of increasing complexity of the partakers module. For the time being, there is no need to comment on which approach is expected to work better: the examples are a way to give more substance to the idea of collecting data on partakers, and also to show how questionnaires can vary widely in practice.

Heterogeneity of approaches
Smith et al. (2014: 32)

Method	Percent of surveys
Data are collected on the presence and/or household meal consumption of non-household members during the recall period	15
Data collected on the number of visitors in the household	11
Data collected on visitors' length of stay	5
Data collected on the number of meals consumed by visitors/guests	10
Data collected by type of meal (breakfast, lunch, dinner)	7
Data collected on the age of visitors/guests	7
Data collected on the sex of visitors/guests	6

Note: N=120 surveys.

A typology of approaches
FAO and WB (2018: 55)

- A. **Food consumer:** count the number of people usually partaking to household's meals, and divide total household consumption by this number.
Limitation: Counting heads of partakers is not precise. The method has difficulties to account for situations in which people do participate only at some meals per day, e.g. employees.
- B. **Meal partakers:** requires an exact accounting of the number of meals taken by household members and non-household members over the same reference period as that for which food data is collected.
Limitation: difficult to implement.

These slides give more structure to the **heterogeneity of approaches** exemplified by the questionnaires shown. Smith et al. (2014) summarize the different pieces of information collected by questionnaires in developing countries; the FAO/World Bank guidelines reduce these different approaches to two main ways of adjusting for partakers – ‘food consumer’ and ‘meal partakers’.

Recommendation # 1
FAO and WB (2018: 55-56)

Information should be collected on the number of meals and the number of individuals (household and non-household members) who participated in each meal.

The addition of an individual **household member-based meal module** should be considered for all surveys that do not yet have it.

...

Recommendation # 3
FAO and WB (2018: 55-56)

If the entire individual household member-based meal module cannot be added, survey design teams should consider **adding questions to a proxy respondent**.

The aim is to capture the **number of meals** were taken at home **by household members and others**, during the reference period.

The discussion on partakers ends with some recommendation from the FAO/World Bank manual. More details can be found there, if the Instructor wishes to elaborate.

5. Timing of visits

Temporal fluctuations

- Fluctuations in consumption and expenditure within the year are common
- Variation **between months**, also called **seasonality**:
 - Agricultural season(s), cyclical food production cycles, festivals and holidays
- But there is also cyclical variation **within months and weeks**:
 - Payday for wage workers, market day, transfer-day¹ for households receiving cash transfers, Friday, Saturday, Sunday (depending on culture) consumption may differ from 'usual'

Next is the issue of determining the **timing of visits** to interviewed households. The choice matters because of fluctuations of consumption over time (seasonality, but also within-month or within-week variation).

Seasonality matters

- Survey objective is usually to mirror **typical consumption throughout the year**
- If variables of interest fluctuate during the year, the **timing of the interview** is not neutral
- Seasonality and higher-frequency fluctuations usually involve:
 - Quantities of food acquired and consumed
 - Dietary patterns
 - Food prices
- These variations are common, although their extent depends on the country

The seasonal effect in Mongolia – III/III

Troubat and Grunberger (2017: 136)

- Peak in December
- It might be associated to celebration of the **Independence Day, 29th of December, or New Year.**
- Lowest in June.

Some evidence is presented on the **magnitude** of seasonal (and higher-frequency) variation in consumption. In Afghanistan, the poverty headcount computed by quarter was found to oscillate widely, due to seasonal variation in both prices and consumption. In Mongolia, Troubat and Grunberger (2017) found that per capita calorie intake varied significantly by day of the week, by day of the month, and by month of the year: the graphs show the percentage change in per capita intakes with respect to a reference day or month (vertical bars indicate confidence intervals for each point estimate). The main point to be made through both examples is that, because consumption is found to fluctuate significantly within the survey year in most countries, the timing of interviews matters for final survey estimates.

Failing to account for seasonality

A survey carried out at **one single time** in the year may be:

- Unrepresentative** of typical consumption across the year
- Not comparable internationally** (what if country A conducts survey in lean season, and country B in harvest season?)
- Not comparable within the same country over time** (what if a major event correlated with consumption patterns moves in or out of the survey period? Think of Ramadan, or harvest periods delayed by weather events)

This point is further stressed by this summary of the negative consequences of failing to adjust for seasonality.

Common approaches to data collection

Smith et al. (2014)

- Repeated visits to the same households** throughout the year. Households are interviewed repeatedly throughout the year (typically 2-4 times, in different seasons)
- Multiple interview rounds distributed by survey subsets.** The sample is split into subsets (usually 12), which are surveyed over 12 months. Subsamples are nationally representative by quarter
- A single interview round, taking place over no more than a few months. **This approach fails to account for seasonality**

Pros and cons of common approaches

- Repeated visits**
 - Pros: seasonal variation captured for all individual households; useful when survey objectives include collecting data on agricultural activities
 - Cons: highest cost, logistical challenges, respondent burden, sample size
- Survey subsets**
 - Pros: cheaper, easier to organize, lower respondent burden than A
 - Cons: seasonal variation captured only on average
- Single round**
 - Pros: easiest
 - Cons: seasonal variation not captured, therefore measurement error

The next slides discuss the **best ways to account for seasonality**. The two most common approaches are introduced, and their frequency in practice is evaluated – together with the frequency of surveys that *do not* account for seasonality at all.

Recommendations

FAO and World Bank (2018: 52-53)

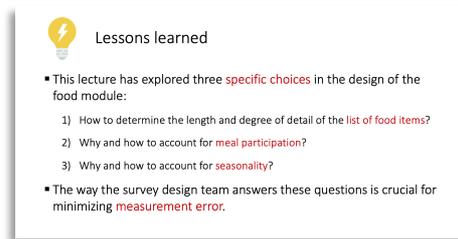
Two options to consider, in order of preference:

- Spread the sample over **12 months** of fieldwork
- Conduct **two visits** per household (e.g. lean period + harvest period)

Whatever the solution chosen:

- Ensure enumeration is equally spread throughout the **days of the week** and the **month**
- Be mindful of changes in timing of **holidays, festivals**, to ensure comparability between survey waves

Recommendations prioritize the ‘survey subsets’ approach, although ‘repeated visits’ are also acceptable.



Lessons learned

- This lecture has explored three **specific choices** in the design of the food module:
 - 1) How to determine the length and degree of detail of the **list of food items**?
 - 2) Why and how to account for **meal participation**?
 - 3) Why and how to account for **seasonality**?
- The way the survey design team answers these questions is crucial for minimizing **measurement error**.

The final slide is used to remind students of the three topics discussed during the lecture. We do not repeat all recommendations – students can refer back to each topic to refresh their memory – but we stress that the ultimate goal of any ‘best practice’ is to minimize measurement error.

Homework

Exercise 1 – Engaging with the literature

See exercise 1, Lecture 1.

Exercise 2 – Food module, international comparisons

This exercise gives students an opportunity to see a practical application of the concepts discussed in the lecture, and become familiar with questionnaires used in different countries. The goal of the exercise is to assess the students’ grasp of the guidelines on the length of the food list, and see whether they can critique questionnaires based on the recommendations learned during the lecture.

Exercise 3 – Meal participation

Example 1 (Somaliland, 2012) collects information on individuals *usually* taking part in the household’s meals. Data collected through this questionnaire would only allow to compute per capita intakes using the ‘food consumer’ approach (estimating the number of people usually partaking during the reference period).

Example 2 (Namibia, 2016) collects daily information on who takes part in each meal (breakfast, lunch, dinner). Data collected through this questionnaire would allow to compute per capita intakes using the ‘meal partakers’ approach (accounting precisely for the number of meals taken by household and non-household members on each day).

Lecture 7

Food away from home (FAFH) and the use of non-standard units for measuring food consumption

Learning objectives

This lecture, together with lectures 5 and 6, has the goal of providing detailed recommendations on the collection of food consumption and expenditure data through household surveys. Lecture 7 focuses on measuring the consumption of food prepared away from home, and on using non-standard measurement units.

Suggested preparation

Much of the material covered in lectures 5, 6, and 7 is adapted from two references: FAO and The World Bank (2018), and Smith, Dupriez and Troubat (2014). The Instructor is assumed to be familiar with both. In addition, Lecture 7 discusses experimental evidence from Farfàn, McGee, Perng, and Vakis (2019) in some detail, and summarizes guidelines from Oseni, Durazo, and McGee (2017). The Instructor is encouraged to review these additional references.

Time allocation

Food away from home	
Concepts	10 min
Empirical evidence	30 min
Recommendations	10 min
Break	15 min
Non-standard measurement units	
Concepts and checklist for implementation	40 min
Recommendations	10 min
Lessons learned	5 min

Annotated lecture

Today's agenda

1. Acquisition vs. consumption
2. Recall vs. diary and length of reference period
3. List of food items
4. Meal participation
5. Timing of visits
6. Food away from home
7. Non-standard measurement units

today

The last lecture on measuring food consumption tackles two very specific topics: **food away from home**, and **non-standard measurement units**.

What is food away from home (FAFH)?

- Difficult to identify a single agreed-upon definition
- General preference: definition based on the **place of preparation** of the food
- FAFH = food **prepared** away from home
- May be consumed at home or not
- **Examples:** prepared meals and snacks that originate from commercial establishments, social programs, schools, other households...

The first few slides define **what exactly is meant by food away from home**: the focus is on food that is *prepared* outside of the home, and therefore does not enter the budget of the household, unless it is explicitly accounted for.

Does FAFH really matter?

- Are these measurement errors **significant** in practice?
- Most likely, **yes**.
- Consumption of food outside the home is **rapidly growing** across the developing world:
 - Percentage of households reporting meals outside increased from 20 to 46% between 1981 and 1998 in **Egypt**, 23 to 39% between 1994 and 2010 in India
 - Household per capita expenditure on FAFH rose at an average annual rate of 9.5% in **China** from 2002 to 2011

Evidence on the importance of FAFH

- **Food security:** Smith (2015) on the "Indian calorie debate", Borlizzi et al. (2017) on Brazil
- **Poverty and inequality:** Farfán, Genoni and Vakis (2017) study Peru and find that, if **FAFH is included**, the poverty rate increases by 1.1 points (18% of baseline), and the Gini index decreases by 1.3 points.
- **Per capita expenditure:** Experimental evidence on questionnaire design by Farfán, McGee, Perng, and Vakis (2019) on Vietnam. We review findings in the next slides.

Some evidence is cited to motivate the emphasis on correctly measuring food away from home. Different studies document that the share of FAFH in total food consumption is increasing in many countries, therefore measuring it correctly is becoming more and more important; several papers estimate the impact of including the expenditure for FAFH in measures of total consumption on food security, poverty and inequality measures. We focus our attention on the findings from a recent experimental study conducted in Hanoi, Vietnam.

Experimental evidence on capturing FAFH

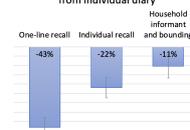
Farfán, McGee, Perng, and Vakis (2019)

- 2,400 households in urban **Hanoi**, Vietnam.
- **Experimental approach:** five different questionnaire designs for capturing FAFH
- What is the impact of different designs on total reported expenditure for FAFH (and therefore on total reported food expenditure)?



Results

% deviation in FAFH expenditure from individual diary



- **One line is not enough:** greatest underestimation of FAFH (-43%), in part because households do not report any FAFH at all

- **Household informant performs well** in this case: likely due to the specific way it was implemented (bounding, worksheet helping to track other hh members)

Evidence from an experimental study by Farfan et al. (2019) is used to showcase different approaches to the measurement of FAFH, and to establish which method performs best. The Instructor is encouraged to review the paper in advance, in order to be able to deliver these slides effectively. Similarly to the studies conducted in Tanzania, and mentioned in previous lectures, this experiment also tests different questionnaire designs, and compares them to a ‘benchmark’ (an individual diary with close supervision, or the “gold standard”). The result that should be most stressed is that **a single question is not enough to adequately capture FAFH**. The issue, then, is how to choose the best among the remaining approaches.

<p>Current practice</p> <ul style="list-style-type: none"> ▪ Serious data gap ▪ Most nationally representative household surveys collect very limited information on food away from home ▪ Dupriez et al. (2014): assessment of 100 surveys from developing countries ▪ 90 out of 100 surveys collect “some” information on FAFH. Of these 90 surveys... 	<p>Main data collection challenges</p> <ul style="list-style-type: none"> ▪ Definition of FAFH: no agreed-upon standards. ▪ Regardless of definition, there should be a protocol to handle ambiguities (consumed out but prepared at home). They can, otherwise, generate confusion or be missed altogether ▪ Accounting for snacks: may seem irrelevant, but are not ▪ A proxy respondent is common, but second-best ▪ Content and quantities: not all meals created equal. Difficult to quantify what is eaten and how much (which matters for nutrition, but also poverty). Recording meal events and mode of acquisition may help
--	---

A review of the **current practice** shows that in many cases, the measurement of FAFH is lacking. Almost one in four surveys capture FAFH using just one question, which, as documented in the literature, is likely to severely underestimate actual consumption; the majority of surveys do not account for snacks; less than one in five surveys collect information on FAFH at the individual level.

<p>Four recommendations FAO and WB (2018: 56-58)</p> <ol style="list-style-type: none"> 1. Design a separate module for FAFH: do not collect FAFH information with just one question 2. Have a clear protocol that specifies whether survey will capture: <ul style="list-style-type: none"> ▪ food prepared at home and consumed outside (jointly with “at-home” module) ▪ food prepared outside and consumed at home (i.e. takeaway) 3. Which pieces of information to collect? <ul style="list-style-type: none"> ▪ Organize data collection around meal events, including snacks and drinks (adapt the meal events list to the local context) ▪ At a minimum, collect info on the value of all meals consumed during each meal event 	<p>Four recommendations FAO and WB (2018: 56-58)</p> <ol style="list-style-type: none"> 4. For whom to collect, and who is the respondent? Two options: <ul style="list-style-type: none"> ▪ Individual level FAFH module. Adults can respond for themselves ▪ Proxy respondents (household level module): <ul style="list-style-type: none"> - Proxy respondent reports incidence of FAFH for all household members, while information on expenditures is collected from each adult - Proxy respondent reports on the incidence and total value of FAFH consumption at the household level, using daily reminder sheet (requires two visits)
---	---

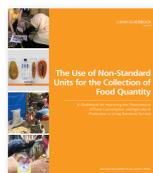
Recommendations adapted from Section 3.5 of the FAO/World Bank guidelines show how to meet the challenges related to the collection of data on FAFH.



Now is a good time for a break.

7. Non-standard measurement units

Main reference for this topic



Oseni, Durazo, and McGee (2017)
Non-standard units LSMS Guidebook

The second part of the lecture, on the use of non-standard measurement units, is modelled closely after the **Guidebook** by Oseni et al. (2017). The Instructor is encouraged to consult sections 3 and 4 of the Guidebook for a detailed discussion of this topic.

What are non-standard units (NSU)?

- **Standard units (SUs)**: standardized across all locations, items. Typically, metric units: grams, kilograms, liters...
- **Nonstandard units (NSUs)**: weights can vary by item, location, customs, type of food preparation... Pieces, bushels, heaps, boxes...

What are non-standard units (NSU)?

Malawi, Third Integrated Household Survey (IHS3) 2010/11

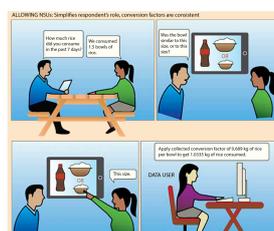
IHS3 Item Code	Item Name and Type	Unit Code	Unit Description	Regional Mean in KG		
				North Region	Central Region	South Region
411	Okra / Thawee	9	Piece	0.08	0.08	0.08
94	Piece (small)	94	Piece (small)	0.08	0.08	0.08
98	Piece (medium)	98	Piece (medium)	0.08	0.08	0.08
10	Heap	10	Heap	0.11	0.21	0.20
10A	Heap (small)	10A	Heap (small)	0.09	0.14	0.13
10B	Heap (medium)	10B	Heap (medium)	0.11	0.21	0.20
10C	Heap (large)	10C	Heap (large)	0.26	0.25	0.32
15	Cup	15	Cup	0.38	0.38	0.38

First, we **define** what is meant by non-standard measurement units, by contrasting them with standard (metric) units. An example from Malawi allows students to form an idea of the kind of NSUs that are used in practice.

Advantages of NSUs

1. May be **more familiar** to respondents
 - SUs may not be used in markets
 - Respondents may not encounter SUs in their daily activities
 - Items may traditionally be consumed in NSUs
 - Evidence: When given the choice, respondents prefer reporting in NSUs. In Malawi IHP5 2013, 73% of consumed item entries are in NSUs
2. Simplified recall: **reduces on-the-spot calculations**, cognitive burden.

In practice
non-standard units



The use of NSUs in a household consumption and expenditure survey has some clear **advantages**: NSUs may be more familiar to respondents, and allowing them to report their consumption in terms of these more everyday units reduces their cognitive burden (they are not forced to make the conversion into standard units in their heads) – a vignette illustrates this point.

Disadvantages of NSUs

1. **Complex to collect**: may vary by item, location...
2. **Not always clearly defined**:
 - A 'piece' of bread?
 - A container can be filled level or heaped
 - A common unit in Ethiopia: chinnet (donkey load)
3. **Need to be converted into SUs** for analysis



However, the use of NSUs has some **disadvantages**, too, mainly linked to the amount of preliminary work that the Statistical Office is required to do in order to make them work properly. The next slides identify a strategy to help meet these challenges.

Compiling and using non-standard units

- The use of NSUs in household surveys requires the construction of a **NSU library**
- Its components are:
 - List of 'allowable' item-NSU combinations
 - **Conversion factors** for each combination (national, regional)
 - **Reference photos** for each combination
 - **Clear protocols** for using conversion factors and reference photos
 - Documentation, documentation, documentation

A checklist for using non-standard units

- Establish a NSU list
- Compile conversion factors
- Take reference pictures
- Integrate NSUs into main survey
- Implement in CAPI (optional)

Making NSUs work requires to build a **NSU library**. To do that, one must go through several steps, summarized here in the form of a **checklist** – this is a condensed version of the guidelines offered by Oseni et al. (2017).

Establish a NSU list

- Plan in advance of main survey, and construct all allowable item-NSU combinations
 - How? National sources, previous surveys, conduct a pilot
- Identify variation:
 - Units: heaps, pails, pieces
 - Sizes: small, medium, large
 - Conditions: shelled, unshelled, threshed, etc.

Conducting a market survey

- Timing matters:
 - **Seasonality** in item availability, harvest time
 - **Ideal scenario:**
 1. Prior to main survey, implement market survey: obtain conversion factors, collect reference photos to use in main survey
 2. During/after main survey: Consider implementing a smaller scale market survey to address gaps in reported item-NSU-(condition) combinations in main survey

First, a **list of NSUs** must be established, in advance of the main survey. For each of the NSU that will be used in the main survey, there must be a **conversion factor**, that allows the analyst to recover metric quantities ex-post. Without conversion factors, data collected using NSUs will be useless: data users will not be able to compute calorie intakes, unit values, and other variables that are crucial for poverty measurement. In practice, the instrument used to collect conversion factors is a **market survey**, to be conducted before the main household expenditure survey. This is a survey collecting data about the nature and weight in grams of NSUs (as well as, in some cases, market prices) from a number of markets where food items are traded.

Both the **coverage and timing** of the market survey need to be planned carefully: in particular, survey specialists must be mindful of the seasonal availability of food items, to make sure all relevant foods are covered. There are several advantages to conducting the market survey in advance of the main survey: reference photos can be taken and then used during fieldwork for the main survey; NSUs that may be missing from any already-existing list can be identified on time and added. However, there are some drawbacks to this choice, too: what if respondents are inclined to report consumption using units that were not covered by the market survey? That would determine conversion factor gaps, which is to be avoided. For this reason, “(...) the ideal plan is to conduct two market surveys—one before and one after the main survey. Both market surveys need not be equally rigorous; one will likely be more comprehensive than the other. For example, the ex-ante survey could be limited, aiming to collect reference photos and weights (for conversion-factor calculations) for the most common NSUs, while the ex-post survey could comprehensively collect weights for all additional NSUs reported during the main survey.” (Oseni et al. 2017: 11).

Take reference photos

- Developed in market survey, to be used in main survey
- Goals: help respondents better specify quantities, provide standard size for NSUs that are not clearly defined (e.g. piece, heap)
- Not as easy as taking a selfie!
 - General guidelines for taking photos
 - Select quality photos for reference photos
 - Properly train enumerators to use NSU materials



Implement on CAPI (optional)

- Market survey**
 - Take a (uniquely-named) photo for each weight measurement within CAPI app
 - Georeference market locations; automatic date and time capture
 - Rigorous data quality checks: range checks, flagging missing observations
- Main survey**
 - Reference photos can be integrated into CAPI app
 - Additional photos at the household-level can be taken
 - Rigorous data quality checks based on:
 - Allowable item-NSU combinations
 - Conversion factor library for flagging potential "outlier" quantities based on checks on the basis of unit values, food consumption and/or caloric intake per capita

The next slides go into details on the other aspects of compiling a NSU list – from taking reference photos to implementing the use of NSUs on CAPI – the Instructor is encouraged to consult the corresponding sections of Oseni et al. (2017) for any clarifications.

Recommendations
FAO and WB (2018: 62-63)

- Decision on whether to allow for NSUs should be addressed in the questionnaire **design phase**.
- Evaluate cost vs. benefit of allowing NSUs. May conduct **pilot survey** to determine to which extent NSUs are actually needed by respondents.
- Complete list** of conversion factors is essential. Market surveys and photo references help.
- Establish NSU database that can be used **across surveys**.

To wrap up the discussion on NSUs, we share a summary of recommendations, taken from the FAO and World Bank guidelines.

 **Lessons learned**

- This lecture has explored **specific topics** in the design of the food module:
 - How to measure **food prepared away from home?**
 - When and how to use **non-standard measurement units?**
- The first issue should be addressed by all surveys: growing relevance of FAFH as a component of food consumption.
- Allowing for non-standard units may be useful in some context, but implementation should carefully follow guidelines.

The last slide recaps the topics of the lecture for students. In just two bullets, we summarize the general takeaways for both FAFH and NSUs.

Homework

Exercise 1 – Engaging with the literature

See exercise 1, Lecture 1.

Exercise 2 – FAFH modules

An in-depth analysis of the two questionnaires may raise some or all of the following points.

Both questionnaires capture the data on food consumed away from home (FAFH) with a recall period of 7 days. Tanzania has a distinct module/section dedicated to FAFH. By

contrast the data is incorporated in the overall food and beverage consumption in Uganda. The major differences between the two questionnaires are twofold:

- 1) FAFH in Uganda is collected for the entire food list, whereas Tanzania concentrates on 7 (somewhat ambiguous) composite items (full meals, barbecued meat chips and snacks, local brews, tea coffee samosa cake and snacks, alcoholic and non-alcoholic beverages).
- 2) Tanzania accounts for FAFH received as gift and differentiate them from other food received for free (collected in another module). In Uganda FAFH is collected for food purchases only (as the phrasing of the question seems to indicate).

The minimum reliability criteria from Smith et (2014: 23) are met for Tanzania only. In both questionnaires the sources of FAFH received in kind are missing.

Exercise 3 – Non-standard measurement units

Picture 1 lacks a familiar reference, like a pen or a bottle, that would allow to gauge the size of the fruits. Picture 2 contains a reference, but it is not very visible, and perspective is skewed: the largest heap is closest to the camera, thus appearing even larger than its actual size. A better way to picture the units would be to frame them straight on against the backdrop. A similar problem occurs with picture 3: this time the bottle is clearly visible, but the baskets are pictured from above, and the largest appears closer to the camera, again skewing perspective and making it difficult to estimate the size of the containers.

Lecture 8

Measuring consumption of non-food non-durable goods

Learning objectives

The goal of this lecture is to provide detailed recommendations on the collection of data on expenditures on non-food, non-durable goods through household surveys. First, the lecture gives a broad description of the analytical procedures usually adopted when analysts construct a proxy of non-food consumption; this knowledge is then used to inform data collection recommendations.

Suggested preparation

The Instructor is assumed to be familiar with the material contained in section 3 (“Constructing the Household Consumption Aggregate”) of Deaton and Zaidi (2002). Guidelines on data collection of non-food non-durable expenditures are from Deaton and Grosh (2000) (general guidelines), Oseni et al. (2018) (education), Gertler, Rose, and Glewwe (2000) (health): this material should be read in advance. The lecture also cites evidence from Heijink et al. (2011) and Lu et al. (2009).

Time allocation

What analysts do, ‘in or out?’	45 min
Break	15 min
What analysts need	
General guidelines	15 min
Education	20 min
Health	20 min
Lessons learned	5 min

Annotated lecture

Where we stand

- In lecture 1, we argued that analysts are interested in **total household consumption expenditure** as a measure of living standards
- In practice, household surveys typically record household **expenditures**, as well as consumption that does not go through the market (in-kind receipts and own-production)
- It is the analyst's job to come up with an estimate of the total value of consumption, starting from this information
- We have covered this process, and how it reflects on questionnaire design and data collection, for what concerns **food items**
- We now turn to **non-food items**

Today's agenda

The lecture is organized as follows:

1. What analysts **do** (research objectives)
2. What analyst **need** (data collection guidelines)

The first slide of this set is meant to help students find their bearings again, after the three lectures on measuring food consumption: it is now time to move on to **non-food consumption**. The outline clarifies the 'grassroots' approach of the lecture: the discussion starts from research objectives (what do welfare analysts do when constructing the non-food component of the consumption aggregate?), and then proceeds to data collection guidelines (how can the questionnaire meet the research objectives?).

1. What analysts do

Specific rules

Include	Controversial	Exclude
Clothing and footwear	Gifts and remittances	Taxes and Levies
...	Health	Purchases of assets, repayment of loans
Education		Means-tested expenditures
Regrettable necessities		Purchase of durables (including houses)
Insurance		Extraordinary (lumpy) expenditures
Utilities		

The takeaway from this first part of the lecture is that **not all expenditures** should be indiscriminately considered to be part of the living standard indicator. The *consumption* aggregate is constructed by analysts based on *expenditure* data, which requires some fine-tuning: some expenditures should not be included in the aggregate, because they do not represent *typical, welfare-enhancing consumption*.

Items in green

- Clothing and footwear
 - ...
 - Education
 - Regrettable necessities (*)
 - Insurance
 - Utilities (*)
- For most items in the list, inclusion is pretty straightforward.
 - For instance, including 'clothing and footwear' means that the more a household consumes on these goods the higher is its standard of living.
 - Is this controversial?
 - No
 - Starred items are deserve a quick discussion.

Utilities

- Expenditure on utilities – water, gas, electricity or telephone – can be, and often are, **problematic**.
- Problems may arise in the presence of
 - subsidies,
 - progressive tariffs, and
 - rationing.
- In these circumstances, welfare comparisons make it necessary to make **corrections** to the reported expenditures.
- A common strategy is **repricing**. See Hentschel and Lanjouw (2000).

Next, the lecture discusses **inclusion and exclusion** 'rules' in further detail. Items to include are generally uncontroversial. Two categories deserve a comment. First, regrettable necessities: these are goods and services that yield no welfare in their own right, but that have to be purchased, for example, in order to earn income, such as work clothes or transport to work. Although one might argue not to include these expenditures in the consumption aggregate (they are not welfare-enhancing), the practical difficulties of telling them apart from ordinary consumption are insurmountable: "Transport to work is a regrettable necessity for someone who has little choice of where to work or where to live, but is consumption for someone who chooses to live in a pleasant suburb. Out-of-pocket

medical expenses are a necessity for some, but a choice for others, as in curative versus cosmetic medicine. It is hard to see how guidelines could be constructed that would allow one and not the other.” (Deaton and Zaidi, p. 20). Ultimately, the recommendation is to include the expenditures that may be considered regrettable, to avoid any ambiguity. Another caveat concerns utilities, which can be problematic when, for different reasons, they are not purchased at market price: in those cases, the price paid by households is not a correct representation of the true market value of the service, and a market price must be estimated by the analyst (Deaton and Zaidi touch on this on p. 34).

Items in red
Expenditures excluded

- Taxes
- Purchases of assets, repayment of loans
- Purchase of durables (including houses)
- Extraordinary (lumpy) expenditures, e.g. marriages and dowries
- Means-tested expenditures (*)

- Taxes are not consumption, but a deduction from income
- Purchases of assets is investment, not consumption
- Durables will be covered in lecture 9 – we touched on them in lecture 2
- "extraordinary", is not "typical" which is what analysts aim at.
- Means-tested expenditure are worth an extra slide

Expenditures to be excluded are only briefly discussed: two upcoming lectures will focus specifically on the two most important categories, durable goods and housing.

*Definition: means testing

- According to the dictionary, 'means testing' is the process of measuring whether a person or a household has the means (income, savings, wealth...) in order to decide if they are eligible for receiving help (benefits, transfers, etc.) from the government.
- For instance, in the US means testing is used to test for eligibility to Medicaid.
- Means testing is used to regulate access to health programs in a number of African countries.

Means-tested expenditures

- "Means-testing" is to determine whether a household is eligible for support from a welfare program, for instance.
- If the income is below a certain threshold, the household is eligible for receiving a certain good or service at a subsidized price.
- If the income is higher than the threshold, the household pays the market price.
- The presence of means-tested expenditures requires a correction to the reported expenditures.

A short discussion is dedicated to means-tested expenditures. First, the Instructor has the option to show the starred slide with the definition of means testing, in case the audience requires it. In terms of construction of the consumption aggregate, the problems related to means-tested expenditures are similar to those arising in the case of subsidized items: these expenditures must be excluded, unless their price can be corrected and equalized across households.

Items in orange

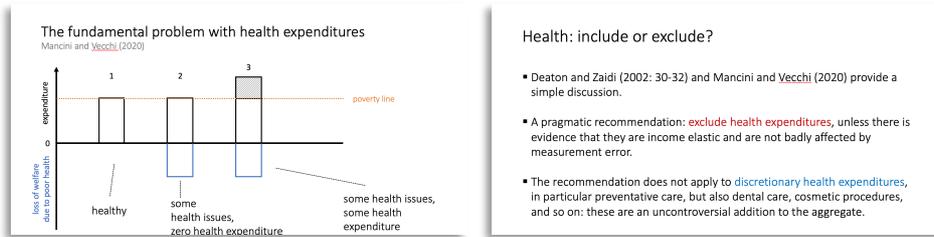
- Gifts, charitable contributions and remittances
- Health (*)

- Gifts, as any other transfers are better excluded to avoid double counting if, as one would expect, the transfer shows up in the consumption of the recipient household.
- Health is tricky category, which deserve a proper discussion (next slide)

Health

- Whether to include or exclude health expenditures is a controversial decision
- By including health expenditures for someone who has fallen sick, we register an increase in welfare when, in fact, the opposite has occurred.
- In principle, we should account for two components:
 - 1) loss of welfare due to illness
 - 2) gain of welfare from health expenditure
- Next slide illustrates.

Finally, the slides touch on mention the controversial topic of **health expenditures**. Why controversial? The fundamental problem is that we can observe and measure the increment in welfare due to receiving (consuming) healthcare, but not the loss in welfare determined by the need for care, arising from a deterioration of one's health status.



A **graph** taken from Mancini and Vecchi (2020) illustrates the point. Observed welfare, measured by consumption expenditure, is represented by the black bars above the horizontal axis. The taller the bar, the higher the expenditure, the better off the household. Below the horizontal axis is what we cannot observe and measure: the loss of welfare due to a compromised health status (in blue). Pictured are three households: household 1 is healthy (there is no bar below the horizontal axis, indicating a loss of zero), and its consumption aggregate puts it right at the poverty line (the bar above the horizontal axis is the same height as the poverty line). Household 2 experiences some health problems (this is represented by the negative blue bar below the horizontal axis), and is also at the poverty line in terms of total expenditure: for the sake of our example, let us assume that household 2 cannot afford any healthcare, and that its expenditure is exactly equal to that of household 1 (bars above the axis for households 1 and 2 are equally sized). Household 3 experiences the same health issues as household 2, but is able to spend some of its resources in healthcare (the shaded portion of the bar), which puts its total expenditure above the poverty line. Comparisons between these households clarify the conundrum related to health expenditures.

Compare households 2 and 3. Their health status is equally compromised, but household 3 is getting healthcare, unlike household 2. If health expenditures were included in the NCA, we would see that 3 is better off than 2, thanks to its ability to consume more of the goods and services it desires. If we were to exclude health expenditures from the NCA, we would lose this distinction: the comprehensiveness of our measure of welfare would deteriorate. Overall, in this situation, inclusion appears to be a better strategy than exclusion.

Now turn to households 1 and 3. With health expenditure included in the NCA, 3 would appear to be better off than 1, despite the fact that it only consumes more goods and services because of a health crisis. As Deaton and Zaidi (2002: 32) put it, “by including health expenditures for someone who has fallen sick, we register an increase in welfare when, in fact, the opposite has occurred”. If we excluded health expenditures, 1 and 3 would appear equally as well off. This is also problematic, because the household whose members are sick is likely to be worse off than the healthy one – but it is also unavoidable, because we cannot measure the loss of welfare due to poor health.

The bottom line is that each choice has its drawbacks: if we exclude health expenditures from the NCA, we miss the welfare-enhancing value of healthcare: keeping health status fixed, we would ideally like to capture the increased welfare of households that enjoy better care. If we include it, we mistakenly attribute higher living standards to households that are actually struggling: health expenditures have the peculiar characteristic of being associated with a reduction in welfare rather than an increase, as is the case with other expenditures.

Deaton and Zaidi (p. 30-32) argue that the first mistake – a loss of coverage of the NCA – is less serious than the second one – a mis-ranking of households.

The 'rules' and the COICOP system



- Classification of Individual Consumption by Purpose
- COICOP is an integral part of the SNA, but it is intended also for use in (...) household expenditure statistics based on household budget surveys (p. 4)

Students may already be aware of the **COICOP classification of expenditures**; if not, the availability of this system is an important takeaway of this lecture. The United Nations COICOP classification provides a basis that each country can extend at will, and it is often used in household consumption and expenditure survey questionnaires. The COICOP classification is used to put the 'rules' discussed in the previous slides to the test. Students should be invited to attempt to explain why certain items should or should not be included in the measure of living standards.

In or out? – I/VI

01	Food and non-alcoholic beverages	✓
01.1	Food	✓
01.2	Non-alcoholic beverages	✓
01.3	Services for processing primary goods for food and non-alcoholic beverages	✓
02	Alcoholic beverages, tobacco and narcotics	✓
02.1	Alcoholic beverages	✓
02.2	Alcohol production services	✓
02.3	Tobacco	✓
02.4	Narcotics	⚠
03	Clothing and footwear	✓
03.1	Clothing	✓
03.2	Footwear	✓
04	Housing, water, electricity, gas and other fuels	✓
04.1	Actual rentals for housing	✓
04.2	Imputed rentals for housing	✓
04.3	Maintenance, repair and security of the dwelling	⚠
04.4	Water supply and miscellaneous services relating to the dwelling	⚠

Food, beverages, clothing, housing and utilities should all be included. The danger sign on narcotics indicates the potential criticism of narcotics not being welfare-enhancing (but Deaton and Zaidi recommend their inclusion, nonetheless).

In or out? – II/VI

05	Furnishings, household equipment and routine household maintenance	✓
05.1	Furniture, furnishings, and loose carpets	✓
05.2	Household textiles	✓
05.3	Household appliances	✓
05.4	Glassware, tableware and household utensils	✓
05.5	Tools and equipment for house and garden	✓
05.6	Goods and services for routine household maintenance	✓
06	Health	⚠
06.1	Medicines and health products	⚠
06.2	Outpatient care services	⚠
06.3	Inpatient care services	⚠
06.4	Other health services	⚠

Expenditures in these categories are to be included, but students should be reminded of the controversial nature of health expenditures.

In or out? – III/VI

07	Transport	✓
07.1	Purchase of vehicles	✓
07.2	Operation of personal transport equipment	✓
07.3	Passenger transport services	✓
07.4	Transport services of goods	✓
08	Information and communication	✓
08.1	Information and communication equipment	✓
08.2	Software excluding games	✓
08.3	Information and communication services	✓

In or out? – IV/VI

09	Recreation, sport and culture	✓
09.1	Recreational durables	✓
09.2	Other recreational goods	✓
09.3	Garden products and pets	✓
09.4	Recreational services	✓
09.5	Cultural goods	✓
09.6	Cultural services	✓
09.7	Newspapers, books and stationery	✓
09.8	Package holidays	✓
10	Education services	✓
10.1	Early childhood and primary education	✓
10.2	Secondary education	✓
10.3	Post-secondary non-tertiary education	✓
10.4	Tertiary education	✓
10.5	Education not defined by level	✓

Vehicles are classified as durable goods. The remaining items are to be included.

In or out? – V/VI

11	Restaurants and accommodation services	✓
11.1	Food and beverage serving services	
11.2	Accommodation services	
12	Insurance and financial services	
12.1	Insurance	✓
12.2	Financial services	✗
13	Personal care, social protection and miscellaneous goods and services	✓
13.1	Personal care	
13.2	Other personal effects	
13.3	Social protection	
13.9	Other services	

If the label ‘financial services’ indicates savings or investment, they should be excluded. Note that insurance services are, in principle, consumption (consumers purchase a service, which is the insurance company’s promise to protect them from a potential bad outcome), but in some countries, insurance is used as a form of investment. In that case, it should also be excluded.

In or out? – VI/VI

14	Individual consumption expenditure of non-profit institutions serving households (NPISHS)	
14.1	Housing	
14.2	Health	
14.3	Recreation and culture	
14.4	Education	
14.5	Social protection	
14.6	Other services	
15	Individual consumption expenditure of general government	
15.1	Housing	
15.2	Health	
15.3	Recreation and culture	
15.4	Education	
15.5	Social protection	

Finally, these categories are used for national accounting, and are typically not needed when designing household surveys.



Now is a good time for a break.

2. What analysts need



Organization of topics: General advice – II/II
Deaton and Grosh (2000: 122)

1. Some expenditures are best collected in a **dedicated module**, together with related non-expenditure information
2. **Housing, durables, health, education** usually have their own module; **employment** module is best place to gather information on household consumption of non-food items provided in kind
3. Create a cross-check to **avoid double-counting** in these instances; make sure questions are worded carefully so analysts are able to exclude duplicate measures of the same item, if needed.

Next, the lecture moves on to consider **data collection** recommendations, on the basis of the analytical priorities discussed so far. It should be acknowledged that guidelines on the collection of non-food, non-durable expenditure data are rather broad, and for many choices that questionnaire design teams will face, there is no single universal answer. This portion of the lecture is based on Deaton and Grosh (2000), which describes the general approach adopted by LSMS surveys.

Education



Discusses best practices for collecting information on **education expenditure** in household surveys

• • •

Five specific recommendations – II/II

- 3) **Age range**
What the best age group of targeted respondents? 5-24?
No, it is best not to restrict the age range.
- 4) **Respondent**
The ideal respondent is the person who is most familiar with education expenditures for individuals in the household. Parents or guardians who make educational payments may be best placed to provide information for children.
- 5) **Recall period**
The baseline recall period should be a full **12 months**. For **recurring** expenditure items (e.g., transportation, school meals), it is appropriate to shorten the recall period to **one month**.

Oseni et al. (2018) provide more specific guidelines on education. The Instructor is referred to the Guidebook for details.

Health



- Health is a critical component of the standard of living
- Analysts need to
 - Assess health problems
 - Identify who receives medical care
 - Household expenditures on health care
- We focus on the last item

• • •

Practical guidelines – II/II
Gertler, Rose and Glewwe (2000: 189)

- Expenditures should include **not only fees** but also any other expenditure incurred by the respondents – more below.
- **Health insurance** poses a problem: data should be collected so as to distinguish between charges paid for (or reimbursed by) the insurance and charges paid by the respondent.

A useful reference for specific guidelines on the health module is Gertler et al. (2000), which draws once again from the LSMS experience.

Evidence

- Heijink et al. (2011) for WHO
- Review of 114 surveys from 1990 to 2010
- Focus on health expenditures



• • •

Question structure and phrasing

- When it comes to designing the questionnaire, **two options**:
 - 1) health expenditure questions can be included in the health module. In that case expenditure questions are preceded by questions on illness and health care use which may help respondents to remember health expenditures.
 - 2) out-of-pocket health expenditures are included in the non-food-non-durable module of the questionnaire. They are surrounded by questions on other services and goods.
- According to Lu et al. (2009), 40% of surveys follow 1), 60% follow 2).

The lecture is wrapped up by a discussion of the **current practice**. The point to be made is that there is some leeway in designing many of the important features of the health module (recall period, disaggregation, structure).

Homework

Exercise 1 – Engaging with the literature

See exercise 1, Lecture 1.

Exercise 2 – ‘In or out?’

Students must read Section 3 of Deaton and Zaidi (2000) in order to engage with this exercise. However, an exact solution is not the goal here. Not all items will be uncontroversially in or out: the exercise can be used to spur a discussion on analytical choices, and their bearing on data collection.

Exercise 3 – The treatment of health expenditures

The goal of the exercise is to let students evaluate the practical repercussions of the concepts discussed during the lecture. Both countries exclude health expenditures from the consumption aggregate; the motivation given by both is in accord with the arguments presented by Deaton and Zaidi (2000), except for the inclusion of purchased medicine in the case of Bhutan, which is essentially left unmotivated.

Exercise 4 – The education module

This exercise gives students an opportunity to see a practical application of the concepts discussed in the lecture, and become familiar with questionnaires used in different countries. The goal of the exercise is to assess the students' grasp of the guidelines on the education module, and see whether they can critique questionnaires based on the recommendations learned during the lecture.

Lecture 9

Durable goods

Learning objectives

The goal of this lecture is to provide detailed recommendations on the collection of data on durable goods. The lecture explains how analysts treat this category of expenditures when constructing a measure of living standards; these analytical requirements guide the design of a survey module on durable goods.

Suggested preparation

The Instructor is assumed to be familiar with Section 3.4 of Deaton and Zaidi (2000) (Consumer Durables), and with Amendola and Vecchi (2014).

Time allocation

What are durables, and why do they require special treatment?	20 min
How to deal with durable goods	
Notation. Acquisition and rental equivalence approaches	40 min
Break	15 min
User cost approach, data requirements	25 min
How to design a dedicated module in the questionnaire?	15 min
Lessons learned	5 min

Annotated lecture

A fundamental presumption

- Long-lived goods (automobiles, appliances, furniture, etc.) have a positive and significant impact on living standards.
- These goods are **special**: measuring the increment in living standards derived from them is not as straightforward as for other goods
- This whole lecture is dedicated to durable goods

Today's four questions

1. What is a durable good?
2. Why do durable goods require **special** treatment?
3. How to deal with durable goods, **analytically**?
4. How to design a dedicated **module** in the questionnaire?

The first few slides of this set make an important point: **durable goods** are not like other consumption goods that contribute to living standards, and they **require special treatment** when we construct a consumption-based measure of welfare. *Why* this is the case and *how*

to treat these goods, both in terms of analysis and if questionnaire design, will be the subject in this lecture.

1. What is a durable good?

What is a durable good? – II/II
Diewert (2009: 447)

- Housing is a durable good.
- Due to its importance, it is customary for analysts to deal with it separately from other durable goods.
- Accordingly, in this lecture we focus on consumer durable goods other than housing

First, we give an exact **definition** of durable goods, using Diewert (2009). Typical examples of durable goods are housing (which requires separate treatment, and will be discussed in the next lecture), cars and other vehicles, household appliances, and so on.

2. Why do durable goods require a special treatment?

Why do durable goods require a special treatment?

- A figure worth a thousands words:



- The durables' service flow exceeds the reference period of the welfare aggregate
- The purchasing price reflects the value of the durable for its entire life
- Need to capture the value of the flow of the service during the reference period

The problem with durable goods

- It is not the purchase of a good that contributes to welfare, but its use.
- This creates a wedge between household expenditure (which we can easily measure) and household consumption (we rarely observe usage directly).
- For non-durable (perishable) goods, it is safe to ignore this wedge: expenditure is a good estimate of consumption expenditure
- But for durable goods, we need to estimate the value of using the good for one year (service flow), and add this value to household consumption expenditure
- How do we estimate the value of owning or having access to durable goods during a given year?

Next, why is it that durable goods require a special treatment? The key concept is that welfare analysts want to measure *consumption* during the reference period, and not expenditure. For most goods, which are purchased and consumed right away, expenditure is a good proxy for consumption. But the use of a durable good exceeds, by definition, the reference period, and its purchase price reflects this long-lasting flow of consumption that the consumer is able to enjoy from the good: if you buy a washing machine today, you are paying for the consumption of it throughout its duration, which is much longer than a year. Instead, welfare analysts want to **measure only the portion of consumption that is enjoyed during the reference period.**

3. How to deal with durable goods, analytically?

Useful reference
Amendola and Vecchi (2014)



- Review of methods and current practice
- Mathematical notation used in the presentation is consistent with this paper

The next question is, how do welfare analysts treat durable goods? Our discussion in the rest of this lecture is a condensed version of the reference cited here: this is required reading for anyone who wishes to understand the details of what is about to follow.

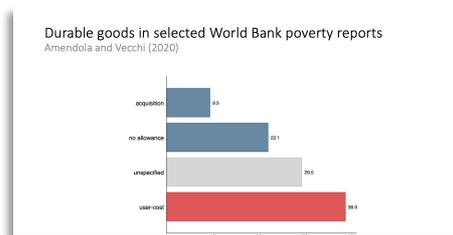
Three approaches

1. **Acquisition Approach**
When the good is purchased its entire value is attributed to the household welfare aggregate
2. **Rental Equivalence**
If a complete set of markets for the services of durables exists, we can use the market rental value of the goods
3. **User Cost**
The annual opportunity cost of holding each durable.

Definition: Opportunity cost

- Opportunity cost = the value of the next-best alternative you give up
 - Choice A has value X, choice B has value Y
 - You can only pick one (scarcity)
 - Opportunity cost of choice A is Y, opportunity cost of choice B is X
- In this case, opportunity cost of holding a durable is the income one would make by selling it and investing the money

There are three main analytical approaches to the treatment of durable goods, and each of them requires different pieces of information – which in turn reflects on how the durable goods module of the questionnaire should be designed. The first approach (**acquisition**) is cited as a baseline: it consists in ignoring the peculiar nature of durable goods and the ‘special treatment’ due to them, and simply using the purchase value as a representation of the value of consuming them. This is obviously incorrect, given the previous discussion. The **rental equivalence** and **user cost** approaches are, instead, two ways of computing the use-value of the durable good during the reference period. One of our ‘definition’ slides with the encyclopedia icon reminds students of the precise definition of the economic concept of opportunity cost, which is crucial for the user-cost approach (the one that, as we will see, is recommended).



The bar chart shows the relative frequency of the approaches in recent poverty assessments around the world: the user-cost approach is most frequent, but the use of the (incorrect) acquisition approach is not rare, and as many as 22% of poverty assessments do not include the value of durable goods in the welfare indicator. Note that the lack of suitable data may be behind such a phenomenon.

Consider your car:
how to calculate its contribution to your standard of living?

Three approaches, one formula

- The **consumption flow** to be included in the consumption aggregate can be calculated by means of a simple formula:

$$CF_t = k_{v,t}^s \times p_{v,t}$$
 - Interpretation: The consumption flow for a generic v-year old durable good purchased s years back in time is a fraction k of the current market value of the good, p_{v,t} (how much the v-year old good is worth on the market at the beginning of the survey period)
 - The coefficient k is typically less than one.
 - This equation should be **memorized**.

How do the different approaches work in practice? We introduce some conventional notation and a basic formula that helps us think more clearly about the problem at hand: in

a nutshell, the **consumption flow** from the durable good during the reference period, CF_t , is a fraction k of its current market price.

Method 1 – Acquisition approach

- A first option consists in adding up reported **purchases** on durable goods (purchase values) and include them in the consumption aggregate
- This would be a **mistake**
- Why?
- Because it would amount to assuming that households that purchased a durable good in the survey period use it all up by the end of the year.
- On the other hand, households that own durable goods purchased before the survey period would be considered “as well off as” households that do not own any durables
- This is in stark contrast with the very definition of durable good: a good that delivers utility for a period longer than the survey year.

The acquisition approach in practice

$$CF_t = k_{v,t}^s \times p_{v,t}$$

$$k_{v,t}^s = \begin{cases} 1 & \text{if } s = 0 \\ 0 & \text{if } s > 0 \end{cases}$$

- If $s > 0$, then $k = 0$, and $CF_t = 0$
Interpretation: items purchased before the survey year ($s > 0$) do not contribute to the household's well-being.
- Does it make economic sense?
- **No**
- If $s = 0$, then $k = 1$, and $CF_t = p_{v,t}$
Interpretation: items purchased during the survey year ($s = 0$) contribute to the household's well-being for their full value (net captures the present value of all services provided by the durable over its entire economic life)
- Does it make economic sense?
- **No**

First, we discuss the acquisition approach – basically, *not* treating durable goods any different than any other good. If a household purchases a durable good, say a car, during the survey period, then they are assumed to have consumed its entire value that same year; if a household did not purchase any cars during the survey period, they are assumed to not enjoy any returns from it, *whether or not they own a car*. This is not correct from a conceptual standpoint: again, it is *use*, not *purchase* that counts for well-being.

Method 2 – Rental equivalence

- Ideally, one could try to estimate the utility that derives from owning (or using) a durable good by collecting information on how much it would cost to **rent it for a year**.
- In **principle**, this is doable - in **practice**, it is not.
- Most countries have no markets for renting most durable goods, and when markets exist it is difficult (impossible?) to control for quality.
- Not recommended

*The rental approach in practice

$$CF_t = k_{v,t}^r \times p_{v,t}$$

$$k_{v,t}^r = \frac{R_{v,t}}{p_{v,t}}$$

- Assume that consumers can rent a car
- Let $R_{v,t}$ denote the **current market rental value** of the v -year-old durable good
- If k is specified as in the formula here (it can be interpreted as the share of the good's value that is consumed in the current period), then $CF_t = R_{v,t}$
- Interpretation: the CF equals the market rental value of the durable owned by the household
- Does it make economic sense?
- **Yes**
- Is it empirically viable?
- Most likely, **no**.

Then, the **rental equivalence** approach. In theory, because rent as a representation of the use-value of a good for a particular time period, it is very sensible to try to measure CF_t as a rental value. The problem is, a **rental market does not exist** for most durable goods, so that this approach is generally unfeasible in practice. We will see that, unsurprisingly, this is not the case for housing. The optional (starred) slide presents the rental equivalence approach in terms of the basic equation introduced at the top of this section.



Now is a good time for a break.

Method 3 – User Cost

- We introduce the user cost approach through a **conceptual experiment**
- Consider a household that owns a durable good.
- Notation: let p_t denote the **market value** of a particular good at the beginning of the survey year t (we forget about the age of the good for a second)
- The household faces two options:
 1. to **sell** the durable good;
 2. to **use** the durable good.

The user cost approach – I/II

sell	use
If the household sells the durable good, and invest the revenue on the financial market, at the end of the year, the household receives	If the household uses the durable good and sells it at the end of the year, the household obtains
$p_t(1 + i_t)$	$p_t(1 + \pi_t)(1 - \delta_t)$
where i_t is the market nominal interest rate.	where π_t is the inflation rate during the year t and δ_t is the annual depreciation rate (due to both physical deterioration and loss of market value).

Finally, we discuss the **user cost method**, which requires a bit more time, because it is the most used in practice. The first step is to recognize that a household that owns a durable good can be seen as facing **two options** at the beginning of the survey period: to sell it right away, or to use it for a year, and then sell it.

The user cost approach – II/II

- The **consumption flow** is the difference between the value of the two options at the end of the year: this is the cost that the household is willing to pay for using the durable good for one year:

$$CF_t = p_t(1 + i_t) - p_t(1 + \pi_t)(1 - \delta_t)$$

which can be approximated by:

$$CF_t = p_t(i_t - \pi_t + \delta_t) = p_t(r_t + \delta_t)$$

CF is the **consumption flow** from durables

The consumption flow, interpreted

$$CF_t = p_t(i_t - \pi_t + \delta_t) = p_t(r_t + \delta_t)$$

- Two cost components:
 - Opportunity cost**
 $p_t r_t$ is the foregone real interest, i.e. the interest one could have earned if one had invested the money in a bank account instead of the consumer good.
 - Depreciation**
 $p_t \delta_t$ is the drop in value of the good during the course of the year.
- Problem: how to estimate the depreciation rate (delta) in practice?

The user cost approach in practice

- Using our formula:
 $CF_t = k_{v,t}^s \times p_{v,t}^s \delta_t$
- Note that if
 $k_{v,t}^s(u) = r_t + \delta_t$
- then
 $CF_t = (r_t + \delta_t) p_{v,t}$
- which is what we have derived through the conceptual experiment seen before.

The difference between the values of the two options is an estimate of CF_t . If a household does not sell the good, but uses it, it is because the good's utility compensates for that difference. Another way to see the final equation is that, by holding the good for a year, the household is 'paying' the *opportunity cost* of selling it, plus the economic depreciation of the good. This can be brought back again to the general formula seen at the start.

Estimating CF_t based on the user cost approach

$$CF_t = p_t(r_t + \delta_t)$$

- Of the two "ingredients" needed to compute CF_t , r_t is the easiest to obtain: it comes from **sources external to the survey**.
- Instead, the **depreciation rate δ_t** , which measures the loss (or gain) in value that durable goods experience with age due to physical deterioration and market value change, must be **estimated**.
- How to estimate δ_t ?
Do **bicycles** depreciate at the same rate as **refrigerators**?

Estimating the depreciation rate – I/II

- We can write: $p_{1,t} = (1 - \delta_1)p_{0,t}$
- And similarly: $p_{2,t} = (1 - \delta_2)p_{1,t}$
- Then: $p_{2,t} = (1 - \delta_2)(1 - \delta_1)p_{0,t}$
- Proceeding iteratively gives: $p_{v,t} = \prod_{i=1}^v (1 - \delta_i)p_{0,t}$

We now have a formula for CF_t , but we still do not know what the depreciation rate **delta** is. If we think about delta as relating the market value of a new good at time t ($p_{0,t}$), with the market value of a 1-year-old good of the same type, again at time t ($p_{1,t}$), then we can proceed iteratively until we reach a general expression. This expression is one step closer to knowing delta (assuming that we know both $p_{v,t}$ and $p_{0,t}$); but still, the expression says that we have a different delta for each year of the good's life, which is highly complex.

Estimating the depreciation rate – II/II

- Given:

$$p_{v,t} = \prod_{i=1}^v (1 - \delta_i)p_{0,t}$$

- The "secret" consists in modelling δ_i . Many options:
 - the **geometric** depreciation model
 - the **straight line** depreciation
 - others not covered here...

The geometric model

- Depreciation rate constant over time: $\delta_i = \delta$
- Market value of age v durable simplifies to: $p_{v,t} = (1 - \delta)^v p_{0,t}$
- Depreciation rate given by: $\delta = 1 - \left(\frac{p_{v,t}}{p_{0,t}}\right)^{\frac{1}{v}}$
- Bottom line: δ can be easily estimated, at least in theory: it only requires information on the market values of homogeneous durable goods of different age, $p_{v,t}$ and $p_{0,t}$.

To overcome this complexity, we need to make assumptions, and to **model delta**. The most common assumption is that of a **geometric depreciation model**: this simply means that we assume delta to be the same every year. This delivers a simple way to estimate delta – at least in theory.

*The straight line model

- Finite economic life. After T years CF falls down to zero. Linear pattern

$$\frac{p_{v,t}}{p_{0,t}} = \begin{cases} \frac{T-v}{T} & \text{if } v \leq T \\ 0 & \text{otherwise} \end{cases}$$

- The depreciation rate increases over time

$$\delta_i = \begin{cases} \frac{1}{T-i} & \text{if } i < T \\ 1 & \text{otherwise} \end{cases}$$

*Depreciation models compared

The optional slides mention some alternative models for estimating delta.

Recap

- User cost is the more appropriate concept to evaluate the consumption flow from durables
- In terms of data requirements, the **geometric depreciation model** is a good compromise
- We need to estimate:
 - Current market value of the durable: $p_{v,t}$
 - Current real interest rate: $r_t = i_t - \pi_t$
 - Depreciation rate: δ

This slide gives an opportunity to recap the user cost approach, and the ‘ingredients’ that are needed to implement it: this is where data comes in. **Questionnaire design must match the analysts’ needs**, and gather enough information to allow for the user cost approach to be feasible in practice.

Data requirements: first best

- Current market value of item of vintage v : $p_{v,t}$
- Current market value of a new item: $p_{0,t}$
- Age v of the durable
- Current nominal interest rate: i_t
- Current yearly inflation rate: π_t

$$CF = (i_t - \pi_t + \delta)p_{v,t} \quad \delta = 1 - \left(\frac{p_{v,t}}{p_{0,t}}\right)^{\frac{1}{v}}$$

Ideally, all of the ‘ingredients’ of the formulas reproduced at the bottom of the slide (the user cost method for estimating the CF from durable goods) would need to be gathered by the questionnaire, except for the nominal interest rate and inflation rate, which are typically available from external sources. If this is the case, the analyst would simply plug survey data into the formula: this is what is meant by **first best**.

Some practical considerations

$$\delta = 1 - \left(\frac{p_{v,t}}{p_{0,t}}\right)^{\frac{1}{v}}$$

- In place of v (age of the durable): **years of ownership** can be a more practical approximation
- In place of $p_{0,t}$ (current market value of a new item): more practical to ask for the **price originally paid for the good when it was purchased**, $p_{v,t-v}$
- The analyst will use an inflation rate π in order to approximate what is needed for estimating delta: $p_{0,t} \approx (1 + \pi)^v p_{v,t-v}$

Realistically, obtaining enough reliable information on all of the ‘first best’ variables is not easy. In practice, a more realistic ‘**second best**’ approach is to ask households about the purchase time, purchase value, and current market value of items they own. The years of ownership of the good can be used as a proxy for the ‘vintage’ or age of the durable, v , if needed. This still allows to apply the user cost method, provided that some changes are made to the formulas: the new expression is reproduced below.

4. How to design a dedicated module in the questionnaire?

- 1) Current market value of item ($p_{i,t}$)
- 2) Price paid in year $t-s$ ($p_{i,t-s}$)
- Age of the durable (v)
- Data requirements for a (practical) first best are met

This is an example of a survey that correctly gathers information needed to implement the second best approach (here called a ‘practical first best’). In this case, we know v , rather than s , which is ok from an analytical standpoint.

Namibia, 2015/16
Household Income and Expenditure Survey (NHIES)

- 1) Current market value of item ($p_{i,t}$)
- 2) Price paid in year $t-s$ ($p_{i,t-s}$)
- 3) Years of ownership (s) or Age of the durable (v)
- We only have the current market value of the item
 - Standard methods cannot be applied

This, on the other hand, is an example of a survey that does not gather enough information to implement even the second best approach.

Some practical considerations

- The years of ownership can be used as a substitute for the age of the durable
- It is uncommon for surveys to collect information on the **current market value of item of vintage v** , and most of the time, what we have instead is the **price paid in $t-s$** . In all these cases, we will need to apply an inflation rate
- When the first best criteria are not fulfilled, alternative methods may exist to achieve a reliable estimation of the durables
- But not always.

Palestine
Expenditure and Consumption Survey, PECS 2011

- Palestine is an **extreme case**
- We only have information about the amount of durables (number of units)
- In those cases, a wise choice is to ignore consumer durable goods and exclude them from the welfare aggregate

Group No	Description of item	Item No.	Total amount last 12 months
501	Furniture		
	Wooden bed	5001	
	Metal bed	5002	
	Wooden tables	5003	
	Wooden chairs	5004	
	Plastic tables	5005	
	Plastic chairs	5006	
	Wooden cupboard	5007	
	Dining room, complete set	5008	
	Kitchen room, complete set	5009	
	Bed room, complete set	5010	

In practice, even if the ‘second best’ approach is not feasible, we can still apply alternative methods, and estimate the necessary variables in an indirect way. In a case like Palestine this is really impossible: the available data is too limited.

Lessons learned

- We are interested in the **use (consumption)** of a durable good, and not in its **value (purchase)**.
- The recommended approach to estimate the value of use is called “the **user cost method**”.
- **Data requirements** depend on the specific method chosen for estimating the so-called consumption flow from durable goods.
- The questionnaire should contain a specific module on ownership of durables.

Lessons learned deserve a few minutes of our time. The takeaways are that durable goods require special treatment when constructing a measure of welfare, because welfare analysts want to proxy consumption; to do that, the user cost method is recommended. The questionnaire should definitely include a module on durable goods, and gather as much information as needed to apply the user cost method.

Homework

Exercise 1 – Engaging with the literature

See exercise 1, Lecture 1.

Exercise 2 – The durable goods module

Ghana, 2017
Ghana Living Standards Survey

SECTION 11: CREDIT, ASSETS AND SAVINGS
PART B: ASSETS AND DURABLE CONSUMER GOODS

ITEM	CODE	Do you own this item?	Year acquired	1. How long ago was this item obtained?		2. What was the purchase price of this item?		3. How much did you pay for this item?	
				Year	Month	Price	Year	Month	Price
Television (color)	101								
Television (not color)	102								
Refrigerator	103								
Washing machine	104								
Washing machine (hand)	105								
Washing machine (auto)	106								
Washing machine (semi)	107								
Washing machine (top load)	108								
Washing machine (front load)	109								
Washing machine (hand)	110								
Washing machine (auto)	111								
Washing machine (semi)	112								
Washing machine (top load)	113								
Washing machine (front load)	114								
Washing machine (hand)	115								
Washing machine (auto)	116								
Washing machine (semi)	117								
Washing machine (top load)	118								
Washing machine (front load)	119								
Washing machine (hand)	120								
Washing machine (auto)	121								
Washing machine (semi)	122								
Washing machine (top load)	123								
Washing machine (front load)	124								
Washing machine (hand)	125								
Washing machine (auto)	126								
Washing machine (semi)	127								
Washing machine (top load)	128								
Washing machine (front load)	129								
Washing machine (hand)	130								
Washing machine (auto)	131								
Washing machine (semi)	132								
Washing machine (top load)	133								
Washing machine (front load)	134								
Washing machine (hand)	135								
Washing machine (auto)	136								
Washing machine (semi)	137								
Washing machine (top load)	138								
Washing machine (front load)	139								
Washing machine (hand)	140								
Washing machine (auto)	141								
Washing machine (semi)	142								
Washing machine (top load)	143								
Washing machine (front load)	144								
Washing machine (hand)	145								
Washing machine (auto)	146								
Washing machine (semi)	147								
Washing machine (top load)	148								
Washing machine (front load)	149								
Washing machine (hand)	150								
Washing machine (auto)	151								
Washing machine (semi)	152								
Washing machine (top load)	153								
Washing machine (front load)	154								
Washing machine (hand)	155								
Washing machine (auto)	156								
Washing machine (semi)	157								
Washing machine (top load)	158								
Washing machine (front load)	159								
Washing machine (hand)	160								
Washing machine (auto)	161								
Washing machine (semi)	162								
Washing machine (top load)	163								
Washing machine (front load)	164								
Washing machine (hand)	165								
Washing machine (auto)	166								
Washing machine (semi)	167								
Washing machine (top load)	168								
Washing machine (front load)	169								
Washing machine (hand)	170								
Washing machine (auto)	171								
Washing machine (semi)	172								
Washing machine (top load)	173								
Washing machine (front load)	174								
Washing machine (hand)	175								
Washing machine (auto)	176								
Washing machine (semi)	177								
Washing machine (top load)	178								
Washing machine (front load)	179								
Washing machine (hand)	180								
Washing machine (auto)	181								
Washing machine (semi)	182								
Washing machine (top load)	183								
Washing machine (front load)	184								
Washing machine (hand)	185								
Washing machine (auto)	186								
Washing machine (semi)	187								
Washing machine (top load)	188								
Washing machine (front load)	189								
Washing machine (hand)	190								
Washing machine (auto)	191								
Washing machine (semi)	192								
Washing machine (top load)	193								
Washing machine (front load)	194								
Washing machine (hand)	195								
Washing machine (auto)	196								
Washing machine (semi)	197								
Washing machine (top load)	198								
Washing machine (front load)	199								
Washing machine (hand)	200								

Data requirements for the second best are met even in the absence of the age v of the durable:

- 1) Current market value of item ($p_{t-s,t}$)
- 2) Price paid in year $t-s$ (p_{t-s})
- 3) Years of ownership (s)

Note a potential problem with the question: “How long ago was [...] obtained”. If many durables were obtained as gifts, then the reported purchase value of goods will often be zero, and analysts may need to impute these values.

Zambia, 2015
Living Conditions Monitoring Survey

Section 7: Household Assets

ITEM	CODE	Do you own this item?	Year acquired	1. How long ago was this item obtained?		2. What was the purchase price of this item?		3. How much did you pay for this item?	
				Year	Month	Price	Year	Month	Price
Television (color)	101								
Television (not color)	102								
Refrigerator	103								
Washing machine	104								
Washing machine (hand)	105								
Washing machine (auto)	106								
Washing machine (top load)	107								
Washing machine (front load)	108								
Washing machine (hand)	109								
Washing machine (auto)	110								
Washing machine (top load)	111								
Washing machine (front load)	112								
Washing machine (hand)	113								
Washing machine (auto)	114								
Washing machine (top load)	115								
Washing machine (front load)	116								
Washing machine (hand)	117								
Washing machine (auto)	118								
Washing machine (top load)	119								
Washing machine (front load)	120								
Washing machine (hand)	121								
Washing machine (auto)	122								
Washing machine (top load)	123								
Washing machine (front load)	124								
Washing machine (hand)	125								
Washing machine (auto)	126								
Washing machine (top load)	127								
Washing machine (front load)	128								
Washing machine (hand)	129								
Washing machine (auto)	130								
Washing machine (top load)	131								
Washing machine (front load)	132								
Washing machine (hand)	133								
Washing machine (auto)	134								
Washing machine (top load)	135								
Washing machine (front load)	136								
Washing machine (hand)	137								
Washing machine (auto)	138								
Washing machine (top load)	139								
Washing machine (front load)	140								
Washing machine (hand)	141								
Washing machine (auto)	142								
Washing machine (top load)	143								
Washing machine (front load)	144								
Washing machine (hand)	145								
Washing machine (auto)	146								
Washing machine (top load)	147								
Washing machine (front load)	148								
Washing machine (hand)	149								
Washing machine (auto)	150								
Washing machine (top load)	151								
Washing machine (front load)	152								
Washing machine (hand)	153								
Washing machine (auto)	154								
Washing machine (top load)	155								
Washing machine (front load)	156								
Washing machine (hand)	157								
Washing machine (auto)	158								
Washing machine (top load)	159								
Washing machine (front load)	160								
Washing machine (hand)	161								
Washing machine (auto)	162								
Washing machine (top load)	163								
Washing machine (front load)	164								
Washing machine (hand)	165								

Here, information is collected on the “most recent” durable. How does this affect the resulting estimate (think about cases in which the household owns more than one durable good for each type)?

Nigeria, 2015/16
General Household Survey

SECTION 5: HOUSEHOLD ASSETS

ITEM CODE	ITEM	NUMBER OF ITEMS	IS CODE	IS CODE	NUMBER OF YEARS	MARKA
1	Refrigerator (17 liter and over)					
2	Refrigerator (12 liter)					
3	Refrigerator (6 liter)					
4	Washing machine					
5	Washing machine					
6	Washing machine					
7	Washing machine					
8	Washing machine					
9	Washing machine					
10	Washing machine					
11	Washing machine					
12	Washing machine					
13	Washing machine					
14	Washing machine					
15	Washing machine					
16	Washing machine					
17	Washing machine					
18	Washing machine					
19	Washing machine					
20	Washing machine					
21	Washing machine					
22	Washing machine					
23	Washing machine					
24	Washing machine					
25	Washing machine					
26	Washing machine					
27	Washing machine					
28	Washing machine					
29	Washing machine					
30	Washing machine					
31	Washing machine					
32	Washing machine					
33	Washing machine					
34	Washing machine					
35	Washing machine					
36	Washing machine					
37	Washing machine					
38	Washing machine					
39	Washing machine					
40	Washing machine					
41	Washing machine					
42	Washing machine					
43	Washing machine					
44	Washing machine					
45	Washing machine					
46	Washing machine					
47	Washing machine					
48	Washing machine					
49	Washing machine					
50	Washing machine					
51	Washing machine					
52	Washing machine					
53	Washing machine					
54	Washing machine					
55	Washing machine					
56	Washing machine					
57	Washing machine					
58	Washing machine					
59	Washing machine					
60	Washing machine					
61	Washing machine					
62	Washing machine					
63	Washing machine					
64	Washing machine					
65	Washing machine					
66	Washing machine					
67	Washing machine					
68	Washing machine					
69	Washing machine					
70	Washing machine					
71	Washing machine					
72	Washing machine					
73	Washing machine					
74	Washing machine					
75	Washing machine					
76	Washing machine					
77	Washing machine					
78	Washing machine					
79	Washing machine					
80	Washing machine					
81	Washing machine					
82	Washing machine					
83	Washing machine					
84	Washing machine					
85	Washing machine					
86	Washing machine					
87	Washing machine					
88	Washing machine					
89	Washing machine					
90	Washing machine					
91	Washing machine					
92	Washing machine					
93	Washing machine					
94	Washing machine					
95	Washing machine					
96	Washing machine					
97	Washing machine					
98	Washing machine					
99	Washing machine					
100	Washing machine					

In this case, standard methods cannot be applied.

- 1) Current market value of item ($p_{t-s,t}$)
- 2) Price paid in year $t-s$ (p_{t-s})
- 3) Years of ownership (s)

Alternative methods exist when we only have the current market value of the durable, and an estimation of the **maximum economic life of the durable**.

Malawi, 2016/17
Integrated household survey

SECTION 4: DURABLE GOODS

ITEM CODE	ITEM	NUMBER OF ITEMS	IS CODE	IS CODE	NUMBER OF YEARS	MARKA
1	Refrigerator (17 liter and over)					
2	Refrigerator (12 liter)					
3	Refrigerator (6 liter)					
4	Washing machine					
5	Washing machine					
6	Washing machine					
7	Washing machine					
8	Washing machine					
9	Washing machine					
10	Washing machine					
11	Washing machine					
12	Washing machine					
13	Washing machine					
14	Washing machine					
15	Washing machine					
16	Washing machine					
17	Washing machine					
18	Washing machine					
19	Washing machine					
20	Washing machine					
21	Washing machine					
22	Washing machine					
23	Washing machine					
24	Washing machine					
25	Washing machine					
26	Washing machine					
27	Washing machine					
28	Washing machine					
29	Washing machine					
30	Washing machine					
31	Washing machine					
32	Washing machine					
33	Washing machine					
34	Washing machine					
35	Washing machine					
36	Washing machine					
37	Washing machine					
38	Washing machine					
39	Washing machine					
40	Washing machine					
41	Washing machine					
42	Washing machine					
43	Washing machine					
44	Washing machine					
45	Washing machine					
46	Washing machine					
47	Washing machine					
48	Washing machine					
49	Washing machine					
50	Washing machine					
51	Washing machine					
52	Washing machine					
53	Washing machine					
54	Washing machine					
55	Washing machine					
56	Washing machine					
57	Washing machine					
58	Washing machine					
59	Washing machine					
60	Washing machine					
61	Washing machine					
62	Washing machine					
63	Washing machine					
64	Washing machine					
65	Washing machine					
66	Washing machine					
67	Washing machine					
68	Washing machine					
69	Washing machine					
70	Washing machine					
71	Washing machine					
72	Washing machine					
73	Washing machine					
74	Washing machine					
75	Washing machine					
76	Washing machine					
77	Washing machine					
78	Washing machine					
79	Washing machine					
80	Washing machine					
81	Washing machine					
82	Washing machine					
83	Washing machine					
84	Washing machine					
85	Washing machine					
86	Washing machine					
87	Washing machine					
88	Washing machine					
89	Washing machine					
90	Washing machine					
91	Washing machine					
92	Washing machine					
93	Washing machine					
94	Washing machine					
95	Washing machine					
96	Washing machine					
97	Washing machine					
98	Washing machine					
99	Washing machine					
100	Washing machine					

The questionnaire only collects information about the price paid if the durable was acquired *in the last 12 months* (in year t !).

- 1) Current market value of item ($p_{t-s,t}$)
- 2) Price paid in year t (p_t)
- 3) Age of the durable (v)

No information is collected on the price paid in year $t-s$; however, note that, in this case, p_t and $p_{t-s,t}$ will probably be very similar value. This case is similar to **Nigeria**, but the fact that there is information on the age of the durable puts the analyst in a better position to use alternative methods for estimating the CF.

Lecture 10

Housing

Learning objectives

The goal of this lecture is to provide detailed recommendations on the collection of data on housing expenditures. The lecture explains how analysts treat this category of expenditures when constructing a measure of living standards; these analytical requirements guide the design of the survey module on housing.

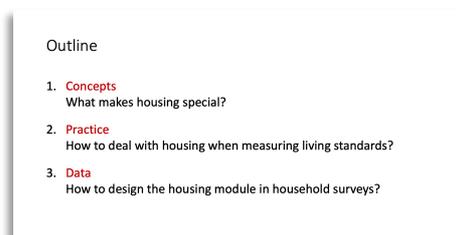
Suggested preparation

The main reading for the lecture is Malpezzi (2002), which focuses on data collection. For a thorough understanding of the analytical requirements underlying the design of the questionnaire, it is also important to be familiar with Deaton and Zaidi (2002), Section 3.5, and Balcázar et al. (2017).

Time allocation

Concepts	30 min
Practice	30 min
Break	15 min
Data	40 min
Lessons learned	5 min

Annotated lecture



Outline

1. **Concepts**
What makes housing special?
2. **Practice**
How to deal with housing when measuring living standards?
3. **Data**
How to design the housing module in household surveys?

The set opens with a roadmap of the lecture, which is organized around three topics: concepts, practice, and data.

1. Concepts

What makes housing special?

12

- Conceptually, a house is a perfect example of a **durable good**.
- Empirically, housing matters. Housing expenditures absorb between 10-30% of total household expenditure.

Conceptually, **housing is a durable good**, and presents the same issues as those discussed in the previous lecture.

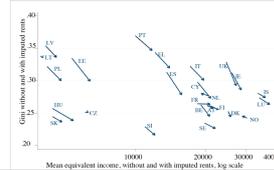
The European experience

Törmälehto and Hannele Sauli (2010)



- Nearly 80% of European households either own their main residence or their rent is below the prevailing market rent
- Imputed rent decreases **inequality** in all (but two) European countries
- A similar result applies to income-based measures of the at-risk-of **poverty** rate

Changes in **income inequality** and **average income** without imputed rent → with imputed rent, 2007



These slides present evidence of the **empirical relevance of housing expenditures**, in the context of welfare measurement: inequality and poverty estimates are significantly affected by housing. The graph from Törmälehto and Sauli (2010) shows, for a selection of European countries, the trajectory of two variables, income inequality (measured by the Gini index, vertical axis) and average income (horizontal axis), before and after including imputed rent in the welfare aggregate: most countries move towards south-east, that is, including imputed rent decreases estimated inequality and increases average incomes.

Why collect data on housing?

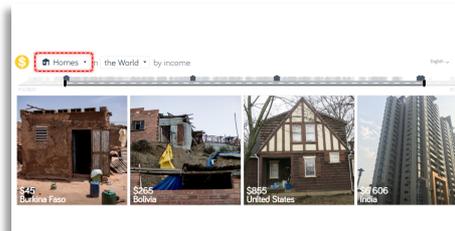
Motivation #1:

Housing characteristics are direct indicators of the household's standard of living.

Dollar street, useful for illustrating

You can check it out here:

<https://www.gapminder.org/dollar-street/matrix>



Data on housing are interesting for several reasons. One is that housing characteristics are themselves an **indicator of living standard**: Dollar Street, a project developed by the Gapminder Foundation, is used to illustrate how dwelling characteristics go hand in hand with well-being. The example of “flying toilets” is a vivid way of communicating the importance of some of these housing features for the standard of living of families.

Why collect data on housing?

Motivation #2:

Housing consumption must be accounted for properly when defining **living standards** and comparing them across households.

Another reason for collecting housing data is that housing is a part of total household consumption, and therefore we need to **account for it in the consumption aggregate**.

Why collect data on housing?

Motivation #3:

- Understanding **housing market** behavior, to help analysts and policymakers understand how housing markets work and how government policies affect housing outcomes.
- See [Malpezzi \(2002: 295\)](#) for more details

▪ In this lecture, the focus is on motivations **1** (housing characteristics) and **2** (housing consumption)

Finally, data on housing expenditures are of interest for analysts that study the **behavior of housing markets**. Of course, our focus is on welfare analysis, therefore motivations 1 and 2 will drive the rest of the discussion.

Some general implications for data collection

- The questionnaire should contain a dedicated **housing module**
- The module should collect data on (at least):
 1. The **characteristics of the household's dwelling**
 2. Expenditures on **utilities**
 3. All pieces of information needed to estimate the **use-value of the dwelling** (the flow of housing services)
- Point 3 requires further elaboration

Motivations 1 and 2 imply that the housing module must contain at least the pieces of information here listed. But what is the **use-value of the dwelling**? The next few slides explain.

A key general principle

- The theory covered for durable goods applies to housing, too
- We are not interested in the **purchase value** of the house: we want the **value of using the dwelling** during the survey period (**flow of housing services**)
- Take three households, **A, B, and C**, living in identical homes. All other things being equal, they should be classified as equally well off (they enjoy the same flow of housing services)
- Imagine **A** pays market rent, **B** owns the home, while **C** pays subsidized rent



How to estimate the flow of housing services?



- Easy
 - In principle, **actual rent** paid is a good proxy for the flow of housing services during the survey period
 - Most surveys collect data on it
- Difficult
 - Owners do not pay rent...
 - Need to estimate the price that owners **would pay** if they had to rent their home
 - This is what we refer to as **imputed rent**
- Difficult
 - If rent is subsidized, it does not represent the actual value of services enjoyed from residing in the dwelling, but something less than that
 - We need **imputed rent**

Theoretical considerations here are very similar to those that apply to durable goods. Welfare analysts are interested in the **value of consumption** of durable goods, including the dwelling, not in their purchase value. What differentiates a house from other durables is the fact that **rent** is usually available, and as mentioned in the previous lecture, rent is an adequate representation of the use-value of a good. For households that rent their dwelling, the case is closed. However, households that own the dwelling do not pay any rent, of course. A similar issue arises when a household pays a rent that is not representative of the true value of their dwelling, but is artificially lowered (by subsidies, for instance). In both of these cases, what we need is the rent that owners or non-market tenants *would* pay, if they were to rent their home on the housing market. This is called **imputed rent**.

The importance of imputed rent

- **Imputed rent** is an estimate of the value of the benefit accruing to the household due to not paying full rent
- It is crucial for **consistent welfare comparisons**: without imputed rent, A, B and C would appear to have different living standards, when in fact they are identical in everything but housing tenure status
- **Homeowners** and **non-market tenants** (households receiving housing free of charge or at rates subsidized by their employers, friends, relatives, the government) require **special attention**

The case of Egypt



- A common situation is the presence of housing **market regulation, rent controls**
- The presence of regulated housing in a market creates a set of tenants who benefit from housing at a subsidized price.
- As we saw in lecture 7, this requires a **correction**

These slides reinforce the concept of imputed rent and its importance for making correct welfare comparisons. Egypt is a real-life case in which rent controls (legally enforced lowered rents for certain tenants) pose this kind of problem to analysts.



Now is a good time for a break.

2. Practice



The challenge of imputing rent

Three main options:

1. **Self-reported rent**
owners are asked the "implicit rental value", that is, how much it would cost to rent their unit on the market
2. **Hedonic housing regression**
Regress actual rents on dwelling characteristics, and predict the rent that owners would pay if they had to rent their unit
3. **Non-hedonic methods**
apply a capitalization rate to the self-reported current value of the unit

Today we focus on 1. and 2.

In practice, computing imputed rent is **not trivial**. Different methods are available. Students are directed to the article by Balcazar et al. for a complete review: this is currently the most comprehensive and updated source on the topic, and it is a fundamental reading for this lecture. In fact, it will be difficult to comprehend the rest of the discussion on housing without having gained at least some familiarity with it. There are three main options for estimating imputed rent, but we will focus on the first two (the last one is less often used in practice).

Self-reported rent

- Respondents are asked to estimate how much it would **cost** to rent their home at full price (willingness to pay, **WTP**)
- Alternative: how much they expect to **receive** as rent for their home (willingness to accept, **WTA**)
- Usually, **WTA > WTP**
- WTP may mitigate misreporting due to 'owner pride' (Ceriani et al 2019: 9)

009. If someone wanted to rent this dwelling (or one exactly like it) today, how much would (s)he rent? (Rent on the open market)

.....

Lesotho, 2017/18
Household Budget Survey (from the Continuous Multi-Purpose Household Survey)
Household Questionnaire (p.61)

Reliability of self-reported rent

- This approach relies on the assumption that owners are **informed and objective** about the value of their dwelling, and the amount they would have to pay to rent a home with similar quality and location attributes
- In practice, this assumption may be unrealistic:
 - **"Thin" rental markets**
No comparable dwellings rented in the area in which respondents live, no information
 - **"Owner pride" factor**
Homeowners may have above-market valuations of their housing, based on subjective reasons, such as special attachment to specific characteristics of their homes

The first approach, **self-reported rent**, relies entirely on the collection of data directly from respondents, by asking homeowners a question of the type: "How much would you pay if you were to rent the dwelling you are currently living in?". This question records the respondent's willingness to pay (WTP) for hypothetically renting her own dwelling;

another approach is to record the respondent’s willingness to accept (WTA) payment in exchange of her dwelling, with a question like “How much would you receive as a rent if you were to lend your apartment?”. The two questions are very similar, and both are used in questionnaires around the world. Ceriani et al (2019) cites some of the literature comparing WTP and WTA approaches, and mentions that usually, WTA questions yield higher reported amounts than the equivalent WTP questions, “due to the fact that a person’s perceived value of an item is often higher when the person owns the item (i.e., the role of the landlord), then when she does not own it (i.e. the role of the tenant). This would suggest that the second version of the question (WTP), may mitigate the possibility of misreporting due to sentimental attachment to the property.” (p. 9).

The self-reported rent question should always be included in the housing module; however, there may be challenges to the quality of the data collected in this way, when respondents are not informed (the housing market in the area does not exist, or is very small, so respondents have no information to base their response on) or not objective (the perceived value of one’s own dwelling is unrealistic).

Hedonic housing regression

- The general idea is to assume that **rent is a function of the characteristics of the dwelling**, including location, structural attributes (e.g. type of construction, number of rooms, age of the building, etc.) and neighborhood characteristics
- Focusing on market tenants, the **relationship** between dwelling characteristics and rent can be estimated (for instance: a house with tile floors goes for a rent that is x% higher than average, all else equal)
- This relationship is then used to **predict** the implicit rental value for households who do not pay (full) rent for their homes, based on their dwelling’s characteristics

The econometric model

- A popular choice is to use a **log-linear functional form**:

$$\ln \text{rent}_i = \beta_0 + \beta_1 x_{1i} + \dots + \beta_k x_{ki} + \epsilon_i$$
 where $\ln \text{rent}_i$ is the natural logarithm of rent (actual and/or self-assessed by owners), and x_{ki} is housing characteristics (number of rooms, roof, floor, wall, type of toilet, location variables...)
- Predict for the rest of the population:

$$\widehat{\text{rent}}_i = \exp(\beta_0 + \beta_1 x_{1i} + \dots + \beta_k x_{ki} + \epsilon_i)$$

An alternative is to estimate rent based on other information, namely the characteristics of the dwelling (a nicer house would go for more on the rental market than a dilapidated one: if the analyst is able to select the characteristics that influence market rent, she can “predict” it). This approach is called **hedonic regression**, usually implemented as a log-linear specification.

* Duan’s estimator

Smearing Estimate: A Nonparametric Retransformation Method

NARAYA DUAN*

The smearing estimate is proposed as a nonparametric estimator of the expected response on the transformed scale after fitting a linear regression model on a transformed scale. The estimate is consistent under mild regularity conditions, and usually causes high efficiency relative to parametric estimates. It can be viewed as a low-parametric bootstrap procedure for estimating the variance of the smearing estimate. A real world example of predicting medical expenditures shows that the smearing estimate can outperform parametric estimates when the parametric assumption is nearly violated.

KEY WORDS: Bootstrap methods; Transformation; Nonparametric; Prediction; Lognormal linear model; Classification; Regression analysis.

The essence of the procedure is to estimate the unknown error distribution by the empirical of the estimated regression residuals, and then take the desired expectations with respect to the estimated error density. In a broader context, the method can be viewed as an application of the bootstrap principle (Efron 1979). In the next section, we present the nonparametric problem, and use an example to demonstrate the possible use of the smearing estimate in the context of the prediction of medical expenditures. We then discuss the smearing estimate as an estimator of the transformed scale expectation from the distributional assumptions on the error distribution.

* The consistency property of the smearing estimate is established in Section 4. © Copyright, 1984, by Springer-Verlag.

*Other methods

- Implicit rent is seen as the **rate of return** that would have been obtained by owners if the home equity had been invested in an interest-bearing account
- Questionnaire asks respondents for a **self-reported current market value** of the dwelling (how much would your home sell for on the market?)
- Different approaches to **estimate capitalization rate** (including user-cost approach, as for other durable goods)

These optional (starred) slides mention two extensions of the topic that was just discussed. **Duan’s estimator** provides more accurate predictions of rents estimated via hedonic regression, for reasons having to do with the logarithmic transformation of the dependent variable. The Instructor is invited to read the paper for more details. The second slide hints at another available method besides hedonic regression, which may be used if all else fails.

Main takeaways

- Estimating an “implicit” rent for all those households who do not pay actual market rent is one of the **main challenges** facing welfare analysts
- **Several estimation approaches**, based on different assumptions: choice depends on context
- Methodology aside, success of estimation rests on the availability of the necessary information from surveys (**data availability**) and its accuracy (**data quality**)
- **Self-reported rent** and **dwelling characteristics** emerge as crucial data requirements

Now is the time to recap the main messages of this section: the last bullet is especially important, and leads to the final part of the lecture.

3. Data

Examples of definitions for the housing module

Malpezzi (2002: 307)

Definitions are **context-specific**: this is just a potential starting point

- **Structure**
“A structure is a physically separate entity such as a house, an apartment building, or a tent. It may contain one or more dwelling units.”
- **Dwelling unit**
“A dwelling is an accommodation unit that contains one or more households. There may be several dwellings in a structure.”
- **Room**
“Whole rooms used for living purposes, such as living rooms, dining rooms, bedrooms (...). Not included are bathrooms (...). If a room is used by occupants of more than one unit, the room is included with the unit from which it is most easily reached.”

These slides offer a few practical recommendations for the **design of the housing module**, mostly from Malpezzi (2002). First, it should be clear that concepts must be defined, however familiar they may seem.

Components of the housing module

Malpezzi (2002: 305-310)

Given our focus on living standards measurement, we summarize recommendations for the following components of the housing module:

1. Dwelling characteristics and housing services
2. Dwelling expenditures

Dwelling characteristics and housing services

- Characteristics of the **structure or dwelling**
 - Building materials (roof, walls, floor...)
 - Age and size of the structure
 - Number of rooms/bedrooms/bathrooms
- Characteristics of the **neighborhood**, availability of **services**
 - Location of the dwelling
 - Availability and distance from services (water, sanitation...)

Next, what should the housing module should ask for, in practice? **Dwelling characteristics** are crucial.

Dwelling expenditures

- a. Questions about expenditures are closely linked with questions about **housing tenure**
- b. Main expenditure to be recorded: **rent** (actual and self-reported)
- c. Other housing expenditures: **utilities, maintenance and repairs**

Housing tenure

- Property rights and tenure vary considerably depending on the country's context: these questions must be **customized**
- At a minimum, questionnaire should differentiate **owning vs. renting**
- **Length of tenure** is important because it often impacts rent paid

Rent

- **Actual rent** (for renters)
 - “How much does the household pay towards rent?”
 - Data must be collected on “**arms-length transactions**”, i.e. transactions between two counterparts who have no special relationship that would suggest that rent paid differs from market prices
 - Crucial to differentiate between households that pay market rent, and households under rent controls or subsidy, related to the landlord, etc.
- **Self-reported rent** (for non-renters)
 - “How much would you charge if you were to rent out this dwelling?”

Utilities, maintenance and repairs

- Distinction between housing expenditures **inclusive of utilities or not** is crucial
- *E.g.* some renters pay for utilities separately, but others pay a monthly rent that includes utility charges
- Questionnaire must be designed to distinguish between these cases
- Expenditure for utilities may be collected elsewhere in the questionnaire (together with other expenditures with same recall period)

In terms of **dwelling expenditures**, the crucial pieces of information are tenure status, rent, and utilities/maintenance/repairs.

The image shows a survey form for the Housing Ghana Living Standards Survey (2012/2013). It is divided into two main sections: 'PART 1: TENURE CHARACTERISTICS' and 'PART 2: HOUSING CHARACTERISTICS'. The form contains various questions and checkboxes related to housing tenure, ownership, and characteristics.

The final slides show **examples** of housing modules in recent HCES. Students are invited to examine them and point out which questions collect the different pieces of information mentioned during the lecture.

 **Lessons learned**

- Analysts need data on **housing characteristics** – they are direct indicators of the household's standard of living.
- Analysts need data on **housing consumption** for inclusion in the consumption aggregate – not an easy task.
- Renters pay rent. For non-renters, we estimate **imputed rent**, the value of the benefit accruing to the household from living in its dwelling.
- Collecting data on **self-reported imputed rent** is key.
- Hedonic regression** is the recommended approach to estimating imputed rent, when self-reported rent is not available, or not reliable.

As usual, the Instructor is encouraged to take a few minutes to recap the main messages of the lecture. In terms of data collection, the main targets for those designing questionnaires are self-reported rent (it is fundamental for welfare analysts) and dwelling characteristics (both because of their use as direct indicators of living standards, and as regressors in hedonic models).

Homework

Exercise 1 - Engaging with the literature

See exercise 1, Lecture 1.

Exercise 2 – Secondary residences

This exercise gives students an opportunity to explore the housing module in a number of different questionnaires, and pushes them to examine their structure in depth, by asking a question that was not covered during the lecture. Each survey will of course collect data on secondary residences in a slightly different way: rather than the answer, what is important here is the students' hands-on engagement with the questionnaire.

Exercise 3 – Housing in theory and practice

Similarly to Exercise 2, this question asks students to comment on questionnaires, this time pushing them to make a connection with the theory. Ultimately, students should be able to conclude (i) whether or not the design of the housing module allows for the implementation of at least one of the recommended methods for the estimation of the use-value of the dwelling, and (ii) what was actually done when constructing the welfare indicator.

Lecture 11

Data validation and diagnostics

Learning objectives

This lecture deals with data validation techniques and data diagnostics. The main goal of the lecture is to provide a gentle introduction to a variety of problems that arise when survey data are collected in the field and processed by the NSI prior to dissemination with final users. The lecture is closely linked to Lecture 12, which focuses on outliers, as the latter are often a specific form of (gross) errors, one of the many covered in this lecture.

Suggested preparation

The single most important reference for this lecture is the manual by de Waal et al. After becoming familiar with the sections indicated in the syllabus, Instructors are encouraged to customize the references with country specific material, as appropriate.

The number of slides in this lecture is intimidating. Many are examples, which only require a few seconds of discussion. Nevertheless, the Instructor is invited to check whether cuts are needed *before* delivering the lecture.

Time allocation

Types of errors	20 min
Data editing	15 min
Missing data	25 min
Break	15 min
Data validation and diagnostics	40 min
Lessons learned	5 min

Annotated lecture

Preamble

- Each step of a **survey** can generate **errors** in the data or the published statistics.
- It has been estimated that national statistical institutes spend some 40% of their resources on identifying errors and fixing them.
- This lecture focuses on **errors**, **missing values** (item nonresponse, or a respondent may give a wrong answer), and **unit nonresponse** (a respondent does not answer at all).
- They are largely responsible for overall **data quality**.

A bird's-eye view of the survey process
today's focus is on steps 3 and 4

1. **Setting survey objectives**
(lecture 1)
2. **Questionnaire design and sampling design**
3. **Data collection and data entry**
the sample is drawn, data are collected from the sampled units and entered into the computer system at the statistical office
4. **Data processing and data analysis**
collected data are edited, missing and erroneous data are imputed, raising weights are determined
5. **Publication and data dissemination**
(lecture 15)

We begin by acknowledging that errors are ‘part of the game’. After showing the first bullet of the first slide we can open the floor and engage with a discussion with students, asking them what kind of errors they can think of. We do not reply, we simply ‘chair’ the discussion. Then, we remind students that this lecture focuses on errors that arise in steps 3 and 4 of the survey process.

Types of errors

Errors can be classified in several ways:

1. **systematic** vs. **random** errors
2. **influential** vs. and **noninfluential** errors
3. **outliers** vs. **nonoutliers**
but outliers are not necessarily errors: more on this in lecture 12

Fatal errors

- There is no definition of ‘fatal error’ in statistics
- We borrow the term ‘fatal error’ from computer science: if you get a fatal error, you generally cannot recover from it, because the computer encounters a problem it cannot resolve.
- So, what could be a ‘fatal error’ in our context, when we begin the survey on day one, with the data collection on the field?

Next, we provide students with a quick overview of different types of errors. The goal is twofold: i) to introduce a few key *concepts* on the different types of error and ii) to build up the students’ statistical vocabulary.

Interviewer falsification of survey data

- A **subtle** form of falsification may consist in surveying a **wrong** household member, or in conducting the survey by telephone when face-to-face interviews are required.
- A **severe** form of falsifying is the **fabrication** of entire interviews without ever contacting the respective household.
- Fabricated interviews can have serious consequences for statistics based on the survey data.

Prevention
3 layers

1. **Study design**
Ensuring reasonable interview length
Creating positive work conditions and avoiding unrealistic production quotas
Interviewer compensations on a per hour, not a per interview basis
2. **Interviewer selection**
Employment of interviewers with personal interest in the data quality (e.g., students)
Little interviewer experience to reduce likelihood of knowledge on cheating opportunities in the system
3. **Subsequent interviewer inspection**
Interview verification via random checks

To liven up this ‘catalogue of errors’, which is useful but potentially a bit dull, we touch on the issue of **interviewer falsification** of survey data. This is a potentially serious threat to data quality.

Ruling out horror stories

Detecting Problems in Survey Data Using Benford’s Law

George Judge
Laure Schechter

ABSTRACT
“Do IOWAs Really Do You Know where your respondents are?”

Judge and Schechter (2009)

- “**Horror stories** are common in which somebody discovers that one (or more) enumerator **answers the survey himself** rather than actually interviewing households (...)”
- It would be useful to become aware of this as early as possible in the data collection process.

Table 6
Correlations (ρ), R^2 in Statistics, Distance d^2 , and χ^2 , χ^2 Tests, and Kraper V_n^2 Tests Between Benford’s Law and CIPP Question Products

Survey	Observations	ρ	R^2	χ^2	χ^2	V_n^2
Handbook (1966)	4,521	0.95	0.90	0.00	0.00	0.00
Chicago (1966)	294	0.99	0.99	0.00	0.00	0.00
Chicago (1967)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1968)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1969)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1970)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1971)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1972)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1973)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1974)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1975)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1976)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1977)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1978)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1979)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1980)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1981)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1982)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1983)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1984)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1985)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1986)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1987)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1988)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1989)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1990)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1991)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1992)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1993)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1994)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1995)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1996)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1997)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1998)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (1999)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2000)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2001)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2002)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2003)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2004)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2005)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2006)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2007)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2008)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2009)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2010)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2011)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2012)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2013)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2014)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2015)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2016)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2017)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2018)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2019)	10,020	0.99	0.99	0.00	0.00	0.00
Chicago (2020)	10,020	0.99	0.99	0.00	0.00	0.00

We provide the Instructor with the option to take a *short* digression on a specific tool, **Benford’s Law**, that can help motivate the audience. Benford’s Law *per se* is a topic that should not take more than a couple of slides in the context of this course – we prepared 6. The reason is that this is an engaging story, it keeps the audience’s attention up and helps to ease into the rest of the discussion. In short, the 6 slides are a rhetoric device, that should be used only by Instructors familiar with the references provided in the syllabus, and confident that the opportunity cost of this longer-than-strictly-needed presentation is not too high.

Data editing and imputation

- The occurrence of nonresponse and errors makes it necessary to carry out an extensive process of checking the collected data, and, when necessary, correcting them.
- This checking and correction process is referred to as "statistical data editing and imputation".



Other errors

- Missing data are a special type of error and must be handled with great care
- Missing data are well-known troublemakers for both data producers (NSIs) and the analysts, and will therefore deserve a bit of our time today

The next bunch of slides discusses the concept of ‘data editing’. Slides are self-explanatory, we believe and can be improved by adding specific examples, if the Instructor wishes so. Our advice is emphasise and elaborate on *over-editing* – it is a common bad practice and the discussion in De Waal et al. deserves some attention.

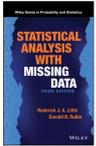
Afghanistan

Living conditions survey 2016-17, Percentage missing values for selected variables

Variable	Base population	Percent missing values
Individual-level variables		
Worked in business, organisation (12.2)	83,788	0.9
Person worked last month (12.6)	83,788	0.9
Days worked in past week (12.14)	34,772	2.1
Industry (12.16)	34,583	2.3
Place of birth (13.2)	195,680	0.7
Literacy (11.2)	121,820	0.9
Attended normal school (11.8)	121,820	0.8
Highest education grade completed (11.8)	92,422	2.1
Currently attending school (11.6)	41,788	1.7
Seeing disability (24.2)	195,680	0.2
Women ever had a live birth (25.7)	179,621	3.2

A useful reference

Little and Rubin (2019)



- Definition**
"Missing data are unobserved values that would be meaningful for analysis if observed; a missing value hides a meaningful value."
- Mechanism**
Why are data missing?
Different mechanisms lead to different strategies for treatment

The discussion of **missing data** is kept within the space of about a dozen slides. Given the relevance and complexity of the topic, the choice of contents is clearly a challenge. We lead with a fundamental idea: the single most important priority when tackling missing data is to understand the *mechanism* driving the missingness itself, that is, why are data missing in the datasets? Different answers imply different coping strategies. The example of the Afghanistan Living Conditions Survey 2016 helps to demonstrate the fact that missing data are ubiquitous, and shows a case of transparency and disclosure of this issue in the official documentation. The reference highlighted here – the manual by Little and Rubin – is a classic, and students are highly encouraged to make use of it.

What causes missing data?

- This is a difficult questions
- While the reality is complex, we can grasp the essence of the discussion by focusing on two different mechanisms
 - "missing completely at random" (MCAR)
 - "missing not at random" (MNAR)

How to decide whether data are MCAR or MNAR?

- Investigating the pattern of missingness in the sample is paramount, before embarking on any action (data editing, dropping records, no action at all).
- Even two-way tables where the distribution of missing values is examined by region, urban-rural areas, and other dimensions are often insightful, despite their simplicity.
- When data are MAR, a useful rule of thumb is that hot-deck techniques are more suitable for elementary expenditure items, while regression-based imputation methods are a better option for large expenditure sub-aggregates.
- If there is evidence that data are MNAR, then the problem is serious, and ad-hoc procedures are required.

Next, we illustrate and discuss the two extremes of the possible missing data **mechanisms**: data could be MCAR, that is, missing completely at random, or MNAR, missing not at random. There is a third case, that of MAR – data missing at random, or data that are MCAR *conditionally* on some observed characteristic – that is not explicitly recalled here, but the Instructor might decide to delve deeper if needed, with the help of the references associated with this lecture. In general, determining which mechanism is behind the observed pattern of missingness is challenging. In practical situations, even simple two-way tables can help shed light on the issue. Making up one’s mind on the mechanism at

work is important, because the optimal course of action to treat missing data differs under MCAR or MNAR.

How to deal with missing data?

- It depends.
- If data are **MCAR**, then observed data can be thought of as a random sample of the complete data, and statistical inference can be carried out based on "complete cases". Simply put, **missing data can be ignored**.
- If data are **MNAR**, the mechanism is referred to as non-ignorable missingness: observed data cannot be treated as if they were a random sample of the complete data. Standard estimation methods would produce biased estimates.

...

Model-based methods
strategy IV/IV

- **Strategy 4** - A broad class of procedures is generated by defining a **model for the complete data** and basing inferences on the likelihood distribution under that model
- Proper treatment of these methods would require an entire course...

The remaining slides deal with the **treatment** of missing data. To impute or not to impute? In general, when data are MCAR, one can effectively ignore the missing observations and perform all analyses on the subset of complete cases: estimates will remain unbiased, albeit less efficient. On the other hand, when data are MNAR, more complex strategies must be deployed. If the analyst were to proceed as if under MCAR, estimates would be biased. We briefly mention the main strategies available to analysts – but the Instructor can move relatively quickly through these slides, as they are not designed to offer a complete presentation of such a complex topic. Again, interested students can be directed to the references. In particular, regarding weighting procedures, a clear and effective introduction can be found in Deaton (1997: section 1.4).

Recap

- **Fatal errors**
forewarned is forearmed
- **Systematic vs. random errors**
randomness is preferable
- **How serious is the incidence of missing values?**
Check and document
- **Is there any pattern in the data missingness?**
MCAR vs. MNAR
- **How to deal with missing values?**
Should we ignore them? Or should we impute them?
It depends on the mechanism (MCAR vs. MNAR)

A recap of the discussion so far leads us to the break.



Now is a good time for a break.

Data validation: what is it?

1. **Range checks**
simplest edit one can think of (details coming shortly)
2. **Internal consistency checks**
combination of edits
3. **Outlier detection**
investigation of extreme values (next lecture)
4. **Other data quality checks**

...

Example
Namibia Household and Expenditure Survey
2015/2016 report

Consistency between "Relationship to the head of the household" and "Marital status":
if the relationship to the head of the household is "Spouse", then the individual must be married or in union; if there is a "Spouse/Partner" in the household, then the head of household must be married or in union.

After the break, we introduce the last topic of the lecture: **data validation and diagnostics**. We start by defining the meaning of these terms, and we provide a selection of examples that can be discussed with students interactively. “Range checks” are usually simple to cover (here a reminder of the discussion on data capture methodologies might be in order), and so are the examples of “consistency checks”. Emphasis here is not much on specific examples, but rather on the fact that there is a need to introduce complex checks involving multiple variables, possibly collected in different sections of the questionnaire.

Other checks

How confident are we about this interpretation of the Whipple index?

- What if, rather than data quality, Whipple's index captures literacy, numeracy, incentives... ultimately, living standards?
- Areas where school enrolment is low, illiteracy is high, income is low, are typically those where age heaping is more pronounced
- Incentives to remember one's exact age may be low: people in poor, isolated, rural areas may have little use for it in their everyday lives
- That doesn't always mean they cannot respond accurately to other questions

The final part of the lecture introduces two tools that are often used as a check of **overall data quality**. The first tool simply consists of a visual inspection of the population pyramid, as estimated by the survey. Does it show suspicious dents or spikes, or any other anomaly? If so ... how to interpret it? If students are not too tired, this is a nice topic for discussion. The second tool focuses on the age variable, and calculates, the so-called Whipple index, a measure of age heaping (the tendency to report one's age as an approximate, 'round' number, namely a multiple of 5). Despite its popularity, the interpretation of this index is controversial. If age heaping is common in the data, this may say less about data quality, and more about the respondents' background, education, habits – ultimately, living standards. The Instructor might wish to discuss pros and cons with students, and end the lecture leaving it as an open issue.

Lessons learned

- Errors are ubiquitous
- Data validation techniques are meant to localize and fix errors
- A special type of error is missing data
- The key issue is whether the reasons for missingness are related to the outcome of interest. When data are MCAR, the impact of missing data is relatively benign. When data are MNAR, then ignoring missing data would lead to biased estimates.
- Imputation of missing values – best method depends on the nonresponse mechanism.
- Data validation and diagnostic should be routine
- A proper documentation of the validation process is an integral part of the metadata to be published

The lecture closes with a plain recap of the main points covered: a) multifaceted and pervasive errors (not just probabilistic, as we have seen, but a long catalogue), b) among errors, missing data deserve special attention. As many problems as there are solutions: this is how we used the second part of the lecture, discussing data validation and imputation methodologies.

We conclude with an important message: the need to produce (and/or ask for) proper and detailed **documentation** of all choices made.

Homework

Exercises 1, 2 and 3 - Engaging with the literature

See exercise 1, Lecture 1.

Exercise 4 – Data imputation

Instructors are encouraged to design and implement this exercise as they see it appropriate. One way is to ask students to write a short essay discussing the pros and cons of the imputation procedures described in the reports in the slide (but others can be chosen).

Lecture 12

Outlier detection and treatment

Learning objectives

This lecture is an extension of the topics discussed in Lecture 11 on data validation and diagnostics. The goal of this lecture is to provide a conceptual framework and some techniques to detect and treat outliers (*i.e.* extreme values), which are commonly found in most datasets. We will focus, in particular, on consumption expenditure.

Suggested preparation

Chapters 1 and 2 of Barnett and Lewis (1994) provide background for the topics of the lecture, and the Instructor is encouraged to review this material.

Time allocation

Definitions	10 min
Do outliers matter?	25 min
How to detect outliers	
Visual inspection	25 min
Break	15 min
Statistical methods	35 min
How to deal with outliers	5 min
Lessons learned	5 min

Annotated lecture

Today is mainly about outliers

- 1) **Definitions**
What do we mean by an outlier, exactly?
- 2) **Motivation**
Do outliers really matter?
- 3) **Detection**
How to detect outliers?
- 4) **Treatment**
How to deal with outliers?

Important premise

- Suggestions shared in this lecture are not a substitute for the protocols that NSOs have in place to ensure data quality
- They are meant to offer **further safeguards** once **"routine"** edits have been completed
- Useful to **analysts**, as well as data producers

The lecture opens with a roadmap: four questions – what are outliers, how influential are they for our statistics of interest, how do we detect them, and what do we do with them – will guide the discussion for the day. Before going into them, though, the Instructor should

share an important disclaimer: none of the recommendations offered during this lecture is intended as a substitute for the routine edits and checks that NSOs have in place to ensure data quality, and to minimize, among other things, the occurrence of extreme and implausible values. Our recommendations are meant to offer further safeguards once “routine” edits and checks – by and large, those discussed in the second part of lecture 11 – have been completed.

What is an outlier?

An outlier is an observation “that appears to deviate markedly from other members of the sample in which it occurs” (Grubbs, 1969)

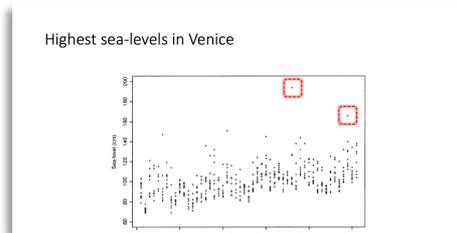


What is an outlier?

- Note: we focus on **univariate** outliers, those found when looking at a distribution of values in a single dimension (e.g. income).
- We use Venice to illustrate



First, the basics: what is an **outlier**? Many definitions exist (a recent study, Aguinis et al. 2013, has documented the use of 14 different definitions, based on a literature review of 28 papers), but the one by Grubbs (1969) is intuitive, and has become a ‘classic’ definition. The note on *univariate* outliers simply means that our focus for the lecture (and the course) is on outliers within the distribution of a *single* variable; a more complex case would be that of *multivariate* outliers, which are defined based on a combination of different variables – but we will not go into that.



As will become apparent in the discussion that follows, **visualizing data** helps with spotting outliers. The distributions shown in the graph (sea levels measured every year from 1930 to 1985) are concentrated around a certain range, and what immediately can be perceived as ‘extreme’ values become visible. Here the Instructor can engage students for a minute on which data points they may want to classify as outliers, exactly, and why. The Venice example should convey the idea that, although the idea of outlier is quite natural and intuitive, there is no hard-and-fast rule to determine what counts as an outlier, and what does not.

What causes outliers?

- Human errors, e.g. data entry errors
- Instrument errors, e.g. measurement errors
- Data processing errors, e.g. data manipulation
- Sampling errors, e.g. extracting data from wrong sources
- Not an error, the value is extreme, just a ‘novelty’ in the data

A dilemma

- Outliers can be genuine values
- The trade-off is between the loss of **accuracy** if we throw away “good” observations, and the **bias** of our estimates if we keep “bad” ones
- The challenge is twofold:
 - to figure out whether an extreme value is good (genuine) or bad (error)
 - to assess its impact on the statistics of interest

Outliers may be the result of various types of error, but they could also be **genuine** values. This is an important message to deliver: outlier is *not* synonym with ‘mistake’. It is not unheard of, for variables like, say, consumption or income, to have extreme values: very rich and very poor people do exist, and they may be the very focus of our analysis. Therefore, data producers and data analysts face a dilemma: assessing whether outliers are genuine values, and, if they are not, understanding their impact on the statistics of interest.

Do outliers matter?

...

Outliers and inequality measures – I
Cowell and Victoria-Feser (1996a)

- This is a beautiful paper
- Explains why outliers (contaminants) are a serious threat to most inequality measures.
- "If the mean has to be estimated from the sample then all scale independent or translation independent and decomposable measures have an unbounded influence function" (p. 89)
- An unbounded IF is a catastrophe.

Why should one care about outliers? The next batch of slides presents some evidence demonstrating that outliers do matter in practice, in terms of their **impact on outcomes** of interest for data users. We summarize the results from at least two, at most three papers (Cowell and Flachaire 2007 may be skipped, as it is quite advanced). First, Cowell and Victoria-Feser (1996): the paper is technical in nature, but its main conclusions are easy to grasp intuitively, and important to understand. The quote from page 82 of the paper is in a language that is inaccessible to most readers, but its meaning will soon be clear.

*The influence function

- F Ideal data, no contaminants
"true" Gini index
- $Gini_{TRUE} = I(F)$
- $G = (1 - \delta)F + \delta H$
 $0 \leq \delta \leq 1$ Real-world data, with $\delta\%$ contaminants
- $Gini_{ESTIMATED} = I(G)$ estimated Gini index
- The influence function, IF: $IF = \lim_{\delta \rightarrow 0} \frac{I(G) - I(F)}{\delta}$

The catastrophe

- Suppose the shape of the income distribution is represented by the continuous frequency distribution in part A
- Suppose that in the sample there are some rogue observations represented by the point mass labelled "contamination".
- Then, according to inequality statistics that are sensitive to the top end of the distribution, the income distribution in A will be indistinguishable from that represented in B (that is, IF is unbounded).

These optional slides elaborate on the concept of **influence function** (IF). The advice to the Instructor is to get into these slides only if she feels comfortable with the contents of the paper. If that is the case, the Instructor can explain that the IF is a measure of the difference between the ‘true’ statistic of interest – in this case, the Gini coefficient – and the same statistic, computed using contaminated data (*i.e.* data containing observations that should not be part of the datasets, but that are present by mistake – they are called ‘contaminants’ in statistics). A **contamination**, that is, one or more ‘wrong’ observations that are far away from most other observations in the distribution, will make the distribution appear much more skewed than it actually is, according to summary statistics and inequality indicators. An unbounded IF means that no matter how few the contaminants (or, in our case, outliers), their impact on the Gini can be infinitely large – that is, the difference between the ‘true’ Gini (what we would observe in the absence of outliers) and the estimated Gini (what we estimate in the presence of outliers) can be very large, infinitely large. This is why an unbounded IF is a problem.

Why an unbounded IF is a catastrophe

- The IF is a measure of the **bias** of an estimator due to the presence of extreme values.
- An unbounded IF means that the bias can be **infinitely large**.
- If the bias of inequality estimators can be infinitely large, **outliers are a priority** for both data producers and data users.

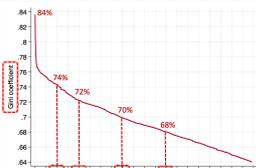


Even if the starred slides have *not* been shown, the implications of Cowell and Victoria-Feser’s findings can still be made intuitively clear, by way of this non-technical slide.

In practice
Hlasny and Verme (2018: 191)

- Many researchers routinely **trim** outliers or problematic observations or apply **top coding** with little consideration of the implications for the measurement of inequality
- One example to illustrate

Sensitivity of the Gini index to extreme values
iterative trimming



A practical **example** can help solidify the idea of the sensitivity of inequality measures to outliers. ‘Trimming’ extreme values in a distribution is common practice: this just means that a researcher may drop a few of the largest observations in the sample – the top 1% households with the highest expenditure, for instance. This is similar to ‘contamination’: the distribution is a certain shape before trimming, and changes after. The graph shows how the estimated Gini coefficient varies when the largest observation in the distribution of per capita consumption is dropped; then the second largest; then the third largest; and so on. The horizontal axis keeps track of how many observations are dropped, as a percentage of the total sample size. As we can see, dropping just a few observations has a huge impact on Gini. There is no limit, in fact, to how large that initial drop may be, regardless of how few observations we are modifying.

***Outliers and inequality measures – II**
Cowell and Flachaire (2007)



- Explains how and why **outliers are a serious threat to most inequality measures**.

***How rapidly the catastrophe occurs**
Rates of increase to infinity of the influence function

- Let us concentrate only on the extremes of the income distribution. Data contamination can occur at very high incomes (say at a point z that approaches infinity) or at very low incomes ($z=0$).

Measure	Generalized entropy, I_α^z	Atkinson, A_α	LogVar	Gini
	$\alpha > 1$	$0 < \alpha < 1$	$\alpha = 0$	$\alpha < 0$
	$\frac{z \rightarrow \infty}{z \rightarrow 0}$			
	$\frac{z \rightarrow \infty}{z \rightarrow 0}$			
	$\frac{z \rightarrow \infty}{z \rightarrow 0}$			

- **Result 1:** GE measures with $\alpha > 1$ are very sensitive to high incomes in the data.
- **Result 2:** GE measures with $\alpha < 0$, and Atkinson measures with $\alpha > 1$ are very sensitive to small incomes in the data.
- We will return on this **catastrophe** in due time, later during this workshop.

These optional slides elaborate on the **impact of outliers on inequality measures**. The Instructor should get familiar with the paper, in case this part is covered in class. The main takeaway is that different statistics of interest (in this case, inequality indicators) show different reactions to outliers.

Outliers and poverty measures
Cowell and Victoria-Feser (1996b)



- Explains why **outliers only rarely** are a serious **threat** to most **poverty measures**.
- Poverty measures are not sensitive to the values (real or contaminated) of the incomes of the rich

Recap

- The answer to the question on whether outliers matter **depends** on the statistic of interest
- Inequality**: both theory (unbounded IF) and practice (incremental truncation) suggest that they matter (tremendously). Not taking this issue into proper account puts inequality comparisons at risk.
- Poverty**: not so much

To conclude this discussion, we mention another paper by Cowell and Victoria-Feser, which investigates outliers and poverty measures, and reaches a rather opposite conclusion with respect to the previous one: **poverty** measures are not sensitive to extreme values.

How to detect outliers?

Visual inspection

- Our procedures are part **graphical**, and part **automatic**. For each commodity, we draw histograms and one-way plots of the logarithms of the unit values, using each to detect the presence of gross outliers for further investigations. [...] [Automatic method] **does not remove the need** for the graphical inspection (Deaton and Tarozzi 2005)

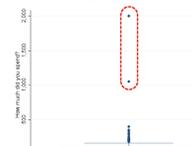
Assuming now that we do care about outliers, as they gravely threaten some of the statistics of interest to the welfare analyst, how should we **detect** them? The question is not trivial, given that there is no single definition of outlier that is valid in every context. Graphical inspection or automatic methods usually work in tandem to detect outliers.

Visual inspection
Malawi IHS3, Cassava tuber expenditure



Visual inspection
Malawi IHS3, Cassava tuber expenditure

- Example 3: use **graphical diagnostic tools**, e.g. the **boxplot graph**



A practical example is used to illustrate **visual inspection** of the distribution of expenditure. We use publicly available microdata from Malawi. Descriptive statistics, graphs of the density of the variable, and box-plot graphs are all tools that help detect any extreme values. In this case, there are two very large, clearly anomalous observations in the expenditure distribution under examination. These methods are effective in pinpointing them – but there are other methods available.



Now is a good time for a break.

Statistical methods

- The literature is rich with methods to identify outliers; in practice, most methods used in empirical work hinge on the underlying *distribution of the data*.
- The idea is simple:
 - **transform** the variable to induce **normality**
 - set **thresholds** to identify extreme values

These are **statistical (automatic) outlier detection techniques**. The idea behind a number of popular outlier detection methods is a simple two-step strategy: (i) transform the consumption aggregate variable to induce *normality*, and (ii) set *thresholds* to identify extreme values. The main benefit of transforming target distributions into approximately normal (Gaussian) ones (step 1), is the ability to work with a well-known, simple distribution, that is computationally straightforward and familiar to virtually all analysts. More specifically, the exact value of the probabilities underlying a normal distribution at different intervals is known. An important consequence of this is that, if we are able to successfully transform the unknown, misshapen, mis-behaved distribution of our variable of interest – whose outliers we want to detect – into a well-known and well-behaved normal, the task of setting specific thresholds beyond which observations are considered ‘extreme’ (step 2) becomes simple, and, perhaps more importantly, easily replicable, whatever the initial distribution. Given that areas underlying the tails of a Normal distribution are known, analysts have a simple decision to make: where in the ‘tails’ of the Normal is an observation become extreme (i.e. unlikely) enough, that we decide to flag it as an outlier? The next slides elaborate.

Transform the variable to induce normality

- The easiest transformation relies on **taking the logarithm** of the variable of interest
- The log “squeezes” large values more, so that skewed distributions become more symmetrical and closer to a Normal distribution.



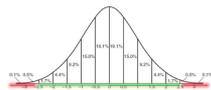
The easiest way to transform the variable of interest into something close to a normal distribution is to take the **logarithm** of the variable, which, as illustrated, changes the shape of a skewed distribution, making it closer to a Normal (in most cases): the logarithm “squeezes” large values more than small ones, so that the logarithm of a skewed distribution becomes more symmetrical.

Set a threshold

- We must specify a **threshold** for deciding whether each observation is ‘too extreme’ (outlier or not?)
- Common ‘thumb-rule’ thresholds : an observation is considered an outlier if it is more than 2.5, 3, 3.5 **standard deviations far from the mean** of the distribution
- In formulas: x is an outlier if $x > \bar{x} + z_{\alpha} s$
where z_{α} equals, say, 2.5.
- We can express the same criterion as $\frac{x - \bar{x}}{s} > z_{\alpha}$
where the left-hand side is called a **z-score** (a variable with mean = 0 and var = 1)

Why 2.5, 3, 3.5, or any other number?

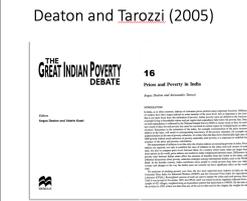
- Under the assumption of normality:



$z_{\alpha} = 2.5$ implies that outliers are in the region where $\alpha = 0.5$ percent of observations normally are.

Next, it is convenient to introduce **z-scores**. By taking a variable and subtracting its mean, we center (translate/shift) the distribution of the variable at zero, and by dividing by the standard deviation, we rescale the variable so that its variance equals one. This amounts to using a standard Normal instead of a Normal, of course. The final step is to set a threshold, after which observations are to be considered outliers. The general rule is: observation x is an outlier if $x > \bar{x} + z_\alpha s$, where \bar{x} is the sample mean of the variable, s is its standard deviation, and z_α is a conventional value, say 2.5. The conventional values of z_α are well known in the context of the Normal distribution: because the Normal is well known, we know that each of these values corresponds to a certain probability left in the tails.

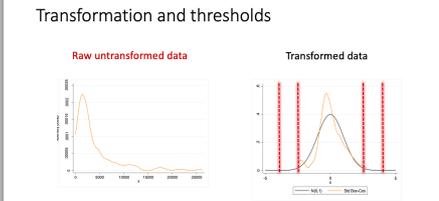
Deaton and Tarozzi (2005)



In the case of **India**, D&T (2005) flagged as outliers prices whose logarithms exceeded the mean of logarithms by more than 2.5 standard deviations:

$$\frac{\ln(x) - E[\ln(x)]}{sd[\ln(x)]} > 2.5$$

Transformation and thresholds



A notable **application** of this simple procedure – normalizing the variable of interest by taking the log, calculating a z-score, and applying a conventional threshold to define an ‘outlier detection region – is in Deaton and Tarozzi (2005). The graph with the orange curves provides a graphical representation of the method. This is simply an example, that should help students consolidate what they have just seen in formulas.

Two questions

- 1) How good is such an approach?
- 2) What to do after flagging outliers?

How good is such an approach?

- Log-transformation is very basic – how to deal with negative values?
- Not recommended when the log-distribution can not be assumed to be a Normal distribution
- Why should we set the threshold using the **mean** and **standard deviation**, which are sensitive to extreme values, if this is exactly what we are worried about?

$$\frac{\ln(x) - E[\ln(x)]}{sd[\ln(x)]} > 2.5$$

- We can do better

The approach can, however, be critiqued: it cannot deal with negative values (because the logarithm only exists for positive numbers), it does not apply to variables that are not normal after the transformation, and – perhaps most importantly – the thresholds that should help to identify outliers are *themselves* sensitive to outliers. As is well known, the sample mean is very much affected by extreme values, and so is the standard deviation. All in all, **we can do better** than the simple ‘take the log and run’ method illustrated so far. How so?

*The Box-Cox transformation



- The Box-Cox transformation:
$$y_h^{(\lambda)} = \begin{cases} (y_h^\lambda - 1)/\lambda & \text{if } \lambda \neq 0 \\ \ln y_h & \text{if } \lambda = 0 \end{cases}$$
- The transformed variable is often remarkably close to a Normal dn.
- Outliers are identified if:
$$y_h > 75\text{th percentile} + 5 \times \text{IQR}$$

*The Box-Cox method: An assessment

- Only applies to **strictly positive variables** (e.g., it does not necessarily work with income)
- Calculation is cumbersome, and often problematic

These optional slides hint at one of these alternatives: it relies on a different ‘normalization’ of the target variable, more complex and more flexible than the logarithm, called the **Box-Cox transformation**. The definition of the threshold is also different, as it relies on the **interquartile range (IQR)** instead of the standard deviation, and on a given percentile, for instance, the 75th, instead of the mean. These statistics are more robust to the presence of extreme values. The main downside of the method is that it is, ultimately, hard to implement, as the Box-Cox transformation cannot always be calculated. Our suggestion is to skip these two slides, unless there are specific questions on methods better than the log to induce normality in the distribution of a variable.

A popular strategy
robustification

- While there is no agreement on the best method, a common solution is to use **robust measures of scale and location** to set the threshold for flagging outliers
- the idea is to replace the sample average \bar{x} with a robust estimator (e.g. the median), and the standard deviation with a robust estimator. A popular option is the median absolute deviation (MAD).

Robustification of the z-score
The median absolute deviation (MAD)

$$z_h = \frac{x_h - \bar{x}}{s}$$

$$z_h = \frac{x_h - \text{med}[x_h]}{\text{MAD}}$$

$$\text{MAD} = b \times \text{med}[|x - \text{med}[x]|]$$

b = 1.4826

if the distribution is **Gaussian**

One solution is to ‘**robustify**’ the z-score. Instead of using a z-score based on the sample mean and the standard deviation, we replace them by plugging in robust estimates of scale and location. Instead of the mean, we use the **median**; instead of the standard deviation, we use the **median absolute deviation (MAD)**. Note that the MAD is defined as the median of all distances between each value of x and the median, times the parameter b (this is just a constant that is required for MAD to be a consistent estimator of the standard deviation under the hypothesis that the distribution is Normal).

We can do even better
Rousseeuw and Croux (1993), IASA

Alternatives to the Median Absolute Deviation
Peter J. Rousseeuw and Christophe Croux*

In robust estimation one frequently needs an initial or auxiliary estimate of scale. For this one usually takes the median absolute deviation (MAD), $\frac{1}{n} \sum_{i=1}^n \text{med}\{|x_i - \text{med}\{x\}|\}$, because it has simple explicit formula, needs little computer time, and is easy either as witnessed by its bootstrap influence function and its 50% breakdown point. But there is still room for improvements in this area: the fast-time MAD, as used in extreme distributions and the fast QN-Gaussian estimator. In this article we use an extreme outlier and 50% breakdown scale estimator that are more efficient. We consider the estimator $S = \frac{1}{n} \sum_{i=1}^n \text{med}\{|x_i - x_j|\}$ and the estimator S_j given by the 25% quantile of the distances $|x_i - x_j|$, $j = 1, \dots, i-1$. Both the fast-time MAD and the new estimator are more efficient than the MAD. We study S and S_j in comparison with the 50% breakdown point estimator. The relative efficiency of S is 48%, whereas S_j gives 62%. We study S and S_j by means of their influence function, that has interest for regression as well as regression, and their breakdown performance. Their influence is also compared at non-Gaussian models, including the negative exponential model where S_j has a lower gross error sensitivity than the MAD.

KEY WORDS: Bias-cov; Breakdown point; Influence function; Robustness; Scale estimation.

Rousseeuw and Croux (1993)

- Rousseeuw and Croux (1993) propose to substitute the MAD with a different estimator:

$$S = \frac{1}{n} \sum_{i=1}^n \text{med}\{|x_i - x_j|\}$$

- For each j we compute the median of $|x_i - x_j|$ ($i = 1, \dots, n$). This yields n numbers, the median of which gives our final estimate S .

$$z_h = \frac{x_h - \text{med}[x_h]}{S}$$

1.1926 at the Gaussian model.

But even this approach can be improved. Rousseeuw and Croux (1993) discuss the pros and cons of the MAD at some length, and end up suggesting that more efficient alternatives exist, especially in the presence of skewed distributions, which is exactly the case of interest here (this is a technical paper, and it is not essential for Instructors or students to master it: the fundamental intuition is enough for the purposes of the lecture). Following this research, we suggest the following, alternative robustification of the z-score: at the numerator of the score, continue using the median, in place of the mean, but at the denominator, switch to the ‘**S-estimator**’ instead of MAD. To calculate the S estimator, we simply need to compute, for each household i in the sample, the expression $\{\text{med}_j |x_i - x_j|\}$ for $j=1, \dots, n$. This gives n numbers, the median of which gives the estimate S (the number 1.1926 in the formula is required for making S a consistent estimator of the standard deviation under the assumption of normality). These are details that are *not important* for students to be confident with – it is enough to explain why those parameter are there. Statistical softwares take care of these details.

Recap

- “take the log and run” is not a recommended practice
- taking the log and **robustifying** the z-score is a better practice
- Belotti and Vecchi (2019) provide [outdetect.ado](#)

Take the log and run vs. robust z-scores

Countries	Year	Outliers (%)					
		log-transformation			best normalization		
		overall	left	right	overall	left	right
		(1)	(2)	(3)	(4)	(5)	(6)
Malawi	2017	0.75	0.14	0.61	0.30	0.22	0.08
Nigeria	2012	1.35	0.11	1.24	0.72	0.32	0.40
India	2012	1.39	0.03	1.36	0.62	0.13	0.49
Pakistan	2014	1.58	0.02	1.56	0.39	0.21	0.18
Guatemala	2014	1.14	0.06	1.08	0.61	0.15	0.46
Peru	2015	0.36	0.09	0.27	0.28	0.16	0.12
Armenia	2013	0.91	0.08	0.83	0.68	0.17	0.51
Georgia	2015	0.75	0.25	0.50	0.73	0.32	0.41

To conclude this topic, we can say that the approach of simply taking the logarithm of the target variable and then using z-scores to detect outliers is not the best available method. Using other statistics to robustify the z-score gives better results in practice. Data from selected countries are used to show the output of the **outdetect.ado** Stata command (Belotti et al. 2020), which performs the Rousseeuw and Croux (1993) robustification.

How to deal with outliers?

(in one slide)

Treatment of outliers

Three main methods of dealing with outliers, apart from removing them from the dataset:

- 1) **reducing the weights** of outliers (trimming weight)
- 2) **changing the values** of outliers (Winsorisation, trimming, imputation)
- 3) **using robust estimation techniques** (M-estimation).

• Documentation, transparency & reproducibility

The final topic for the lecture is **outlier treatment**, which is discussed only briefly. It bears repeating that checking for “gross mistakes” and other context-specific sources of measurement error is a necessary first step when dealing with extreme values – indeed with any data points that data producers deem “implausible”. Given the impact of outliers on inequality estimates, however, taking *some* action, once these routine checks are completed and if outliers are still present, is recommended. To explain, the Instructor can recall a prescription from Huber’s classical manual on Robust Statistics (Huber 1981: 4): “Any reasonable, formal or informal, procedure for rejecting outliers will prevent the worst.”. Provided choices are motivated and documented by the analyst, such procedures may fall into the broad categories recalled here.



Lessons learned

- Outliers can be **genuine** observations... be gentle to the data and document each and every step of the data processing
- As far as inequality is concerned, outliers are the worst enemy (**unbounded IF**)
- Outlier detection:
 - go beyond the “take the log and run” strategy. It works well only if you can describe the data with a Gaussian distribution. Typically, however, distributions are skewed.
 - Use a “take the log, **robustify** the z-score and run”, strategy.
- **Outlier treatment**: it depends. Quantile regression is a good candidate.

As usual, it is good practice to take some time to end the lecture, and summarize the main takeaways. The first point should emerge strongly: documentation of each of the steps discussed during the lecture is key for any type of data cleaning or analysis.

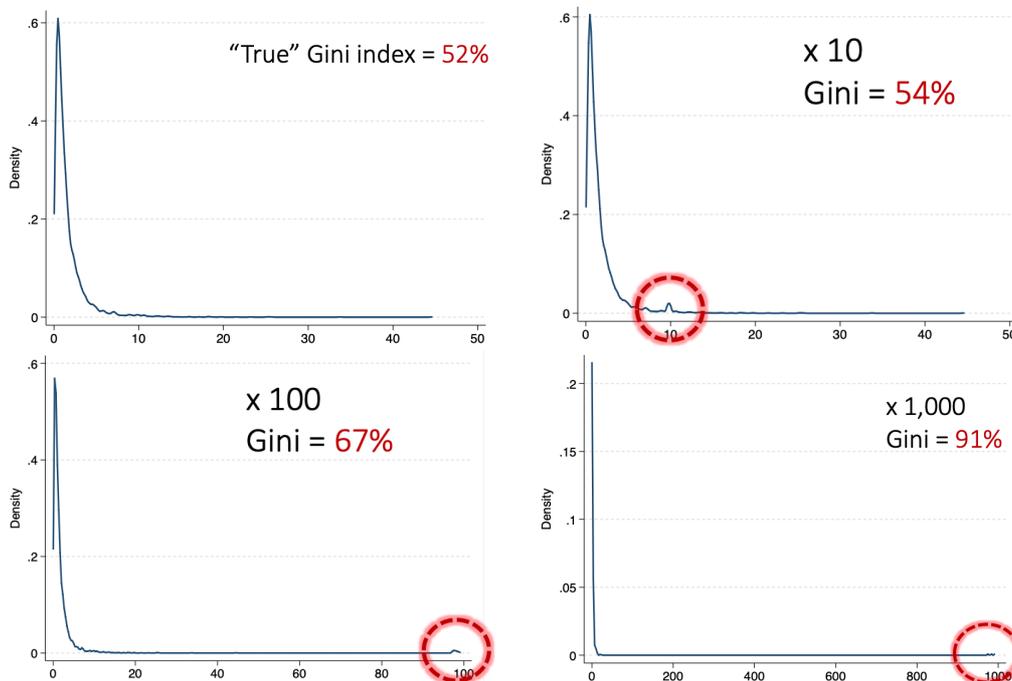
Homework

Exercise 1 - Engaging with the literature

See exercise 1, Lecture 1.

Exercise 2 – Do-it-yourself...

This exercise is meant to give students an opportunity to do a hands-on experiment on the impact of outliers on a statistic of interest (in this case, the Gini index), and to get a sense of how these extreme values may arise in practice (wrong placement of a decimal separator is a common data entry mistake). The code pictured in the slide is written for Stata, but the same operations can be performed with any statistical software. Results should look similar (not necessarily identical) to those pictured below:



Exercise 3 – Inequality measures

The table referenced by the exercise is an example of how strong the impact of the method chosen to treat outliers can be on final estimates. It should be stressed that none of the columns in the table report the 'correct' estimates, necessarily: they are simply the results of different ways to deal with extreme values. The importance of documenting outlier detection methods should also be emphasized.

Lecture 13

Measuring inequality

Learning objectives

The goal of this lecture is to provide an overview of both concepts and methods underlying inequality measurement. The task of estimating inequality is greatly facilitated by the availability of suitable data – thus the lectures provides the theoretical framework required for data producers and users to share a common view of what is meant by ‘inequality’ and how to measure it.

Suggested preparation

The lecture follows chapters 1 and 2 of Cowell (2011). The Instructor is assumed to be familiar with this material.

Time allocation

Tables and graphs	45 min
Break	15 min
Indicators	55 min
Lessons learned	5 min

Annotated lecture

Outline for final lectures

- Once datasets have been finalized, it is time to produce results, with the aim of representing the patterns emerging from the data.
- In practice?
 - Inequality this lecture
 - Poverty next lecture
- Basic summary statistics on household demographics, education, access to services, etc.
- Average expenditures and incomes final lecture

Inequality and poverty measurement

- a measure of living standards
- high-quality data on households' living standards
- a distribution of living standards (inequality)
- a critical level (a poverty line) below which individuals are classified as "poor"
- one or more poverty measures

The lecture opens with two introductory slides, where we provide a short recap of where we stand, and discuss the next steps. The main message is that today's class is the first of a twin set of *theoretical* lectures on **inequality** (today) and **poverty measurement** (tomorrow).

Cowell (2011)



99.9% of this lecture is explained with better words in Cowell's work: this book and other (countless) journal articles

Warning

- During the course we stressed the distinction between the concepts of *living standard*, *income*, *expenditure*, *consumption*, etc.
- In this lecture we make an **exception**, and use these terms **interchangeably**
- Similarly, I will **not** make a distinction between income per *household*, per *capita*, or per *adult equivalent*
- For once, and for today only, we will be (occasionally) inconsistent

This lecture is modelled after the reference cited: students interested in inequality measurement should certainly aim to get familiar with Cowell's book. We also share a disclaimer about **terminology**: in contrast to the rest of the course, this lecture will move beyond the construction of a welfare indicator, into analysis of its distribution. For this reason, we will not be as specific or precise about what the welfare indicator actually *is*: our attention here will be on inequality *measures*, not on data, per se. Data, however, still continue to be important, of course: learning about analysis will help students be aware of how data are ultimately put to use.

*Functional vs. personal income distribution

- Economists distinguish between:
 - Functional** distribution of income distribution among *factors of production* land (*rent*), labor (*wages*), and capital (*profits*)
 - Personal (or size)** distribution of income distribution among *persons*, irrespective of their economic function
- We focus on the latter.

*The functional distribution of income
Average factor shares in Indian Economy, 1960-61 to 1991-1992



Sector	Labour	Land	Capital
Primary sector	56.42	30.30	13.28
Secondary sector	67.68	3.47	28.85
Tertiary sector	61.57	3.74	34.69
All sectors	60.69	15.21	24.10
Public sector	86.15	0.83	13.02
Private sector	56.53	17.76	25.71

These slides are optional, and the Instructor is encouraged to skip them, unless the audience demonstrates a specific interest on the difference between **functional and personal inequality**. Our focus in this class is on the latter: functional inequality is typically more of interest to macroeconomists than to welfare economists.

Focus on the term 'inequality'

- "When we say **income inequality**, we mean simply **differences in income**, without regard to their desirability as a system of reward or undesirability as a scheme running counter to some ideal of equality" (Kuznets 1953: xxvii)
- In practice, how can we appraise the inequality of a given income distribution? Three main options:
 - Tables
 - Graphs
 - Summary statistics

Next, we spend a few words on **what we mean by 'inequality'**. The point here is to clarify that in this lecture we aim at defining and measuring inequality from a purely statistical standpoint, that is, we do *not* discuss a number of interesting issues/questions related to fairness and equity – just 'differences in income' as in the Kuznets' quote. The bottom part of the slide is the outline for the rest of the lecture. The plan is to answer the question: what is the **best way of describing (measuring) inequality**? Using tables? Graphs? Indicators? Each option will be examined and evaluated, so that ultimately, we will be able to motivate the use of well-known inequality measures. The spoiler – the Instructor may share it in advance – is that tables are *not* a good idea, the Lorenz curve 'wins' among graphical tools, and Gini and GEI 'win' among inequality measures.

Tables: an assessment

- In general, tables are **not recommended** when the focus is inequality
- Difficult to get a clue of the extent of inequality in the distribution by looking at a table.
- To illustrate, let us look at the the distribution of income and taxable income in South Africa.



Can you tell how high or low inequality is?

Income group	Number of taxpayers	Income before deductions (R million)	Deductions allowed (R million)	Taxable income (R million)
<= 0	119 998	-15 096	11	-15 107
1 – 70 000	401 447	14 835	565	14 270
70 001 – 350 000	2 689 263	543 389	58 380	485 009
350 001 – 600 000	764 197	317 905	44 483	273 423
600 000 +	926 000	906 546	109 690	796 856
Total	4 898 565	1 759 939	213 128	1 546 811

Source: 2018 Tax statistics, National Treasury and the South African Revenue Service

Let us quickly evaluate our first candidate: **tables**. The Instructor may show the example from South Africa to the audience, and ask the question: is this table effective in conveying the extent of inequality? We *briefly* engage with students (a couple of minutes are enough) and listen to their opinions. After weighing pros and cons, we argue that tables are *not* the best way to summarize inequality.

Another candidate: graphs

Do graphs (diagrams) **help** represent and understand inequality?



If tables are not good enough to summarize inequality, what about **graphs**? Again, we briefly engage with students. Which graphs would they use? Why? We listen, we do not reply and/or comment on each question. Then we take students through the next slides, where we assess a selection of popular graphs. We start with the **histogram**.

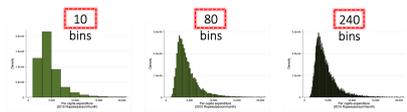
Histograms

- Let the interval $[x^-, x^+]$ denote the **range** of the data.
- Partition $[x^-, x^+]$ into m non-overlapping **bins** (intervals of equal width $h = (x^+ - x^-)/m$).
- A histogram estimate of the **density** $f(x)$ is the fraction of observations falling in the bin containing x , divided by the bin width h :

$$\hat{f}(x) = \frac{\text{(fraction of sample obs. in same bin as } x)}{h}$$
- The area of each bar ($= h \times \hat{f}(x)$) is interpreted as the fraction of sample observations within the bin. All bar areas sum up to unity.

Histograms?

Note: extreme values (top 1%) have been removed



Histograms: an assessment

- The position and number of bins is **arbitrary**
- Inherently **lumpy**: discontinuities at the edge of each bin
- Can provide very **different pictures** of the **same distribution**
- Look of graphs highly dependent on treatment of extreme values
- Read Cowell, Jenkins and Litchfield (1996) for more.

The aim of first slide of this bunch is to check with the students whether they understand histograms as we think they should. Would they be able to explain how to construct a histogram from scratch? Faced with a histogram, would they be able to answer technical question (e.g. how to interpret the height of a bar, how many bins to choose, etc.)? The message: take this opportunity to refresh your knowledge of these basic tools. Only by understanding *the details* can they answer the question of which is the ‘best’ histogram out

of the three shown in the slides. The point that the Instructor should make here is that the number of bins, which greatly influences the appearance of the graph, and therefore our perception of the distribution and of inequality, is ultimately arbitrary. In addition to that, the note says that “the top 1% observations have been removed”: this is common when drawing graphs of highly skewed distributions, given that plotting the top values often makes the graph impossible to read. But in our context, this robs us of crucial information: top values are precisely the interesting ones – how can we gauge inequality without them? After discussing pros and cons of histograms we conclude that, despite their popularity, they are certainly not the optimal way of summarizing inequality.

Beyond histograms

- A **probability density function (PDF)** is the ‘continuous version’ of a histogram
- A convenient way to introduce the PDF is by starting from the **cumulative distribution function (CDF)**

The next candidates are the **probability density function (PDF)**, and the **cumulative distribution function (CDF)**. If these statistical tools are already well-known to students, as they likely are, there is no need to cover them in great detail; if not, the following slides can be used to refresh their understanding of some of the most fundamental tools in statistics. Familiarity with these concepts is key when working with inequality measures,

The cumulative distribution function (CDF)

- The **cumulative distribution function (CDF)** of a random variable X is defined as follows:

$$F(x^*) = \int_0^{x^*} f(x) dx$$

- $F(x^*)$ is the **proportion of individuals having X less than or equal to x^*** .
- If X is income and, say, $x^* = 2,000$ Rps., then $F(x^*) = \Pr(X < 2,000)$, that is the fraction of people with less than 2,000 Rps.

The empirical cumulative distribution function (ECDF)

Iraq IHSES 2007/2012., Cumulative distribution of welfare aggregate (p.21)

- Pick any income level on the x -axis, and the curve $F(x)$ will tell you the percentage of individuals in the population having a level of income **lower than x** .
- The CDF is more effective in telling you about the incidence of poverty than about inequality.

IRAQ: The Unfulfilled Promise of Oil and Growth
Poverty, Inclusion and Welfare in Iraq, 2007-2012

We follow the same logic as we did for histograms, and define the **CDF** first; then, we show an example from a recent report, and wonder whether the graph provides useful information about inequality. The CDF returns the share of the population with incomes lower than any given value: a very handy piece of information indeed, when you are concerned about poverty. Plug the poverty line into the CDF, and you get the poverty headcount rate back. This is not really what we need when we care about inequality. After arguing this, we move on to the PDF.

The probability density function (PDF)

- The **probability distribution function (PDF)** is the derivative of the CDF:

$$f(x) = \frac{dF(x)}{dx}$$

- Simply, the pdf describes the likelihood that a random variable X takes on a given value.
- Note: for continuous random variables, the probability that X takes on any particular value x is 0, the probability being defined only over (a,b) intervals.

Iraq, 2007/12

PDF of welfare aggregate (p.21)

- Does the PDF work any better than the CDF when describing inequality?
- Probably, yes. What do you think?

IRAQ: The Unfulfilled Promise of Oil and Growth
Poverty, Inclusion and Welfare in Iraq, 2007-2012

We introduce the **PDF** (two hidden slides give the Instructor the options to delve into the details, but should not be needed with students who have a background in statistics).

***Interpretation**

Two extremes:

- perfect equality:** everyone is concentrated at one particular income value
- uniform density:** income is spread uniformly from the poorest to the richest individual – significant inequality
- in-between, typical case.**

PDFs: an assessment

- Similar drawbacks as histograms
- The bandwidth (how “smooth” the graph is) is **arbitrary**
- In most cases, PDFs require some **trimming** of top values to avoid looking “squished” and being unreadable
- In general, the PDF is not very effective in showing what is going on in the **upper tail**, but that is important information when focusing on inequality

The starred slide titled “Interpretation” may help to explain the way one might approach the PDF as a graphical tool to evaluate inequality. The two extreme scenarios (a and b) are clear enough; the problem is that in any practical situation, we find ourselves in c, and visually discriminating between more or less unequal PDFs is not easy. We conclude that the PDF ultimately suffers from the same limitations as the histogram.

The quantile function

- Let $p = F(x)$ be the proportion of people in the population with income **lower than** x .
- The **quantile function** $Q(p)$ is defined as:
 - $F(Q(p)) = p$ or $Q(p) = F^{-1}(p)$
- $Q(p)$ is the **income level** below which we find a proportion p of the population.

***The Parade of Dwarfs**
Pen (1971)

- Assume that everyone in the population has **height** proportional to **income**.
- Line people up in order of height, and let them **march**.
- After some time, the shape of such a parade will be represented by the curve called **Parade of Dwarfs (and a Few Giants)**.

Next, we introduce the **quantile function**, the inverse function of the CDF. We concede that perhaps it is not totally ineffective at summarizing inequality, and if time permits, we explain why by means of the ‘parade of dwarfs’ (see the excellent explanation in Cowell 2011: the famous metaphor describes the quantile function as a parade of people, that are as tall as their incomes – the parade is made up by many dwarfs and a few giants). We conclude that ultimately, we hope to find better graphical tools. It is finally time to introduce *the* graph that – so we argue – succeeds in representing inequality: the Lorenz curve.

The Lorenz Curve (1905)
Picture and intuition

- Horizontal axis:** cumulative % of **population** (individuals ordered poorest to the richest)
- Vertical axis:** cumulative % of **income** received by each cumulative % of population.
- 45-degree line:** Lorenz curve if perfect equality.
- The overall distance between the 45-degree line and the Lorenz curve is indicative of the amount of **inequality** present in the population.

The Lorenz Curve (1905)
Mathematically

- The **Lorenz curve** $L(p)$ is defined as follows:
$$L(p) = \frac{\int_0^p Q(q) dq}{\int_0^1 Q(q) dq}$$
- The **numerator** sums the incomes of the poorest $p\%$ of the population;
- The **denominator** sums the incomes of all.
- The **ratio** $L(p)$ indicates the cumulative % of total income held by a cumulative proportion p of the population.
- Example:** if $L(0.5) = 0.3$, then we know that the **50% poorest individuals hold 30%** of the total income in the population.

We discuss the **Lorenz curve** at some length, and conclude that this is a useful tool, one that helps welfare analysts tremendously. This must become part of every student’s analytical toolbox.

Quantile function and Lorenz curve: an assessment

- These graphical tools emphasize the **ranking** of shares of the population on the basis of income
- The Lorenz curve clearly shows how far the distribution is from perfect equality
- Still, no graph is as straightforward and easily comparable as a **scalar measure** of inequality

Recap and next steps

- Not all **graphs** are OK to represent inequality
- **Lorenz curve** is the most popular
- A better conceptual understanding comes from constructing inequality **measures** from first principles.
- The most straightforward approach: inequality measures as **pure statistical measures of dispersion**.

With the Lorenz Curve, we conclude the review of graphical tools for analysing inequality, and this is when we are likely to need a break.



Now is a good time for a break.

The second part of the lecture is devoted to **inequality indicators**: scalar measures of inequality, that summarize our phenomenon of interest with a single number. Again, we follow the same logic: we evaluate and dismiss candidates, until we arrive at our ‘winner’.

Measures of dispersion

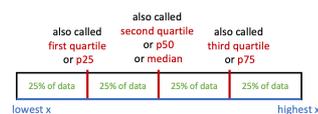
- **Range** $R = x_{\max} - x_{\min}$
What do you think?
- **Variance** $\sigma^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \mu)^2$
- **Coefficient of Variation** $CV = \frac{\sigma}{\mu}$
- ...

There is an obvious solution to the problem of measuring inequality using a single number: what about familiar **measures of dispersion**, like the variance? No need to spend much time discussing this: each measure has some advantages, and some (many more) disadvantages – the point is that we have many candidates, and we don’t really know how to choose. We would like *one* measure that we can trust (or even a few of them), not a multitude of different ones that we are not very confident in. Ultimately, we can do much better.

Quantiles, Quintiles, Quartiles, ...

- Take a group of observations for variable x (e.g. x = expenditure)
- Order observations from smallest x to largest x (poor to rich)
- Partition the set of observations into n subsets of equal size
- Q_1, Q_2, \dots, Q_{n-1} are values of x that identify the cut points between the n subsamples, and are called **quantiles**.
- Some quantiles have special names

Example with $n = 4$

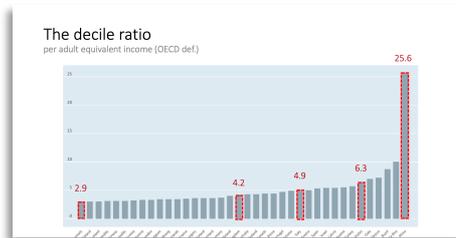


Note that p indicates **percentiles**, that is, quantiles when $n = 100$

These slides review the definition of **quantiles**: again, this is likely to be well known to students, so it may be covered quickly. Quantiles are fundamental building blocks for some widely used inequality indicators.

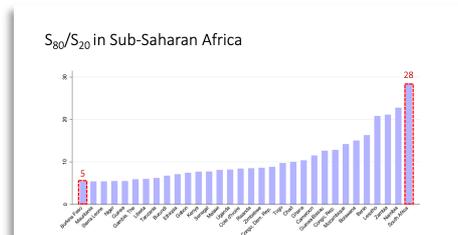
Quantile ratios

- A **quantile ratio** measures the **gap** between the rich and the poor.
- It is defined as the ratio of two quantiles, $Q(p2)/Q(p1)$ using percentiles $p1$ and $p2$.
- Three popular indices are:
 - The **quintile ratio** ($p2 = 80$ and $p1 = 20$):
 $QR = Q(p80)/Q(p20)$
 - the **decile ratio** ($p2 = 90$ and $p1 = 10$):
 $DR = Q(p90)/Q(p10)$



Quantile share ratios

- Let S_{20} denote the share of (equivalised disposable) income received by the **bottom 20%** of the population, and S_{80} the income share received by the **top 20%** of the population.
- The **quintile share ratio** is defined as follows:
 $S_{80-20} = S_{80}/S_{20}$
- The quintile share ratio is the level-1 Laeken indicator, chosen by the EU to monitor income distribution.



These are **quantile ratios** and **quantile share ratios**, which are defined and ‘shown in action’ with a few examples. The graphs help students form an idea of what is the expected range of each of these indicators. Both measures are commonly encountered in practice, but we move on to an even more useful one.

The Gini Coefficient
A definition

- Yitzhaki (1997) counts more than a dozen formulas available for the Gini index.
- A classic definition of the Gini coefficient:

$$G = \frac{1}{2n\mu} \sum_{i=1}^n \sum_{j=1}^n |x_i - x_j|$$
- The Gini coefficient ranges from 0 (all recipients have the same income: **maximum equality**), to 100 (all income is received by one recipient: **maximum inequality**).

The Gini Coefficient
Interpretation – Farris 2010: 857

- Consider the following experiment.
- Pick **two households at random** and record the lower of their two incomes: call the result y . Let μ denote the mean income. It turns out that:

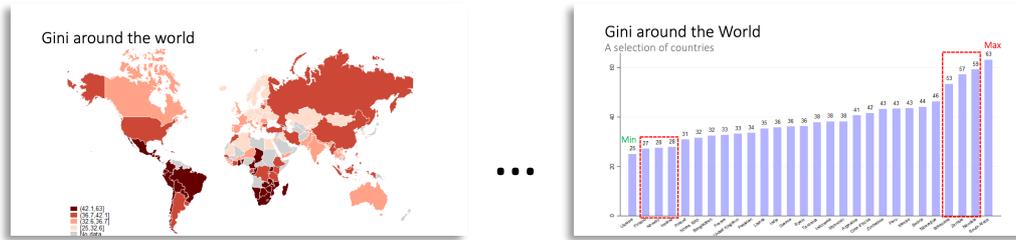
$$\frac{y}{\mu} = 1 - G$$
- Assuming that the Gini index is **40%**, we can conclude that the lower of two family incomes, chosen at random, is about **60%** (=100-40%) of the national mean: on average, the poorest of any two families earns only about 60% of the national mean.

The Gini Coefficient
Graphical interpretation

- The Gini index is **two times the area A** between the Lorenz curve and the equality diagonal:

$$Gini = \frac{A}{(A+B)} = 2A = 2\left(\frac{1}{2} - B\right) = 1 - 2B$$

We are talking, of course, about the **Gini coefficient**, arguably the single most used inequality indicator. We give a mathematical definition, and then a few interpretations, among the many available (the one by Farris, about average income differences, and a graphical one based on the Lorenz curve).



A few more graphs help once again with grasping the usual order of magnitude of the Gini index across countries in the world. We invite students to get as familiar with the Gini coefficient as much as possible, due to its wide practical application. We emphasize that, after the Lorenz curve, this is the second analytical tool that everybody should have in her toolbox.

*Atkinson's paper

The paper

Tony Atkinson (1944-2017)

This optional slide delves deeper on the issue of **inequality comparisons**. The approach is to simply direct students to one of the foundational papers on the topic.

Recap

- Quantile ratios, quantile share ratios, Gini, are all popular inequality measures
- They do a fine job at representing inequality with a number
- **Problem** they do not always have all the **properties** that we would want for an inequality measure
- **Solution** solve the problem backwards. First lay out some desirable properties, then construct a measure that complies with them

A short recap of the main takeaways leads us to the next and final topic of the lecture. We have told students that some measures of inequality are better or more popular than others: but why is that, exactly? One way to think about the choice of the ‘best’ indicator more rigorously is in terms of desirable properties of an inequality measure. As we might do when we buy a car or a phone, it seems logical to first identify some characteristics we care about (good reception, nice design, a certain amount of memory) and then see which phones possess these characteristics, and select them based on that. The same can be done with inequality measures: this logic is called the **axiomatic approach**.

Deriving inequality measures from axioms

- **Axiom**: a statement accepted as true as the basis for argument or inference.
- The **axiomatic approach** allows us to “custom-build” inequality measures that fit our needs:
 1. We define a set of elementary properties (axioms) that we think inequality measures ought to have
 2. We obtain a mathematical formula that delivers a class of inequality measures satisfying the axioms

Five axioms of inequality measures

- A. **Anonymity (or Symmetry)**
Who is earning the income does not matter
- B. **The Population Principle**
Population size does not matter
- C. **Scale Invariance (or Relative Income Principle)**
Income levels do not matter
- D. **The (Pigou-Dalton) Principle of Transfers**
Rank-preserving rich-to-poor transfers reduce inequality
- E. **Decomposability (or Subgroup Consistency)**
The measure is additively decomposable

***Five axioms of inequality measures**

(A) **Anonymity (or Symmetry)**

- If income distribution X is any permutation of income distribution Y, then $I(X) = I(Y)$.
- In short, it does not matter who is earning the income.

(P) **The Population Principle**

- When one income distribution is an n-fold replication of another, the two are distributionally equivalent.
- The population size does not matter: all that matters are the proportions of the population who earn different levels of income.

...

***Five axioms of inequality measures**

(D) **Decomposability (or Subgroup Consistency)**

- An **additively decomposable** inequality measure is one which can be expressed as a **weighted sum** of the inequality values calculated for population groups **plus** the contribution of differences between group means.

$$I = \sum_{k=1}^K \omega_k I_k + I(\bar{y}_1, \dots, \bar{y}_K), \quad \sum_{k=1}^K \omega_k = 1$$

where I_k is the inequality index calculated within the k-th group, and ω_k are the population shares.

This approach has revolutionized the way inequality is measured, by laying out five axioms that a good inequality measure must comply with. The axioms are simply stated in one slide; a series of five optional slides goes into more detail as to what each axiom exactly implies. The instructor is advised to cover them only if time permits.

Generalized Entropy Indices (GEI)
Shorrocks (1980)

- Inequality measures that satisfy all axioms (A to E), **must** have the following form:

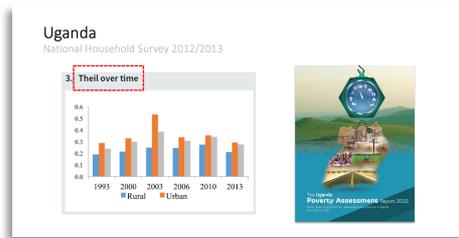
$$GE(\theta) = \frac{1}{\theta^2 - \theta} \left[\frac{1}{n} \sum_{i=1}^n \left(\frac{z_i}{\bar{z}} \right)^\theta - 1 \right]$$

where θ is a parameter that may be given any value (positive, zero or negative).

- Depending on the value assigned to θ , different indices are obtained.

The family of Generalized Entropy Indices

- If $\theta = 0$
Mean Logarithmic Deviation: $GE(0) = MLD = \frac{1}{n} \sum_{i=1}^n \log \left(\frac{z_i}{\bar{z}} \right)$
- If $\theta = 1$
Theil Index: $GE(1) = THEIL = \frac{1}{n} \sum_{i=1}^n \frac{z_i}{\bar{z}} \log \left(\frac{z_i}{\bar{z}} \right)$
- If $\theta = 2$
Half Coeff. of Variation Squared: $GE(2) = \frac{CV^2}{2}$
- The case of Uganda illustrates



The axiomatic approach leads to the **Generalized Entropy class of Indices (GEI)**, which comply with all axioms. Gini and the GEI indices are the most commonly used by scholars the world over (note that Gini does not comply with all axioms, but it is so popular that welfare analysts are, in some way, ‘stuck’ with it: it is still a useful indicator, regardless). The instructor is advised to read Cowell’s highly accessible account of the axiomatic approach, which underlies the slides. The take-away for students is: GEIs are important, and in particular the Mean Log Deviation (MLD).

Inequality decomposition

- Inequality decompositions are typically used to estimate the extent to which the **heterogeneity** of the population affects overall inequality. Two popular techniques are:
 - Decomposition by **population sub-group**
 - Decomposition by **income source**
- We focus on the former:
 - Societies can often be partitioned into groups (e.g. North-South). We would like to be able to decompose total inequality into two components, namely the inequality **within** the constituent groups, and inequality **between** the groups:

$$I_{TOTAL} = I_{WITHIN} + I_{BETWEEN}$$

***Inequality decomposition**

- The most popular additively decomposable inequality index is the **Mean Logarithmic Deviation**.
- Partition the population into $k = 1, \dots, K$ groups. Then:

$$MLD = \underbrace{\sum_{k=1}^K v_k MLD_k}_{\text{WITHIN}} + \underbrace{\sum_{k=1}^K v_k \log\left(\frac{\bar{x}}{\bar{x}_k}\right)}_{\text{BETWEEN}}$$

where v_k are population shares.

Rwanda 2016/17
Fifth Integrated Household Living Conditions Survey, EICV5 (2016/17)

	2014	2017	Change
Total	0.442	0.374	-0.068
Urban	0.533	0.417	-0.116
Rural	0.225	0.198	-0.027
Decomposition:			
within inequality	0.335	0.277	-0.058
between inequality	0.107	0.097	-0.010
Minus: 'between' inequality as % of total inequality	24%	26%	

***Panama 2008**
Living Standards Measurement Surveys (LSMS):
Decomposition of inequality by regions

	1997			2003			2008		
	GE(0)	GE(1)	GE(2)	GE(0)	GE(1)	GE(2)	GE(0)	GE(1)	GE(2)
Total	44.0	41.8	61.1	40.2	38.7	54.6	42.3	42.1	67.3
Between-group inequality	7.7	7.1	6.7	11.0	8.0	6.8	10.0	6.9	5.7
Between as a share of total	17.5	16.9	11.0	27.3	20.7	12.4	23.6	16.4	8.4
Within-group inequality	36.3	34.8	54.3	29.2	30.7	47.9	32.3	35.2	61.7

If time permits, the Instructor can mention one final topic that may be of interest: the last inequality axiom has to do with decomposability of the measure, and these slides explain the benefits of such a property. Once inequality is estimated at the national level, if we care about population subgroups (such as urban and rural areas, or different regions) a decomposable measure can be additively decomposed into a within and a between component (this is explained non-technically in one slide, and a hidden, optional slide may be covered if the Instructor believes showing the mathematical formula for the decomposition can help). One or more example of inequality decomposition using an index of the GEI family can be shown (note that the Gini index is not additively decomposable).

 **Lessons learned**

- Many ways to **describe** inequality, some more effective than others
- Graphs**: most notable are quantile functions and Lorenz curves
- Measures**: different inequality measures lead to different results
 - The **Gini** index is an extremely popular measure - we recommend its use.
 - The **GEI** (generalized entropy indices), and in particular the **MLD** (mean log deviation), are recommended because of their theoretical properties.

Lessons learned can be customized, as required. Here we summarize a few key messages, namely that there are many ways to describe inequality, but only few of them ‘work’, that is only few have sound theoretical foundations and a wide practical application. In essence, our recommendation is to focus on the Lorenz curve (and possibly the quantile function), the Gini coefficient, and the Mean Log Deviation.

Homework

Exercise 1 - Engaging with the literature

See exercise 1, Lecture 1.

Exercise 2 – Inequality in South Asia

The main point here is that it is not a good idea to mix expenditure- and income-based inequality estimates.

Exercise 3 – Functional vs. Personal Income distribution

The exercise is for students interest in this specific topic (functional vs. personal distribution), and as a way to check whether their understanding of the report has benefitted from attending the lecture and read Cowell's reference.

Lecture 14

Measuring poverty

Learning objectives

The goal of this lecture is to provide a conceptual framework for measuring poverty. Given the dimension and the technical contents of the literature, we limit the discussion to a short review of how analysts define a poverty line, and how they measure the incidence, depth and severity of poverty.

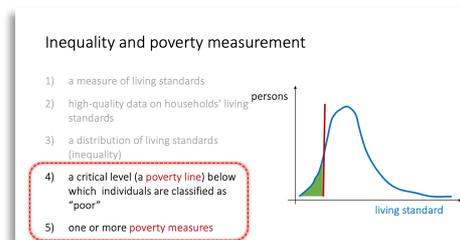
Suggested preparation

The lecture follows the sections indicated in the syllabus from Ravallion (2016). The Instructor might find it useful to become familiar with the material in Chapters 2-4 in Haughton and Khandker (2009).

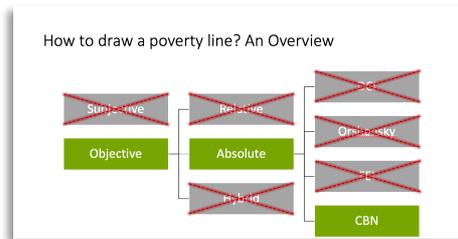
Time allocation

Poverty lines	55 min
Break	15 min
Poverty measures	45 min
Lessons learned	5 min

Annotated lecture



The lecture begins, as usual, with the outline of the topics that will be covered. Today is a simple two-part structure, where we explain poverty lines in part 1, and poverty measures in part 2.



The opening slide argues that there is a **variety of approaches** that have been proposed to ‘draw the line’: we have chosen one, out of the many possible, which is called the Cost-of-basic-needs (CBN) method.

The illustration of the CBN method will take place after dismissing all other methods. We share our strategy with the audience. We announce from the onset, for instance, that we will define and explain what is meant by a ‘subjective’ poverty line and an ‘objective’ one, and explain that we will criticize and dismiss the former and focus on the latter. Next, we define and explain what is meant by ‘relative’, ‘absolute’ and ‘hybrid’ poverty lines, and explain that we after discussing pros and cons associated to each option we will focus on absolute poverty lines. And so on, until we reach the ‘CBN’.

Subjective poverty lines – I/III

- Poverty lines are inherently **subjective judgments** people make about what constitutes a socially acceptable **minimum standard of living** in a particular society at a given time (Ravallion 1994: 42).
- The subjective poverty approach is based on the **self-assessed** adequacy of a family’s food, housing, and clothing.
- How are poverty lines estimated, **in practice?**

Objective poverty lines

- An objective poverty line is one based on some **objective metric**, such as consumption or income.



The first set of slides is on the **subjective vs. objective** poverty debate. The discussion is only supposed to share the *ideas* underlying these methods, to conclude that there is revival in the interest of the subjective approach, but currently, it is not the one implemented in national studies of poverty in most countries in the world. We leave it to interested students to read more on the topic.

Absolute poverty lines

- An **absolute poverty line** is one which is **fixed** in terms of the average standard of living (or welfare).
- Example: cost of a bundle containing “basic commodities”, **however defined**.
- Note 1: ‘**fixed**’ is a false friend. An absolute poverty is defined in a specific context and time, that is fully historically determined. Fixed ≠ unchanging.
- Note 2: ‘**absolute**’ is not a synonym of ‘low’ – an absolute poverty line can be as generous as the analyst or the society wishes.

Relative poverty lines

- A **relative poverty line** is one which **varies** with the average standard of living.
- Example: half the mean of per capita income.
- The EU definition of **relative poverty line**: “Low income rate after transfers with low-income **threshold set at 60% of median [equivalized] income**, with breakdowns by gender, age (...)”
- Question: **why 60%**? Why this specific number?
- Answer: **who knows?** Indicator 11 (“Dispersion around the low income threshold”). **Three thresholds: 40, 50 and 70% of the median income.**

The discussion on **absolute vs. relative poverty** is an important one. First, we define both concepts. The meaning of ‘absolute’ in this context should be made very clear: it does not mean that it is fixed and immovable. On the contrary, an absolute poverty line may change over time, if the country decides to re-calculate it in order to make it consistent with the consumption patterns that are currently prevalent in the country. ‘Absolute’ also does not mean ‘extreme’: on the contrary, it might indicate a level of consumption that is well above

starving. The deeper meaning of these caveats will be clearer when the absolute approach is contrasted with the relative one.

The problem with relative poverty: the richer...the poorer?

x1	x2	x3	x4	x5	total	mean	poverty line (50% of the mean)	poor
2	2	16	20	60	100	20	10	40%
3	3	24	170	300	500	100	50	60%

An awkward feature of relative poverty lines is that a policy which raises the living standards of all, but **proportionally more** those of the rich, will increase poverty, notwithstanding the fact that the absolute living standard of the poor has increased!

Relative poverty lines
an assessment

- In short, relative poverty = inequality
- In lecture 13 we discussed inequality measures at length – we have better tools for measuring inequality than ‘relative poverty’ measures
- However... World Bank (2017):

Recommendation 16: The World Bank should introduce a “**societal**” **headcount ratio** measure of global consumption poverty that takes account, above an appropriate level, of the standard of living in the country in question, thus **combining fixed and relative** elements of poverty

The reason why we ultimately dismiss the relative poverty approach is explained by means of a numerical ‘toy’ example, which leads us to a paradoxical conclusion regarding the level of poverty in two countries, or two moments in time (the two rows of the table). This should illuminate the fundamental defects of relative poverty. The key message in these slides is that relative poverty is – de facto – a measure of inequality. This message can be softened and nuanced, but this is the first bring-home conclusion from this set of slides.

Absolute poverty lines
many popular methods but one key idea: food is the anchor

- 1) Direct Calorie Intake (DCI)
Kakwani (2003)
- 2) Food Energy Intake (FEI)
Dandekar & Bath (1971) + Greer & Thorbecke (1986)
- 3) Food-share
Orshansky (1963, 1965)
- 4) Cost of Basic Needs (CBN)
Rowntree (1901) + Ravallion (1994)

Now that our focus on objective, absolute poverty has been motivated, we focus on the **CBN approach** to setting a poverty line. Our account of the CBN method is masterly explained in Martin Ravallion’s work, from his 1994 book (check the references) all the way to his recent 2016 book. There is no need to repeat all details here. The instructor is encouraged to customize these slides the way she/he thinks best for the audience. In our experience, the most effective way is first to explain the idea in words, next to use the support of one or more graphs, and finally to double check the understanding of the audience by illustrating the formulae (we provide the non-parametric version of the method, but the Instructor might wish to replace it with its parametric, regression-based counterpart. This is what students find in the Appendix of Ravallion 1994). This is the structure that the Instructor will find in the following slides. In the next pages of this Manual, we give a quick description of the main steps of our argument.

The Cost of Basic Needs (CBN) method

- In a nutshell: estimate the cost of a **consumption bundle** adequate to meet **basic** consumption needs.
- Question
What constitutes a ‘**basic need**’ and what does not?
- Constraint
The choice of the basic-needs bundle should reflect local perceptions of what constitutes poverty (**specificity**).
- Solution
A safe start consists in including **foodstuffs** among the basic needs. After, we’ll think of how to add an **allowance** for consumption of **non-food goods/services**.

The CBN method: A strategy

- Three steps:
 - 1) Estimate the **cost of a ‘basic food bundle’**: this gives the **food poverty line**
 - 2) Estimate the **allowance for ‘basic non-food goods’**
 - 3) Add 2) to 1): this gives the **(total) poverty line**

The basic premise of the CBN method is rather intuitive. The goal is to estimate the **cost of a basic consumption bundle**: this will give us a threshold, beyond which individuals

are considered to be poor, i.e. not able to satisfy basic needs. We start from the cost of a basic food bundle, and then we add the cost of any non-food items that are considered ‘basic’ on top of it. The whole essence of the CBN method is in how these two problems are solved.

The food poverty line (FPL)

- How to define a ‘basic food bundle’?
- The key idea, which does not require any arbitrary assumption on consumption patterns, is to:
 - 1) estimate the **minimum energy requirement** for the average individual in the target population (say 2,000 kcal/person/day)
 - 2) **price** that amount of calories, using the **average cost of one kcal** which is computed using the survey data.
- A monetary amount is obtained, and that is the food poverty line (FPL)
- Note that 3) takes account for local tastes (preferences)

Determining the **food poverty line** – the cost of a basic food bundle – is perhaps the easiest of the two steps. The most ‘neutral’ way of doing it (that is, a way that does not require any assumptions on which food items are essential and which are not) relies on calories. First, we estimate the minimum caloric requirement for the average individual in the population (to do that, we rely on parameters that are disseminated by medical authorities, and on the demographic structure of the population). Then, we estimate the average cost of one calorie in our sample of household-level data (usually, by restricting the sample to the bottom part of the expenditure distribution, so that we target the right group of people). Multiplying the minimum energy requirement by the average cost of a calorie gives us a monetary amount: the FPL.

The non-food allowance (NFA)

- How much is the minimum for non-food necessities?
- We start by asking the data
- Focus on a subset of people that are **most likely poor**, and see how much they spend on non-food
- Two ways to define that target population:
 - 1) people whose total expenditure is about as much as the food poverty line (**lower bound**)
 - 2) People whose food expenditure is about as much as the food poverty line (**upper bound**)

Allowance for non-food goods

Next is the determination of a **non-food allowance**. Again, the approach is to try and assume as little as possible about what the poor need: our preference always goes to objective estimates (‘asking the data’). The key idea is to define a target group of people in the population, who are most likely poor, and to calculate how much they spend on non-food items on average. According to the way we define this target, we can obtain two different non-food allowances. The method is explained by way of a graph first.

*Lower bound CBN poverty line

- $PL = FPL + NFA_L$
- $NFA_L = E_n(x_n^{nonfood} | x_n \approx FPL)$

*Upper bound CBN poverty line

- $PL = FPL + NFA_U$
- $NFA_U = E_n(x_n^{nonfood} | x_n^{food} \approx FPL)$

Optionally, and if deemed useful, the Instructor can use formulas, as well. The gist is that the lower bound non-food allowance (NFA) is the average non-food expenditure among

those households whose *total* expenditure is close to the food poverty line (they are extremely poor, cannot even afford the FPL, therefore whatever they spend on non-food must be essential for them); the upper-bound NFA is the average non-food expenditure among those households whose *food* expenditure is close to the food poverty line (they are still poor, although not quite as much, and obviously consume non-food items as well as food items).

Lower and Upper Bound CBN Poverty Lines
Recap

- $LBPL = FPL + E_h(x_h^{nonfood} | x_h \approx FPL)$ (lower bound PL)
- $UBPL = FPL + E_h(x_h^{nonfood} | x_h^{food} \approx FPL)$ (upper bound PL)
- Which one to choose?
- It is customary to report results on them all (FPL, LBPL, UBPL), but if there needs to be one number, it is often based on UBPL

Finally, the FPL and NFAs are added together, to obtain **two poverty lines, lower and upper**. Conventionally, both poverty lines are used, although some analysts may prefer to focus on the UBPL.

Important remark

- The CBN method hinges on the **food poverty line**
- A good food poverty line requires good estimates of **calorie intake**
- Good estimates of calorie intake require a well designed **questionnaire** (lectures 5-7)

The Instructor is encouraged to stress this point: data producers should focus on the collection of data that allow for the estimation of **calorie intakes** because, among other uses, calories are the foundation of the CBN method for estimating poverty lines.

Zambia, 2015
Living Conditions Monitoring Survey

Food item	Unit	Quantity	Unit price	Cost
Cooking oil local	l	1	28	28
Dried beans	1kg	2	13	27
Dried bean	1kg	1	85	85
Dried legume	1kg	2	104	207
Fresh milk	50ml	4	5	20
Onion	1kg	4	10	40
Shelled groundnuts	1kg	3	13	39
Table salt	1kg	1	5	5
Tomatoes	1kg	4	5	21
White onlon	20kg	3.6	54	194
Vegetables	1kg	7.5	4	29

Total per family (six people or 4.52 AE) 685

Total per AE 152

Food Poverty Line

Zambia, 2015
Living Conditions Monitoring Survey

The non-food allowance was determined as the average non-food consumption of households whose total consumption was close to the food poverty line.

$LBPL = FPL + E_h(x_h^{nonfood} | x_h \approx FPL)$

	2015
Total	LBPL 214
Food	FPL 152
Nonfood	Non-food Allowance 62

Note: At average national prices of April/May 2015. Source: CSO World Bank estimations.

One example concludes the presentation on poverty lines.



Now is a good time for a break.

Poverty measures
Basic ideas

- Poverty measures **aggregate** information.
- A **poverty measure** is a function of individual incomes $x = (x_1, \dots, x_N)$ and the poverty line z :
 $P : \mathbb{R}^N \rightarrow \mathbb{R}_+$
- The literature on poverty measures is huge and technical in nature. It deals with the choice of the **functional form** of a **suitable** poverty index.
- In practice, three indices have taken center stage:
 - 1) the **headcount ratio**
 - 2) the **poverty gap index**
 - 3) the **poverty gap squared index**

After the break, the lecture resumes by the question: after estimating a poverty line, how do we summarize the extent of poverty in the society? We engage a short discussion with students. For instance, we can claim that ‘the best way for measuring poverty is to count the number of households (or individuals?) whose expenditure (or income?) falls short of the poverty line. The we ask: can anyone think of a ‘better’ way to measure poverty? What’s wrong with counting the poor? We keep the discussion short by not replying to suggestions – the point of this discussion is to deliver a first message: there are many **measures** available, different indicators deliver different results, not clear how to identify the best indicator. Hopefully, students should be motivated to attend the second part of the lecture. For our purposes, the complexity of the choice can be reduced, in practice, to the discussion of only three indices, which we denote here by **H, PG and PG2**.

The rest of the lecture is devoted to introducing and assessing the three indices.

The poverty headcount ratio (H)
Mongolia HSES 2016, Cumulative distribution of per capita consumption (p.10)

- The **headcount ratio** is the proportion of the population that is classified as poor.
- $H = \frac{\alpha}{N} = \frac{1}{N} \sum_{h=1}^N I(x_h \leq z)$
- $I(\cdot)$ is an **indicator function** that is 1 if its argument is true, 0 otherwise.
- Interpretation: **incidence** of poverty

The Poverty Gap Squared
Definition

- The **squared poverty gap** index attributes more weight to the poorest among the poor:

$$PG2 = \frac{1}{N} \sum_{i=1}^N \left(1 - \frac{x_i}{z}\right)^2 I(x_i \leq z) = \frac{1}{N} \sum_{i=1}^N \left(1 - \frac{x_i}{z}\right)^2$$

- The contribution of the i -th individual to PG2 is larger the poorer she is, that is, the larger is her poverty gap $(z - x_i)/z$:

$$PG2 = \frac{1}{N} \sum_{i=1}^N \left(1 - \frac{x_i}{z}\right) \times \left(1 - \frac{x_i}{z}\right)$$

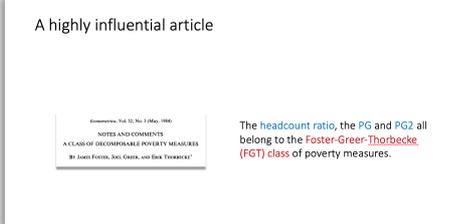
weight gap

For each index, the strategy proposed in the slide is as follows. First, we provide a definition in words (“The headcount poverty rate is defined as...”). Next, we give the same definition using a formula. The notation is consistent for the three indices and the use of the **indicator function** (defined in the first slide) is introduced and described as a useful device to filter out the ‘rich’, that is, individuals whose expenditure is greater than the poverty line. We help students to read the math, exactly as if they would read plain English – this is a challenge for the Instructor, sometimes, but a high-return investment. In our experience students will appreciate mastering the formulas that we have selected. Third, we provide a visual illustration of each index (with the exception of the poverty gap squared index, for which none is available). So, the headcount ratio is illustrated by means of the CDF, while

the PG by means of the quantile function. Finally, we provide a short assessment of each index, a phase that lends to engaging with students (time permitting).

The material covered in these slides is standard, and the Instructor should feel free to amend it by adding numerical examples, for instance.

A highly influential article



FGT (1984)
Definition

The FGT class of poverty measures:

$$P_{\alpha} = \frac{1}{N} \sum_{h=1}^M \left(\frac{z - x_h}{z} \right)^{\alpha} I(x_h \leq z), \quad \alpha \geq 0$$

α	P_{α}	Index
0	$P_0 = H$	HEADCOUNT RATIO
1	$P_1 = PG$	POVERTY GAP INDEX
2	$P_2 = PG2$	POVERTY GAP SQUARED
...		
∞	P_{∞}	weights the poorest person

The last two slides introduce the **FGT class of poverty measures**. It is important that students familiarize with this formula. The paper is a famous and classical one, but it is *not* a recommended reading for our students. Too technical and not worth the effort, for the intended audience of this course. It is important however that students understand that i) H, PG and PG2 belong to the same family of indices (the so called ‘FGT class of poverty indices’), ii) different indices assign different weights to different people (this is what we explain in the slide for PG2), 3) the FGT class of poverty indices has become the most popular among welfare analysts thanks to its theoretical axiom-derived properties: the logic is similar to that introduced in the lecture on inequality.

 Lessons learned

- 1) We argued in favour of objective, absolute, CBN poverty lines.
- 2) Regarding poverty measures:
 - The headcount ratio is a crude and ‘theoretically inferior’ poverty index. H is useful, but **should not be used exclusively**.
 - The Poverty Gap Index and the Squared Poverty Gap Index are complements to H; poverty analysis should **combine the three measures**. We recommend FGT (1984).
 - The **axiomatic approach** does not succeed in identifying the “best” poverty measure. Yet, it is **useful**, as it reveals the principles underlying the poverty measures.

The last slide summarizes the main points discussed during the lecture. We have been very parsimonious in summarizing the discussion on poverty lines, for instance. The Instructor can consider adding a second slide with additional conclusions/considerations, should she wish so.

Homework

Exercise 1 - Engaging with the literature

See exercise 1, Lecture 1.

Exercises 2 and 3 – DASP and ADePT

Students who have basic command of Stata can take advantage of exploring DASP and ADePT, two popular tools in use among welfare analysts. The Instructor can devise and assign specific exercises after preparing suitable small-size datasets.

Lecture 15 Describing data

Learning objectives

This final lecture closes the course by discussing the final stage of the survey process: the dissemination of final results. It presents general principles that guide the presentation of findings from a typical household survey, with emphasis on how to design effective tables and graphs, and how to present inequality and poverty estimates.

Suggested preparation

Glewwe and Levin (2005) is a useful reference on the presentation of descriptive statistics from household survey data, and the Instructor is assumed to be familiar with it. Schwabish (2014) is focused on the visualization of economic data, and is the source for some of the material presented in the lecture.

Time allocation

Background information	20 min
Descriptive statistics	40 min
Break	15 min
Poverty and inequality measures	45 min
Lessons learned	5 min

Annotated lecture

The dissemination phase

- Last step of the journey
- Typically, after survey and data processing are concluded, a report is released, aimed at describing **main findings** from survey



The subject of this lecture is the construction of the **final report** that is typically released in the final stages of the survey process, to disseminate the main findings from the survey, and potentially accompany the dissemination of microdata files for public use. An example of this kind of report is reproduced here (cover of the Rwanda ‘Main indicators report’).

Dissemination of what?

Topics vary according to contents of the survey, **target audience**, etc., but there are some common elements:

1. **Background information** on sampling
2. **Descriptive statistics** (roughly corresponding to survey modules)
3. In the case of income and expenditure surveys, measures of **inequality and poverty**

Next slides cover these 3 points, with tips for effective presentation and examples.

A note on the target audience

- Addressing a **technical or academic audience** is different than informing **laypeople or the media**
- NSOs often use different communication instruments for these different audiences, e.g. technical reports vs. press releases
- But ‘**technical**’ is not a synonym for ‘**obscure**’
- This lecture covers **general principles** that are useful for different audiences

The rest of the lecture will present some guidelines that can be applied to the construction of the final report. Context matters, of course, but it is possible to make some general points. We focus on three main components of the report listed here: background information on sampling, descriptive statistics, and, more to the point of this course, the presentation of statistics on poverty and inequality. One important point that the Instructor should stress is that, of course, the **target audience** matters: the best strategy for effective communication changes depending on whether the NSO wants to reach ‘technical’ people or the general public. However, one of the takeaways of the lecture is that a ‘technical’ report does not need to be obscure: certain general rules for clarity and rigor apply to all types of dissemination.

1. Background information

...

What not to miss

- Reports should document at least the following **survey design features** and **processing choices**:
 - a. **Sampling design**
Sample size, stratification, representativeness...
 - b. **Data collection and processing**
Fieldwork, outlier detection and treatment, data imputation...
 - c. **Definitions of economic concepts and aggregates used**
E.g. disposable income, total household consumption, imputed rent... May be presented as a glossary

Some **background information** on the survey process should be included in any report. These slides present a ‘checklist’ for the pieces of information that should never be missed.

Documentation on sampling design
what to include

- Sampling design report with
 - Allocation of sample into strata and indication of **excluded strata**, if any
 - Estimation **formulas** (selection probabilities and weights)
- Household listings forms
- Sample frames
 - For the first sampling stage/s: list of all **sampling units**
 - For the last sampling stage: list of all **households** in each sample point
- Non-response rates
- On the survey datasets
 - Sampling **weights**

...

Other documentation
What to include

- Organizations included in preparation of work
- List of data sets and contents
- Codes not found in the questionnaire
 - Occupation codes
 - Industry codes
- Other information
 - Exchange rates
 - Consumer Price Index
- Supervisor control forms
- Questionnaire control forms
- Maps
- Abstract

We now go into further detail, by providing some concrete examples. These slides illustrate the type of information on **sampling design, data collection, and the definition of economic concepts and aggregates** that may appear in a typical report: there is no need to dwell on each detail, usually experts will be tasked with drafting these technical sections. However, it is important for students to understand what is meant by ‘documentation’, in practice. Examples from recent reports – Kenya, Uganda, and South Africa – have the same illustrative purpose.

2. Descriptive statistics

...

Text

- **Effective writing** complements good tables and graphs
- This lecture will focus on the latter: writing deserves a separate discussion
- A useful reference



Next, we delve into the topic of presenting **descriptive statistics**. Text, tables, and graphs are our main tools when it comes to describing data; text, that is, writing about data, is a topic in and of itself, and is not going to be discussed here. Interested students are encouraged to check the cited reference.

Tables

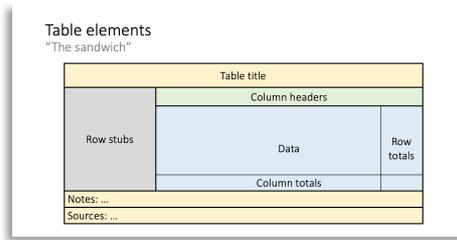
- Tables are **omnipresent** in data dissemination reports
- Often used when describing two variables jointly (two-way tables), e.g. income by region, population by age...

What do you think of this table?

Income inequalities by area of residence and geopolitical Zone for 2001 and 2010

		2001	2010	% change from 2004 to 2010
National		0.4296	0.447	4.1
Area of Residence	Rural	0.4230	0.4334	2.2
	Urban	0.4154		4.2
Geo-political zones				
	1 South South	0.3849	0.434	12.8
	2 South East	0.376	0.4442	18.1
	3 South West	0.4088	0.4097	0.2
	4 North Central	0.4059	0.422	5.4
	5 North East	0.4114	0.4468	8.6
	6 North West	0.4028	0.4056	0.7

We quickly move on to **tables**. The example shown here is simply a device to be used – if the Instructor feels like it may be useful – to engage the audience. No long debates are needed: the discussion should simply highlight a few points that students may deem important (using clear definition, consistent formatting, alignment, etc.) which the Instructor will soon organize in a few clear ‘rules’.



This slide is meant to give students some **terminology**, and to highlight the importance of tables being self-contained, thanks to notes and sources.

What makes a good table
Golden rule #1

Express contents clearly

1. The **table title** should answer the questions "what", "where" and "when", but still be concise
2. Tables should be **self-contained**: use notes to clarify definitions, abbreviations, etc.
3. Percentage distributions of discrete variables should be clearly identified as either **percentages of households or percentages of the population**
4. **Row and column totals** should be reported, when they identify a marginal distribution

What makes a good table
Golden rule #2

Reduce clutter

1. Avoid unnecessary **colors, repetitions** (e.g. use % or \$ just once, in the title, rather than throughout the table)
2. **Precision of numbers**: do not present too many significant digits. Percentages: one decimal digit is usually enough. Numbers with four or more digits: no decimals at all. Large numbers: express them in thousands or millions
3. Be mindful of **spacing and alignment**

We offer a few '**golden rules**' on how to present effective tables. These boil down to two simple but crucial principles: explaining contents clearly, and reducing clutter. We elaborate on each 'rule' by offering specific tips.

What's wrong with this table?

Final energy consumption by sector - Percentages	1980	1985	1990	1995	2000	2002	2003
Transport	27.81	27.92	28.74	31.12	35.53	39.48	39.13
Residential	38.81	38.98	39.44	37.60	35.93	33.71	29.07
Industry	31.47	27.21	23.86	22.11	21.43	19.53	18.79
Buildings	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Services	9.95	10.96	13.98	15.45	16.33	18.37	19.3

Source: UNECE (2009: 12)

A possible table redesign
UNECE (2009: 12)

Share of total energy consumption, by sector (in percent) Ireland, 1980-2003	1980	1985	1990	1995	2000	2002	2003
Transport	27.8	27.9	28.2	31.1	35.8	39.5	39.1
Residential	31.1	29.9	28.4	27.6	24.3	23.7	24.0
Industry	31.5	27.2	23.9	22.1	21.4	19.5	18.8
Agriculture	n/a ¹	n/a ¹	3.5	3.7	3.1	2.9	2.8
Services	9.6	11.0	14.0	15.5	14.4	14.4	15.3
Total	100.0	100.0	100.0	100.0	100.0	100.0	100.0

¹Data on energy consumption for the agricultural sector was not collected until 1996.
Source: Department of Public Expenditure, Ireland.

Note clarifies meaning of "n/a":

An example clarifies even further. Students can be engaged directly in commenting the first table. The 'redesigned' table is an example, and is by no means the only way to put together an effective table, but it incorporates a number of improvements following the 'golden rules': a more informative title, alignment and formatting that facilitate legibility, only one decimal digit, column totals to clarify the interpretation of cells, notes and sources.

Graphs

- In many cases, presentation of data can be made more interesting and intuitive by using **graphs** or charts rather than **tables**
- Many of the "golden rules" that help make better tables also apply to graphs

What makes a good graph
Golden rule #3

Express contents clearly

1. A good **graph title** answers the same questions as a good table title
2. Graphs should be **self-contained** too (use notes)
3. **Explain encoding**: always label axes and data series clearly
4. **Avoid visualizations that mislead the eye**: two notorious "sins" are bar charts with a nonzero baseline, and 3D pie charts

We adopt a similar approach for the case of **graphs**. The two 'golden rules' are the same in spirit, but they are obviously adapted to the context of data visualization.

Bar charts with nonzero baseline

- Bar charts rely on **bar length** to show data: compare lengths to compare values
- Shifting the baseline **distorts the visual**: a value twice as high no longer corresponds to a bar twice as long
- Graphs on the right show the same data, but appear very different

3D pie charts

- Pie charts encode data in the **area** of each slice: larger slice equals higher share
- A 3D pie chart **distorts angles**, making the slice that is "closer" to the viewer appear larger than it actually is
- This visualization can **mislead** viewers, and should be avoided

Here are two common examples of graphics that do not express their message in the clearest way. Bar charts with nonzero baseline and 3D pie charts are misleading, and should be avoided.

What makes a good graph

Golden rule #4

Reduce clutter

- Again, avoid unnecessary **colors and decorative elements** that obfuscate the message of the graph
- Precision of numbers**: same recommendations as for tables
- Do not crowd graph with **too many data points**: viewer should be able to understand the message of the graph easily, without having to parse too much visual information (if that is the issue, select a subset of relevant values, or consider using a table instead)

The golden rule of 'reducing clutter' applies to graphs as well.

On pie charts

A. Pie Chart

B. A Pie Chart, Rotated

"Because **pie charts force readers to make comparisons using the areas of the slices or the angles formed by the slices—something that our visual perception does not accurately support** — they are not an effective way to communicate information" Schwabish (2014: 223)

Graph redesign

A pie chart, labeled

Pie chart alternative: a bar or column chart

Source: Schwabish (2014: 223)

Pie charts, even when not 3D, can be a less-than-optimal way of presenting information, as mentioned by Schwabish (2014).

What's wrong with this graph?

Percentage of Employed Who Are Active Managers, by Sex, 2008

Source: Schwabish (2014: 223)

Graph redesign

Percentage of Employed Who Are Active Managers, by Country, 2008

Source: Schwabish (2014: 223)

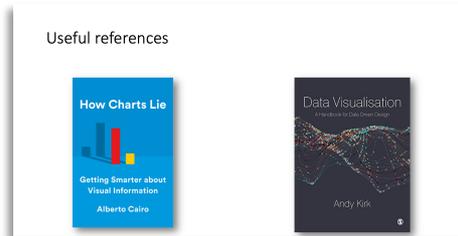
What's wrong with this graph?

Zambia, Living conditions monitoring survey report 1996 and 2015

Incidence of poverty rural/urban, 1996

Incidence of poverty rural/urban, 2015

These slides show some examples of cluttered visualizations, and of some ways in which they can be modified to bring the data back into the spotlight, eliminating distractions. The ‘redesigned’ graphs are not supposed to be the only good solution, just an example of visualization that takes advantage of the ‘golden rules’.



Interested students can take advantage of the references cited here: these are accessible resources on data visualization that do not focus specifically on poverty or inequality, but can certainly prove useful to Statistical Officers.



Now is a good time for a break.

3. Inequality and poverty

Finally, some guidelines on the **presentation of results on inequality and poverty**, one of the main outcomes of consumption and expenditure surveys.

Overview

- Tips for presentation of generic summary statistics still apply
- There are a few additional points to be made specifically about presenting results on poverty and inequality:
 - a. Popular **measures** and graphics (from lectures 13 and 14)
 - b. Best practices for making **comparisons**

...

Suggested poverty measures
2015/16 Kenya Integrated Household Budget Survey (KIHBS)

▪ **FGT**, the rest is extra credit

Table 4.3: Overall Poverty Estimates (Individual) by Residence and County, 2015/16

Residence / County	Headcount Rate (%)	Distribution of the Poor (%)	Poverty Gap (%)	Severity of Poverty (%)	Population ('000)	Number of Poor ('000)
National	36.1	100.0	10.4	4.3	45,371	16,401
Rural	40.1	71.3	11.5	5.0	29,127	11,687
Peri-Urban	27.5	5.6	6.9	2.6	3,340	920
Core-Urban	29.4	23.1	8.9	3.9	12,905	3,795

Poverty and inequality measurement have been the subject of previous lectures. At this stage – the dissemination of results – it is useful to know which of the available measures have come to be expected by the international community of readers, practitioners, and academics, who are likely in the target audience of the report. The **Gini index and FGT**

poverty measures are essential; more measures are a welcome addition. This group of slides includes examples from recent reports.

Making comparisons

- Many audiences (policy makers, general public) are especially interested in comparisons of poverty and inequality, over time or across regions
- Poverty and inequality **trends** are among the most visible and impactful results to emerge during dissemination
- Comparability** of underlying data and methods is key: if processes that led up to estimates differ, comparison is **invalid**
- Being **transparent** on comparability is key!

Tanzania, 2012

Poverty Assessment, HBS 2007 and 2011/12 recall modules

Consumption and expenditure categories	HBS 2011/12 Recall period (months)		HBS 2007 Recall period (months)	
	1	12	1	12
Clothing and footwear (COFOP 3)		X		X
Housing and utilities (COFOP 14 - selected other)				
- Sewer	X		X	
- Utilities		X		X
- Energy	X	X	X	X
- Building maintenance		X		
Housing equipment (COFOP 16)				
- Household durables, furniture and furnishings		X		X
- Small household appliances		X		X
- Expenditure on transport services				
Health expenditures (COFOP 16)				

Another important point is that of **comparability**. Methodological changes are a threat to consistent comparisons over time and across countries, and it is crucial that they are thoroughly documented when results are released, to allow for a critical assessment of trends.

The importance of uncertainty

- Poverty calculations are based on a **sample** of households, and samples carry a margin of error in representing the population
- Standard errors** should always be estimated along with poverty point estimates
- Crucial when making **comparisons** (over time, across regions): poverty changes should not be taken at the face value
- Note: probability weighting, clustering, and stratification, are **survey design features** which must be taken into account when estimating standard errors.

Sensitivity Analysis

Conclusions

Table 6.2 Key characteristics of poverty and its robustness to measurement assumptions.

Characteristics of poverty	Baseline, consumption per capita	GC1 I scale	GC2 II scale	Higher poverty line	Lower poverty line	Expenditure per capita
Mixed (semi-urban) municipalities in RS	yes	yes	yes	yes	yes	yes
Rural municipalities in FRH	yes	yes	no	no	yes	no
IDPs and Refugees	yes	yes	yes	yes	yes	yes
Households headed by persons with low education (primary or less)	yes	yes	yes	yes	yes	yes
Households headed by persons with education above secondary	no	no	no	no	no	no
Unemployed (ILO) and inactive adults	yes	yes	yes	yes	yes	yes
Employed according to registration	no	no	no	no	no	no
Registered according to identity	no	no	no	no	no	no
Household headed by elderly	no	no	no	no	no	no
Large households	yes	yes	yes	yes	yes	yes

Source: Staff estimates based on BHPLDHS 2007.

A final point concerns **uncertainty**. One of the recurring topics of the course has been the impact of survey design features – as well as analytical choices – on final estimates. These slides emphasize the importance of being transparent in conveying the unavoidable uncertainty of estimates, and recommends two tools, which are too rarely used in standard reports: standard errors, and sensitivity analysis. Examples from recent publications illustrate.

Homework

Exercise 1 - Engaging with the literature

See exercise 1, Lecture 1.

Exercise 2 - Standard Errors

This exercise asks students to elaborate on the meaning and importance of reporting standard errors. One way to discuss the tables is the following. The point estimate plus or minus 1.96 times its standard error gives the upper and lower bounds a 95% confidence interval. For instance, the confidence interval for the poverty headcount rate in urban areas is (27.5, 39.3), obtained as $33.4 \pm 1.96 \cdot 3$. If one were to repeat this calculation for all reported point estimates, one would discover that at least some of the confidence intervals for strata overlap, which would imply that, for instance, poverty rankings of geographical areas may need to be taken with caution.

Exercise 3 - Sensitivity analysis

This exercise is aimed at conveying the meaning of sensitivity in a practical setting. Changing the minimum calorie requirement from 2100 to 2200 kcal per person per day causes a large change in the poverty line (green to blue line), and, consequently, in estimated poverty. Adding just 23 more calories to the calorie norm generates a jump in the line that is almost as large (blue to grey line), while increasing the norm to 2300 only moves the line minimally. The poverty line is much more sensitive to the calorie norm at around 2200, while it becomes more robust over 2238.

References

- A'Hearn, B., Amendola, N., and Vecchi, G. (2016). On historical household budgets. *Rivista di Storia Economica*, 32(2), 137-176.
- Alkire, S. (2013). Towards frequent and accurate poverty data.
- Alkire, S., and Samman, E. (2014). Mobilising the household data required to progress toward the SDGs.
- Alvarez, E., Garcia-Fernández, R. M., Blanco-Encomienda, F. J., and Munoz, J. F. (2014). The effect of outliers on the economic and social survey on income and living conditions. *World Acad. Sci., Eng. Technol., Int. J. Soc., Behav., Educ., Econ., Bus. Ind. Eng.*, 8, 3276-3280.
- Amendola, N. and G. Vecchi (2014), Durable goods and poverty measurement, World Bank Policy Research Working Paper no. 7105.
- Atkinson, A. B. (1970). On the measurement of inequality. *Journal of economic theory*, 2(3), 244-263.
- Atkinson, A. B. (2015). *Inequality: What Can Be Done?* Harvard University Press.
- Atkinson A. B., WBG (2017). *Monitoring Global Poverty, Report of the Commission on Global Poverty.*
- Backiny-Yetna, P., Steele, D., & Dijma, I. (2017). The impact of household food consumption data collection methods on poverty and inequality measures in Niger. *Food Policy*, 72, 7-19.
- Balcázar, C. F., Ceriani, L., Olivieri, S. and Ranzani, M. (2017), Rent-Imputation for Welfare Measurement: A Review of Methodologies and Empirical Findings. *Review of Income and Wealth*, 63: 881-898.
- Barnett, V., and Lewis T. (1994). *Outliers in Statistical Data*. 3rd edition. J. Wiley and Sons.
- Barton, A. J. (1958). Asking the embarrassing question. *Public Opinion Quarterly*, 22, 67-68.
- Beegle, K., De Weerdt, J., Friedman, J., and Gibson, J. (2012). Methods of household consumption measurement through surveys: Experimental results from Tanzania. *Journal of Development Economics*, 98, 3-18.
- Bouis, H., Haddad, L., and Kennedy, E. (1992). Does it matter how we survey demand for food?: Evidence from Kenya and the Philippines. *Food Policy*, 17(5), 349-360.
- Borlizzi, A., Delgrossi, M. E., and Cafiero, C. (2017). National food security assessment through the analysis of food consumption data from Household Consumption and Expenditure Surveys: The case of Brazil's Pesquisa de Orçamento Familiares 2008/09. *Food policy*, 72, 20-26.
- Brzozowski, M., Crossley, T. F., and Winter, J. K. (2017). A comparison of recall and diary food expenditure data. *Food Policy*, 72, 53-61.

- Burgess, T. F. (2001). A general introduction to the design of questionnaires for survey research. Leeds: University of Leeds.
- Chen, S., and Ravallion, M. (1996). Data in transition: Assessing rural living standards in southern China. *China economic review*, 7(1), 23-56.
- Cowell, F. (2011). *Measuring inequality*. Oxford University Press.
- Cowell, F., and Flachaire, E. (2007). Income distribution and inequality measurement: The problem of extreme values. *Journal of Econometrics*, 141(2), 1044-1072.
- Cowell, F.A., Jenkins, S.P., Litchfield, J. (1996): The Changing Shape of the U.K. Income Distribution: Kernel Density Estimates. In: Hills, J. (Ed.): *New Inequalities. The Changing Distribution of Income and Wealth in the United Kingdom*. Cambridge University Press, Cambridge.
- Cowell, F., and Victoria-Feser, M. (1996). Robustness Properties of Inequality Measures. *Econometrica*, 64(1), 77-101.
- Cowell, F., and Victoria-Feser, M. (1996). Poverty measurement with contaminated data: A robust approach. *European Economic Review*, 40(9), 1761-1771.
- Dang, H. A., Jolliffe, D., and Carletto, C. (2018). Data Gaps, Data Incomparability, and Data Imputation. *Ecineq WP*, 456.
- De Waal, T., Pannekoek, J., and Scholtus, S. (2011). *Handbook of Statistical Data Editing and Imputation*. New York: John Wiley and Sons.
- De Weerd, J., Beegle, K., Friedman, J., and Gibson, J. (2016). The challenge of measuring hunger through survey. *Economic Development and Cultural Change*, 64(4), 727-758.
- Deaton, A. (1997). *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*. Washington, D.C.: World Bank.
- Deaton, A. , and Grosh., M. (2000), "Consumption." In M. Grosh, and P. Glewwe eds., *Designing Household Survey Questionnaires for Developing Countries: Lessons from Ten Years of LSMS Experience*. Washington, DC: World Bank.
- Deaton, A., and Muellbauer, J. (1980). *Economics and consumer behavior*. Cambridge University Press.
- Deaton, A., and Tarozzi, A. (2000). *Prices and poverty in India*. Princeton, July.
- Deaton, A. and Zaidi, S. (2002). *Guidelines for Constructing Consumption Aggregates for Welfare Analysis*. LSMS Working Paper No. 135. Washington, DC: The World Bank.
- Diewert, W. E. (2004), "Durables and User Costs" in ILO, *Consumer Price Index Manual: Theory and Practice*, chapter 23, ILO / IMF / OECD / UNECE / Eurostat / World Bank.
- Diewert, W. E. (2009), "Durables and Owner-Occupied Housing in a Consumer Price Index" in W. E. Diewert , J.S. Greenlees and C.R. Hulten (eds.), *Price Index Concepts and Measurements*, University of Chicago Press.
- Dupriez, O. (2007). *Building a household consumption database for the calculation of poverty PPPs*. Technical note. Available at: <http://go.worldbank.org/4YG7I5RGT0>.

- FAO and The World Bank. 2018. Food data collection in Household Consumption and Expenditure Surveys. Guidelines for low- and middle-income countries. Rome.
- Farfan, G., McGee, K. R., Perng, J., and Vakis, R. (2019). Poverty Measurement in the Era of Food Away from Home: Testing Alternative Approaches in Vietnam. Policy Research Working Paper Series 8692, The World Bank.
- Farris, F. A. (2010). The Gini index and measures of inequality. *The American Mathematical Monthly*, 117(10), 851-864.
- Fellegi, I. P., and Holt, D. (1976). A systematic approach to automatic edit and imputation. *Journal of the American Statistical Association*, 71(353), 17-35.
- Fiedler, J. L. and Mwangi, D. M. (2016). Improving household consumption and expenditure surveys' food consumption metrics: developing a strategic approach to the unfinished agenda. IFPRI
- Friedman, J., Beegle, K., De Weerd, J. and Gibson, J. (2017). Decomposing response error in food consumption measurement: implications for survey design from a randomized survey experiment in Tanzania. *Food Policy*, 72: 94–11.
- Finn, A. and Ranchhod, V. (2017). *Genuine Fakes : The Prevalence and Implications of Data Fabrication in a Large South African Survey*. Published by Oxford University Press on behalf of the World Bank.
- Foster, J., J. Greer, and E. Thorbecke (1984). A Class of Decomposable Poverty Measures. *Econometrica*, 52, 3: 761–65.
- Gaddis, I. (2016). Prices for poverty analysis in Africa. The World Bank.
- Gertler, Paul J., Elaina Rose, and Paul Glewwe. (2000), "Health." In M. Grosh, and P. Glewwe eds., *Designing Household Survey Questionnaires for Developing Countries: Lessons from 15 Years of the Living Standards Measurement Study*. Washington, D.C.: World Bank.
- Gibson, J. (2007). A guide to using prices in poverty analysis. World Bank, Washington, DC.
- Gibson, J., Beegle, K., De Weerd, J., and Friedman, J. (2013). What does variation in survey design reveal about the nature of measurement errors in household consumption?. The World Bank.
- Gibson, J., and Rozelle, S. (2002). How elastic is calorie demand? Parametric, nonparametric, and semiparametric results for urban Papua New Guinea. *Journal of Development Studies*, 38(6), 23-46.
- Gibson, R. S. (2005). *Principles of nutritional assessment*. Oxford university press, USA.
- Glewwe, P. (2005). Chapter III: Overview of questionnaire design for household surveys in developing countries. In United Nations Statistical Division, United Nations Department of Economic and Social Affairs (Eds.), *Household surveys in developing and transition countries*. New York, NY: United Nations.
- Glewwe, P., and Levin, M. (2005). Presenting simple descriptive statistics from household survey data. In UN, *Household Sample Surveys in Developing and Transition Countries*. Studies in Methods Series F No. 96.

- Grosh, M., and Glewwe, P. (1998). Data Watch: The World Bank's Living Standards Measurement Study Household Surveys. *The Journal of Economic Perspectives*, 12(1), 187-196.
- Grosh, M. and Glewwe, P. (2000). *Designing Household Questionnaires for Developing Countries, Lessons from 15 years of Living Standards Measurement Study, Volume One: World Bank*.
- Grubbs, F. E. (1969). Procedures for detecting outlying observations in samples. *Technometrics*, 11(1), 1-21.
- Harrison, D. E., and Krauss, S. I. (2002). Interviewer cheating: Implications for research on entrepreneurship in Africa. *Journal of Developmental Entrepreneurship*, 7(3), 319.
- Haughton, J. and Khandker, S. R. (2009). *Handbook on poverty and inequality*. Washington, DC: World Bank.
- Heijink, R., Xu, K., Saksena, P., and Evans, D. (2011), Validity and comparability of out-of-pocket health expenditure from household surveys: a review of the literature and current survey instruments. Geneva: World Health Organization, 28.
- Hentschel and Lanjouw (2000), "Household welfare measurement and the pricing of basic services", *Journal of International Development*, 12: 13-27.
- Heston, A. and A.O. Nakamura (2009), Questions about the equivalence of market rents and user costs for owner occupied housing, *Journal of Housing Economics*, 18, 273—279.
- Hlasny, V., and Verme, P. (2018). Top Incomes and Inequality Measurement: A Comparative Analysis of Correction Methods Using the EU SILC Data. *Econometrics*, 6(2), 30.
- Iarossi, G. (2006). *The power of survey design - a user's guide for managing surveys, interpreting results, and influencing respondents*. Washington, DC: World Bank.
- Ibarra, G. L., Mendiratta, V., and Vishwanath, T. (2017). Rental regulation and its consequences on measures of well-being in the Arab Republic of Egypt. The World Bank.
- Jolliffe, D. (2001). Measuring absolute and relative poverty: the sensitivity of estimated household consumption to survey design. *Journal of Economic and Social Measurement*, 27(1, 2), 1-23.
- Judge, G., and Schechter, L. (2009). Detecting problems in survey data using Benford's Law. *Journal of Human Resources*, 44(1), 1-24.
- Krosnick, J. A., Et Presser, S. (2010). Question and questionnaire design. In J. D. Wright Et P. V. Marsden (Eds.), *Handbook of survey research (second edition)* (pp. 263-313). Bingley, UK: Emerald Group.
- Lanjouw, P. (2012), "Consumption-Based Measures in Developing Nations. Lessons from Brazil", in Besharov and Couch (eds.), *Counting the Poor*. New York: Oxford University Press. Ch. 13.

- Lipton, Michael and Ravallion, Martin (1995). "Poverty and policy," Handbook of Development Economics, in: Hollis Chenery and T.N. Srinivasan (ed.), Handbook of Development Economics, edition 1, volume 3, chapter 41, pages 2551-2657 Elsevier.
- Little, R. J., and Rubin, D. B. (2019). Statistical analysis with missing data (Vol. 793). Wiley.
- Lohr, S. L. (2009). Sampling: design and analysis. Nelson Education.
- Lu C, Chin B, Li G, and Murray CJ. (2009) Limitations of methods for measuring out-of-pocket and catastrophic private health expenditures. Bull World Health Organ;87(3):238-44, 244A-244D.
- MacDonald, L., Macpherson, D. A., Sirmans, G. S., and Zietz, E. N. (2006). The value of housing characteristics: a meta analysis. The Journal of Real Estate Finance and Economics, 33(3), 215-240.
- Malpezzi, S. (2002). "Housing". In Grosh, M. and Glewwe, P. (eds.). Designing Household Questionnaires for Developing Countries, Lessons from 15 years of Living Standards Measurement Study, Volume One: World Bank.
- Mancini, G. and Vecchi, G. (2019), On the Construction of a Welfare Indicator for Inequality and Poverty Analysis, mimeo.
- Meyer, B., and Sullivan, J. (2003) "Measuring the Well-Being of the Poor Using Income and Consumption", The Journal of Human Resources, 38.
- Meyer, B. D., and Sullivan, J. (2009). "Five decades of consumption and income poverty". National Bureau of Economic Research.
- Meyer, B. D., and Sullivan, J. (2011). "Further results on measuring the well-being of the poor using income and consumption." Canadian Journal of Economics, 44(1), 52-87.
- OECD (2013). OECD Guidelines for Micro Statistics on Household Wealth.
- Oseni, G., Durazo, J., and McGee, K. (2017). The Use of Non-Standard Units for the Collection of Food Quantity. LSMS guidebook.
- Oseni, G., Huebler, F., McGee, K., Amankwah, A., Legault, E., and Rakotonarivo, A. (2018), Measuring Household Expenditure on Education: a new guidebook on measurement. LSMS guidebook.
- Pape, Utz Johann and Mistiaen, Johan A.. (2018). Household expenditure and poverty measures in 60 minutes : a new approach with results from Mogadishu.
- Pradhan, M. (2009). Welfare analysis with a proxy consumption measure: evidence from a repeated experiment in Indonesia. Fiscal Studies, 30(3-4), 391-417.
- Pyatt, G. (1976). On the interpretation and disaggregation of Gini coefficients. The Economic Journal, 86(342), 243-255.
- Ravallion M. (1994). Poverty Comparisons.
- Ravallion M. (2008) Poverty Lines. In: Durlauf S.N., Blume L.E. (eds) The New Palgrave Dictionary of Economics. Palgrave Macmillan, London.

- Ravallion, M. (2016). *The Economics of Poverty History, Measurement, and Policy*. Oxford: Oxford University Press.
- Ravallion, M. and B. Bidani (1994). How Robust is a Poverty Profile?, *World Bank Economic Review*, 8: 75-102.
- Rousseeuw, P. J., and Croux, C. (1993). Alternatives to the median absolute deviation. *Journal of the American Statistical Association*, 88(424), 1273-1283.
- Schwabish, J. A. (2014). An economist's guide to visualizing data. *Journal of Economic Perspectives*, 28(1), 209-34.
- Sen, A. (1976). Poverty: An Ordinal Approach to Measurement, *Econometrica*, 44(2): 219-31.
- Sen, A. (1987). *The Standard of Living*. Cambridge: Cambridge University Press.
- Sen, A. (1987). *Commodities and Capabilities*. New Delhi: Oxford University Press.
- Shorrocks, A. F. (1980). The class of additively decomposable inequality measures. *Econometrica: Journal of the Econometric Society*, 613-625.
- Smith, L. C. (2015). The great Indian calorie debate: Explaining rising undernourishment during India's rapid economic growth. *Food Policy*, 50, 53-67.
- Smith, L. C., Dupriez, O., and Troubat, N. (2014). Assessment of the reliability and relevance of the food data collected in national household consumption and expenditure surveys. *International Household Survey Network*.
- Stigler, G. J. (1954). The early history of empirical studies of consumer behavior. *Journal of Political Economy*, 62(2), 95-113.
- Stiglitz, J. E., Sen, A., and Fitoussi, J. P. (2009). *Measurement of economic performance and social progress*.
- Tourangeau, R., Groves, R. M., & Redline, C. D. (2010). Sensitive topics and reluctant respondents: Demonstrating a link between nonresponse bias and measurement error. *Public Opinion Quarterly*, 74,413–432.
- Troubat, N. and Grünberger, K. (2017). Impact of survey design in the estimation of habitual food consumption. The case of the 2007/08 Socio Economic Survey of Mongolia applied to urban households. *Food Policy*, 72(C): 132–145.
- UNECE (2009). *Making Data Meaningful, Part 1: A Guide to writing stories about numbers*. United Nations, Geneva.
- UNECE (2009). *Making Data Meaningful, Part 2: A Guide to presenting statistics*. United Nations, Geneva.
- World Bank Group. (2015). *A Measured Approach to Ending Poverty and Boosting Shared Prosperity: Concepts, Data, and the Twin Goals*. Policy Research Report. Washington, DC: World Bank.
- World Bank. (2017). *Monitoring Global Poverty: Report of the Commission on Global Poverty*. Washington, DC: World Bank.

Xu K, Ravndal F, Evans DB, and Carrin G. (2009), Assessing the reliability of household expenditure data: results of the World Health Survey. *Health Policy*;91(3):297-305.

Zheng, B. (1997). Aggregate Poverty Measures, *Journal of Economic Surveys*, 11(2): 123-62.