H i C N Households in Conflict Network
The Institute of Development Studies - at the University of Sussex - Falmer - Brighton - BN1 9RE
www.hicn.org

Estimating Poverty in a Fragile Context
The High Frequency Survey in South Sudan

Utz Pape* and Luca Parisotto

HiCN Working Paper 305

May 2019

Abstract: The High Frequency South Sudan Survey, implemented by the South Sudan National Bureau of Statistics in collaboration with the World Bank, conducted several waves of representative surveys across seven of the ten former states between 2015 and 2017. These surveys provided a long overdue update to poverty numbers in South Sudan, with the previous national poverty estimates dating as far back as 2009. The escalation and expansion of the civil conflict posed severe challenges to the planning and implementation of fieldwork. The surveys therefore capitalized on several technological and methodological innovations to establish a reliable system of data collection and obtain valid poverty estimates. Focusing on the 2016 urban-rural wave, this paper describes the design and analysis of the survey to arrive at reliable poverty estimates for South Sudan, utilizing the Rapid Consumption Methodology combined with geo-spatial data for inaccessible survey areas.

Keywords: Consumption Measurement, Poverty, Questionnaire Design

JEL classification: C83, D63, I32

Acknowledgements: The findings, interpretations and conclusions expressed in this paper are entirely those of the authors, and do not necessarily represent the views of the World Bank, its Executive Directors, or the governments of the countries they represent. The authors would like to thank Kristen Himelein, Syedah Iqbal and Ambika Sharma for discussions. In addition, the authors thank Véronique Lefebvre, Sarchil Qadar, Amy Nineman and Tom Bird from Flowminder and WorldPop for modelling and imputing poverty from spatial data, in collaboration with the authors.

* Authors in alphabetical order. Corresponding author: Utz Pape (upape@worldbank.org).

1. Introduction

Civil war broke out across The Republic of South Sudan in December 2013 only two years after gaining independence on the 9th of July 2011. The South Sudanese conflict has since continued to escalate, resulting in a large-scale humanitarian crisis where more than a third of the population has been forcibly displaced (World Bank, 2018).[2] Given the extremely difficult context, very little was known about welfare and livelihoods during the early years of the country's independence in 2011.[3] The last nationally representative household survey measuring consumption and poverty was conducted as far back as 2009. To fill this data gap, the High Frequency South Sudan Survey (HFS), implemented by the National Bureau of Statistics (NBS) in collaboration with the World Bank and funded by the U.K. Department for International Development, conducted several waves of representative surveys across seven of the ten former states between 2015 and 2017 (Appendix A). In the period prior to and during the first wave of the HFS in 2015, conflict had primarily been concentrated in the Greater Upper Nile region (Figure 12in Appendix D).[4] This period of relative stability across the remaining Greater Equatoria and Greater Bah El-Ghazal regions allowed the preparation and relatively calm implementation of Waves 1 and 2 of the country in 2015 and early 2016.

In summer 2016, clashes broke out in Juba. The escalation of the conflict coincided with the beginning of the implementation of Wave 3 of the HFS, a second urban-rural representative wave measuring consumption and poverty. The third wave of the HFS provides a relatively rare and extremely valuable glimpse of trends in welfare, consumption, and poverty in a country going through a period of upheaval. Indeed, the South Sudanese economy has since displayed all the characteristics of a war economy, including severe output contraction, rapid currency devaluation, and soaring inflation (FAO and WFP, 2017; International Monetary Fund, 2016). Unsurprisingly, driven by these powerful shocks the incidence of poverty has risen to extremely high levels. In 2016, the HFS estimated that more than 4 in 5 people across seven of the ten former states were living under the international poverty line of US$ 1.90 PPP 2011 (82 percent). Such high levels of deprivation are not merely a direct result of the crisis but also reflect a history of instability, characterized by a poorly functioning state and a lack of institutional services provision (de Vries and Schomerus, 2017; de Waal, 2014; World Bank, 2017). In 2017 South Sudan ranked 187 of 189 countries in the Human Development Index, with a life expectancy of merely 57 years.[5]

The HFS was designed with the expectation of potential instability and thus capitalized on recent technological and methodological innovations to obtain reliable national poverty statistics in difficult contexts.[6] Closely monitoring fieldwork is key to implementing such a large project in a risky context. The HFS leveraged the expansion of cellular networks across South Sudan to build a near real-time monitoring system, whereby the data could be uploaded daily to a dedicated server and checked for consistency. Computer Assisted Personal Interviewing (CAPI) also allowed built-in consistency checks, eliminating the need for expensive and potentially dangerous re-visits. Adherence to the sample design can be closely monitored with GPS software, tracking enumerators inside and outside areas with mobile phone coverage. The HFS also leveraged innovations in questionnaire design which permitted reducing the number of consumption items asked to the respondents while still obtaining unbiased poverty estimates through

---

[2] See, UNOCHA: https://www.unocha.org/south-sudan & UNHCR: http://data.unhcr.org/SouthSudan/regional.php

[3] Not only has insecurity made fieldwork dangerous, but much of the South Sudanese population lives in isolated and hard to access areas. More than 85 percent of the 12 million South Sudanese reside in sparsely populated rural areas connected by a mere 200 km of paved roadways – about 2 percent of all roads – spanning an area of 650,000 square kilometers, approximately the size of France (Pape et al. 2017; African Development Bank 2013). The poor state of infrastructure combined with the size of the country means nationally representative surveys are expensive and time-consuming.

[4] The Greater Upper Nile region was where the opposition forces, the SPLM-IO, kept their stronghold and were thus contested in the fighting. In Appendix Figure 11 this region corresponds to the non-HFS states, where the number of conflict events in non-HFS covered states is much greater throughout 2014. During the year 2015 the conflict lost some of its intensity. Especially in the HFS states, where although conflict events continued to be recorded most of the violence remained concentrated, particularly in a few select areas which were relatively close to the border with the non-HFS states (Figure 12).

[5] UNDP Human development index, available at: http://hdr.undp.org/en/composite/HDI; and World Development Indicators.

[6] For a comprehensive review of issues in data collection in fragile and conflict situations see Mneimneh et al., (2016).

within-survey multiple imputation (Pape and Mistiaen, 2018, 2015). The lower amount of time spent collecting consumption data allowed the HFS to devote more time to collecting complementary data. Indeed, the HFS questionnaires contained additional modules covering asset ownership, education, labor market outcomes, perceptions of government performance and provision of public goods and services, psychological well-being, perceptions of violence and safety, allowing a well-rounded depiction of welfare and livelihoods.

The rapid escalation of the conflict in the summer of 2016, including several violent incidents affecting international humanitarian and development staff, led to the closure of the World Bank Office in South Sudan, disrupting the implementation of the third wave of the HFS. Therefore, the NBS implemented the third wave of the survey more independently relying mainly on remote support. A multitude of challenges had to be met, including large inflation, fuel unavailability, electricity shutdowns, insecurity, delay in payment of staff salaries, high NBS staff volatility, and cash flow limitations. Even though the NBS and the World Bank project team managed to mitigate a number of those challenges, the final sample reached only about 50 percent of the intended sample size. Nevertheless, this paper will argue that despite the enormity of challenges faced during fieldwork and the slight methodological departures from established approaches to poverty estimation (e.g. Deaton and Zaidi, 2002), the data collected by the HFS provide the best-possible insights on welfare and livelihoods during a critical period of the country's history.

The data from the HFS are complemented by video testimonials providing a glimpse of the lives of the people of South Sudan. At the end of the interviews, respondents are offered to provide a short video testimonial where they can share their views and give a sense of their lives. The testimonials capture the dire situation on the ground and provide a much richer qualitative picture that accompanies and complements the quantitative data. While the data may help the government fine tune its policies, the videos may reach a broader audience and depict the sense of powerlessness, the pain of hunger, the stress of hopelessness and the feelings of disappointment that express people's experiences. Overall, this helps to create a more rounded perception of the situation on the ground in South Sudan.[7]

The levels of deprivation documented by the HFS are staggering. As mentioned above, more than 4 in 5 people across the seven states covered in 2016 were living under the international poverty line of $1.90 USD PPP (83 percent). Such breadth of poverty places South Sudan among some of the poorest countries in the world. The depth of poverty is just as important as its breadth, with the average poor household consuming about one-half of the international poverty line (a poverty gap index of 47 percent). The incidence of poverty is much more widespread in rural areas compared to urban areas, with the rural poverty headcount reaching up to 86 percent compared to 65 percent in urban areas ($p<0.001$). The rural poor also to experience a deeper poverty than urban residents, with a poverty gap equal to 50 percent compared to 31 percent in urban areas ($p<0.001$). Widespread fighting and large-scale displacement over several consecutive planting seasons have disrupted many households' normal agricultural activities, resulting in increasingly large production deficits each year and widespread food insecurity. This has had a devastating effect on livelihoods, given that except for a few oil enclaves the productive structure of South Sudan is one of a rural pastoralist society where more than 4 in 5 people practice subsistence agriculture (World Bank, 2018, 2016).

Despite initial intentions to expand the HFS across the entire country, continued insecurity made it impossible to reach the former states of Jonglei, Unity, and Upper Nile. To account for this gap in coverage and obtain countrywide poverty rates, a statistical model imputes poverty in inaccessible areas. The resulting poverty predictions are intended as supplemental to the survey estimates and serve as a proof-of-concept for using geo-spatial information alongside on-the-ground data collection. A growing body of research has emerged leveraging the increasing availability of alternative data sources such as satellite imagery and other geo-spatial characteristics. The estimates are derived by exploring the potential correlations between existing spatial data sets as well as custom-derived spatial data with geo-referenced poverty estimates obtained in the HFS. Once a set of spatial correlates were

[7] The translated testimonials are available at: http://www.thepulseofsouthsudan.com.

selected several models were trained and evaluated using a cross-validation approach. The final model was used to predict poverty rates at the 100m*100m level into all settled areas of the country including where survey data were not available. To aggregate the estimates at the state and county level, the 100m*100m level are weighted using a newly developed data set of human settlements across South Sudan constructed by combining a variety of publicly available data sources.

This paper describes the design and analysis of the third wave of the HFS in 2016.[8] The paper is focused on Wave 3 of the HFS, conducted between mid-2016 and early 2017, representing the most recent wave covering both urban and rural areas. Furthermore, the period between late 2016 and early 2017 was a critical period in South Sudan's history, when the conflict and refugee crises were reaching their peak. In Section 2, the paper describes the survey design and implementation, including the deviations from the original sample frame presenting consistency-checks used to evaluate potential selection issues that affect representativeness. Section 3 will detail the process of calculating consumption aggregates and estimating poverty using within-survey multiple imputations, including calculating durables consumption flow and spatial-time deflators. Section 4 gives a brief overview of the results of the poverty estimation, while a comprehensive assessment of poverty trends is available elsewhere (World Bank, 2018). Section 5 describes the estimation of poverty rates using satellite data as a proof-of-concept while Section 6 concludes the paper with a short discussion of main limitations.

2. Survey Design and Implementation

Sample Design

The 2016 Wave of the HFS was conducted between mid-2016 to early 2017 and consisted of the second nationally representative survey wave of the HFS. The survey covered rural and urban areas across 7 of the 10 former states of South Sudan. The regions covered include Greater Equatoria, Greater Bahr el Ghazal, and Lakes. The 10 former states are used in planning for the HFS instead of the 28 more recent ones because the sample was constructed based on the sampling frame derived from the 5th Sudan Population and Housing Census from 2008.[9] The survey was designed to be representative at the state level and employs a stratified two-stage clustered sample design. Within each state the primary sampling units are enumeration areas (EAs) that were drawn randomly proportional to size. The EAs were drawn by the NBS for the 2008 Census (Southern Sudan Center for Census, Statistics, and Evaluation, 2010).[10] The number of EAs and households was equalized across states in order to balance the fieldwork across teams. Within the EAs, 12 households were drawn randomly as the unit of observation based on a listing exercise.[11]

The EAs were allocated across urban and rural areas within each state to minimize the variance of indicators of interest across the strata while explicitly taking into consideration the design effect. The data used for the sample size calculations came from the NBHS 2009, and the indicator used for the sample size calculations was the real total per capita household expenditure.[12] While this variable is one of several that are of interest in the HFSSS, consumption/expenditure is generally strongly positively correlated with other indicators of interest. For the purposes of comparison, the relative standard error (complex standard error / mean) is used. The allocation was

[8] The data from Wave 3 (2016) of the HFS and the code used to process these data can be downloaded from the World Bank MicroData Library at the following link: http://microdatalib.worldbank.org/index.php/catalog/9584/
[9] The more recent states have largely been drawn based on the counties subdivision of the former states, the geographical boundaries have therefore largely remained intact.
[10] Urban EAs were drawn to contain approximately 100 to 150 households, while urban EAs would generally contain between 200 to 300 households.
[11] The number of households per EA was determined to be 12 to allow an equal split into 4 groups per EA to facilitate the implementation of the Rapid Consumption Methodology. The specific options of 8, 12, and 16 were considered. Eight households per cluster was deemed as too small as the number of EAs necessary and the associated travel time could not be done within the fieldwork calendar. Sixteen resulted in very high design effects, over 3 in most cases and as high as 5 for some strata, and was therefore deemed too large. Twelve households per EA was therefore selected as the ideal cluster size.
[12] The top and bottom 1 percent of outlier observations were trimmed for the sample size calculations.

done so as to ensure a minimum of 10 EAs per combination of urban-rural and state distinction, according to the following rule:

$$n_u = \begin{cases} if\ n_u \geq 10 & n_u \quad n_r \\ if\ n_u < 10 & 10 \quad 40 \end{cases}, n\left(\frac{N_u S_u * deff}{N_u S_{u*deff} + N_r S_r * deff}\right), n_r = n\left(\frac{N_r S_r * deff}{N_u S_u * deff + N_r S_r * deff}\right), n = n_u + n_r = 50$$

where n_h is the sample size in stratum *h, n* is the total sample size, *H* is the total number of strata, N_h is the total population of stratum *h*, *N* is the total overall population, and S_h is the standard deviation in stratum *h.* The results from the sample size calculations are shown in Appendix B, Table 3. The chosen sample allocation provides estimates that are representative at the national, urban/rural, and state level. Sampling weights were calculated on the basis of the 5th Sudan Population and Housing Census from 2008 (Appendix B). In cases where fewer than 12 households were interviewed in an EA, the sampling weights were adjusted at the EA level to reflect this.

Data collection was intended to be implemented in two phases, by randomly splitting each stratum into two equal-sized parts, where each phase of data collection would cover half of the sample. The advantage of a two-phased approach was early availability of representative data after half of the survey was implemented. The two-phased approach reduces the risk that an eruption of violence during field work invalidates the representativeness of the survey. However, such an approach is not guaranteed to maintain representativeness if some areas remain inaccessible throughout the entirety of fieldwork. It also comes at the cost of optimizing the organization of fieldwork by reducing the enumerators' ability to sweep over their intended area.

Survey Implementation

The survey was implemented using tablets as survey devices. The data collection system consisted of Samsung Galaxy Tablet computers equipped with SIM cards, mobile data plans, microSD cards (16 GB capacity), and external battery packs.[13] Teams of four enumerators and one supervisor were provided with a mobile generator using fuel to ensure that tablets could be charged overnight. Computer Assisted Personal Interviewing (CAPI) data collection can be used to improve data quality by imposing sophisticated systems of constraints on the enumerators' entries. This was particularly relevant for consumption and price data, which need to be measured precisely as a prerequisite for a reliable poverty analysis. Indeed, CAPI has been experimentally shown to improve data collection while minimizing the potential for enumerator error (Caeyers et al., 2012; Fafchamps et al., 2010). Furthermore, it can be used to create more sophisticated questionnaires, with elaborate conditional skipping patterns that are much easier to implement (De Leeuw et al., 1995).

The rapidly expanding cellular network in South Sudan meant that the data could be transmitted via mobile networks and made available quickly to data analysts (Pape and Mistiaen, 2014). The near real-time transmission of data to a cloud enabled the implementation of a monitoring system including a dashboard tracking the cumulative number of interviews, the fraction of missing variables, as well as additional quality indicators at any level of disaggregation.[14] This helped to identify challenges in the field work as well as weak enumerators early on and mitigate their impact on data quality, e.g. by providing individually tailored extra trainings for selected enumerators. In addition, the real-time analysis code calculates core indicators of the survey, e.g. consumption, educational attainment, and unemployment, to check incoming data while field work is still ongoing. This head-start on building the analysis code ensures that swiftly after the end of data collection the cleaned data can be made available, which considerably accelerated the process from data collection to the publication of results.

The availability of real-time data facilitated monitoring by allowing much closer tracking of the geographic progression of fieldwork. The GPS coordinates for each interview were recorded and uploaded along with the data,

[13] The Android application AirDroid was used to remotely manage devices, this remote management software meant that errors in the tablet configuration were detected and could be solved by updating the tablets remotely in cases where enumerators may have needed help from the survey analysts.

[14] In areas without 3G activities, enumerators saved conducted interviews on the tablet and submitted data once they had 3G connectivity.

allowing tracking enumerators and ensuring the sampling design was implemented. Furthermore, GPS tracking software helped to track devices at all times using a web interface (www.gps-server.net), the exact path of the devices was recorded even retrospectively and uploaded to the server once they entered areas with 3G/WIFI connection. Given the frequent disruptions and slow rate of data collection their combination provided a useful reference to understand where field teams were at any time, and could be cross-checked with reports of conflict activity etc. Overall, this system allowed close supervision of the implementation of the sampling design (Pape and Mistiaen, 2014).

Fieldwork and Insecurity

Sporadic eruptions of fighting meant that teams of enumerators were at times forced to remain idle and wait for the situation to deescalate before reaching certain areas. A few areas that had been subjected to heavy fighting and that may have experienced mass displacement could not be reached at all. Therefore, fieldwork was delayed and the quality of documentation was negatively affected. In the end, despite the relatively long duration of data collection, the final sample fell short of the intended sample. Fortunately, the two-phased approach described above implies that representative data are already available after the first half of the survey implementation. Indeed, the final sample that was collected during Wave 3 only reaches only about 50 percent of the intended sample size, i.e. the first of the two phases. This was true across all states (Table 4 in Appendix E).

Nevertheless, many of the selected EAs had to be replaced when security rendered field work unfeasible.[15] One hundred EAs were surveyed of the 350 EAs in the original sample, the rest of the 64 EAs were replacement EAs. Replacements were done in three batches where each time new enumeration areas had to be drawn from the master sample frame. The replacement sequence was defined by assigning enumeration areas randomly to the original enumeration areas, maintaining the order of the original enumeration areas as in the original sample. The large number of replacements was concerning given fear of selection bias. Therefore, the team ran checks to ensure that the set of EAs surveyed do not systematically differ from a random sample as best as it could. It is important to keep in mind that assessing representativeness is a difficult task, generally due to the lack of a counterfactual or a point of reference to compare estimates. Despite these checks, it is plausible that selection bias in favor of less conflict-affected areas leads to an under-estimation of poverty. The resulting estimates are therefore interpreted as lower-bound estimates.

The checks are based on comparisons of Wave 3 data from 2016 with the nearest available reference point, Wave 1 data from 2015. Specific outcomes were compared across the two waves as well as at lower levels of aggregation and within specific regions (Table 8 in Appendix E shows an example). This process was severely complicated by the magnitude of the South Sudanese crisis. The conflict, displacement crisis, and near-hyperinflationary price increases are powerful shocks, which are expected to cause severe disruption even in a relatively short amount of time.[16] The checks therefore concentrated on outcomes that are less likely to be affected by the crises and are relatively time-invariant.

Adults' educational outcomes is one such indicator which is expected to remain relatively stable from one year to the next assuming only small demographic changes. In South Sudan, the adult literacy rate (18+), the proportion of adults with no education, and the proportion of adults with only primary education were comparable between 2015 and 2016 (Table 8). Similarly, cultural norms should be expected to remain stable, such as the prevalence of polygamy and the gender of the household head, both of which are again seemingly unchanged. Some types of infrastructure can provide good indicators if they are not susceptible to be destroyed in the fighting. Mobile

[15] Replacement EAs were approved by the project manager. Replacement of households were approved by the supervisor after a total of three unsuccessful visits of the household.

[16] At the very start of Wave 3 data collection year-on-year inflation was equal to almost 650 percent. The CPI between the start of Wave 1 and the end of Wave 3 had increased by almost 1,600 percent. Similarly, more than a third of the population was displaced by mid-2018.

telephone networks are a good example, since they generally comprise relatively heavy infrastructure that is not easily destroyed through the type of warfare occurring in South Sudan. This is also a good indicator of sample selection favoring wealthier areas, especially in the context of South Sudan where only one in four households is covered. Access to electricity is a similar indicator given that it is exclusive to a few selected areas of South Sudan. Again, the latter two indicators do not seem different from 2015 and 2016. Finally, the share of households living far from schools, health centers, and markets, did not change significantly – this generally holds for various thresholds.

More importantly, the path of enumerators and geographic coverage of Wave 3 data was closely inspected to ensure that it remained broadly comparable to that of previous HFS waves and other sources of population data. This helped to control that entire areas were not systematically excluded. As an exception, the city of Yei was not surveyed at all in Wave 3 because it was the site of several large battles during fieldwork and subsequently experienced a massive wave of displacement. This was likely the most severe case, and in many other instances where fighting affected specific areas enumerators simply delayed fieldwork until it was safe to continue. This explains to some extent the prolonged duration of fieldwork relative to the low number of interviews conducted in total.

3. Measuring Poverty in a Fragile Context

Calculating Consumption Aggregates

Poverty in the HFS was measured according to a standardized methodology best described in the seminal contribution by Deaton and Zaidi (2002). Poverty analysis consists of comparing a welfare measure to a predetermined poverty line. Therefore, the first step is to calculate a measure of welfare. The measure chosen for the HFS is the households' consumption expenditure per capita.[17] The nominal household consumption aggregate consists of the sum of consumption expenditure per person on three primary components, i) total expenditures on food items, ii) total expenditures on non-food items, and iii) the value of the consumption flow from the durable goods owned by the household.[18] The consumption aggregate is then deflated to reflect spatial and temporal cost of living differences.

$$y_i = y_i^f + y_i^n + y_i^d \quad (1)$$

Accurately measuring consumption in highly volatile environments is a complex task, primarily because insecurity and uncertainty severely restrict the time that can safely be spent by enumerators in certain areas and the time spent conducting each interview. Consumption modules tend to be bulky and take a long time to administer. At the very least, it requires asking information on quantities consumed, quantities purchased, and prices of purchase – including additional information on home production in a context such as South Sudan – for what is often upwards of 300 to 400 consumption items (Beegle et al., 2012). Reducing the length of the questionnaire is therefore a key strategy when designing surveys for fragile contexts. For example, it is common to remove rarely consumed items or to combine categories of items (e.g. vegetables). However, Beegle et al. (2012) and Olson Lanjouw and Lanjouw (2001) show that such approaches tend to result in underestimated consumption levels, and hence overestimate the poverty rate.

[17] In the context of South Sudan using consumption as a measure of welfare is preferable to a measure of income for two main reasons: (i) there exists no real reliable information on income given poor administrative record keeping, and (ii) employment is primarily irregular and informal in nature, with subsistence agriculture accounting for about two-thirds of employment, non-farm business ownership for one-eighth, and salaried labor also only about one-eighth (World Bank, 2018).

[18] In some cases, housing is included in the consumption aggregate. However, calculating the consumption flow obtained from housing requires estimating rental values from the open market (Balcazar et al., 2014). Unfortunately, the housing market in South Sudan is highly underdeveloped, making such estimations impossible in any sort of accurate manner. Indeed, in the 2016 HFS, 91 percent of households were owned by the residents and fewer than 4 percent were rented. Thus, housing was excluded from the consumption aggregate.

Another set of approaches for obtaining poverty estimates in a fragile context consists of modeling consumption, or poverty, based on a set of observable covariates and then projecting estimates using cross-survey imputation. In this manner, infrequent bulky consumption surveys can be combined with more frequent surveys that collect information on the covariates necessary for imputing poverty (for example labor force surveys as in Douidich et al., 2013; or SWIFT[19]). However, this methodology is problematic in contexts where there is no consumption survey to underlie the estimation, or where there may have been deep structural change that changes the relationship between covariates and poverty across time (Beegle et al., 2016; Christiaensen et al., 2010). This is most likely the case in South Sudan, where the last full consumption survey was conducted in 2009 and which had experienced a period of rapid development leading up to independence in 2011 and until the breakout of the current conflict in 2013.

Within-survey imputation can alleviate some of these concerns because the assumption of similar covariate distribution between the data used to estimate poverty and that used to project is more likely to hold, or the differences may not be as great. One approach consists of administering a full consumption module to a subset of respondents, generally in more secure areas where time-constraints are not binding, and then impute consumption for less secure areas based on a smaller set of covariates (Fujii and Van der Weide 2013). However, safer areas where the full consumption module can be administered may still systematically differ from insecure areas where only the covariates are collected, thus violating the assumption of equally distributed covariates.

The HFS employed a method of within-survey imputation, but instead of imputing the totality of consumption in certain areas based on data from other areas it imputed a randomly different fraction of consumption across all enumeration areas covered in the survey (Pape and Mistiaen, 2018, 2015). Food and non-food consumption items were first into a core and multiple optional modules. Each household was then asked only about the core items and those items in one of the optional modules, and consumption of items in the remaining optional modules was estimated through multiple imputation. The imputation does not suffer from bias caused by different covariate distributions, since data on every one of the optional consumption modules are collected within each enumeration area. Furthermore, because a majority of consumption is accounted for by a relatively small set of items collected for each household, additional variance introduced by the imputation is minimized.

This section will describe the rapid survey consumption methodology, a more detailed treatment and simulations can be found in (Pape and Mistiaen, 2018, 2015). First, food and non-food consumption for household *i* are estimated by the sum of expenditures for a set of items

$$y_i^f = \sum_{j=1}^{m} y_{ij}^f \text{ and } y_i^n = \sum_{j=1}^{m} y_{ij}^n \quad (2)$$

where y_i^f and y_i^f denote the food and non-food consumption of item *j* in household *i*.[20] Previous consumption surveys in the same country or consumption surveys in neighboring / similar countries can be used to estimate food shares.[21] In South Sudan, the item assignment could draw from the NBHS 2009 survey.[22] The list of items was partitioned into 1 core and 4-optional modules each with m_k items:

[19] Survey of Well-being via Instant and Frequent Tracking.

[20] As the estimation for food and non-food consumption follows the same principles, we neglect the upper index *f* and *n* in the remainder of this section.

[21] In a case where a previous survey is not available the items can be randomly assigned to the module. This would result in larger standard errors but would not introduce bias.

[22] With manual modifications to treat 'other' items correctly. Items 'other' are often found to capture remaining items for a food category. Using the Rapid Consumption Methodology, this creates problems as 'other' will include different items depending on which optional module is administered. This can lead to double-counting after the imputation. Therefore, 'other' items are re-formulated and carefully assigned so that double counting cannot occur.

(3) $$y_i = \sum_{k=0}^{4} y_i^{(k)} \text{ with } y_i^{(k)} = \sum_{j=1}^{m_k} y_{ikj}$$

The core module was designed to maximize its consumption share based on NBHS 2009 consumption, and therefore contains all the most commonly consumed items. This includes staple foods such as dura, maize, onions, okra, common types of flour (e.g. millet, maize, cassava, and groundnut flour), common types of meat (e.g. goat, sheep, poultry, beef), and some fruits. The nonfood core module similarly captures common expenditures including fees for education, common types of transportation, common medicines and health related expenditures, and clothing. Optional modules were constructed using an algorithm to assign items iteratively to optional modules so that items are orthogonal within modules and correlated between modules.[23]

This step is followed by the actual data collection. Conceptual division into core and optional items is not reflected in the layout of the questionnaire. Rather, all items per household are grouped into categories of consumption items (like cereals, meats, etc.). Using CAPI, it is straight-forward to hide the modular structure from the enumerator. For each household, only the core module $y_i^{(0)}$ and one additional optional module $y_i^{(k^*)}$ are collected. In each enumeration area, 12 households were interviewed with an ideal partition of three households per optional module.[24] The assignment of optional modules was stratified per EA to ensure that an equal number of households are assigned to each optional module. This served to minimize potential EA effects during the imputation process.

Household consumption was then estimated using the core module, the assigned module and estimates for the remaining optional modules:

(4) $$\hat{y}_i = y_i^{(0)} + y_i^{(k^*)} + \sum_{k \in K^*} \hat{y}_i^{(k)}$$

where $K^* := \{1, \dots, k^* - 1, k^* + 1, \dots, M\}$ denotes the set of non-assigned optional modules. Consumption of non-assigned optional modules is estimated using multiple imputation techniques taking into account the variation absorbed in the residual term.

Multiple imputation was implemented using multi-variate normal regression based on an EM-like algorithm to iteratively estimate model parameters and missing data. This technique is guaranteed to converge in distribution to the optimal values. An EM algorithm draws missing data from a prior (often non-informative) distribution and runs an OLS to estimate the coefficients.[25] Iteratively, the coefficients are updated based on re-estimation using imputed values for missing data drawn from the posterior distribution of the model. The implemented technique employs a Data-Augmentation (DA) algorithm, which is similar to an EM algorithm but updates parameters in a non-

[23] In each step, an unassigned item with the highest consumption share was selected. For each module, total per capita consumption was regressed on household size, the consumption of all assigned items to this module as well as the new unassigned item. The item in questions was then assigned to the module with the highest increase in the R2 relative to the regression excluding the new unassigned item. The sequenced assignment of items based on their consumption share can lead to considerable differences in the captured consumption share across optional modules. Therefore, a parameter is introduced ensuring that in each step of the assignment procedure the difference in the number of assigned items per module does not exceed *d*. Using *d=1* assigns items to modules (almost) maximizing equal consumption share across modules. Increasing *d* puts increasing weight on orthogonality within and correlation between modules. The parameter was set to *d=3* balancing the two objectives.

[24] Field work implementation aimed to achieve a balanced partition among optional modules but due to challenges in following the protocol exactly some enumeration areas are not completely balanced.

[25] The model employed in the HFS was constructed using the following indicators: demographics variables including household size, the fraction of children, the fraction of elderly persons, the sex of the household head, and the employment status of the household head; indicators of access to amenities including the water source, whether the household had electricity to power its lighting, the number of sleeping rooms, and whether the household had access to a toilet; geographic indicators including an urban-rural dummy and state fixed effects; finally, the model included dummies for each quartile of consumption of food and non-food per capita. One hundred imputations were run for the consumption imputation process to maximize the accuracy of results.

deterministic fashion unlike the EM algorithm. Thus, coefficients are drawn from the parameter posterior distribution rather than chosen by likelihood maximization. Hence, the iterative process is a Monte-Carlo Markov Chain (MCMC) in the parameter space with convergence to the stationary distribution that averages over the missing data. The distribution for the missing data stabilizes at the exact distribution to be drawn from to retrieve model estimates averaging over the missing value distribution. The DA algorithm usually converges considerably faster than using standard EM algorithms:

$$\hat{y}_i^{(k)} = \beta_0^{(k)} y_i^{(0)} + x_i^T \beta^{(k)} + u_i^{(k)} \tag{5}$$

The performance of the estimation technique was assessed based on an *ex post* simulation using the NBHS 2009 data and mimicking the Rapid Consumption methodology by masking consumption of items that were not administered to households. The results of the simulation were compared with the estimates using the full consumption from NBHS 2009 as reference. The simulation results distinguish between different levels of aggregation to estimate consumption.[26] The methodology generally does not perform well at the household level (HH) but improves considerably already at the enumeration area level (EA) where the average of 12 households is estimated. At the national aggregation level, the Rapid Consumption methodology slightly over-estimates poverty by 1.6 percent. Assessing the standard poverty measures including poverty headcount (FGT0), poverty depth (FGT1) and poverty severity (FGT2), the simulation results show that the Rapid Consumption methodology retrieves almost unbiased estimates. Generally, the estimates are robust as suggested by the low standard errors.[27]

The assumption that the imputed components of consumption follow a joint normal distribution might provide an explanation as to why poverty is slightly overestimated. This would be due to the imputed means of consumption of the imputed items being slightly lower than the actual means since their true distributions are generally skewed to the right. This possibility was explored by assuming a non-parametric error term in the imputation procedure through the use of chained equations, which performed almost indistinguishably as well as the multivariate-normal approximation.

[26] The performance of the estimation techniques is presented using the relative bias (mean of the error distribution) and the relative standard error. The relative error is defined as the percentage difference of the estimated consumption and the reference consumption (based on the full consumption module, averaged over all imputations). The relative bias is the average of the relative error. The relative standard error is the standard deviation of the relative error. The simulation is run over different household-module assignments while ensuring that each optional module is assigned equally often to a household per enumeration. The relative bias and the relative standard error are reported across all simulations.

[27] These standard errors are estimated empirically using a bootstrap approach taking into account intra-cluster correlation within enumeration areas.

Figure 1: Relative bias of simulation results using Rapid Consumption estimation.

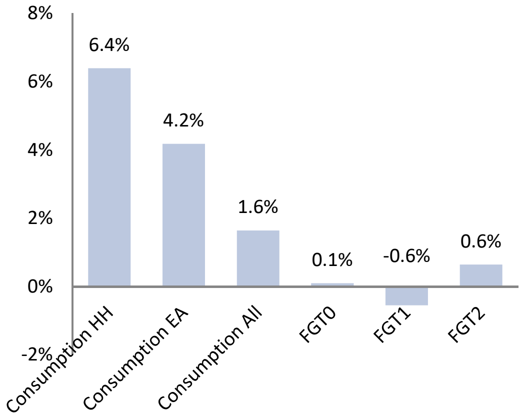

Figure 2: Relative standard error of simulation results using Rapid Consumption estimation.

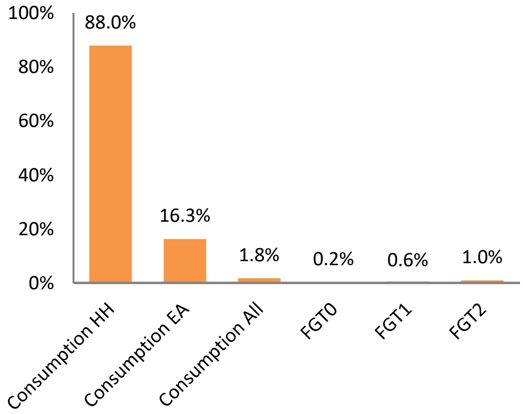

Source: Authors' own calculations based on NBHS 2009 data.

Durable Consumption Flow

The consumption aggregate includes the consumption flow of durables calculated based on the user-cost approach, which distributes the consumption value of the durable over multiple years (Amendola and Vecchi, 2014). The user-cost principle defines the consumption flow of an item as the difference of selling the asset at the beginning and the end of the year as, this is the opportunity cost of the household for keeping the item. The opportunity cost is composed of the difference in the sales price and the forgone earnings on interest if the asset is sold at the beginning of the year. The current price of the durable is p_t. If the durable item would have been sold one year ago, the household would have received the market price for the item twelve months ago plus the interest on the revenue for one year. The market price from 12 months ago is calculated by adjusting for inflation π_t and annual physical or technological depreciation rate δ arriving at[28]

$$\frac{p_t(1+i_t)}{(1+\pi_t)(1-\delta)} \tag{6}$$

with the nominal interest rate denoted as i_t. Alternatively, the household can use the durable and sell it after one year of usage for the current market price p_t. The difference between these two values is the cost that the household is willing to pay for using the durable good for one year. Hence, the consumption flow is:

$$\mathrm{y}^d = \frac{p_t(1+i_t)}{(1+\pi_t)(1-\delta)} - p_t \tag{7}$$

By assuming that $\delta \times \pi_t \cong 0$, the equation simplifies to

$$\mathrm{y}^d = \frac{p_t(r_t+\delta)}{(1+\pi_t-\delta)} \tag{8}$$

where r_t is the real market interest rate $i_t - \pi_t$ in period t. Therefore, the consumption flow of an item can be estimated by the current market value p_t, the current real interest rate r_t, the inflation rate π_t and the depreciation rate δ. Assuming an average annual inflation rate π, the depreciation rates δ can be estimated utilizing its relationship to the market price[29]:

[28] Assuming a constant depreciation rate is equivalent to assuming a "radioactive decay" of durable goods (see Deaton and Zaidi, 2002).

[29] In particular π solves the equation $\prod_{i=t-k}^{t}(1+\pi_i) = (1+\pi)^k$.

(9) $$p_t = p_{t-k}(1+\pi)^k(1-\delta)^k$$

The equation can be solved for δ obtaining:

(10) $$\delta = 1 - \left(\frac{p_t}{p_{t-k}}\right)^{\frac{1}{k}} \frac{1}{(1+\pi)}$$

The depreciation rates estimated the 2015 HFS wave were used to calculate the consumption flow in the 2016 wave. The reason being that estimating depreciation rates is much more prone to errors in a context of high and unstable inflation such as that observed in South Sudan in 2016.[30] Furthermore, there are few reasons to expect depreciation rates to drastically change over such a short period of time. In 2015, based on equation (10), item-specific median depreciation rates are estimated assuming an inflation rate of 0.5 percent, a nominal interest rate of 5.5 percent and, thus, a real interest rate of 5 percent (Table 6). For all households owning a durable but that did not report the current value of the durable, the item-specific median consumption flow is used. For households that own more than one durable, the consumption flow of the newest item is added to the item-specific median of the consumption flow times the number of those items without counting the newest item.[31]

Spatial and Temporal Price Deflators

Prices fluctuated considerably in South Sudan in 2016 (Pape and Dihel, 2017; World Bank, 2018). Prices therefore need to be adjusted to make consumption comparable across the several months of fieldwork. Furthermore, there are important differences in the cost of living between urban and rural areas. This is particularly marked in South Sudan given the sheer isolation of rural areas and state of poor market linkages across the country (African Development Bank, 2013; Pape et al., 2017). A Laspeyres deflator was chosen to calculate price differences across urban and rural areas and months of data collection, due to its relatively light data requirements. The base period for deflating prices was chosen as October 2016 in urban areas. Urban areas were chosen as a reference because the national CPI calculated by the NBS is based on prices in urban markets across some of the largest cities in South Sudan, and hence would facilitate deflating consumption across the frequent waves of data collection in the HFS.

The Laspeyres index reflects the item-weighted relative price differences across products. Item weights are estimated as household-weighted average consumption share across all households before imputation. Based on the democratic approach, consumption shares are calculated at the household level. Core items use total household core consumption as reference while items from optional modules use the total assigned optional module household consumption as reference. The shares are aggregated at the national level (using household weights) and then calibrated by average consumption per module to arrive at item-weights summing to 1. The item-weights are applied to the relative differences of median item prices for each urban/rural and month pair. Missing prices are replaced by the item-specific median over all households. The reference strata was chosen as the urban strata for one specific month of data collection. The month with the most data points was generally chosen for the reference time period. The Laspeyres deflator can be expressed as such:

(11) $$L_{i,t} = \sum_{k=1}^{k} w_{i,k,m}\left(\frac{p_{i,k}}{p_{0,k}}\right)$$

The Laspeyres $L_{i,t}$ for strata i and month t is equal to the sum of, over all items k: $w_{i,k,m}$, the national budget share of item k in optional module m, times the ratio of $p_{i,k,m,t}$, the median price of item k in strata i at month t, and

[30] One potential source of bias being that the value placed by respondents on durable goods may be inflated given high levels of uncertainty regarding the future of the currency. Another is that the volatility of inflation across time periods is problematic given the formula assuming one inflation rate prevailing across the different years.

[31] The 2016 HFSSS questionnaire provides information on a) the year of purchase and b) the purchasing price only for the most recent durable owned by the household.

$p_{0,k,m,0}$, the median price of item k in the reference strata in the reference month. Two sets of price deflators were calculated, one for food and another for nonfood items, the nonfood price deflator was used to deflate the consumption flow of durable goods.

Poverty Line

Determining a household's poverty status requires a poverty line against which to compare consumption. A poverty line serves as a reference point for what might be an acceptable minimum standard of well-being, below which one could be considered deprived, or living in poverty (Ravallion, 2017, 1998). The choice of the poverty line considers what might constitute an acceptable minimum standard of living and the potential impact of resulting poverty estimates on policy decisions. Once a poverty line has been chosen, poverty analysis is then typically based on comparing the first three poverty measures of the Foster-Green-Thorbecke (FGT) class of poverty indicators. FGT measures consist essentially of variations of specification 0, where the parameter α takes the value of 0 for the poverty headcount, 1 for the poverty gap, and 2 for poverty severity (Foster et al., 1984).

$$FGT(\alpha) = \frac{1}{n}\sum_{i=1}^{p}\left[\frac{z-y_i}{z}\right]^{\alpha}$$

Where y_i denotes the consumption y of individual i, n denotes the total population, and z the poverty line.

Theoretically, a national poverty line could have been estimated for South Sudan in the year 2016 using the survey data. However, the international poverty line of US$1.90 PPP was chosen.[32] Given that the international poverty line was based on the predicted poverty line for the world's 15 poorest countries, combined with the expectation that poverty in South Sudan was to be relatively high, the international poverty line was considered an appropriate metric, also offering the ability to make international comparisons. Hence, the $1.90 USD PPP (2011) poverty line was first converted into current SSP and adjusted to reflect South Sudanese purchasing power using the South Sudan PPP conversion factor for 2011. It was then adjusted for inflation up to October 2016 using the national CPI calculated by the National Bureau of Statistics, resulting in a value of approximately 65 SSP (October 2016).

4. Results from the HFS

In 2016, more than 4 in 5 South Sudanese people in the seven states covered in the HFS lived under the international poverty line of US$1.90 PPP (2011) per capita per day. The poverty headcount was equal to 83 percent in 2016, with a 95 percent confidence interval from 81 to 85 percent. These levels of poverty place South Sudan among some of the poorest countries in the world. South Sudan's poverty headcount ratio is much higher than the average estimates of other countries at similar levels of development (Figure 3). The estimated poverty headcount ratio is not particularly sensitive to the choice of poverty line, since average consumption levels are so low that the poverty line lies at a point where the slope of the cumulative distribution of consumption tapers off (Figure 5). The deterioration of economic conditions has driven many poor households down to hardship conditions (Figure). The poverty gap, which measures poor households' average deficit in consumption relative to the poverty line, is equal to 47 percent in 2016. The average poor household is therefore consuming about one-half of the poverty line in 2016 (US$ 1.00

[32] The international poverty line was first introduced in the 1990 World Bank World Development Report with the intent of measuring poverty across countries in a consistent manner. This international poverty line used data on 33 national poverty lines for the 1970s and 1980s and represented the predicted poverty line for the poorest country in the sample, equal to about $0.76 USD PPP (1985). The international poverty line was subsequently adjusted for inflation as new sets of PPP were made available through the International Comparison Program. The computation of the current international poverty line of $1.90 USD PPP per day was obtained as the unweighted average of the poverty line for the 15 poorest countries, as such: i) by adjusting the national poverty lines of the 15 poorest countries for inflation up to 2011; ii) then converting the national poverty lines to real USD using the 2011 PPPs; and iii) then computing the simple average of the 15 national poverty lines. The resulting average poverty line is equal to $1.88 USD PPP (2011) per person per day, which was rounded up to $1.90 USD PPP (2011).

2011 PPP). The poverty severity index, which is the square of the poverty gap and thus places more weight on people with consumption levels further below the poverty line, was equal to 0.31 ($p<0.001$).

Figure 3: Poverty headcount in low and lower middle-income countries.[33]

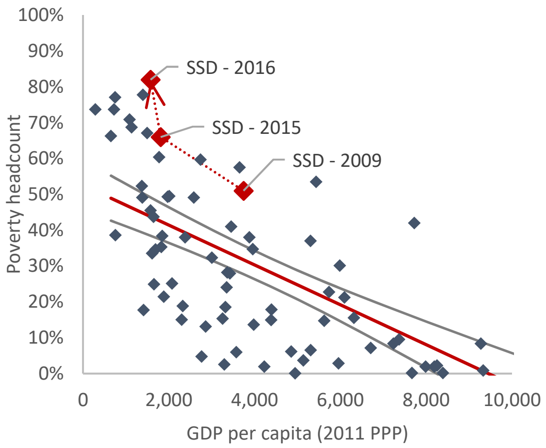

Figure 4: Gini index in SSA countries.

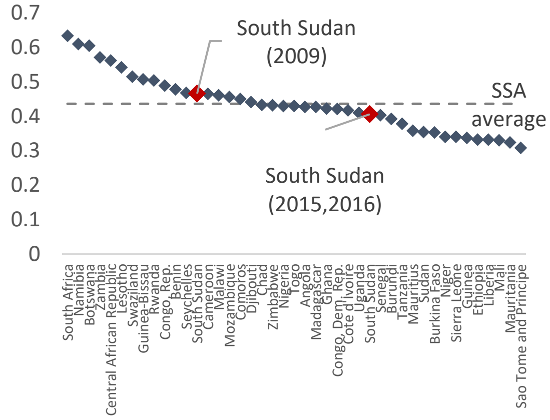

Figure 5: Cumulative consumption distribution.

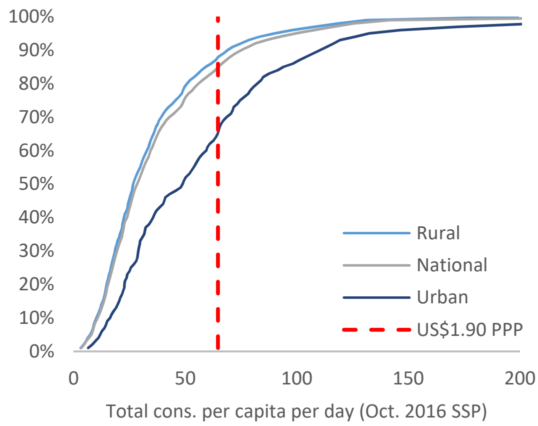

Figure 6: Consumption distribution, 2016.

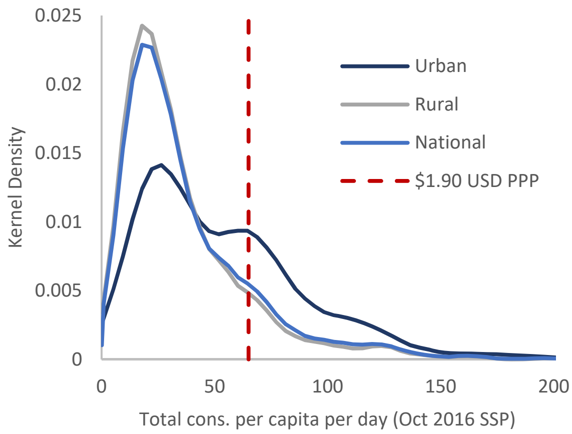

Note: Figure 3 includes low income and lower middle-income countries with poverty data post-2008. All data for South Sudan refers to the seven states covered by the HFS.

Source: Authors' own calculations based on NBHS 2009, HFS 2015-2017, World Bank WDI, and IMF World Economic Outlook.

Such high levels of deprivation translate into widespread hunger and food insecurity. Disruptions to agricultural production and the near hyperinflationary increases in prices of most staple foodstuffs have left most households struggling to find enough food to sustain themselves (World Bank, 2018). Widespread fighting and large-scale displacement over several consecutive planting seasons have disrupted many households' normal agricultural activities, resulting in increasingly large production deficits each year and widespread food insecurity (FAO and WFP, 2017). This has had a devastating effect on livelihoods, given that except for a few oil enclaves the productive structure of South Sudan is one of a rural pastoralist society where more than 4 in 5 people practice subsistence agriculture (World Bank, 2018, 2016). Food security has continuously deteriorated since late 2012, sometimes even reaching famine conditions in certain vulnerable counties. During the most recent harvest season in 2017, a time

[33] Data for real GDP per capita in 2011 PPP for South Sudan were obtained from the IMF World Development Outlook Database.

when food should be abundant, as many as 4.8 million people were severely food insecure (FAO and WFP, 2017). By mid-2018, the number of severely food insecure people is expected to rise to 6.2 million, reaching more than half of the total population.[34]

Table 1: Poverty headcount and average consumption per strata for the seven HFS covered states, 2016.

	Poverty headcount ratio				Mean consumption				
	Mean	**Standard Error**	**[95% CI]**		**Mean**	**Standard Error**	**[95% CI]**		**N**
National	0.83	0.01	0.80	0.86	73.30	2.68	67.99	78.60	1,848
Rural	0.86	0.02	0.83	0.89	67.36	2.70	62.03	72.70	1,281
Urban	0.65	0.02	0.60	0.70	113.99	5.59	102.94	125.05	567
Warrap	0.86	0.05	0.77	0.95	63.98	7.13	49.88	78.08	135
Northern Bahr El Ghazal	0.90	0.03	0.84	0.95	62.63	5.64	51.49	73.77	299
Western Bahr El Ghazal	0.90	0.02	0.87	0.94	60.17	6.33	47.66	72.68	310
Lakes	0.84	0.02	0.80	0.88	71.22	3.46	64.38	78.06	232
Western Equatoria	0.53	0.04	0.46	0.61	130.51	7.45	115.79	145.23	300
Central Equatoria	0.80	0.05	0.70	0.90	86.53	8.27	70.18	102.88	311
Eastern Equatoria	0.95	0.01	0.93	0.98	43.88	3.58	36.80	50.96	261

Note: Standard errors estimated through linear regressions; all estimates weighted using population weights.

Source: Authors' own calculations based on HFS 2016-2017 data.

The incidence of poverty is much more widespread in rural areas compared to urban areas. Rural poverty was equal to 86 percent in 2016 compared to 65 percent in urban areas ($p<0.001$, Figure 5). The rural poor also experience deeper poverty than urban residents, with a higher poverty gap and poverty severity. In 2016, the urban poverty gap was equal to 31 percent compared to 50 percent for the rural poverty gap ($p<0.001$, Figure 5). A similar pattern can be observed for poverty severity, the urban severity index was equal to 19 percent and the rural index equal to 33 percent ($p<0.001$). A stochastic dominance analysis based on a comparison of the cumulative consumption expenditure distribution across rural and urban areas reveals that this is not due to the chosen poverty line but that at any point along the distribution the urban consumption expenditure curve lies consistently below the rural curve (Figure 5). The isolated nature of many rural areas contributes to these observed poverty rates, given that they are often cut off from public services as well as humanitarian assistance.

Measuring inequality, the Gini index in South Sudan declined from 2009 to 2016, from about 0.47 in 2009 to 0.41 in 2016 (Figure 4).[35] The average Gini index for countries in Sub-Saharan Africa is approximately 0.44, with South Sudan at 0.41 indicating slightly lower inequality but higher inequality compared to the global average Gini index of 0.38. While all households suffered consumption losses because of the conflict and macroeconomic crises, the consumption losses experienced by better off households were larger than those of the poorer households (World Bank, 2018). Thus, the driver of the reduction in inequality was not pro-poor growth but rather a greater decline in welfare for wealthier households relative to poorer households. This is not entirely unexpected since the poorer households already experienced extreme deprivation, and thus could not fall much further even as the crisis worsened. Inequality remains nevertheless greater in urban areas than in rural areas though only slightly, at 0.41

[34] FEWSNET Food Security Outlook, February to September 2018.

[35] The Gini index is calculated from the area under the Lorenz curve, which plots the cumulative percentage of consumption expenditure against the cumulative percentage of the population, with perfect equality lying along the 45-degree line.

and 0.39 respectively. Indeed, many of the households with the highest consumption levels reside in urban areas, with better access to markets and opportunities.

Figure 7: Cumulative consumption distribution by state.

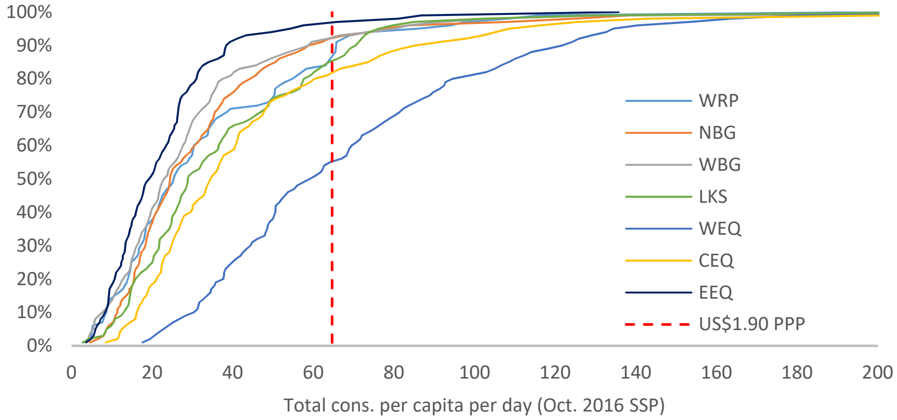

Source: Authors' own calculations based on HFS 2016 data.

Poverty in 2016 is generally high but it is higher in former states that were more exposed to the conflict. The incidence of poverty reached extremely high levels in the former states of Eastern Equatoria, Northern Bahr el Ghazal, and Western Bahr el Ghazal, where about 9 in 10 people live under the international poverty line (95, 90, and 90 percent, respectively). In the former states of Lakes and Central Equatoria, the poverty headcount is slightly lower at about 8 in 10 people, though still extremely high by international standards (84 and 80 percent, respectively). One notable exception is the former state of Western Equatoria, as it was less affected by the conflict compared to the other states and has benefitted from high fertility and favorable weather conditions. Indeed, Western Equatoria, in the "green belt" of South Sudan, was the only state to record a consistent cereal production surplus in the years from 2014 to 2016 (FAO and WFP, 2017). Accordingly, the residents of Western Equatoria were much more likely to be able to sustain their livelihoods through own production compared to those in other states and thus maintain better standards of living (World Bank, 2018).

5. Imputing Poverty Using Geo-Spatial Data

Extending Poverty Estimates to Non-Covered Areas

Despite initial intentions to expand the HFS across the entire country, continued insecurity made it impossible to extend the survey to the former North-Eastern states of Jonglei, Unity, and Upper Nile. To account for this gap in coverage and obtain countrywide poverty rates, a statistical model was developed to impute poverty in non-covered areas leveraging the growing availability of satellite imagery and geo-spatial data. Recent advances in the processing and availability of satellite imagery and geo-spatial data have led to a growing field of research on predicting a range of outcomes based on diverse such data sources.[36] Indeed, there is a growing body of evidence indicating that household-survey derived indices of poverty correlate strongly with many geographic features that can be observed from space or derived from ground-based data (Engstrom et al., 2017; Jean et al., 2016; Krizhevsky et al., 2012; Sedda et al., 2015).

[36] An organization called Planet currently operates more satellites than even the U.S. and Russian governments. Planet recently launched 88 additional satellites, allowing almost daily coverage of the entire globe with a resolution of 3 to 5 meters per pixel (Engstrom et al., 2017a).

One of the earlier applications of the use of satellite and geo-spatial data to predict outcomes was the use of night-time lights to predict GDP. Night-time lights are well-suited to predicting cross-country levels of GDP (Henderson et al., 2012; Pinkovskiy and Sala-i-Martin, 2016). However, at the within-country level they are much better suited to predicting population density than welfare, and the correlation of night-time lights with local wages and local poverty rates has typically been found to be weak (Engstrom et al., 2017a; Mellander et al., 2015). Night-time lights may therefore not be very well suited to uses akin to small-area estimation, particularly in a place such as South Sudan where only about 3 percent of households have access to electricity (World Bank, 2018). More recent research has focused on training deep-learning algorithms to extract a diverse range of features from high resolution satellite imagery, for example counting the number of cars on a street, distinguishing road types, recognizing materials roofs are made of, tree coverage, the contrast and number of jagged edges, etc. (Engstrom et al., 2017). This allows making poverty predictions at a much higher level of disaggregation (Jean et al., 2016; Krizhevsky et al., 2012; Sedda et al., 2015). Engstrom et al. (2017) provide a useful overview of the current state of the literature and show the predictive power of a range of indicators constructed from satellite data in estimating poverty at the village-level.

In the case of the HFS in South Sudan, predictions from a set of linear models were used to project poverty estimates to inaccessible areas based on already extracted satellite features and geo-spatial data, given the objective of creating reliable and transparent poverty measures. The poverty imputation follows a process that is relatively similar to small area estimation, though only the point estimates were estimated and not higher moments of the outcome distribution (see for example: Elbers Chris et al., 2003; Guadarrama et al., 2016; Haslett, 2016). Poverty as measured in the 2016 wave of the HFS is regressed on a range of geo-spatial characteristics such as distance to urban centers, distance to the electricity grid, annual rainfall, annual temperatures, urban-rural status, IPC phase classification, and others. The estimated model is then used to calculate expected poverty rates across regions where the household survey data are not available, but where the geo-spatial data are available. Poverty rates are predicted at the 100m*100m level across South Sudan. The poverty estimates then need to be weighted by local population counts to eliminate potential bias caused by vast uninhabited areas. Given the lack of reliable administrative data on settlements or population counts, local populations were in turn estimated using a set of covariates derived from geo-referenced data such as urbanicity, roads, clinics, and buildings.

Estimating Settlements Data

The aggregation of poverty estimates to the county and state levels needs to be calibrated against suitable population estimates. Naively aggregating poverty rates across broad geographic regions would result in extremely high poverty rates given the vast uninhabited expanses isolated from the rest of the country, in which a model would likely predict high poverty rates. Indeed, South Sudan is sparsely populated relative even to most other large African countries, in 2008 South Sudan had a population density of approximately 13 persons per kilometer squared compared to the Sub-Saharan Africa average of 35.[37] Because an accurate high-resolution map of population density is not available for South Sudan, the spatial distribution of settlements was used as a proxy for population density in order to calculate weights with which to weight poverty estimates.

The methodology had to employ a novel process to generate estimates of settlements given the absence of more recent and up to date population data since the 2008 Census. This process was based on a wide variety of data sources and variables associated with population density, leveraging varied sources of data such as open source data from Open Street Maps on residential areas, roads, health facilities, schools, data from the Global Urban Footprint project, as well as data form the survey itself. The map of settled areas in South Sudan was built by processing and regrouping the data sets (Table 9 in Appendix E). The map of settled areas was created as a binary map (1=settled, 0=not settled) at 100m resolution. While drawing the map, the data sets were manually checked against Google Satellite imagery for the presence of settlements. One advantage of this system of estimation for settlements is that each component can be updated independently as new data become available or the situation within the country

[37] World Development Indicators.

changes. Finally, the map of settlements was adjusted for displacement and for the locations of IDP camps, given extreme rates of displacement in South Sudan.

Other variables were tested but not used for the creation of the map of settled areas (Table 9). This includes night-time lights, which are commonly used in studies predicting outcomes from satellite data (Mellander et al., 2015; Pinkovskiy and Sala-i-Martin, 2016). However, given that only about 3 percent of households in South Sudan have access to a stable source of electricity, there is very little variation to exploit in trying to identify within-country correlations between deprivation and electric light (World Bank, 2018). Indeed, night time lights would only really predict small industrial enclaves such as oil fields and did not accurately capture where the population actually lives.

Figure 8: Urban (red) and rural (blue) settlements.

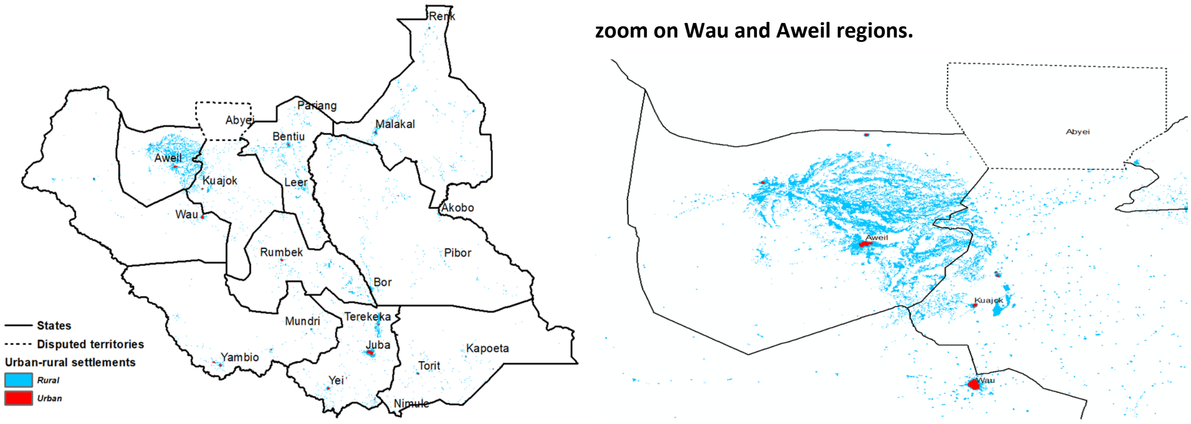

Source: Flowminder / WorldPop.

An ‘urban gradient’ variable was also derived from the map of settled areas. This estimation was based in large part on the distance to major roads and the wave 1 and wave 3 survey points labelled as ‘urban’, i.e. the urban classification of enumeration areas based on the 2008 Census exercise. Each 100x100m pixel was classified as a city, city extent, town, town extent, large village, small village, villages far from major roads and unsettled. Distinction between villages and towns was primarily based on the presence of major road intersection and settlement size. A simpler urban/rural settlements map was also produced with only 3 classes: unsettled, rural, urban (towns and cities). All HFS survey points labelled as ‘urban’ fall in the urban category. Finally, a map of ‘distance to urban centers’ was created based on the generated urban/rural settlements map.

Variable Selection and Model Estimation

Many variables were tested for correlation against each household’s probability of being poor averaged per EA. Given that the variance of the probability of being poor was greater across EAs than within EAs, the choice was made to average the probability of poverty per EA. In this manner, a greater degree of spatial variation could be observed, thus increasing the potential to observe meaningful correlations between the probability of poverty and the predictors, i.e. the geo-spatial variables. The variables tested included more traditional geo-spatial characteristics that are commonly used in such applications, such as average temperatures, average rainfall, annual cloud cover variation and annual cloud cover (Table 11). It also tested determinants of public services provision and proxies for distance to economic activity, such as distances to different types of roads, urban centers, the electricity grid cultivated areas, schools, and water bodies. Finally, a set of variables indicative of the crisis were used, such as the number of people in need as calculated by OCHA, the IPC phase classification, and the number of conflict fatalities as collected by the Armed Conflict Location Events Data between 2011-16 and between 2014-2016. Finally, the various urban gradients calculated in the previous step were also tested for correlation with poverty rates.

A dummy variable was added for the capital city, Juba, and the former state of Western Equatoria because no variable tested alone could explain the lower levels of poverty observed in Western Equatoria or Juba. The urban gradient alone provided little predictive power as other large towns such as Wau had very high average poverty rates. Therefore, a spatial variable indicating Western Equatoria and Juba was created, with its values smoothed for 200km across the WEQ border and smoothed 2km around the city center of Juba. The resulting map takes the value of 1 in Western Equatoria and in the Juba center, the value of 0 outside these two regions, and a gradient of values between 0 and 1 across its border. This variable does not help explain variation in poverty, but merely reflects observations from the survey and helps to account for chance correlations in the prediction. In other words, this avoids predicting low poverty in the entire western part of the country based on the low poverty rates observed around Western Equatoria and Juba. Of the variables having a relatively large correlation with poverty, some are redundant, some are due to 'chance' as explained above – and some show a trend both within Western Equatoria / Juba and in the rest of the country and hence are deemed as reliable correlations.

Figure 9: Example maps of variables used in the estimation.

IPC phase classification in January 2017.

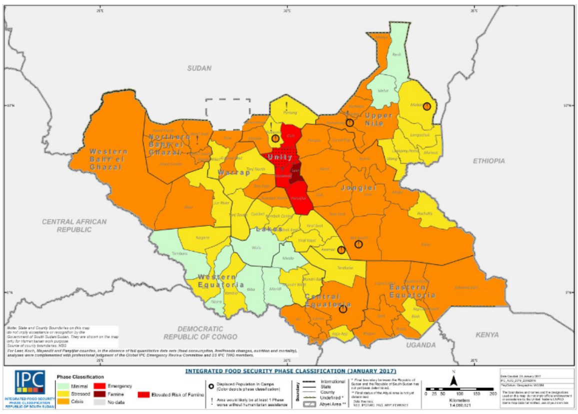

Distance to electricity grid.

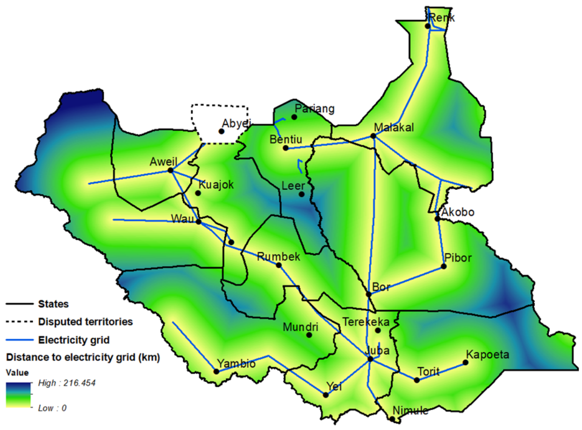

Annual average temperatures.

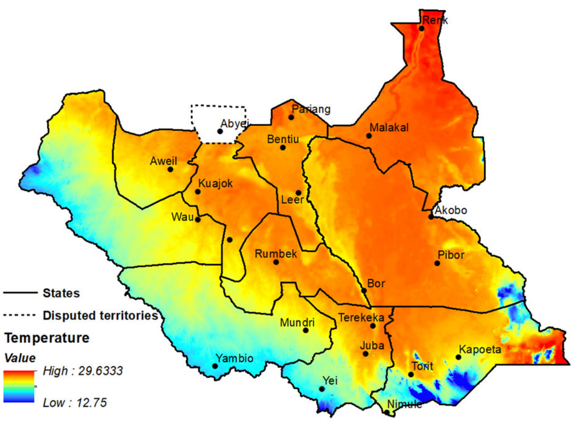

Annual average precipitation.

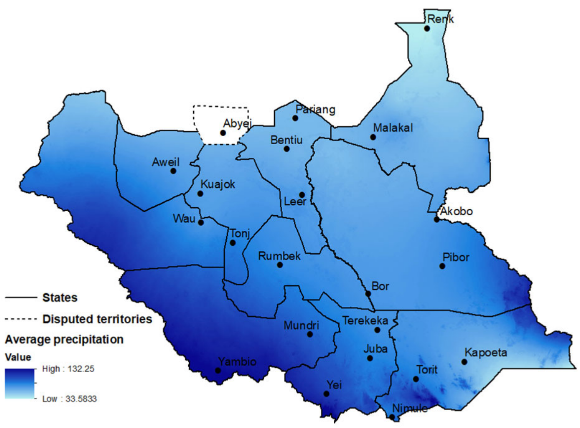

Source: Flowminder/WorldPop using data from IPC Info, WorldClim, and Africa Infrastructure Country Diagnostic (AICD).

While each of the covariates described above provides some level of predictive power for poverty, a combination of non-orthogonal variables is more likely to better predict poverty. Because of the relatively small number of enumeration areas used in this study (156), focus was placed on a simple linear model. Furthermore, comparisons against polynomial and more complex models indicated that a linear model retained the largest R^2 (=0.7). The level of predictive power was confirmed using an out-of-sample cross validation. In the cross-validation exercise the model was first built using 75 percent of the survey data. Then, the remaining 25 percent was used to predict EA-level poverty values and check the predictive power of the model, therefore confirming the efficiency and validity of the results. The cross-validation approach was performed 10 times and the average predictive power was used. The following variables were selected in the final model: the IPC phase classification, distance to urban centers, annual average temperature, distance to the electricity grid, annual average precipitation, an urban/rural/unsettled dummy, and a dummy for Juba and Western Equatoria (Table 12).

Results

Imputing poverty headcount ratios in the states not covered by the HFS based on satellite and geo-spatial data indicate potentially extremely high levels of poverty in those regions as well. Estimating poverty for every kilometer squared across South Sudan results in the map shown in Figure 10. The poverty map obtained reflects the variations of the in WEQ or Juba variable (lower poverty in WEQ and Juba), and variations of the IPC phase (e.g. North East). The influence of the Distance to urban centers can be seen e.g. around Raga (North West town), and the distance to the electricity grid can also be seen but to a lesser extent. Influence of temperature and precipitations can be seen along the Nile and in the South East. At a smaller geographic scale predicted poverty follows the urban/rural/unsettled classification (Figure 8). The weighted poverty rates indicate extremely high poverty rates in the Greater Upper Nile regions, which is expected given the predominantly rural nature of the region and its state of instability. The poverty headcount across almost all the non-covered states reaches upwards of 9 in 10. Therefore, based on the trends depicted in Table 13, the extent of deprivation has reached extremely high levels throughout almost the entire country except for Western Equatoria.

Figure 10: Poverty maps, headcount FGT(0) in 2016.

Aggregate per state – imputation in non-HFS states **Poverty predictions at 100mx100m level**

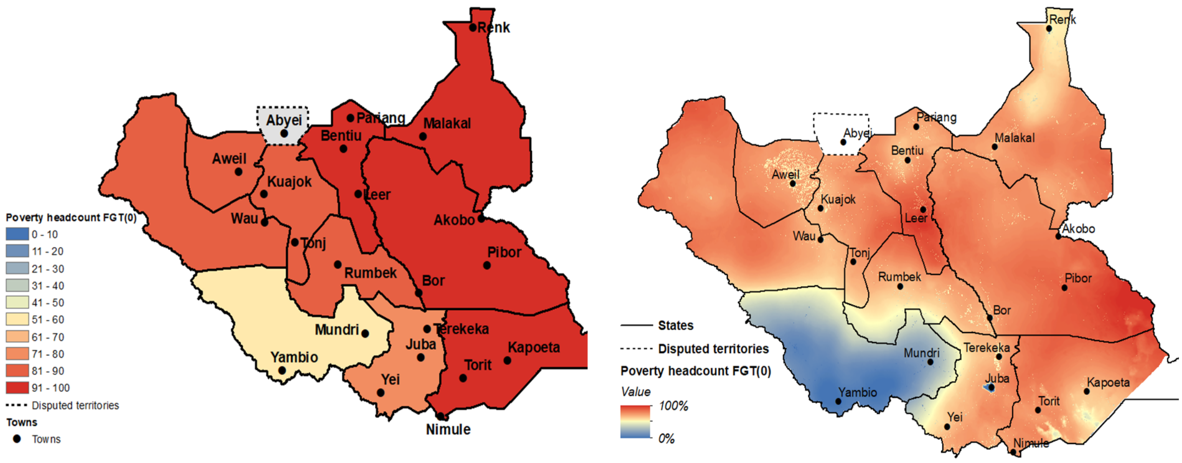

Source: Authors' own calculations based on HFS 2016 data and Flowminder / WorldPop computations.

Limitations

The results presented here are an attempt to make the best use of available data given a number of limitations. Firstly, no spatial random effect was used in the present model largely due to the fact that EAs were mostly sampling in a North-West / South-East gradient, with little information available on the East-West spatial structure. In the present case, geographic covariates have provided sufficient predictive power that this lack of spatial autocorrelation is not necessarily an issue. However, further data from other regions in the country would provide significant advantages for defining this spatial random component. A related issue is the use of spatially smooth predictors, for example the distance to urban centers and the distance to the electricity grid. Such variables are informative especially with respect to their impact on poverty and can be better predictors than binary variables indicating access based on a cutoff might be. However, they also can have difficulty predicting "pockets of poverty" sitting in otherwise wealthier areas, for example slums in urban areas, or the converse. This could exacerbate the spatially smooth predictions already introduced by the assumption of constant coefficients from the linear regression. Unfortunately, the impact this may have had on the estimation is difficult to test using cross-validation with survey data that was designed to be widely distributed geographically. Therefore, it is impossible to test what the share of variation in welfare across EAs is dampened by the use of these spatially smooth predictors. This is an area which warrants future research given the predictive power of such variables and could be better tested using data collected more finely over large areas such as a Census.

Secondly, there is a very poor understanding of the population distribution in South Sudan and no reliable sampling frame against which to extrapolate our predictions. The implications of this are that while the model can predict into geographic pixels based on the existing data, it is difficult to aggregate by county without knowing how to weight each pixel according to the population present within it. Thus, poverty maps aggregated by area are likely to over-estimate poverty rates as most areas within each county are likely to have lower population density and high poverty. The solution to this problem is to define a new sampling frame for the country, then re-calculate county-level predictions based on this sampling frame. This was attempted in this study by estimating an urban gradient based on multiple data sources and their relationship with urbanicity. However, some of these data are likely to be out of date for many of the same reasons that a traditional Census exercise is complicated. The rapid and enormous movement of people caused by the conflict is likely to have compounded this problem. Building newer and more up to date population sample frames should be a priority for researchers interested in South Sudan. This could be achieved either by conducting a traditional census, or by leveraging the recently available satellite imagery using and machine-learning based methods to extract features. Such extract could be used to help define settled areas and their associated population density to create a predictive population surface (Engstrom et al., 2017a; Jean et al., 2016; Pasquale et al., 2017). Based on this, new sample frames can be built to use for future data collection work, which is badly needed in the context of South Sudan.

The model structure was voluntarily kept simple (linear combination) to ease its interpretation given that it was constructed as a proof of concept to show the potential of spatial data for imputing poverty to supplement poverty survey estimates. Furthermore, although where these techniques may have the most value, which is where there might have been a crisis or emergency or where safety is a concern, these techniques are also the most difficult to apply. Indeed, the link between poverty and such variables is much more likely to be structural than transient across much of South Sudan. Indeed, a set of issues that arise in this estimation method is the difficulty of modeling the dynamics of poverty and shocks. Many of the areas where the enumerators could not go were inaccessible because of recent conflict and it is difficult to account for this in a cross-sectional model as such, given the potentially endogenous nature of conflict and poverty whereby some conflict events are concentrated around wealthier areas. One of the areas for future research might be to leverage the time series that area available for various types of geo-spatial data to try to account for some of these dynamics relating poverty rates to shocks and imbalances.

6. Conclusion

The HFS conducted several rounds of data collection at a time of upheaval in the short history of South Sudan. In particular, Wave 3 of the HFS consisted of a major data collection effort during what effectively became one of the deepest humanitarian crises in recent history. The HFS was conceived within the context of the crisis and was therefore designed to leverage new technologies for monitoring and implementation as well as methodological innovations in survey design. This allowed the HFS team to monitor closely the survey and facilitate the implementation while facing a multitude of challenges induced by the escalating crisis. Unfortunately, the growing intensity of the conflict eventually led to a shortened survey with deprived sample size. In the end, after almost 9 months of fieldwork, only about one-half of the intended sample of households was interviewed. While the disruptions caused by the conflict have had impact on the data collected, consistency checks suggest that this impact was relatively small. In addition, any introduced sample selection bias due to the conflict is likely to be a downward bias leading to under-estimation of poverty.

The HFS presents a rare data point in a fragile setting. Only very few similar surveys have managed to collect comprehensive data on welfare and livelihoods in such a complicated and volatile context. Indeed, the HFS documents some staggering levels of deprivation, which are also corroborated by accounts from a multitude of organizations operating in the country. The methodology employed to estimate poverty in the HFS is based on the best available methodologies specifically adapted to the context of fragility. The estimation is also entirely reproducible through the publicly available code and data published in the World Bank MicroData Library.[38] Overall, the HFS provides an extremely detailed picture of welfare and livelihoods for the South Sudanese population between 2015 and 2017. This is especially true when combined with the other three waves conducted between 2015 and 2017, as in the South Sudan Poverty Assessment (World Bank, 2018).

Finally, the satellite imputation, although limited in scope and means, provides an additional glimpse of livelihoods across the country. Although the results are only a proof-of-concept, it remains a useful exercise to complement the survey-based data rather than assuming a national average for inaccessible areas. Much research has already gone into the field of small area estimation, which is likely to benefit enormously from the recent availability of cheaper and more encompassing – geospatial – data sets. Although such models are not likely to replace survey data, as these are needed to train the models, they can be used to supplement data collection and provide information either at more frequent intervals or for hard-to-reach areas. One particular area for future research that might be especially relevant would be to explore how such sources of data can be leveraged to estimate outcomes during rapidly evolving and dynamic events, exactly when representative surveys and other traditional data collection exercises are especially difficult to implement.

[38] See: http://microdata.worldbank.org/index.php/catalog/2914

References

African Development Bank, 2013. South Sudan Infrastructure Action Plan - A Program for Sustained Strong Economic Growth.

Amendola, N., Vecchi, G., 2014. Durable Goods and Poverty Measurement, Policy Research Working Papers. The World Bank. https://doi.org/10.1596/1813-9450-7105

Balcazar, C.F., Ceriani, L., Olivieri, S., Ranzani, M., 2014. Rent Imputation for Welfare Measurement: A Review of Methodologies and Empirical Findings. Policy Research Working Paper, World Bank 40.

Beegle, K., Christiaensen, L., Dabalen, A., Gaddis, I., 2016. Poverty in a Rising Africa. The World Bank. https://doi.org/10.1596/978-1-4648-0723-7

Beegle, K., De Weerdt, J., Friedman, J., Gibson, J., 2012. Methods of household consumption measurement through surveys: Experimental results from Tanzania. Journal of Development Economics 98, 3–18.

Caeyers, B., Chalmers, N., De Weerdt, J., 2012. Improving consumption measurement and other survey data through CAPI: Evidence from a randomized experiment. Journal of Development Economics 98, 19–33. https://doi.org/10.1016/j.jdeveco.2011.12.001

Christiaensen, L., Lanjouw, P., Luoto, J., Stifel, D., 2010. The Reliability of Small Area Estimation Prediction Methods to Track Poverty.

De Leeuw, E.D., Hox, J.J., Snijkers, G., 1995. The effect of computer-assisted interviewing on data quality. A review. Market Research Society. Journal. 37, 1–19.

de Vries, L., Schomerus, M., 2017. Fettered Self-determination: South Sudan's Narrowed Path to Secession. Civil Wars 19, 26–45.

de Waal, A., 2014. When kleptocracy becomes insolvent: Brute causes of the civil war in South Sudan. Afr Aff (Lond) 113, 347–369. https://doi.org/10.1093/afraf/adu028

Deaton, A., Zaidi, S., 2002. Guidelines for Constructing Consumption Aggregates For Welfare Analysis 108.

Douidich, M., Ezzrari, A., Van der Weide, R., Verme, P., 2013. Estimating Quarterly Poverty Rates Using Labor Force Surveys: A Primer, Policy Research Working Paper. World Bank.

Elbers Chris, Lanjouw Jean O., Lanjouw Peter, 2003. Micro–Level Estimation of Poverty and Inequality. Econometrica 71, 355–364. https://doi.org/10.1111/1468-0262.00399

Engstrom, R., Hersh, J., Newhouse, D., 2017a. Poverty from space: using high-resolution satellite imagery for estimating economic well-being. The World Bank.

Engstrom, R., Newhouse, D., Haldavanekar, V., Copenhaver, A., Hersh, J., 2017b. Evaluating the relationship between spatial and spectral features derived from high spatial resolution satellite data and urban poverty in Colombo, Sri Lanka, in: 2017 Joint Urban Remote Sensing Event (JURSE). Presented at the 2017 Joint Urban Remote Sensing Event (JURSE), pp. 1–4. https://doi.org/10.1109/JURSE.2017.7924590

Fafchamps, M., Mckenzie, D., Quinn, S., Woodruff, C., 2010. Using PDA consistency checks to increase the precision of profits and sales measurement in panels. https://doi.org/10.1016/j.jdeveco.2010.06.004

FAO, WFP, 2017. Crop and Food Security Assessment.

Foster, J., Greer, J., Thorbecke, E., 1984. A Class of Decomposable Poverty Measures. Econometrica: journal of the econometric society 761–766.

Fujii, T., Van der Weide, R., 2013. Cost-effective estimation of the population mean using prediction estimators. The World Bank.

Guadarrama, M., Molina, I., Rao, J.N.K., 2016. A COMPARISON OF SMALL AREA ESTIMATION METHODS FOR POVERTY MAPPING. Statistics in Transition. New Series 17, 41–66. https://doi.org/10.21307/stattrans-2016-005

Haslett, S.J., 2016. Small Area Estimation Using Both Survey and Census Unit Record Data. Analysis of Poverty Data by Small Area Estimation 325–348.

Henderson, J.V., Storeygard, A., Weil, D.N., 2012. Measuring economic growth from outer space. American economic review 102, 994–1028.

International Monetary Fund (IMF), 2016. Article IV Consultation Report, February 2017 (Art. IV). World Bank Group.

Jean, N., Burke, M., Xie, M., Davis, W.M., Lobell, D.B., Ermon, S., 2016. Combining satellite imagery and machine learning to predict poverty. Science 353, 790. https://doi.org/10.1126/science.aaf7894

Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. Imagenet classification with deep convolutional neural networks, in: Advances in Neural Information Processing Systems. pp. 1097–1105.

Mellander, C., Lobo, J., Stolarick, K., Matheson, Z., 2015. Night-Time Light Data: A Good Proxy Measure for Economic Activity? PLOS ONE 10, e0139779. https://doi.org/10.1371/journal.pone.0139779

Mneimneh, Z.N., Pennell, B.-E., Kelley, J., Hibben, K.C., 2016. Surveys in societies in turmoil. The SAGE Handbook of Survey Methodology 178.

Olson Lanjouw, J., Lanjouw, P., 2001. How to compare apples and oranges: Poverty measurement based on different definitions of consumption. Review of Income and Wealth 47, 25–42.

Pape, U., Benson, M., Ebrahim, M., Lole, J., 2017. Reducing poverty through improved agro-logistics in a fragile country: findings from a trader survey in South Sudan. World Bank Group, Washington, D.C.

Pape, U., Mistiaen, J., 2018. Household Expenditure and Poverty Measures in 60 Minutes: A New Approach with Results from Mogadishu. World Bank Policy Research Working Paper 25.

Pape, U., Mistiaen, J., 2015. Measuring household consumption and poverty in 60 minutes: The Mogadishu high frequency survey. World Bank.

Pape, U., Mistiaen, J., 2014. Utilizing Mobile Technology to Innovate CAPI Data Collection in Fragile Contexts 14.

Pape, U.J., Dihel, N.C., 2017. South Sudan - economic update : taming the tides of high inflation - policy options. World Bank, Washington, D.C.

Pasquale, A.D., McCann, R.S., Maire, N., 2017. Assessing the population coverage of a health demographic surveillance system using satellite imagery and crowd-sourcing. PLOS ONE 12, e0183661. https://doi.org/10.1371/journal.pone.0183661

Pinkovskiy, M., Sala-i-Martin, X., 2016. Lights, Camera… Income! Illuminating the national accounts-household surveys debate. The Quarterly Journal of Economics 131, 579–631.

Ravallion, M., 2017. Poverty comparisons. Routledge.

Ravallion, M., 1998. Poverty lines in theory and practice. The World Bank.

Sedda, L., Tatem, A.J., Morley, D.W., Atkinson, P.M., Wardrop, N.A., Pezzulo, C., Sorichetta, A., Kuleszo, J., Rogers, D.J., 2015. Poverty, health and satellite-derived vegetation indices: their inter-spatial relationship in West Africa. International Health 7, 99–106. https://doi.org/10.1093/inthealth/ihv005

Southern Sudan Center for Census, Statistics, and Evaluation, 2010. Southern Sudan Counts: Tables from the 5th Sudan Population and Housing Census, 2008.

World Bank, 2018. Impact of Conflict and Shocks on Poverty: South Sudan Poverty Assessment, 2017.

World Bank, 2017. South Sudan - Governance analysis : building sustainable public sector capacity in a challenging context.

World Bank, 2016. South Sudan Poverty Profile: Findings from the 2015 wave of the High Frequency South Sudan Survey.

APPENDICES

APPENDIX A

The High Frequency Survey in South Sudan

The High Frequency Survey conducted waves of almost nationally representative surveys across South Sudan between 2015 and 2017. The HFS was based on a pilot which collected six waves of panel data across 4 of the largest urban centers between 2012 and 2014. The pilot was then scaled up in 2015 to a representative wave covering 6 of the 10 former states of South Sudan. Between 2015 and 2017, the HFS was expanded to a seventh state and conducted three more waves. Waves 2 and 4 were limited to urban areas but included a panel component. The HFS was accompanied by market price surveys which collected weekly price data and daily exchange rate data in 17 locations across the entire country.

Table 2: Dates and sample for data collection for all four waves of the HFS

	Wave 1 Feb.-Oct.2015			Wave 2 Feb.-Apr.2016	Wave 3 Sep.2016-Feb.2017			Wave 4 May-Jul.2017
EAs/HH	**Rural**	**Urban**	**Total**	**Urban**	**Rural**	**Urban**	**Total**	**Urban**
Warrap		-		**15/173**	8/95	5/40	**13/135**	**15/144**
Northern Bahr El Ghazal	40/480	10/120	**50/600**	**15/177**	20/239	5/60	**25/299**	**15/126**
Western Bahr El Ghazal	20/225	30/360	**50/585**	**11/126**	14/166	12/144	**26/310**	**15/137**
Lakes	40/478	10/120	**50/598**	**15/180**	19/172	5/60	**24/232**	**15/133**
Western Equatoria	34/406	16/192	**50/598**	**15/176**	18/216	7/84	**25/300**	**15/156**
Central Equatoria	16/192	34/408	**50/600**	**15/177**	16/192	10/119	**26/311**	**15/95**
Eastern Equatoria	40/453	10/116	**50/569**	**15/180**	20/201	5/60	**25/261**	**15/153**
Total	190/2,234	110/1,316	**300/3,550**	**101/1,189**	115/1,281	49/567	**164/1,848**	**105/944**

Source: HFS 2015-17.

The fourth wave of the HFS was accompanied by the Crisis Recovery Survey (CRS), a representative survey of four of the largest IDP camps in South Sudan. The CRS was conducted simultaneously to Wave 4 of the HFS in mid-2017. It covered the four largest protection of civilian (PoC) camps with well-defined boundaries accessible to enumerators. The camps include Bentiu PoC, Bor PoC, Juba PoC1 and 3, and Wau PoC. Although the CRS covers PoCs, where only 12 percent of South Sudan's IDPs are located, the detailed microdata fill important information and knowledge gaps for IDP-focused programming.

The HFS and CRS questionnaires cover a large range of topics and draw a well-rounded picture of socio-economic livelihoods of people in South Sudan. The HFS questionnaire covers topics including demographics, employment, education, consumption, as well as perceptions of well-being and of the effectiveness of public institutions. Consumption is measured using the newly developed rapid consumption methodology. The CRS and Wave 4 HFS questionnaires, designed to be exactly comparable, also collected details on displacement-specific outcomes guided by the IASC framework.[39] These were developed to understand the motivations for displacement, return intentions, sense of security, relations with the surrounding community, social capital, and pre-displacement outcomes in the standard of living, education and labor.

[39] The Inter Agency Standing Committee Framework on Durable Solutions for Internally Displaced Persons aims to provide guidance for achieving durable solutions following internal displacement in the context of armed conflict, situations of generalized violence, violations of human rights and natural or human-made disasters. The Framework primarily aims to help international and non-governmental actors to better assist governments dealing with humanitarian and development challenges resulting from internal displacement. The Framework is also designed so that it can be used to assist those in the field in determining whether a durable solution to internal displacement has been found, depending on the context of the local environment.

Figure 11: High Frequency Survey coverage, 2015-2017.

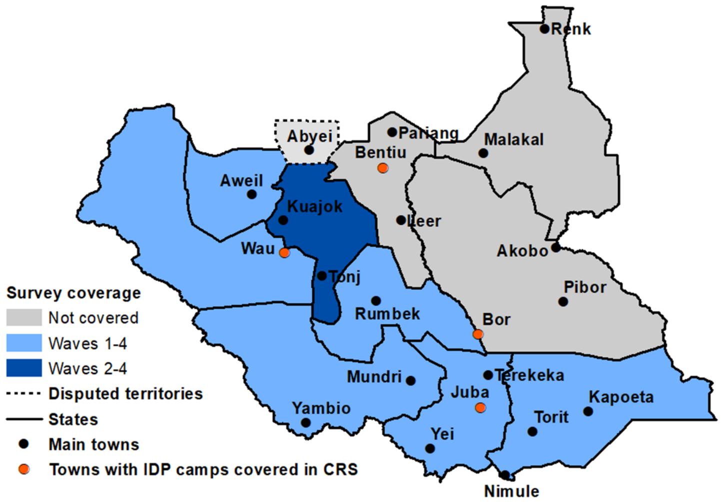

Source: HFS 2015-17 and CRS 2017.

APPENDIX B

Table 3: Sample design calculations.

	No. HH (Census)	Urban (%)	Mean (Cons.)	std dev	equal/optimal (10 min) Urban EAs	Rural EAs	rel. err.
Central Equatoria	175,962	31.2%	133.0	90.0	13	37	0.032
Eastern Equatoria	151,199	9.9%	107.3	80.2	10	40	0.045
Western Equatoria	115,595	17.1%	126.1	99.9	13	37	0.028
Warrap	167,654	7.6%	73.3	49.8	12	38	0.043
Western Bahr El Ghazal	57,487	44.7%	122.1	144.6	33	17	0.029
Northern Bahr El Ghazal	130,832	6.3%	61.1	52.1	10	40	0.049
Lakes	90,315	7.2%	119.3	119.0	10	40	0.019
Rural	746,136	--	94.3	74.0	--	249	0.003
Urban	142,908	--	152.4	155.1	101	--	0.098
Total	889,044	16.1%	103.5	90.1	101	249	0.026

Source: Authors' own calculations based on NBHS 2009 data.

- **Sampling weights**

Sampling weights are used to make survey observations representative for the sample. The sampling weight is the inverse probability of selection. The selection probability P for a household can be decomposed into the selection probability P_1 of the EA and the selection probability P_2 of the household within the EA:

(1) $$P = P_1 P_2$$

The selection probability P_1 of an EA k is calculated as the number of households within the EA divided by the number of households within the stratum multiplied by the number of selected EAs in the stratum:

(2) $$P_1 = \frac{|K|\hat{n}_k}{\sum_{k' \in K} \hat{n}_{k'}}$$

where $\hat{n}_k$ denotes the number of households in EA k estimated using the Census 2008 data and K is the set of EAs selected in the corresponding stratum. Replacement enumeration areas were assigned the sampling weight of the enumeration area that they were replacing. In Wave 3, the number of enumeration areas surveyed in each stratum differed from the original sample. The weights were therefore scaled to correct for the change in the value of K.

The selection probability P_2 for a household within an EA k is constant across households and can be expressed as:

(3) $$P_2 = \frac{|H|}{n_k}$$

where $|H|$ is the number of households selected in the EA and n_k denoting the number of listed households in EA k. Usually, the number of households per EA is 12 while a few exceptions exist due to invalid interviews.

Sampling weights were scaled to equal the number of households per strata using the Census 2008 data. Thus, the sampling weight W can be written as:

(4) $$W = \frac{c}{P} \text{ with } c = \frac{\sum_{k \in K} \hat{n}_k}{\sum_{k \in K} n_k}$$

Table 4: No. of enumeration areas per strata, 2016.[40]

	Intended			Actual		
	Rural	**Urban**	**Total**	**Rural**	**Urban**	**Total**
Warrap	37	13	50	8	5	13
Northern Bahr El Ghazal	40	10	50	20	5	25
Western Bahr El Ghazal	37	13	50	14	12	26
Lakes	38	12	50	19	5	24
Western Equatoria	17	33	50	18	7	25
Central Equatoria	40	10	50	16	10	26
Eastern Equatoria	40	10	50	20	5	25
Total	**37**	**13**	**350**	**115**	**49**	**164**

Source: Authors' own calculations based on HFS 2016 data.

[40] Note that the date of data collection refers to the period when most of the interviews were collected. In some cases, a few interviews were conducted in the month after the end of fieldwork as part of follow-ups to improve data quality.

APPENDIX C

- **Cleaning consumption data**

Food expenditure data are cleaned in a three-step process. First, units for reported quantities of consumption and purchase are corrected. Second, quantities consumed and purchased converted into kilograms are cleaned, where potential data entry errors and outliers are detected and corrected. Third, prices per kilogram calculated using the cleaned quantities are corrected in a similar manner. The cleaning rules were maintained across the 4 survey waves to ensure comparability. More details on the specific cleaning rules are provided below:

- Rule 1 (data entry errors for units): For records that have the same figure in quantity purchased and consumed but have different units, it is assumed that the correct unit is the one that takes the quantity (consumed or purchased, converted into kilograms) closer to the weighted median value for the same item.

	N	%
Not-tagged	14,818	99.5
Tagged	70	0.5
Total	14,888	100

- Rule 2 (mistakes in reported units): Items that are likely to be reported in the wrong unit are corrected following generic rules. An example of a typical mistake is to report consumption of 100 kilograms of a product (like salt) where the supposed correct unit is grams. In this case, all quantities given in kilograms that exceed 10s0 would be corrected so as to be given in grams instead.

Cons. Q.	N	%
Not-tagged	14,871	99.9
Tagged	17	0.1
Total	14,888	100

Purc. Q.	N	%
Not-tagged	14,507	97.4
Tagged	381	2.6
Total	14,888	100

- Rule 3 (missing quantities): Items that were consumed but have a missing quantity, consumed or purchased, are replaced with the item-specific median quantity.

Cons. Q	N	%
Not-tagged	12,851	86.3
Tagged	2,037	13.7
Total	14,888	100

Purc. Q.	N	%
Not-tagged	13,211	88.7
Tagged	1,677	11.3
Total	14,888	100

- Rule 4: (quantities beyond 'hard' constraints): Quantities consumed and purchased that are below or above the item-unit quantity constraints are replaced with the item-specific median.

NONE

- Rule 5 (data entry errors for quantities or prices): Records with the same value for quantity consumed or quantity purchased and price, or with the same value for all three, are assumed to have a data entry error in the price or quantity. They are replaced with the item-specific medians.

	N	%
Not-tagged	14,859	99.8
Tagged	29	0.2
Total	14,888	100

- Rule 6 (quantities per capita too high): For items consumed by more than 300 households, quantities that were 3 standard deviations above the mean value per capita were replaced with item-specific medians.

Cons. Q	**N**	**%**
Not-tagged	14,757	99.1
Tagged	131	0.9
Total	14,888	100

Purc. Q.	**N**	**%**
Not-tagged	14,780	99.3
Tagged	108	0.7
Total	14,888	100

- Rule 7 (missing prices): Items that were consumed but have zero or missing prices are replaced with the item-specific median prices. The reason why this is so high is because many households obtained much of the food consumed from home production, and thus could not answer when asked the price at which they purchased these goods.

	N	**%**
Not-tagged	11,715	78.7
Tagged	3,173	21.3
Total	14,888	100

- Rule 7 (price outliers): Prices in the item-specific price per kilogram distribution that lie above the 95th percentile are replaced with item-specific medians, so are prices for items consumed by more than 300 households that lie above 3 standard deviations above the mean.

Hard constraints	**N**	**%**
Not-tagged	14,531	97.6
Tagged	357	2.4
Total	14,888	100

3 sd	**N**	**%**
Not-tagged	13,885	93.3
Tagged	1,003	6.7
Total	14,888	100

All medians are estimated at the EA level if a minimum of 5 observations are available. If the minimum number of observations is not met, weighted medians are estimated at the strata-level requiring a minimum number of 10 observations before proceeding to the item level. Medians are estimated excluding zero values and tagged values so as not to replace reported values with zeroes or invalid values.

The non-food data set only contains price values without quantities and units, the cleaning process was therefore much simpler. Two cleaning rules are applied and tagged observations are replaced with item-specific medians at the EA, state, and survey level as is done for food consumption. The cleaning rules are the following:

- Rule 1 (price outliers): Prices that are beyond the hard constraints, above or below, are replaced with item-specific medians. Given the high inflation over the subsequent HFS waves, the value of the hard constraints used in Wave 1 were adjusted for inflation using the national NBS CPI.

Max	**N**	**%**
Not-tagged	10,864	94
Tagged	689	6
Total	11,553	100

Min	**N**	**%**
Not-tagged	10,969	94.9
Tagged	584	5.1
Total	11,553	100

- Rule 2 (zero or missing prices): Zero and missing prices for consumed items are replaced with item-specific medians.

Zero	N	%
Not-tagged	11,310	97.9
Tagged	243	2.1
Total	11,553	100

Missing	N	%
Not-tagged	10,862	94
Tagged	691	6
Total	11,553	100

The medians are calculated following exactly the same process as in food cleaning. All medians are estimated at the EA level if a minimum of 5 observations are available. If the minimum number of observations is not met, weighted medians are estimated at the strata-level requiring a minimum number of 10 observations before proceeding to the item level. Medians are calculated excluding zero values and tagged values so as not to replace reported values with zeroes or invalid values.

For durables, the cleaning process involved cleaning ownership statistics as well as the calculated depreciation rates. The quantity of an item is replaced by the item-specific survey median (due to paucity of data) if the reported quantity is unrealistically high assessed by manual inspection. The purchase value of durables is recorded in the year and currency of purchase. Outliers of purchase values in the reported currency are identified by hard constraints and replaced by the item-specific survey median. Items with at least 3 observations purchased in the same year are replaced by the respective item-year specific median. Alternatively, the item-state-level median prices are used if at least 5 observations are given. Hypothetical selling prices are replaced by the item-state level median if at least 5 observations are available. Without the minimum number of observations available, the item-specific median is used. All prices reported in foreign currencies are converted into SSP through conversion to USD.

- Rule 1 (quantity outliers): Quantities above 100 units of an asset are replaced with the item-specific median.

	N	%
Not-tagged	5,007	99.9
Tagged	5	0.1
Total	5,012	100

- Rule 2 (price outliers): (i) Prices above hard constraints are replaced with the item-specific median. (ii) For specific assets where outliers are identified that fall below the hard constraints and for which we have enough observations to estimate a distribution, the top 5 percent of observations are replaced with item-specific medians.

Selling Above	N	%
Not-tagged	5,004	99.8
Tagged	8	0.2
Total	5,012	100

Purchase Above	N	%
Not-tagged	4,759	95
Tagged	253	5
Total	5,012	100

Selling Below	N	%
Not-tagged	4,851	96.8
Tagged	161	3.2
Total	5,012	100

Purchase Below	N	%
Not-tagged	4,654	92.9
Tagged	358	7.1
Total	5,012	100

- Rule 3 (missing prices): Missing prices are replaced with the item-specific median.

Missing Purchase	N	%
Not-tagged	3,713	74.1

Missing Selling	N	%
Not-tagged	2,569	51.3

Tagged	1299	25.9	Tagged	2443	48.7
Total	5,012	100	Total	5,012	100

- Rule 4 (missing vintages): Items with missing vintages are replaced with the item-specific median.

	N	%
Not-tagged	4,602	91.8
Tagged	410	8.2
Total	5,012	100

Table 5: Core vs. module shares[41]

	Food Consumption				Non-Food Consumption			
	Number of items	Share NBHS 2009	Share HFS 2016 (collected)	Share HFS 2016 (imputed)	Number of items	Share NBHS 2009	Share HFS 2016 (collected)	Share HFS 2016 (imputed)
Core	33	80%	92%	73%	26	65%	89%	61%
Module 1	27	5%	3%	12%	21	8%	2%	8%
Module 2	26	5%	2%	6%	20	9%	4%	14%
Module 3	26	5%	2%	6%	18	7%	3%	10%
Module 4	28	5%	1%	3%	25	11%	2%	7%
Total	**140**	**100**	**100**	**100**	**110**	**100**	**100**	**100**

Source: Authors' own calculations based on NBHS 2009 and HFS 2015 data

Table 6: Estimated median depreciation rates. [42]

Asset	**Depreciation rate**	**Asset**	**Depreciation rate**
Cars	0.05	Radio or transistor	0.17
Trucks	0.02	Mobile phone	0.21
Motorcycle/motor	0.12	Computer or laptop	0.03
Rickshaw	0.12	Refrigerator	0.05
Bicycle	0.04	Fan	0.16
Canoe or boat	0.04	Mattress or bed	0.10
Plough	0.21	Mosquito net	0.11
Television	0.04	Electric ironer	0.07
Satellite dish	0.12	Hoe, spade or axe	0.12
DVD or CD player	0.16		

Source: Authors' own calculations based on HFS 2015.

[41] The share of module 4 is missing in the HFS 2015 data due to a technical glitch. See footnote 21.
[42] Washing machines and Air conditioners were not bought.

Table 7: Urban and rural Laspeyres deflators, 2016.

	Food		Non-Food	
	Rural	**Urban**	**Rural**	**Urban**
Sep-16	1.09	1.15	0.83	1.08
Oct-16	1.18	1.00	0.88	1.00
Nov-16	1.21	1.23	1.08	1.67
Dec-16	1.05	1.23	0.86	1.67
Jan-17	1.11	0.99	0.95	1.25
Feb-17	1.07	1.37	1.07	1.43
Mar-17	1.25		1.54	
Apr-17	1.46		1.43	

Refence strata and period is Urban areas in October 2016.

Source: Authors' own calculations based on HFS 2016.

APPENDIX D

Figure 12: Heatmap of conflict fatalities, Dec. 2013-Oct. 2017.

Dec. 2013-Dec. 2014.

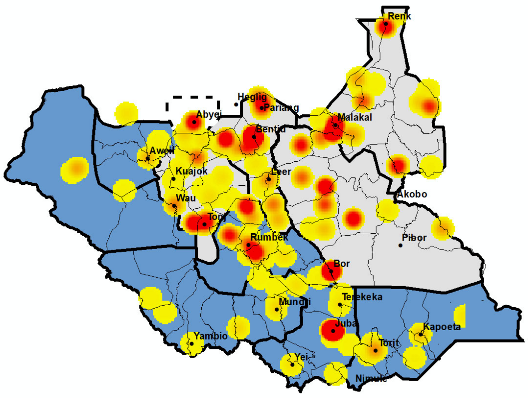

Jan.-Dec. 2015

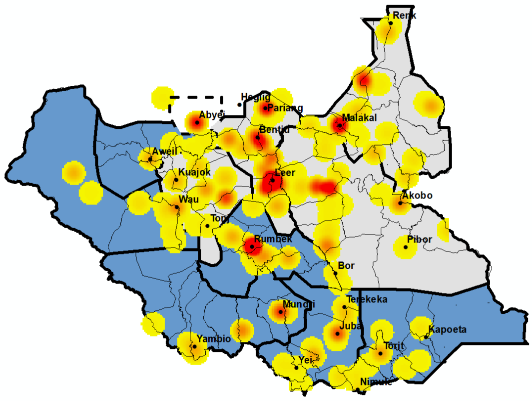

Jan.-Dec. 2016

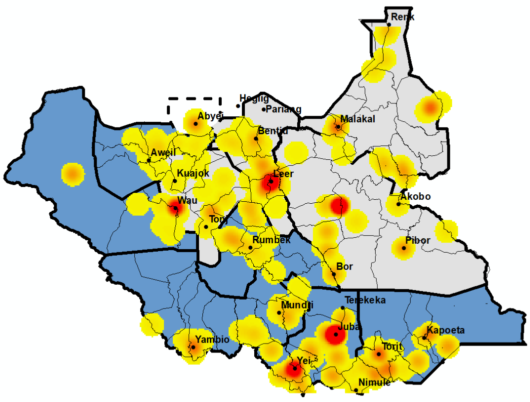

Jan.-Oct. 2017

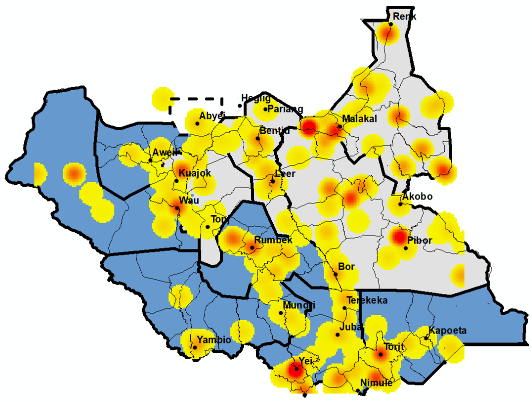

Note: all densities in maps above are color-labelled on the same scale; counties lying outside of the state boundaries are disputed territories.

Source: Authors' own calculations based on ACLED data.

Table 8: Difference in means between selected variables in Wave 1 (2015) and Wave 3 (2016) of the HFS.

	(1) Wave 3 - 2016 Mean/SE	(2) Wave 1 - 2015 Mean/SE	t-test (1)-(2) Difference
Household owns its property	0.905	0.916	-0.011
	[0.009]	[0.009]	
Phone network available at household	0.236	0.244	-0.008
	[0.012]	[0.013]	
Household is more than two hours walking from a health center	0.299	0.27	0.029
	[0.016]	[0.012]	
Household is more than two hours walking from a school	0.103	0.124	-0.021
	[0.011]	[0.008]	
Household is more than two hours walking from a market	0.327	0.32	0.008
	[0.016]	[0.012]	
Household has access to electricity	0.014	0.028	-0.014
	[0.003]	[0.008]	
Adult literacy rate (18+)	0.376	0.353	0.023
	[0.009]	[0.008]	
Adults with no education (18+)	0.564	0.55	0.014
	[0.009]	[0.008]	
Adults with only primary education (18+)	0.242	0.233	0.009
	[0.008]	[0.007]	
Household head practices polygamy	0.333	0.338	-0.005
	[0.015]	[0.012]	
Household head is male	0.579	0.603	-0.024
	[0.016]	[0.013]	
Household head is employed	0.732	0.715	0.018
	[0.014]	[0.012]	
Average age	19.311	19.026	0.285
	[0.217]	[0.186]	

*** and * indicate significance at the 1 and 5 percent level.*

Source: Authors' own calculations based on HFS 2016-2017 data.

APPENDIX E

Table 9: Variables used to create a map of settled areas.

	Variables used	
Variable name	**Description**	**Processing step**
Global Urban Footprint	Infrared-based raster of predicted presence or absence of buildings.	Dilated 100m
OSM residential areas	Volunteer-reported residential locations	Rasterized
OSM buildings	Volunteer-reported point locations of buildings	Rasterized, dilated 100m
OSM residential roads	Volunteer-reported vector of road locations	Rasterized, dilated 100m, then eroded to identify blob-like structures (residential areas)
OSM road intersection	Volunteer-reported point locations of road intersections	Rasterized, Dilated 100m
OSM health sites	Volunteer-reported locations of health facilities	Rasterized, Dilated 200m
WB health facilities	Point locations of health facilities reported in WB Points of interest database.	Rasterized, Dilated 200m
Schools	Point locations of schools reported in WB Points of interest database.	Rasterized, Dilated 200m
Household survey interviews, HFS Wave 3.	Data points from the HFS Wave 3 survey	Rasterized, Dilated 100m
	Variables rejected	
Variable name	**Description**	**Reason not used**
Night time lights DMSP	Satellite-detected intensity of night-time visible light radiance.	Brightest for power plants and oil fields in the north. These data do not bring more information on settled areas
Night time lights VIIRS	Satellite-detected intensity of night-time visible light radiance.	Often high level in areas that do not appear to be settled on satellite imagery
Waterpoints	GPS coordinates of reported water points	Many water points were not in settlements - perhaps because dataset is dated (<2012)

Note: 'rasterized' means that point or vector data were converted to gridded data at100m. 'Dilated' means that pixels were added around 'on' pixels, expanding shapes or points by a constant radius. 'Eroded' means that outer pixels of shapes were removed, suppressing linear structures and keeping only the core of blob like structures.

Source: Flowminder / WorldPop using data from NASA, NOAA, South Sudan Ministry of Water Resources and Irrigation, Global Urban Footprint, Open Street Maps, HFS 2016.

Table 10: Summary Statistics of Geo-Spatial variables

	Mean country	Min country	Max country	Mean settle	Min settle	Max settle	Mean sample	Min sample	Max sample
Distance to electricity grid (km)	114.01	0.00	505.08	27.67	0.00	213.40	12.80	0.00	72.65
Distance to schools (km)	86.42	0.00	459.24	6.33	0.00	200.77	3.59	0.00	24.83
Distance to waterpoints (km)	91.96	0.00	470.21	7.18	0.00	119.63	1.78	0.00	21.99
Distance to national roads (km)	127.22	0.00	491.87	17.03	0.00	235.94	6.96	0.00	66.82
In WEQ or Juba	0.06	0.00	1.05	0.09	0.00	1.05	0.26	0.00	1.04
IPC phase Jan. 2017 smoothed 50km	2.54	0.98	4.58	2.70	0.99	4.57	2.36	1.01	3.00
MODCF intra annual SD 100mres	1763.08	449.00	3005.00	1906.11	650.00	2933.0	1622.45	980.00	2249.00
MODCF mean annual 100mres	5338.89	3014.00	9199.00	5115.11	3084.0	8195.0	5600.55	4316.00	7945.00
SSD conflicts 2011 2016	21.46	0.00	4130.17	239.32	0.00	4130.1	378.19	0.04	3051.54
SSD conflicts 2014 2016	12.34	0.00	1608.22	128.57	0.00	1608.2	181.64	0.04	960.01
Distance to major roads (km) 100mres	83.87	0.00	470.11	5.56	0.00	131.34	2.40	0.00	26.39
Distance to plantations in 2014 100mres	71.66	0.00	458.39	2.22	0.00	96.44	2.13	0.00	23.62
Distance to urban centres (km) 100mres	158.95	0.00	543.57	38.77	0.00	256.28	14.74	0.00	74.80
Precipitations 100mres	959.98	405.54	1586.70	892.39	411.00	1584.6	1015.16	752.43	1538.01
OCHA percent people in need, 2016	24.63	0.00	252.70	51.67	0.00	252.70	36.70	33.01	115.18
Temperature 100mres	26.90	12.81	28.63	27.24	18.13	28.59	26.85	23.53	27.86
Urban gradient	0.03	0.00	8.00	3.55	1.00	8.00	4.92	2.00	8.00
Urban-rural settlements	0.01	0.00	2.00	1.04	1.00	2.00	1.35	1.00	2.00

Source: Flowminder / WorldPop.

Table 11: Variables tested for correlation with poverty.

Variable	Correlation with poverty
IPC phase (01/2017)	0.34
Seasonal cloud cover variations	0.28
Annual cloud cover	-0.37
OCHA nb people in need	0.02
Mean conflict fatalities 2011-2016	-0.49
Mean conflict fatalities 2014-2016	-0.51
Distance to 1,2,3 roads	0.02
Distance to cultivated areas 2014	0.17
Distance to urban centres	0.5
Annual temperature	0.41
Distance to electricity grid	0.36
Distance to schools	0.25
Distance to water bodies	0.10
Distance to national roads	0.25
Annual precipitation	-0.61
Urban gradient	-0.41
Urban-rural-unsettled	-0.45
In WEQ	-0.62
In Juba	-0.44
In WEQ or Juba	-0.81

Source: Flowminder / WorldPop.

Table 12: Estimated coefficients for best-fit linear model.

Variable name	Coefficient Estimate
(Intercept)	0
IPC phase	0.04
Distance to urban centers	4.7e-4
Annual temperature	0.03
Distance to electricity grid	3.6e-4
Annual precipitation	2.0e-4
urban/rural/unsettled	-0.13
In WEQ or Juba	-0.46

Source: Flowminder / WorldPop.

Table 13: State-level predictions of poverty headcount (percent).

	Poverty (survey)	*Poverty (predicted)*	*Poverty Rural (survey)*	*Poverty Rural (predicted)*	*Poverty Urban (survey)*	*Poverty Urban (predicted)*
Central Equatoria	*80*	*76*	*84*	*84*	*17*	*63*
Eastern Equatoria	*95*	*91*	*97*	*94*	*28*	*42*
Jonglei		*92*		*95*		*17*
Lakes	*84*	*86*	*86*	*89*	*29*	*47*
Northern Bahr el Ghazal	*90*	*90*	*91*	*93*	*12*	*68*
Unity		*92*		*95*		*17*
Upper Nile		*92*		*95*		*36*
Warrap	*86*	*89*	*90*	*92*	*43*	*65*
Western Bahr el Ghazal	*90*	*88*	*53*	*92*	*38*	*70*
Western Equatoria	*53*	*68*	*61*	*74*	*39*	*31*
Total	***83***	***92***	***86***	***92***	***66***	***77***

Source: Authors' own calculations and Flowminder / WorldPop (predictions).