Advances in hunger measurement

Traditional FAO methods and recent innovations

Carlo Cafiero
ESS

1. Introduction. Assessing food insecurity: a complex task

Food security has come to be customarily defined as the situation that occurs "when all people, at all times, have physical, social and economic access to sufficient safe and nutritious food that meets their dietary needs and food preferences for an active and healthy life" (FAO, 1996). Food security is thus more than simply "freedom from hunger"; more dimensions are included to highlight that it is a condition that applies at the individual level on a continued basis, that health and nutritional aspects associated with food consumption and individual tastes and preferences are as important as the mere fulfillment of basic dietary energy needs, and that the right to food extends well beyond mere survival, being the basis for a healthy and productive life.

Recognition of the complexity of the problem presents obvious challenges for monitoring it. Despite the fact that in recent policy forums achieving food security has been described as "a measurable and monitorable goal" (FAO, 2001), and as unpalatable to many as it may be, it should be clearly recognized that, if reference is made to the broad definition involving "all people at all times" and to both "dietary needs and food preferences," *no direct measure of the state of food insecurity in the World will ever be possible*. Such a measure, in facts, would imply the ability to continuously monitor all the dimensions that constitute food security at the level of individuals in a population. It would mean, for example, to frequently record individual food consumption of single food items, convert it into nutrient intakes, and compare both quantity and quality of nutrients to supposedly known individual requirements and preferences. Such an endeavor is clearly impossible.

Real progress in our ability to inform the international community and to guide policy can be achieved if scope and limitations of any attempt at measuring hunger and food security are properly recognized. Our opinion, and the philosophy underlying this article, is that to give a sensible meaning to the task of monitoring the achievement of food security, several points must be considered. First, and foremost, as departures from the food security ideal pictured in the definition above include many different situations (inadequate dietary energy intake, inability to satisfy food preferences, uncertainty about the future ability to access food, etc.), hardly any single indicator can ever be deemed sufficient to respond to the need of adequate monitoring. More likely, a number of appropriately chosen indicators, each focusing on one key dimension of the problem, should be considered as element of a coherent suite.

Second, “to measure” in this context must be taken as to mean “to estimate” which implies that the question is to be addressed in probabilistic terms. Consequently, due attention needs to be devoted to the statistical aspects of the inference that can be drawn from available data. This is particularly problematic when monitoring is needed in real-time fashion and must be based on scattered, and often rather imprecise, data.

Many indicators have been proposed, used and sometimes unfairly criticized over the years for lack of consideration of the many statistical problems involved. In order to properly address the qualities of an indicator and of the methods used to estimate it, in fact, a clear definition of the concept that the indicator is meant to capture must be provided and understood. Failure to do so might contribute to create a gap between the statistical measurement and the public perception of the problem. That such a gap exists with reference to food security is not surprising, given the attention that the problem receives in the public. In analyzing a similar state-of-affairs with respect to the measurement of economic progress, Stiglitz, Sen and Fitoussi (2009, pp. 7-8) have pointed out that the gap may be created either because the statistical concept may be correct, but the actual measurement process is imperfect, or because there are questions on what the right concepts are, and what the appropriate use of different concepts is. Though criticisms can be raised and attention should be devoted to both aspects in discussions about measuring social phenomena, one main message in this article is that it is important to avoid confusion between two levels: the appropriateness of a given concept to capture aspects of a socially relevant phenomenon, and the possible problems in the measurement process.

This paper elaborates on some of the methodological aspects linked to the issues introduced above, with special reference to the practice that FAO has been following in monitoring the state of food insecurity in the world. The FAO indicator of the prevalence of undernourishment is described with the aim of clarifying the statistical concept that informs it and the way it is estimated. The objective is to shed light on some of the aspects that have contributed to make the debate on food security assessment over the past decade less productive of what might have been, while advancing suggestions for possible improvements in our collective ability to effectively monitor food security.

2. The FAO methodology

Since its establishment, FAO has been charged with the responsibility of monitoring the state of the world food situation to enable the international community to appropriately direct actions aimed at promoting the universal achievement of the right to adequate food. FAO's statistics division has been at the forefront of such an effort by developing methods and tools for data and information dissemination to respond to the demand for effective food security monitoring.

This section discusses the current state of the FAO's food security monitoring effort as performed through estimation of the "Prevalence of Undernourishment" (PoU) indicator, routinely published in the State of Food Insecurity (SOFI). It does so by clarifying a) which aspects of food insecurity are captured by the indicator, b) what is the statistical concept informing the methodology, and c) how the available data are used in the inferential process leading to the estimates.

2.1. The operational definition of "undernourishment" embedded in the FAO indicator.

FAO has received a mandate from the international community to monitor progress towards achievement of the objectives set by the World Food Summit and the UN Millennium Development Goal.[1] The terms "undernourishment" and "hunger", as used in describing the two targets, have been usually interpreted as referring to a situation of *continued inability to obtain enough food*, i.e., a quantity of food sufficient to conduct a healthy and active life. The meaning of terms as "undernourishment" or "hunger" is clearly narrower than that of food insecurity as implied in the definition reported above. Even once established that undernourishment is not synonym with food insecurity[2], the definition as "inability to obtain enough food" is still too vague to lead to practical monitoring. To reach a valid operational definition of undernourishment several issues need to be addressed.

First, considering the complexity of the process of human nutrition, and the fact that there are both quantity and quality dimensions associated with food, the expression '*enough food*' needs to be qualified. The FAO method has been traditionally based on the assumption that the most relevant aspect to be monitored is *dietary energy*

[1] The 1996 World Food Summit pledged to "... to eradicate hunger in all countries, with an immediate view to reducing the number of **undernourished** people to half their present level no later than 2015" (FAO, 1996) while the MDG Target 1.C is defined as to "halve, between 1990 and 2015, the proportion of people who suffer from **hunger**." (UN, 2000)

[2] In a sense, FAO's "undernourishment" can be considered as the extreme form of food insecurity, arising when even the mere caloric supply is inadequate to cover basic needs.

intake, and that 'enough' ought to be evaluated with reference to a normative benchmark described as *dietary energy requirement* as established by nutritionists. According to such a definition, a human being is considered undernourished if the level of his or her habitual dietary energy intake is below the minimum level that nutritionists would deem appropriate.

Second, the definition calls for a *continued* inability to access enough food over a certain period of time that must be defined. The question of what is the appropriate time span to assess undernourishment is not a trivial one. If our interest is towards highlighting deep, chronic undernourishment, the reference period should be long enough for the consequences of low food intake to be detrimental to health. Though temporary food shortage may be stressing, the FAO definition of the indicator is based on a year, and the relevant average consumption of food over that period is referred to as the *habitual* level.

Next, although the proper comparison between caloric intake and caloric requirement ought to be conducted, in principle, at the *individual* level, this is still deemed too difficult to be operational on a broad scale, as food access data is usually collected only at the household level. Most of the methods proposed so far for the assessment of countries' food insecurity must thus be recognized as referring to households, rather than to individuals, and this is a clear limitation in trying to assess the relevance of food and nutritional disparities due, for example, to problems with intra-household allocation of food.[3]

It should thus be clear that the FAO indicator is designed to capture a clearly (and somehow narrowly) defined concept of undernourished, namely a *state of food deprivation lasting over an extended period of time*. As such, it is certainly not sufficient to assess the overall welfare cost associated with food and nutrition problems. It does not capture, for example, costs associated with food procurement that do not result in reduced food consumption which may nevertheless have strong impacts on the quality of life of people being forced to strive to maintain adequate caloric intake levels. Equally importantly, the FAO indicator is not meant to capture short-lived effects of temporary crises, or to distinguish the roles and impacts of external causes (i.e., production or trade shocks) from those of the possible inadequacies of coping strategies (i.e., savings, changes in overall consumption patterns, food item substitution, etc.) One conclusion of all this is that, rather than aiming at substituting the PoU indicator with alternative single food security

[3] As will be made clear below, the FAO has made an effort to provide proper inference based on the individual state of undernourishment, even when lacking ideal data, through proper statistical treatment of the available ones, something that critics of the method have failed to recognize.

measures, we need to continue discuss how to broaden the set of indicators to monitor food security in its various dimensions. As will be hopefully clear after reading the following pages, the FAO indicator on the prevalence of undernourishment remains an indispensable component of any such a set of indicators.

2.2. The inferential process

As it has been abundantly documented over the years (see for example FAO 1996, Appendix 3, Naiken 2003), the FAO method is defined with reference to a probability distribution for the individuals' yearly habitual Dietary Energy Consumption in a population, x, and a threshold level, called Minimum Dietary Energy Requirement (*MDER*) relevant for the same population.

The Prevalence of Undernourishment (*PoU*) is then defined as:

$$PU \equiv \int_{x<MDER} f(x)dx \qquad (1)$$

that is, the probability that consumption falls below the threshold.

In such a framework, the distribution of yearly habitual dietary energy consumption $f(.)$ across individuals is intended to capture both the overall level and the distribution of food consumption in the population, thus capturing two of the recognized dimensions of food security, namely *availability*, through the location parameter (i.e., the *mean*), and differential *access, through* the higher moments (*dispersion, skewness* and *kurtosis*).

In evaluating the merits of this estimator, one of the most common sources of misinterpretation is the fact that the probability distribution in (1) has tended to be interpreted as the empirical distribution of the actual food consumption in the population, that is, the distribution that could be obtained, for example, through a food consumption census of the population, but such interpretation is largely misleading. Under such an interpretation, in fact, it would be very difficult for example to make sense of a unique threshold level to be applied to all individuals, as it is obvious that energy requirements vary among individuals. If reference were to be made to the empirical distribution of food consumption in the population, than the proper measure of the prevalence of undernourishment ought to be:

$$PU_2 = \iint_{(x,r)\in\{x<r\}} f(x,r)drdx \qquad (2)$$

where the possibility of a *joint distribution* of dietary energy consumption (x) and requirements (r) is explicitly recognized.[4] The attractiveness of an approach as in (2) is that it gives the impression that it may be possible to classify *individuals* in the population as being undernourished or not, based on the comparison between the individual intake x_i and the individual requirement r_i. Estimating the prevalence of undernourishment in the population would then amount to simply head-count those who are classified as undernourished.

The major obstacle to the proper application of a joint distribution framework as in (2) is that *individual dietary energy requirement is practically unobservable.* It is in fact widely recognized that individual dietary energy requirement proper depends not only on clearly identifiable individual characteristics such as body mass and level of physical activity, but also on a rather elusive individual degree of efficiency in the metabolism of food. As an important practical consequence of this fact, normative food requirement standards can only be given as *ranges valid for groups of individuals* (usually defined by age, sex, and physical activity) in recognition of the many unobservable factors affecting the individual requirement.[5] When the only information available on an individual is the combination of sex, age and level of physical activity, only a range of energy requirement levels that are compatible with good health can be given, and the FAO/WHO/UNU experts repeatedly make a point that the norm corresponding to the average of the range provided should not be used at the individual level, lacking a more comprehensive assessment of other individual characteristics.

Contrary to what seems to have been implied by some critics of the FAO methodology, the difficulty to precisely assess individual energy requirement

---

[4] In the past, the FAO approach has been described with reference to a joint distribution, most notably by Svedberg (2000), who claimed that under such an approach, the choice of a single threshold level would necessarily lead to estimation errors. As pointed out by Naiken (2007) and Cafiero and Gennari (2011), Svedberg criticism is vitiated by a fundamental misrepresentation of the FAO methodology (see also below.)

[5] "Estimates of energy requirements are derived from measurements of individuals. Measurements of a collection of individuals of the same gender and similar age, body size and physical activity are grouped together to give the average energy requirement - or recommended level of dietary intake - for a class of people or a population group. These requirements are then used to predict the requirements and recommended levels of energy intake for other individuals with similar characteristics, but on whom measurements have not been made. Although individuals in a given class have been matched for characteristics that may affect requirements, such as gender, age, body size, body composition and lifestyle, *there remain unknown factors that produce variations among individuals*. Consequently, there is a *distribution of requirements within the class or population group*." (FAO/WHO/UNU, 2002, p. 5, emphasis added).

thresholds (which led to the suggestion of monitoring anthropometrics as an indirect way to assess undernutrition) does not exclude the possibility of conducting a valid inference at the *population* level, based on a probabilistic statement and a proper understanding of the concepts involved. To facilitate such an understanding, we suggest that the distribution in (1) is interpreted as the probability distribution of the possible levels of habitual dietary energy consumptions for the population's representative individual (that is, the "average" individual in terms of all the observed and unobserved characteristics that may affect energy requirements), and the threshold level be interpreted with reference to that 'special' individual. In other words, x can be interpreted as the level of dietary energy consumption that would be observed on a randomly-selected individual in the population. The PoU is then a statement on the probability that a randomly selected individual would found to be undernourished.

Admittedly, the procedure may appear as too convoluted to readers who are not familiar with the principles of statistical inference. Unfortunately, conceptually simpler procedures (such as for example the one proposed by Smith, 2003 and Smith *et al.,* 2006) would need to confront such fundamental issues of data availability that they cannot be considered valid alternative for the same objective of assessing country level undernourishment at the global level (see further below).

Once the foundation of the method described by equation (1) is understood, the question arises on how to set the caloric threshold level to obtain the best estimate, an issue that has raised lots of controversies in the past.[6] Most of the discussion so far has been concerned with the roles of *intra-* and *inter-individual* variation in energy requirement, with unresolved issues among nutritionists and others on whether or not human beings can effectively adapt their dietary energy requirements to the environment (See for example the review of the issues in Osmani, 1992). Notice that, in view of the interpretation of the model in equation (1) as referring to the representative individual in the population, the question is of no relevance here. Irrespective of whether it is due to individual adaptation to changing conditions, or to differences in the metabolism among different individuals, the only thing that matters here is the fact that there will be not a single value, but rather a whole *range* of energy requirements that is compatible with good health and nutrition for the *representative* individual in the population.

While keeping the consideration in mind, there should be no controversy on the fact that the starting point must be found in the normative standards set by nutritionists,

[6] Osmani (1992) provides an excellent collection of articles contributing to the debate in the 1980's.

such as those produced by the joint FAO/WHO/UNU group of experts in various occasions.[7] At a superficial consideration, one may be tempted to refer to the average requirement given as "recommended level" of dietary intake for a group of individuals of a certain age and sex (see fn. 3) as the most suited threshold to determine whether an individual belonging to that group is likely to be undernourished or not. After all, lacking any other information, why not considering any individual as the "average" of the group of same sex and age in terms of caloric requirements? Unfortunately, such a simplistic reasoning is faulty and would lead to grossly mistaken inference.[8] The reason is simply that, even if a group of individuals were composed only of well-nourished people, it would be expected that roughly 50% of them have intakes below the average requirement. Using the average requirement as a threshold in any situation where less than half of the population is undernourishment would produce an overestimate (a point effectively made as early as 1960 by P.V. Sukhatme, and subsequently recognized, among others, by Srinivasan (1983)). The obvious corollary of recognizing that there is variability of intakes also among well-nourished, similar people is that the proper threshold should be set at a level below the average, so that the probability of overestimating the prevalence of undernourishment in the group is reduced.[9]

By choosing as a threshold the level that corresponds to the *minimum* of the range of dietary energy requirement indicated by nutritionists as compatible with good health and normal physical activity for that group (as FAO does in setting the MDER as the threshold), the probability of overestimating the prevalence of undernourishment in a group is minimized. Of course, as lucidly pointed out, among others, by Osmani (1999), there is still the possibility that the measure thus obtained is an *under* estimate of the actual prevalence of undernourishment, as there would be individuals who, despite consuming at a level higher than the MDER, are undernourished. In choosing the threshold level, a trade-off thus arises between the risk of overestimation and the risk of underestimation. By choosing a threshold that is higher than the MDER, the probability of underestimation of the prevalence of undernourishment is reduced, but at the cost of increasing that of overestimation.

---

[7] The latest expert consultation was held in 2001. Results are published in UNU/WHO/FAO (2004)

[8] Several researchers in the past have fallen in such a mistake, including Smith *et al.* (2006),

[9] The use of a single threshold value in estimating the prevalence of undernourishment has been criticized by many others (e.g., Anand and Harris, 1992)on the account that it may never lead to correct identification of all the undernourished in a population. It is worth to remind the reader that the objective of the FAO methodology discussed here is not, and has never been, that of identifying individuals as being undernourished, but only to estimate the likely prevalence in the population.

How to choose the most appropriate threshold depends therefore on the cost that is associated with the two estimation errors, which is likely to be a function of the expected *size* of estimation errors. It is conceivable that over- or underestimating by a "little" is less costly than doing it by a "lot". In judging what the proper threshold should be, therefore, not only the mere possibility of underestimating or of overestimating should be considered, but also by how much.

Seen from a statistical inferential perspective, this is the aspect that has probably been lying underneath the debate that, over the years, has surrounded the FAO estimator. One aspect that has emerged clearly in the debate is that the extent of possible errors in setting a given threshold crucially depends on the supposed relationship between the (observable) level of intake and the (unobservable) level of requirements. The conclusion Osmani draws from the review of the elements in the debate is that: "a value-neutral way of tackling [the problem of estimating the prevalence of undernourishment] would be to shun the 'cut-off approach' altogether and use instead a joint distribution of intake and requirement" (Osmani, 1999, p. 142).

There are two main issues that can be raised with such a conclusion and that, if properly recognized, could contribute to move forward in the debate and in making proper use of the evidence provided by an indicator such as the PoU. The first one is to realize that the existence of a pure "value-neutral" way of tackling any estimation problem is essentially an illusion created by what could be termed an "idealistic" view of statistical inference. As best put by the late Arnold Zellner in his pioneering and valuable contribution to the theory and practice of statistical inference, no estimator can ever be deemed void of value judgment. Whatever estimator can be conceived, the most that can be hoped is that it is proven optimal *under a specific loss function*. The value judgment that is supposedly avoided by claiming the "optimality" of the estimator resurfaces, if only implicitly, in the choice of the loss function.

With reference to the problems related to the use of a single cut-off point in estimating undernourishment, use of the MDER as the threshold level can be then easily justified by a consideration that the cost of overestimating undernourishment vastly dominates that of underestimating it, something that is in line with the definition of "undernourishment" as the *extreme* form of food insecurity. This is not to deny the possibility of a risk of underestimation. However, the possible *extent* of such underestimation has been deemed small by the proponents of the FAO methods, based on the presumption that the probability of undernourishment conditional on intakes being larger than MDER should become very small as soon as intakes become larger than the MDER (see FAO, 1996, Appendix 3, Figure 1). If that is true, a threshold at, or very close to, the MDER is an appropriate one, unless one associates a disproportionately higher cost to the risk of underestimating compared with that of overestimating it. Nothing would prevent that, other considerations deemed relevant or other perspectives on the food deprivation problem being considered, a different (higher) threshold may be used, resulting in an estimator that would implicitly give more weight to the cost of underestimating it.[10]

The second problem with the conclusion drawn by Osmani is the idea that use of an approach based on the joint distribution of intake and requirement might have paved the way towards a better estimator. As elaborated by Naiken (2007), and hinted to in Cafiero and Gennari (2011), the possibility of defining a joint probability distribution for dietary energy intake and requirement is highly problematic, to say the least. It has to confront the very definition of the concept of "adequate nourishment" from a statistical point of view. If we define the condition ($x < r$) as "undernourishment" and, by symmetry, ($x > r$) as "overnourishment", it is clear that the event ($x = r$) should indicate adequate nourishment. Obvious considerations suggest that any credible continuous joint density function would assign zero probability to such an event, and therefore should not be used. One possibility is to define adequate nourishment as an approximation, ($x \approx r$) by recognizing that small differences between x and r may exist due to measurement errors and therefore, even though ($x = r$) gets assigned zero probability, a joint distribution could be defined that assigns the proper mass to the event "adequate nourishment." However, it should be self-evident that, even if such a distribution could be conceived, it is very unlikely that it will have elliptical iso-density contours and therefore it ought to be fundamentally different from joint normal or similar distributions.[11]

Ruling out the possibility of using a joint distribution framework (see Naiken, 2007 for an extended treatment), the question seems to revolve around the problem of whether or not it is possible at all to conduct meaningful inference based only on the (observed) variability in intakes when this in part reflects (unobservable) variability in requirements. It should be clear from what has been said thus far that the FAO method provides indeed one such possibility, which is fully consistent with the principles of statistical inferences, even though the "optimality" of the proposed estimation method stands on an unproven – though perfectly reasonable - assumption on the probability of being undernourished conditional on various levels of observed intake.

---

[10] See also below for further discussion of the possibility of defining alternative estimators by the choice of higher cut-off points.

[11] A double normal distribution, with various degrees of correlations, has been used by Svedberg (2000, ch. 9) in his criticism of the FAO methodology. In light of the discussion being conducted here, the whole argument raised by Svedberg appears virtually irrelevant.

There is however another argument that can be made to demonstrate that inference on undernourishment can properly be drawn on observing just the distribution of food consumption, and that appears to have eluded the attention of researchers so far. Any population for which the issue of measuring undernourishment is meaningful, should be properly viewed as composed of three different groups: that of the undernourished individuals, that of the adequately nourished, and that made of people with *excess* food consumption. Sampling from such a population, therefore, would correspond to what has been termed in the statistics literature as a problem of conducting inference on a finite *mixture* of distributions (Everitt, 1985), and there is ample literature on how to conduct such inference.[12] An interesting avenue for further research on how to improve the practice of estimating the extent of undernourishment (which would bring the important addition of being able to estimate at once also the extent of *over*-nourishment) resides in the application of statistical procedures developed to deal with mixture distributions on datasets of individuals' or households' food consumption obtained from representative samples of the population. The aim would be that of classifying each observed case in one of the three underlying sub-populations, thus effectively providing an estimate of the relative size of the groups in the overall population.

2.3. Are there viable alternatives?

Admittedly, the need for a relatively complicated estimation procedure such as the one just described is not obvious, and many commentators have been attracted by a simpler estimation procedure based on direct use of survey data, as described for example in Smith and Subandoro (2005)[13] and usually referred to as the non-parametric approach.

The idea behind the method is rather simple. Inference is drawn by measuring the proportion of households in a representative sample that are classified as food insecure based on the comparison of an estimated level of habitual food consumption in the household and a household specific food requirement threshold. If the estimated total food consumption in the household is found to be lower than

[12] Principles of how to conduct inference from finite mixture distributions are given in detail by Titterington, Smith and Makov (1985), using maximum likelihood principles and the E-M algorithm, and by Diebolt and Robert (1994), using Bayesian procedures based on Markov chain Monte Carlo (MCMC) methods.

[13] This alternative procedure is reminiscent of the one typically used for poverty assessments based on data from income and expenditure surveys, pioneered at the poverty group of the World Bank, where the incidence of poverty is estimated by counting the proportion of households whose average income is estimated to be below a reference value (typically 1.25 USD/day or 2.0 USD/day).

its total food requirements, the household is classified as undernourished. The prevalence of undernourishment in the sample is then taken an estimate of the one in the population.

Attractiveness of the non parametric approach resides in the fact that, if it were indeed possible to correctly classify each single household in the sample, it could be possible to make fuller use of the rich information contained in the data that are usually collected through household surveys. This would provide important elements to answer the question of "who the food insecure are", rather than just estimating their number in the population. In practice, however, the possibility of correctly classifying sampled households has to confront two problems related to the appropriateness of the data collected for estimating the two needed variables, namely the habitual food consumption and the dietary energy requirements at the individual household level.

We already discussed why the dietary requirement norms given by nutrition experts should only be safely applied to groups, and not to individuals. Lacking information on the individuals' actual body masses and levels of physical activity, in fact, there is a high probability of mismatching individual to their requirements, and thus wrongly classifying "tiny" persons as undernourished while failing to recognize "big" undernourished ones. In order to reduce that risk, and the consequent error in classifying households, it seems that anthropometrics (better if height and weight separately) and occupation status of all members of the household should be collected, together with food consumption, if a survey is to be used within a non-parametric approach.

As for food access data, based on the household surveys' data available to FAO at that time, Naiken (2001, p.11) noted that "the household per capita dietary energy consumption figures derived from the food consumption data collected at the household level are imprecise and, in many cases, found to be unrealistically high or low." Few years later Smith *et al.* (2006), in preparing a report on Sub Saharan Africa, had to exclude 50 out of the 76 available surveys because they did not satisfy a minimal set of requirements needed to be able to produce an estimate. Even for the surveys that were used to conduct the assessment on the remaining twelve countries, the authors recognized that household level food consumption data referred to food *acquisition* rather than food *consumption*, and that large differences could be expected between the two. Unfortunately, even if the estimated mean household calorie consumption in the sample might have been unbiased because "households in a large population group are equally likely to be drawing down on food stocks as they are to be accumulating them" (Smith *et al.*, 2006, p.), each *individual household* value would likely be biased whenever household level storage of food is relevant. This shall have two serious consequences on the reliability of the results of a household food security analysis conducted without attempting to control for the difference between acquisition and consumption. First, the number of households classified as undernourished will be larger than it should, as households who are reporting acquisitions lower than actual consumption (i.e., those drawing down from own stocks), will likely be found with higher frequency at the lower end of the sample distribution, whereas those reporting large acquisition will be on the upper tail. The point is that, conditional on being undernourished, the probability of having acquisition larger than consumption in any given period is likely to be much lower than the opposite (for example, because the very poor could simply not afford to acquire more than they consume.) Therefore, mistakenly referring to acquisition instead of consumption in a head count approach is likely to yield *ceteris paribus*, an overestimate of the prevalence of undernourishment.

Second, and perhaps more worrying, if the individual status of being undernourished is going to be used to conduct disaggregated analysis by subpopulation groups (the possibility of which constitutes one of the major attractiveness of household survey data), the risk exists that the analysis would yield inconsistent results if the difference between acquisition and consumption happens to be correlated to the grouping variable.[14]

To conclude on the non-parametric approach, we can say that in principle, if a specialized survey could be designed to capture precisely enough the variables of interest, there is no doubt it would be the natural choice. A minimal set of requirements for such a survey should include features that would allow:

a) to discern actual food consumption of the surveyed household from food acquisition over the surveyed period, recognizing that the latter may include food that is acquired for other uses (storage, food given to guests, etc.);
b) to have the means to verify the presence and, in case, to control for seasonal variation in food consumption;
c) to precisely assess the household's member dietary requirements, which would call for data on height and physical activity level, in addition to sex and age.

If those features are included, the quality and reliability of the estimation through a head count of the households in the sample would depend mostly on the size and

[14] As an example, consider the case in which most households build up food stocks in the period after the harvest. If that is not being taken into account when defining the sampling plan, it may happen that one area of the country is surveyed in that period, with the result that there will be a bias in the data correlated with the geographic location of the household.

representativeness of the latter, with no need for a parametric model for the distribution of food access in the population. If instead, as it is often the case, surveys do not satisfy the minimal requirements above, survey data are still best used within a model based approach, where the parametric model is a guarantee against the excess, uncontrollable variability present in individual household data.

3. From theory to implementation: modeling assumptions and data.

The discussion so far has been concerned with the theoretical foundation and the methodological principles underlying the FAO method to assess the prevalence of undernourishment. In this section, the question of how even a perfectly sound statistical inferential model is translated in practice is discussed, by touching on the fundamental issues of what 'ancillary' assumptions are needed and on the important issue of the availability and quality of the data used to inform the estimate.

3.1.1. Specific ancillary assumptions

Implementing the procedure for estimating undernourishment as synthesized in equation (1) above requires a set of *ancillary assumptions* that will give a concrete practical meaning to the symbols x, $f(\cdot)$ and MDER. The current practice at FAO is based on a set of operationalizing assumptions regarding the operational definition of dietary energy *consumption*, x, the choice of the parametric model for $f(\cdot)$, and the way to estimate the MDER, which will be briefly discussed in turn.[15]

Individual dietary energy consumption, x

Though the proper comparison in terms of human energy balance ought to be conducted between energy intakes and requirements the FAO method is defined in terms of the *quantities available for consumption at the household level*, that is, inclusive of possible household-level food waste. The variable x in equation (1) thus refers to the amount of dietary energy contained in the food that reaches the household, not the one effectively *ingested*. While an effort could be made in determining the level of household level waste, there are two considerations that may question the desirability of referring to actual food intake rather than food consumption in assessing food insecurity. The first is that, especially among food insecure households, for which food is presumably very precious, such wastes are conceivably very limited. Second, even admitting that household level waste may be significant, one has to ponder whether it would be morally proper to classify as food

---

[15] A thorough description of how the FAO methodology for assessing undernourishment at national level has been implemented so far is presented in Naiken (2003).

insecure households that waste food to which they have gained access, as it may happen by applying the FAO method to the distribution of food intake. In practice, this contributes to make the FAO a *conservative estimate* of extreme food insecurity.

Notice that the point raised here does not concern the food that is possibly wasted *before* it reaches the households, which should be properly accounted when measuring x, a point that raises some concerns with respect to the practice of referring to the per capita dietary energy supply obtained from aggregate food balance sheets as an estimate of the per capita dietary energy consumption at the household level. Recent evidence (Gustavsson *et al.*, 2011) points to the fact that food wasted during distribution at the retail level may amount to up to five percent of the quantities of grains available at the retail level, with even larger percentages for perishable products such as fresh fruits and vegetables.

Also related to the definition of dietary energy consumption is the assumption on how food is distributed among households' members. Even though the method is devised in terms of the representative *individual,* not the household, in the population, extensive data on food consumption is only available at the household level. In the FAO estimation practice, food available at the household level is simply divided by the number of individuals in the household. While this does not induce a bias in the estimated average consumption[16], it means that implications of possible unequal distribution of food within the household are neglected. Unfortunately, lacking reliable data on *individual* consumption level, there is no alternative at the moment. It is hoped that, in the future, more dataset would be made available containing truly individual level consumption data. If successful, this would set the possibility also for conducting proper analyses, for example, of possible gender and age discrimination in access to food.

The functional form of $f(\cdot)$

One other operational assumption needed to give practical content to the method regards the form of the distribution of x in the population. In principle, if reliable sufficient data on habitual food consumption of individuals in a population were available, one could avoid choosing a parametric distribution model, and rely on non-parametric representations of the empirical distribution of the available data, for example through kernel approximations to the frequency histogram.

[16] Recall that, as what is aimed at is a measure of the food consumption of the representative individual in the population, averaging across household members does not induce a bias, to the extent that the surveyed households are representative of the population in terms of sex-age composition

Unfortunately, such data are inexistent. Nationally representative surveys that provide information on food consumption are not usually designed to collect the average individual food consumption over the year. The data provided – on households' total expenditure on, or total acquisition of, food – must be heavily processed in order to infer on the needed information.[17] Besides the fact that the data collected refer to the whole household, and not to individuals, there are three major problems with existing surveys that suggest not relying on the empirical distribution of recorded household data as an approximation to the representative individual's distribution of habitual food consumption.

First, surveys often do not report the actual quantities of food being consumed over the reference period. Quantities may need to be estimated by converting monetary values on expenditures into quantities, which then are converted into calorie equivalents. Second, even when quantities are reported, to increase reliability of the data these are recorded with reference to a shorter period (usually a week or a fortnight) than the one to which the assessment of the undernourishment condition refers to (usually one year). Third, the quantities or expenditures being reported often refer to food acquired during the reference period but not necessarily consumed. All of these factors make the estimate of individual food consumption from the data collected on the single household very shaky, and suggests that it is wise to use the data on household level food consumption from surveys in an indirect way, as the FAO does, to estimate parameters of an assumed distribution of individual habitual food consumption in the population. This will avoid biases – especially in estimation of the coefficient of variation – induced by spurious variability due to seasonality and to the difference between household level acquisition and individual consumption.

While it is hoped that in the near future more nationally representative household surveys will explicitly collect average quantities of food consumed over the year by the individual members in a household, use of a parametric distribution seems to be the best way to protect against mistaken inference. During the preparatory work for the 1996 World Food Survey, the Log Normal model[18] was adopted and it has

---

[17] The quality of the household level information pertaining to food consumption contained in household income or expenditure surveys may be so low that they are of no use for food security assessments. As reported by the authors, in their analysis of food security in Sub-Saharan countries Smith *et al.* (2006) were able to use only 13 out of more than 100 available household surveys, having to disregard the others for various problems related with data quality.

[18] A random variable x is said to follow a Log Normal distribution with parameters μ and σ if the variable $y = log(x)$ is normally distributed with mean μ and standard deviation σ.

maintained since. The distribution was chosen due to some desirable characteristics[19] and to the parsimony in the number of parameters (only two) needed for its characterization. At the time it was chosen, the model was tested against two possible alternatives (the Normal and the Beta distribution) showing a markedly superior fit of the data from a limited number of available *individual* food consumption surveys.

The major drawback of the Log Normal model is that it may not be flexible enough to capture possible changes in the degree of skewness associated with food consumption. As mean food consumption in a population increases, in fact, it is conceivable that levels of food consumption below the mean would increase proportionally more than do levels of consumption above the mean. Such changes may make the distribution less positively skewed than what is imposed by the Log Normal model.[20] While the matter of whether or not a Log Normal distribution still provides a good representation of the distribution of habitual food consumption in a population remains an empirical one, some reservation may be legitimate after more than fifteen years since it was introduced, especially for cases where the mean of the distribution has increased considerably as demonstrated by larger values of DES. In those cases, a distribution with zero or even negative skewness might be more appropriate, considering that – contrary, for example, to what happens to income or to the value of total consumption – caloric consumption cannot increase indefinitely.

The caloric threshold

As discussed in section 2 above, the threshold used to estimate the proportion of undernourished based on the distribution of caloric consumption has been one of the most controversial issues in discussions of the FAO method to assess undernourishment, being linked to the unresolved question of whether or not human organisms are able to adapt their food requirements to varying intake levels. When properly interpreted as the minimum dietary energy requirement compatible with good health for the population's representative individual, the point related to whether variability of requirements in a group of individual of the same sex, age and physical activity is due to intra- or to inter- individual variability becomes irrelevant. In any case, it must be undisputed that there will not a single, but a whole range of values of dietary energy requirements that are consistent with adequate

[19] Mainly the fact that it is assigns positive probability to positive values only, and that it is positively skewed.

[20] The degree of skewness of a lognormal distribution cannot vary independently from the standard deviation and can never be negative, as the coefficient of skewness is a positive monotonic function of the standard deviation, σ.

nourishment, and that the minimum of such a range should be used as the proper threshold for estimating the PoU.

As already hinted to above, in the FAO method the cutoff point is derived by aggregating sex-age-specific minimum energy requirements norms as defined by nutritionists, using the proportion of the population in the different sex-age groups as weights. The sex-age-specific *minimum* energy requirements for adults and adolescents are based on the energy expenditure corresponding to the lower limit of the range of acceptable body weight for a given height and the lowest acceptable PAL index. The lowest acceptable body weight for a given height has been estimated as the one corresponding to the fifth percentile of the distribution of body mass indexes (BMI) (WHO, 1995), and a PAL index corresponding to light activity (1.55 for males and 1.56 for females) has been taken to reflect the lowest acceptable activity level.

Use of the PAL index corresponding to light physical activity in defining the threshold, should be reiterated, neither implies nor should be meant as to suggest for example, that people working hard are not going to be recognized as undernourished when their intake level is not sufficient to cover the higher requirement associated with their physical activity. It is simply the analytical device needed to avoid overestimation of undernourishment, as the extreme form of food insecurity, when referring to the distribution of observed food consumption whose variation reflects, in part, the co-movement of food consumption and requirements in the component of the population who is adequately nourished.

3.1.2. Data problems

The discussion thus far should have helped clarifying that the FAO method is founded on sound theoretical bases and that the needed ancillary assumptions have been always suggested by pragmatic considerations with careful attention to avoid inducing any bias. Nevertheless, all this is not sufficient to ensure high reliability of the estimate. A crucial role, in this sense, is played by the quality and reliability of the data used to inform the estimate of the parameters of the distribution and of the cut-off point.

The practice at FAO thus far has been to base the estimates on three major sources of data:[21]

[21] It may perhaps help to remind here that the FAO has never been directly involved in the collection of data used to inform the production of the indicator. The major endeavor of the FAO statistics division

1. Demographic data on population size and composition, as provided by official sources (UN Population Division, Demographic and Health Surveys – DHS, etc.)
2. Official data on production and trade as well as on estimates of food and non-food uses of the major commodities, used to assess the aggregate availability of food at country level, as reported by the FAO's Food Balance Sheets (FBS) (http://faostat.fao.org/site/368/default.aspx) compiled for a large number of countries in the world.
3. Data from nationally representative household income and expenditure surveys (HIES) containing information on expenditure on food.

The demographic data are used both to calculate the food availability per capita and the MDER. Estimates of the sex-age structure of the populations of all countries in the world are released every other year by the UN Population Division, and promptly used to revise, when needed, the estimates of the PoU. As it is used to assess the average per capita level of food availability, and given the sensitivity of the estimated PoU to such value, the estimated total population in a country is a crucial variable in the entire procedure. FAO refers to official figures as published by the UN Population Division, though it has not been unusual that those figures diverged from the official population numbers reported by countries, thus giving rise to controversies.[22]

FAO's Food Balance Sheets are used as a source of information on the total availability of food in a country in a given year. In principle, they should include all food sources: produced, imported or otherwise made available (i.e., through food aid). The total supply of each food source is converted into dietary energy, and an estimate of the per capita Dietary Energy Supply (DES) is obtained by dividing it through the country's population size (FAO, 2001).

DES numbers obtained from FBS have been typically used by FAO as the preferred source for estimating the average dietary energy consumption. Advantages of this choice are that the FBS are routinely produced by FAO for a large number of countries using a common methodology. Problems are associated with the reliability of the underlying official data on production, trade and the extent of various non-food uses. While concerns can be (and have been) raised on the reliability of the various elements of the FBS, the resulting estimate of the mean total dietary energy supply is arguably going to be more precise than each of the individual components, owing to the fact that errors in the various different components may cancel out with

---

has been that of collecting, validating and publishing data provided directly by countries, assisting them in developing their capacity in production and analysis of food and agricultural data.

[22] One of the reasons why the UN assessment of *de facto* population living in a country may differ from official records is the presence of temporary or illegal migrants.

aggregation. One issue that raises particular concerns pertains to storage. Storable commodities such as grains, in fact, contribute towards a large share of total food availability in many countries. Failing to accurately capture inventory changes will undoubtedly affect the precision of DES in any single year. This is the fundamental reason why official estimates of the prevalence of undernourishment have been published in SOFI only as three year average, on the account that errors in the measurement of stock variation would be eliminated through averaging.

In principle, data from Household Income and Expenditure Surveys (HIES) or other surveys could also be used to estimate the mean per capita caloric consumption, as it has been suggested especially by those who are critical of the reliability of FBS data. Some have even suggested that FBS data should be abandoned altogether, and estimate of undernourishment should be fully based on data from household surveys.

While it is obvious that household surveys are a precious source of data, there are several problems with complete reliance on surveys to estimate undernourishment. First, even when there are doubts on the quality of FBS, one should have elements to believe that for the same countries, the quality of the data collected through surveys is higher than that of data on aggregate food production, trade and uses. Second, the concrete possibility exists that major components of households' food consumption could be missed or wrongly captured in the surveys. A known problem is that of food consumed away from home, for which usually only expenditures are collected, with no indication both of the kind of food being acquired and of whether the expenditure also includes payment for services associated with food consumption. It is very difficult in those cases to assess the caloric consumption that can be associated with food consumed away from home.[23]

The truth is that, in general, sizeable discrepancies have been known to exist between the per capita availability of food as estimated from FBS and from HIES on the same country in the same year, but that is not a sufficient reason to believe that the average food consumption obtained from survey data is a better estimate. While problems can certainly be found in the approximations and assumptions that are required in the compilation of FBS, there are several reasons why the two measures may be expected to be different, including the fact that the reference population to which the two measures apply is different. While HIES record the food available to the population residing within households, the FBS measures food available to the broader population, including citizen residing in public facilities such as hospitals,

[23] For example, analyzing a set of Indian household consumption surveys, Smith (2011) reveals a systematic negative bias in mean household food consumption due to the way in which data on food consumed away from home is collected.

military compounds and jails, unregistered migrants and tourists. On the other hand, food supply as recorded in the FBS may overestimate the quantity available for household consumption when food losses at the level of retail distribution are not negligible, as recently suggested (Gustavsson et al. 2011), and discrepancies between average food supply from the two sources may help some light on the possible extent of such losses.

Nationally representative HIES are an indispensable source of data for another reason, namely that they are virtually *the only* source of available information to assess the other parameter of the distribution of food consumption, that is, the coefficient of variation.[24] Even for that purpose, however, the type and quality of the data collected is far from ideal. To assess the caloric content of the food consumed, for example, one needs to identify as precisely as possible the actual food items being consumed. Most surveys instead still only record acquisition of broadly defined categories of food, thus making the conversion into calories prone to error. Moreover, the collected data refers to food acquired over short periods of time (usually one week) and includes that being served to guest. Lacking information on the actual number of partakers to the acquired food, it is difficult to precisely estimate the individual habitual food consumption over a year. All these problems prevent use of the collected data for direct estimate of the parameters of the distribution of habitual food consumption through the corresponding sample values. Short term variability (due for example to seasonality) in addition to the differences between acquisition and consumption, makes the variability present in the survey data excessive. To control for excess variability when estimating the CV, the FAO has devised an indirect estimation method, based on tabulation of the food consumption data by classes of household income, and intended to control for excess variability. The procedure consists in calculating the variability in caloric consumption *between* income classes, and eliminating all variability that is observed within each class, thus estimating what has been indicated as the coefficient of variation "due to income", $CV(x|v)$.[25] By doing so, however, all variation in food consumption that is uncorrelated with income levels is unaccounted for, including the variability that is due to sex, age and body sizes which, instead, should be considered. For that reason,

[24] For countries for which no household surveys were available, FAO had to devise indirect means to estimate the parameters of the distribution, by referring, for example, to tabulations of income distribution and to other indicative data on food insecurity. These indirect methods are obviously fated to generate more unreliable estimates of the distribution of food intake and will be discontinued as soon as suitable household survey data will be available.

[25] Choice of the terminology may be considered unfortunate has it has created the impression that a *causal* relationship between income and food consumption was assumed, which is not the case. Classification by income, and then taking the

an additional component due to other factors, $CV(x|r)$, is added to the estimated total CV, calculated as $CV(x) = \sqrt{CV(x|v)^2 + CV(x|r)^2}$. (See Naiken, 2003, for further details).

4. The way forward

Almost since it has been introduced, the FAO method for monitoring undernourished has been subject of critical comments, mostly originated by the recognized inadequacy of the indicator in capturing various aspects of the complex phenomenon broadly indicated as food insecurity. One of the most critical aspects has been related to the timeliness of its reporting. Even though published every year, the PoU indicator can only be properly calculated with reference to the years for which the latest data on food supply and consumption is available, which has usually meant two to three years earlier than the publication date. To try and overcome this problem, preliminary estimates have been produced by FAO for the most recent years, based on forecasting models that, instead of actual data, used projections of food supply and demand, similar to those used by USDA in their long term projections of food insecurity.[26] These projections are obviously much less reliable than the actual estimates based on data, and have proven particularly problematic when the uncertainty surrounding the international food economy has been large, as it has been the case following the 2007/08 food price crisis.

The need to thoroughly review the methodology and to explore ways to improve on the FAO ability to timely monitoring the State of Food Insecurity emerged clearly during the 36th session of the Committee on World Food Security (CFS) and culminated in a Round Table to review the methods for estimating hunger organized in Rome in September 2011.

The review process of the methodology for estimating undernourishment following the recommendations of the CFS has confirmed the solidity of the general foundation of the method, as discussed in the previous sections of this article, but has also revealed two of the most relevant critical aspects in the overall endeavor: the importance of the quality of the underlying data, and the need to integrate the indicator of PoU with other measures, to capture the various dimensions of food insecurity that cannot be measured simply through the level of caloric consumption.

Actions to improve the quality of the agricultural data produced by developing countries have been initiated, most notably with launching of the Global Strategy to

---

[26] See http://www.ers.usda.gov/topics/international-markets-trade/global-food-security.aspx

improve Agricultural and Rural Statistics, a joint initiative of FAO and the World Bank (see http://www.fao.org/docrep/015/am082e/am082e00.pdf), which should bring about significant improvements in the medium term.

The review process conducted also revealed the possibility for improvements that could be introduced in the FAO methodology in the shorter term, and that will be discussed below.

4.1. Recent improvements in the implementation of the FAO method

One of the most critical aspects in the implementation of the PoU estimator has been the lack of recent reliable survey data on which to update the estimates of the CV which, for almost all countries, have been kept constant for many years. During the same period, the mean caloric availability has been updated (often increased) due to the annual releases of the FBS.

Increasing the mean, while keeping the CV constant, however, has the consequence of increasing the skewness of the distribution if the Log Normal assumption is maintained, as the model is not flexible enough to reflect changes in the *skewness* and possibly *kurtosis* that may have resulted from changes in food availability and access occurring over the years.

To address this issue, two initiatives have been taken by the Statistics Division at FAO. First, a significant number of new household surveys have been processed, to the aim revising the estimate of the CV and to verify whether there is evidence of changing skewness in the distribution of food consumption. At the same time, new, more flexible functional forms (such as the three-parameter Skew-Normal and Skew-Log Normal, or the four-parameter Skew –T models) are currently being considered as a viable alternative to the previous Log Normal.

Efforts are also underway to improve estimates of the various components of the FBS, including levels and changes in food stocks and better assessment of non food uses.

The combined result of these initiatives will result in improved estimates of undernourishment for many countries in the World, already in the next issue of the SOFI publication.

4.2. Other indicators that may be produced given the already available data

As noted in the first section, an indicator of chronic hunger as the prevalence of undernourished in a population is by no means sufficient to provide a comprehensive picture of the many dimensions related to lack of adequate nutrition, both in terms of the causes and of the consequences.

In the following we put up for debate the proposal for three additional indicators, which could be easily produced given the available information, and which would be fully consistent with the theoretical underpinnings of the current methodology. While this list does not exhaust the set of needed additional indicators, it is certainly a starting point for a constructive debate.

4.2.1. An indicator of the prevalence of continued over-nourishment

As the methodology already provides for estimation of the distribution of DEC and (albeit indirectly) of the range of normal DER in the population, it would be natural to also measure the proportion of the population which is over nourished, an important indicator to dig into what has been termed the "double burden" of malnutrition. To do so, the area below the marginal distribution of DEC and above a threshold equals to the estimated maximum of the range of the average individual's requirement can be easily calculated. The increased flexibility of the new functional form for the distribution of DEC should ensure reliable estimates for the countries for which recent surveys are available.

4.2.2. An additional indicator of food deprivation

The cut-off point currently used to evaluate chronic hunger is defined, as we saw, by referencing to the physical activity level associated with a sedentary lifestyle. Consideration of a higher PAL coefficient, corresponding to a moderate activity, would lead to a higher threshold, say MDER* > MDER, that would allow estimation of the prevalence of "*economically significant*" hunger.

4.2.3. A measure of the depth of food deficit

The average of the individual's dietary energy requirement, ADER, is a proper *normative* reference for adequate nutrition in the population. While it would be mistaken to take the value ADER as the cutoff point to determine the prevalence of undernourishment (as some of the critics have suggested), its value could be used to calculate the ***depth of the food deficit*** (FD), that is the amount of dietary energy that would be needed to ensure that, if properly distributed, hunger would be eliminated. Such an index could be calculated as:

$$FD \equiv \int_0^{L_r} (ADER - x) f_x(x) dx \qquad (3)$$

5. Conclusions

The discussion presented in this article points to the conclusion that, to provide the international community with a comparable, worldwide, periodic assessment of the likely number of people suffering from deep food deprivation, a model based estimation procedure is still needed. At the moment, there is likely no viable alternative that is superior, or even just equivalent to the FAO method for estimating the prevalence of undernourishment.

If we ever reach the point that nationally representative surveys collecting reliable data on habitual food consumption are conducted every year and processed in a timely and consistent manner throughout the World, then a simpler head-count method based on classification of individuals could be used.

Several considerations, however, suggest that such type of food consumption surveys will never become routine, and therefore there is no alternative to the proper use of a model based estimation procedure, soundly rooted in statistical inference principles, that uses all the information that is reasonable to be expected being collected in all countries, including that from continuously improved household consumption surveys.

In any case, the information on food availability and access, as condensed in the FAO PoU estimator, while still fundamental to monitor the extreme form of food deprivation throughout the world, is going to be increasingly insufficient to provide the needed guidance to policy actions, as there are other dimensions of food and nutrition insecurity that are emerging as relevant in the contemporary debate on the food problem. In particular, there may be important welfare costs associated with increased food prices that do not get reflected in reduced caloric consumption, as households are forced to sacrifice other consumptions and needed investments, simply to maintain their food consumption at acceptable levels.

To cover the whole spectrum of biological, economic, and psychological stress that is associated with food insecurity, more research and different kind of data is needed to develop a broader set of indicators. One very promising avenue in this respect is represented by the growing number of applications of measures along *scales of food insecurity* obtained thorough surveys of peoples experiences, through questionnaires that ask people to report on their worries, behaviors and adopted coping strategy when facing food insecurity. Surveys of the type currently used in the United States to assess national food insecurity, and that have been adapted to various other context, such as for example with the Escala Latinoamericana Y Caribena de Seguridad Alimentaria (ELCSA), could relatively easily be scaled up to provide timey monitoring worldwide of a broader concept of food insecurity than caloric deprivation, thus responding to a pressing demand from the international community.

References

Anand, Sudhir and Christopher J. Harris (1992) "Issues in the Measurement of Undernutrition". Chapter 7 in Osmani (ed.) cit.

Cafiero, Carlo and Pietro Gennari (2011) "The FAO indicator of the prevalence of undernourishment". Presented at the Workshop "Measuring Food Insecurity and Assessing the Sustainability of Global Food Systems", The National Academies of Science, Washington, DC, February 16-17, 2011. Available at: http://www.fao.org/fileadmin/templates/ess/ess_test_folder/Workshops_Events/Food_Security_for_All_FEB2011/Background_paper.pdf http://www.fao.org/fileadmin/templates/ess/ess_test_folder/Workshops_Events/Food_Security_for_All_FEB2011/Presentation_SFA_2.pdf

Diebolt, J. and Robert, C. P. (1994). "Estimation of finite mixture distributions through Bayesian sampling." *Journal of the Royal Statistical Society*, Series B 56, 363–375.

Everitt, B. S. (1985). "Mixture distributions." In *Encyclopedia of Statistical Sciences*, S. Kotz, N. L. Johnson, and C. B. Read, eds. Wiley, New York, Vol. 5, pp. 559–569.

FAO (1996) *The Sixth World Food Survey*. Rome, FAO

FAO (2001) "The World Food Summit Goal and the Millennium Development Goals". Committee on World Food Security, 27th session, Rome, 28 May - 1 June 2001. (Full text available at: http://www.fao.org/DOCREP/MEETING/003/Y0688E.HTM)

Gustavsson, Jenny, Christel Cederberg, Ulf Sonesson, Robert van Otterdijk and Alexandre Meybeck (2011) "Global food losses and food waste. Extent, causes and prevention." Study conducted for the International Congress SAVE FOOD! at Interpack 2011,Düsseldorf, Germany. Rome:FAO

Naiken, Loganaden (2003) "FAO methodology for estimating the prevalence of undernourishment". In *Measurement and Assessment of Food Deprivation and Undernutrition*. Proceedings of the International Scientific Symposium convened by the Agriculture and Economic Development Analysis Division in Rome, 26-28 June 2002. (HTML version available at: http://www.fao.org/docrep/005/Y4249E/y4249e00.htm)

Naiken, Loganaden (2007) “The Probability Distribution Framework for Estimating the Prevalence of Undernourishment: Exploding the Myth of the Bivariate Distribution.” FAO Statistics Division Working Paper Series No ESS/ESSG/009e FAO Rome, August. (Full text available at: http://www.fao.org/fileadmin/templates/ess/documents/food_security_statistics/working_paper_series/WP009e.pdf)

Osmani, Siddiqur R. (1992) *Nutrition and Poverty.* UNU-WIDER Studies in Development Economics. Clarendon Press, Oxford.

Ravaillon, Martin (2011) “On Multidimensional Indices of Poverty.” Policy Research Working Paper N° 5580, The World Bank, Washington, DC, February 2011. (Full text available at: http://www-wds.worldbank.org/external/default/WDSContentServer/WDSP/IB/2011/02/28/000158349_20110228142416/Rendered/PDF/WPS5580.pdf

Sibrian, Ricardo , Loganaden Naiken and Jorge Mernies (2007) “The FAO Parametric Versus the IFPRI Non-Parametric Approach in Estimating the Prevalence of Undernourishment: Issues Relating to the Use of Household Level Data from National Household Surveys.” FAO Statistics Division Working Paper Series No ESS/ESSG/12e FAO Rome, August. (Full text available at: http://www.fao.org/fileadmin/templates/ess/documents/food_security_statistics/working_paper_series/WP012e.pdf)

Smith, Lisa C. (2011) “The great Indian calorie debate: An investigation of divergent trends in poverty and undernourishment during India’s rapid economic growth.” Unpublished Report presented to the Food and Agricultural Organization of the United Nations.

Smith, Lisa C. and Ali Subandoro (2005) “Measuring food security using household expenditure surveys. Food Security in Practice series. International Food Policy Research Institute, Washington, D.C.

Smith, Lisa C., Harold Alderman and Dede Aduayom (2006) “Food Insecurity in Sub-Saharan Africa. New Estimates from Household Expenditure Surveys.” Research report 146. International Food Policy Research Institute, Washington, D.C.

Svedberg, Peter (2000) *Poverty and Undernutrition. Theory, Measurement and Policy*. UNU-WIDER Studies in Development Economics. Oxford University Press.

UNU/WHO/FAO (2004) *Human energy requirements*. Report of a Joint FAO/WHO/UNU Expert Consultation. Rome, 17-24 October 2001. (Full text vailable at ftp://ftp.fao.org/docrep/fao/007/y5686e/y5686e00.pdf)

WHO (1995) “Physical status: the use and interpretation of anthropometry”. WHO Technical Report Series No. 854. Geneva.