

Measuring Disaster Crop Production Losses Using Survey Microdata

Evidence from Sub-Saharan Africa

Yannick Markhof

Giulia Ponzini

Philip Wollburg



WORLD BANK GROUP

Development Economics

Development Data Group

March 2022

Abstract

Every year, disasters account for billions of dollars in crop production losses in low- and middle-income countries and particularly threaten the lives and livelihoods of those depending on agriculture. With climate change accelerating, this burden will likely increase in the future and accurate, micro-level measurement of crop losses will be important to understand disasters' implications for livelihoods, prevent humanitarian crises, and build future resilience. Survey data present a large, rich, highly disaggregated information source that is trialed and tested to the specifications of smallholder agriculture common in low- and middle-income countries. However, to tap into this potential,

a thorough understanding of and robust methodology for measuring disaster crop production losses in survey micro-data is essential. This paper exploits plot-level panel data for almost 20,000 plots on 8,000 farms in three Sub-Saharan African countries with information on harvest, input use, and different proxies of losses; household and community-level data; as well data from other sources such as crop cutting and survey experiments, to provide new insights into the reliability of survey-based crop loss estimates and their attribution to disasters. The paper concludes with concrete recommendations for methodology and survey design and identifies key avenues for further research.

This paper is a product of the Development Data Group, Development Economics. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at <http://www.worldbank.org/prwp>. The authors may be contacted at pwoollburg@worldbank.org.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

Measuring Disaster Crop Production Losses Using Survey Microdata: Evidence from Sub-Saharan Africa

Yannick Markhof * †

Giulia Ponzini * ‡

Philip Wollburg * § 1

JEL Codes: C81, C83, O12, O13, Q15, Q54.

Keywords: crop losses, disasters, crop production, agriculture, surveys, survey methods.

¹ Authors listed in alphabetical order. † ymarkhof@worldbank.org, UNU-MERIT, United Nations University and Development Data Group, World Bank. ‡ gponzini@worldbank.org, Development Data Group, World Bank. § pwillburg@worldbank.org, Development Data Group World Bank. The authors would like to thank Shukri Ahmed, Joe Alpuerto, Piero Conforti, and Wirya Khim at FAO and Gero Carletto, Sydney Gourlay, Adriana Paolantonio, and Alberto Zezza at the World Bank for their comments and suggestions. This paper was produced with financial and technical support from the 50x2030 Initiative to Close the Agricultural Data Gap, a multi-partner program that seeks to bridge the global agricultural data gap by transforming data systems in 50 countries in Africa, Asia, the Middle East and Latin America by 2030.

1. Introduction

Between 2008 and 2018, disasters accounted for declines in crop and livestock production totaling USD 108.5 billion in low- and lower-middle income countries (LMICs), with the impact expected to further exacerbate with climate change (FAO, 2021). This toll is not only significant from an aggregate economic perspective, but its consequences are particularly dire in the poorest countries where over two-thirds of the population directly depend on agriculture for their livelihoods. Recent estimates by the FAO (2021) suggest that as much as 6.9 trillion kilocalories, equivalent to the amount it would take to end undernourishment among half of the undernourished population in low-income countries,² are lost to disasters diminishing crop and livestock production in the developing world.

With the impacts large in magnitude, widespread throughout LMICs, and likely to exacerbate, the accurate measurement of disaster-induced damages and losses³ has been identified as a key component for economic development⁴ and enshrined in Sustainable Development Goal 1.5.2 and Sendai indicator C. A large share of the disaster impact in LMICs accrues in agriculture and particularly due to reductions in pre-harvest crop production (FAO, 2018). Therefore, their accurate measurement at a micro-level is paramount to track progress toward the SDG and Sendai goals and inform strong disaster risk reduction strategies.

In a recent effort to provide a standardized methodology for measuring damages and losses in agriculture, Conforti, Markova and Tochkov (2020) propose the use of survey data to measure pre-harvest losses in crop production. However, there is yet scant evidence applying the method for micro-level loss measurement and a need for refinement to make full use of a variety of data sources including agricultural household survey data and geospatial data (FAO, 2018). This paper conducts a thorough descriptive assessment of the suitability of existing survey data to measure pre-harvest production losses of annual crops in LMICs at the micro-level. For this, it draws on a diverse set of data sources, rich plot-level data on harvest, inputs and different losses proxies along with household- and community-level data from the Living Standards Measurement Study – Integrated Surveys on Agriculture (LSMS-ISA), geospatial data, as well as experimental data from crop cut and survey experiments, in three SSA countries.

The paper’s approach is guided by a set of six distinct questions that we attempt to empirically test and mold into concrete recommendations for survey design and measurement. These questions fit into two broad issues for the measurement of disaster crop losses, the **quantification** of losses (i.e.

² Authors’ calculations based on (FAO et al., 2020; Our World In Data, no date; World Bank, no date)

³ Throughout this paper, we follow the definitions employed in (FAO, 2021, pp. 198–200):

“Disaster: A serious disruption of the functioning of a community or a society at any scale due to hazardous events interacting with conditions of exposure, vulnerability and capacity, leading to one or more of the following: human, material, economic and environmental loss and impacts (UNDRR Terminology).”

“Damage: The monetary value of total or partial destruction of physical assets and infrastructure in disaster-affected areas, expressed as replacement and/or repair costs.” In agriculture, this comprises, for example, the destruction of stored agricultural inputs or assets used in the production process.

“Loss: *The change in economic flows occurring as a result of a disaster.*” In the context of this paper, the focus will be on losses due to declines in crop production.

⁴ Specifically, accurate measurement of damages and losses matters for: 1. disaster risk reduction; 2. poverty eradication; 3. food security; 4. climate change action; and 5. sustainable development (FAO, 2020).

the size of crop losses) and their **attribution** to specific disasters. For the quantification of losses, the paper assesses the suitability of proxies based on (i) land area and crop damage, (ii) expert assessments, (iii) past harvest data or data from similar but unaffected plots, and (iv) farmers' post-planting expectations to inform a micro-level crop production loss estimate. Subsequently, it (v) studies the reliability of self-reports for linking losses to specific disasters and (vi) compares them to geospatial data on disaster exposure.

Based on these questions, the paper makes a contribution in three ways. First, it operationalizes the methodology proposed in Conforti, Markova and Tochkov (2020) for use with state-of-the-art, crop production survey data at the plot level based on the format of the LSMS-ISA questionnaire. Second, the paper provides an in-depth investigation into the feasibility of measuring pre-harvest crop production losses due to disasters at scale and yet a highly disaggregated level and provides concrete recommendations for survey design. As such, it informs a new generation of survey data that facilitates the integration of a diverse set of data sources for measurement and analysis. Lastly, the paper identifies crucial remaining knowledge gaps and delineates relevant avenues for future research. Importantly, the methodology explored lends itself primarily to small- to medium-sized disasters and those with slow onset (e.g. droughts). For large scale disasters and swift policy action requiring an immediate appraisal of damages and losses, post-disaster needs assessments (PDNAs) and aggregate-level analyses may present complementary vehicles that are more readily deployable.

The remainder of this paper is structured as follows. Section 2 contextualizes the measurement of disaster crop losses in the literature and depicts the key approaches and challenges associated with their accurate quantification and attribution to disasters. Section 3 introduces the data sources the paper draws on and the key questions that motivate its empirical strategy. Section 4 presents the paper's analysis for each of these questions while Section 5 closes with our headline findings and highlights key takeaways as well as open questions.

2. Crop Production Losses Measurement in the Literature – Approaches and Challenges

To accurately measure crop production losses due to disasters, this paper distinguishes between two issues, first, the quantification of crop production losses, that is determining their magnitude, and second, attributing losses to specific disasters.

The task of measuring crop production losses relates to a large body of research broadly concerned with crop yield prediction and the explanation of geographical and intertemporal variation in yields. Part of this is an extensive agronomic literature measuring yield potential, that is the optimum achievable yield in a certain location without limitations from pests, diseases, nutrition, or suboptimal crop management apart from uncontrollable factors such as temperature or precipitation in rainfed systems (Lobell, Cassman and Field, 2009; van Ittersum *et al.*, 2013;

Grassini *et al.*, 2015; van Loon *et al.*, 2019).⁵ These studies typically employ so-called “process-based crop models”, mathematical models using daily data on a host of biophysical conditions to simulate the processes associated with plant growth. Models differ in their inclusion of different processes (such as crop phenology, light interception and utilization, or evapotranspiration) and their respective parametrization (Rosenzweig *et al.*, 2014). They can thus yield different results and are often combined in an ensemble to reduce variance and improve accuracy (Asseng *et al.*, 2013; Rosenzweig *et al.*, 2014; Zhao *et al.*, 2017).

Yield gaps, however, do not provide estimates of the counterfactual harvest in the absence of disaster impact. Rather, they measure productivity gaps and thus reflect gaps in resources used, technology employed, and crop management efficiency as are prevalent in smallholder agriculture in LMICs (Assefa *et al.*, 2020; Beza *et al.*, 2017; Mueller *et al.*, 2012). Furthermore, achieving yield potential may not even be desirable from an efficiency perspective due to diminishing returns to crop management effort and input use (van Ittersum *et al.*, 2013). While yield gaps are based on the yield potential under optimal management but *conditional* on uncontrollable factors such as weather shocks, this paper’s interest lies on the exact opposite: *Given* the farming practices observed but *in the absence* of shocks, what is the attainable yield and consequently the disaster-related loss?

Another stream of literature studies the sources of crop yield variability across years and different farms, regions, or even countries (Restuccia and Rogerson, 2013; Gollin, Lagakos and Waugh, 2014; Adamopoulos and Restuccia, 2018; Gollin and Udry, 2021). This literature finds that, among other factors, the stochastic nature of the production process (e.g. due to weather shocks) accounts for a significant share in the variability of yields (Ray *et al.*, 2015; Amare *et al.*, 2018; Gollin and Udry, 2021).

Related to this is a literature analyzing farmers’ decision making in the face of uncertainty over future shock realizations. In the context of LMICs, future shock expectations and importantly also observed behavior have been found to vary according to differences in risk preferences, e.g. due to past shock exposure (Gloede, Menkhoff and Waibel, 2015; Said, Afzal and Turner, 2015; Cassar, Healy and von Kessler, 2017; Brown *et al.*, 2018; Hanaoka, Shigeoka and Watanabe, 2018), access to information, e.g. weather forecasts (Karlan *et al.*, 2014; Rosenzweig and Udry, 2019), and access to insurance markets, e.g. rainfall index insurance (Giné, Townsend and Vickery, 2009; Karlan *et al.*, 2014). Furthermore, large differences in idiosyncratic expectation building have been found (Freudenreich and Kebede, 2019).

These findings matter for this study in that they suggest that not only current-season shock realizations but also past shock experiences and future expectations of shocks account for variability in crop yields. While this paper’s main interest is in measuring losses due to actual disaster exposure, disaster crop losses themselves may be endogenous. This is because the aforementioned factors matter for farmers’ expected returns to their farming decisions and their resulting behavior, such as input decisions, sets an effective ceiling for the attainable harvest.

⁵ This literature is commonly concerned with the measurement of *yield gaps*, i.e. the difference between potential and realized yields.

However, it is important to note that this paper is primarily concerned with the *measurement* of disaster crop losses as opposed to identifying their causal drivers. In as much as our identification strategy for *measuring* losses is thus unaffected by these factors, they do not pose a major challenge for this study.

On the methodological side, a strand of literature relevant for this paper is concerned with estimating and forecasting the climate change impacts on agricultural yields. For this, three approaches are common (Challinor *et al.*, 2014; Liu *et al.*, 2016; Zhao *et al.*, 2017). The first two approaches use process-based crop models in conjunction with meteorological models and can be split into grid-based models (e.g. Rosenzweig *et al.*, 2014; Webber *et al.*, 2018) and less coarse but more geographically scattered site-based models (e.g. Asseng *et al.*, 2013).⁶ These models are also becoming increasingly popular as in-season, real-time monitoring tools for early crop loss prediction and swift relief (e.g. the World Food Program's LEAP tool employed in Ethiopia and documented in Drechsler and Soer (2016)).

As discussed previously, crop models vary in their data requirements but usually require a rich set of sensor data at high frequency and sufficient spatial resolution. They are thus demanding for accurate, large-scale, and micro-level measurement in LMICs where smallholder agriculture with small and irregular plot sizes, intercropping, and large idiosyncratic variations in yields are the norm. Despite recent advancements in high resolution yield estimation (e.g. Burke and Lobell, 2017), survey data may thus provide complementary information. Survey data are typically the data source for a third approach in the climate change and crop yield literature that employs statistical modeling based on observed harvest data at varying levels of aggregation (e.g. Lobell, Schlenker and Costa-Roberts, 2011; Imran, Zurita-Milla and Stein, 2013; Liu *et al.*, 2016).

Survey data may be a valuable source of complementary information for several reasons.

- 1) They can provide information down to the plot-crop level and thus possibly higher micro-level accuracy than most grid-based models to date. This opens up opportunities for better understanding of micro-level dynamics that can inform disaster risk reduction policies.
- 2) They provide high data density. Relative to site-based models, this reduces the need for strong assumptions about the generalizability of the sample for larger aggregates (though it may not eliminate it entirely) (van Wart *et al.*, 2013; van Bussel *et al.*, 2015).
- 3) They provide an approach that may be more suitable to the peculiarities of smallholder farming in LMICs (see above).
- 4) They are able to pick up non-meteorological sources of crop losses such as pests and diseases (Liu *et al.*, 2016).
- 5) Their information can be effortlessly integrated with other information collected in surveys that make understanding, for example, behavioral components of crop management

⁶ In comparison to gridded models that rely on remote-sensed data of sufficiently high resolution, site-based models typically use sensor data collected at several distinct experimental sites, e.g. farms or even single plots. While the former can thus cover a larger geographic area and may be more representative for aggregate measurement, the latter has the advantage of more fine-grained measurement based on data observed “on the ground”.

possible. This facilitates causal investigations into variation in crop losses or the micro-level impacts of disasters (Hallegatte and Rozenberg, 2017; Markhvida *et al.*, 2020).

Importantly, this paper does not study nor does it propose the use of survey data as a substitute for the other approaches presented. Rather, it highlights how survey data may yield additional insights and ultimately contribute to the development of more accurate, integrated approaches suitable for the context of LMICs. For example, high-quality survey data may be useful to train remote-sensing algorithms (Jain *et al.*, 2016) and survey data, geospatial data, and/or crop model data can be combined in statistical models to improve accuracy and analyze causal relationships (Imran, Zurita-Milla and Stein, 2013; Hill and Porter, 2017).

Next to the issue of determining the magnitude of crop losses (quantification), accurately attributing losses to disasters is an essential challenge for this paper to study. As the most common disasters in our study region are rainfall-related (FAO, 2018, 2021), this task poses some challenges akin to those covered in the literature on agricultural index insurance (Miranda and Farrin, 2012; Cole *et al.*, 2013; Karlan *et al.*, 2014; Tadesse, Shiferaw and Erenstein, 2015).

Self-reported insurance claims may be hard to verify without significant administrative effort and may thus be subject to moral hazard and adverse selection (Tadesse, Shiferaw and Erenstein, 2015). To overcome issues of information asymmetry and reduce administrative costs, insurance contracts that tie payouts to a local index proxying for disaster impact intensity, e.g. a rainfall gauge placed at a representative location for a small geographic area, have received considerable attention in the literature. However, while “measuring” shock exposure through an index provides an objective proxy, it leaves farmers with a certain “basis risk”, the possibility that the index does not capture idiosyncratic deviations at the farm or even plot level.

Basis risk is a substantial impediment to the uptake of index insurance (Giné, Townsend and Vickery, 2008; Miranda and Farrin, 2012; Cole *et al.*, 2013; Karlan *et al.*, 2014; Jensen, Mude and Barrett, 2018) and of particular concern in heterogenous regions with many micro-climates (Barnett, Barrett and Skees, 2008). The extent of basis risk furthermore depends on shock intensity with more severe shocks reducing the idiosyncrasy in incidence (Barnett, Barrett and Skees, 2008; Hazell and Hess, 2010; Jensen, Barrett and Mude, 2016). Consequently, most insurance payouts have been found to constitute large, infrequent payouts in response to severe disasters (Giné, Townsend and Vickery, 2007). Yet, idiosyncratic risks have been found to be substantial in the case of smallholder livestock farming and inadequately covered by index insurance (Jensen, Barrett and Mude, 2016). This suggests that crop losses could be quite localized and require fine-grained measurement to accurately attribute losses to disasters. Similarly, a number of studies have found large variation in drought exposure between different measures. Notably, farmers’ self-reported shock exposure commonly deviates from measured data collected from weather stations or satellite data (Meze-Hausken, 2004; Sohnesen, 2020). Determining the degree of idiosyncrasy in farmers’ self-reported shock exposure and disentangling variation due to differences in perception from differences in actual exposure is thus also a relevant issue for this paper.

Lastly, a number of factors commonly studied in the behavioral literature may matter for the measurement of crop production losses through survey data (Schilbach, Schofield and

Mullainathan, 2016; Waldman *et al.*, 2020). Examples include availability bias⁷ (Tversky and Kahneman, 1973; Karlan *et al.*, 2014; Brown *et al.*, 2018), scarcity⁸ (Shah, Mullainathan and Shafir, 2012; Mani *et al.*, 2013; Lichand and Mani, 2020), satisficing⁹ (Krosnick and Presser, 2010), and a large body of methodological research on the accuracy of farmers' self-reports (Carletto, Savastano and Zezza, 2013; Carletto, Gourlay and Winters, 2015; Arthi *et al.*, 2018; Gourlay, Kilic and Lobell, 2019; Wollburg, Tiberti and Zezza, 2021).

3. Empirical Strategy and Data

In exploring the suitability of survey data for disaster crop loss measurement, this paper will be guided by the methodology for damages and losses assessment in agriculture that was recently developed in Conforti, Markova and Tochkov (2020) for the FAO (henceforth also “the FAO methodology”). Complementing previously developed disaster impact assessment approaches by the United Nations Economic Commission for Latin America and the Caribbean (UN ECLAC, 2003) and the multi-agency Post Disaster Needs Assessment methodology (European Union, UN Development Group and World Bank, 2013), they provide a standardized methodology for measuring damages and losses due to disasters across the agricultural sub-sectors of livestock, forestry, aquaculture, fishery, and (this paper's focus) crops. To measure the impact in the crop sector, the FAO methodology distinguishes between annual and perennial crops as well as damage (e.g. the destruction of stored inputs or stored crops) and loss (e.g. reductions in crop yields).

$$DL(Crops) = \text{Annual crop prod. damage} + \text{Perennial crop prod. damage} + \text{Annual crop prod. loss} \\ + \text{Perennial crop prod. loss} + \text{Crop assets damage}$$

where $DL(Crops)$ is the total monetary value of damages and losses in the crop sub-sector. This paper's focus will be crop production losses of annual crops.¹⁰

In keeping with Conforti, Markova and Tochkov (2020), we can distinguish between two types of crop losses. The first case describes losses due to reductions in harvested area. It occurs when parts of the planted area are fully destroyed and thus all harvest was lost in these areas.¹¹ The second type of crop loss arises due to reductions in productivity of the plants harvested. This case occurs

⁷ For example, Karlan *et al.* (2014) find that farmers overweight recent adverse events in their investment decisions.

⁸ Lichand and Mani (2020), for instance, find that the mere prospect of adverse rainfall and thus income reductions reduced farmers cognitive abilities in a range of survey experiments.

⁹ For example, providing an estimate of disaster crop losses, in theory, involves estimating the counterfactual harvest in the absence of shocks which is complex and may lead to the use of heuristics and inaccuracy in self-reports. Such complex questions have been found to be error prone without careful survey design and the provision of appropriate aids that simplify the task (Krosnick, Narayan and Smith, 1996; Krosnick and Presser, 2010; Delavande, Giné and McKenzie, 2011).

¹⁰ We expect our findings to similarly apply to perennial crops. However, they bring added complexity due to the possible cumulation of shocks across seasons.

¹¹ For example, part of a plot may have been flooded completely wiping out the harvest in the flooded area.

when a disaster only partially affects an area, reducing yields but leaving something to be harvested.¹² Equation 1 formalizes this relationship.

$$L_{i,j,t} = \underbrace{[p_{i,g,t-1} \times y_{i,j,t}^c \times (ha_{i,j,t}^c - ha_{i,j,t}^r)]}_{\text{Losses due to reductions in the area harvested in fully destroyed areas}} + \underbrace{[p_{i,g,t-1} \times (y_{i,j,t}^c - y_{i,j,t}^r) \times ha_{i,j,t}^r]}_{\text{Losses due to reductions in productivity in partially affected areas}} + \gamma \quad (1)$$

Where L_i is the quantity of annual crop i lost; j is the unit of observation (e.g. the plot or the farm); t is the season where the disaster occurred; $p_{i,g,t}$ is the price of crop i in geographic zone g in the time period immediately before the disaster¹³; y^r is the realized (disaster-reduced) yield and y^c the counterfactual yield that would have been achieved in the absence of a disaster; ha^r is the area that was harvested and ha^c the full planted area (the area that would have been harvested in the absence of a disaster); γ are short-run post disaster maintenance costs.¹⁴

Conceptually, measuring crop production losses due to disaster thus boils down to comparing the realized yield y^r to a counterfactual yield y^c . More precisely, the counterfactual yield is the yield that would have been realized in the absence of any disasters holding all other factors such as input decisions, farming practices, or non-disaster shocks constant. In the following, we refer to this as the “attainable yield” (Loïc *et al.*, 2018). In partially affected areas where plants are less productive, losses are then calculated by multiplying the difference between attainable and realized yield by the size of the harvested area ha^r . For the remainder of this paper, we will refer to these as “plant productivity losses”.¹⁵ Analogously, in completely destroyed areas, losses are obtained by multiplying the attainable yield by the size of the completely destroyed area (the difference between the planted area ha^c and harvested area ha^r). Such losses will be referred to as “planted area losses”.

Three important points are evident from the distinction between realized and attainable yield. First, whenever losses due to disaster occur, the attainable harvest will be unobserved and thus becomes the key unknown to determine.¹⁶ Second, in econometric terms, the challenge is to single-out the adverse effect of disasters on yields and clean identification relies on the ability of any methodology to adequately control for potential confounding factors. Third, measurement crucially depends on the choice of an accurate benchmark for the counterfactual yield. For this, the FAO methodology suggests the use of past yield data, e.g. based on a multi-year trend. Other

¹² For instance, in some areas, pests may damage the plants, reducing their productivity but not completely destroying them.

¹³ For the remainder of this paper, we focus on the quantification of losses as this poses the key methodological challenge to measurement and leave their valuation for future research that may also look into the feasibility of creating aggregates from micro-level crop loss measurement.

¹⁴ For example, this could be the costs associated with removing uprooted trees and branches from a field or clearing a field of waterlogging. As these costs are comparatively straightforward to assess by valuing the time required and equipment hired to do so, we will not consider them further for the sake of this paper.

¹⁵ In keeping with the phrasing in the surveys used for this study, we will use the term “damage on the crop” interchangeably. Note that these are distinct from “damages” as defined in Footnote 2 and rather a special case of “losses” as defined in the same footnote.

¹⁶ Trivially, therefore, only in the complete absence of any disaster losses will realized harvest equal attainable harvest and the latter thus be observable.

proposed options include the use of average yield data from comparable districts unaffected by any shocks or yield forecasts based on meteorological and remote-sensed data (FAO, 2020).

As the accuracy of different approaches for micro-level measurement of the attainable yield through surveys likely varies, their empirical assessment constitutes the first goal of this paper. We call this the “quantification of losses”. Concretely, the paper explores a number of survey-based ways to measure or proxy for the attainable harvest using plot-level data¹⁷ including

- (i) explicitly measuring planted area losses and plant productivity losses based on farmers’ self-reports;
- (ii) expert assessments of the share of the attainable harvest that was lost;
- (iii) obtaining attainable yields from past data or comparable but disaster-unaffected plots;
- (iv) using farmers’ post-planting harvest expectations to proxy for the attainable yield.

Importantly, these approaches need not be mutually exclusive, and the paper will aim to highlight the respective strengths and weaknesses of each approach in an aim to combine them for best results.

In addition to the *quantification* of losses, the second crucial focus of this paper is the *attribution* of losses to specific disasters. One possible approach relies on self-reported shock identification, either from the farmer, a community congregation, or an ‘expert’. Here, the paper analyzes the consistency of shock identification on plots of the same crop in close proximity and thus assesses the suitability of self-reports to confidently attribute losses to disasters. Alternatively, the paper compares self-reported drought exposure to geospatial rainfall data to obtain a similar “objective measure” as in the agricultural index insurance literature. Through this, it aims to disentangle variation in self-reports due to different shock exposure from variation due to differences in farmers’ perceptions. Notably again, the paper also explores the scope for integrated approaches that combine self-reported and sensor data to harness the strength of each.

For our analysis, we draw on a diverse set of data sources. Our main data source is rich plot-level household survey panel data¹⁸ for two Sub-Saharan African countries, Ethiopia and Malawi, supported by the Living Standards Measurement Study – Integrated Surveys on Agriculture (LSMS-ISA). The LSMS-ISA data have a distinct focus on agriculture and contain, among others, detailed data on harvest, input use, crop management, as well as a large set of community, household, and individual characteristics. The agricultural modules of the surveys feature a two-visit structure in which each household was visited after planting and again after the harvest had been completed. At the first visit, among others, self-reported information on input use, plot characteristics, and expected production are elicited while the second visit asks farmers about their

¹⁷ This is in contrast to other approaches including the “default” angle of the FAO methodology which aim to measure crop losses at the farm-level or higher.

¹⁸ The LSMS-ISA data we use are *farm-level* panels but not necessarily *plot-level* panels due to constant changes in plot dimensions and the crops cultivated on them (Carletto, Gourlay and Winters, 2015). In this version of the paper, we therefore do not exploit the panel dimension of our data and treat each wave as a separate cross section. However, this does not mean that the fact that we are dealing with panel data does not carry advantages for our analysis. For example, it should attenuate concerns about changes in the sample when comparing descriptives across rounds. An exception is ESS18 in Ethiopia which received a new baseline sample in this round and is thus not strictly comparable (from a panel perspective) to the sample from previous rounds.

realized production and any crop losses incurred. In the case of Ethiopia, we furthermore have crop cut data for a sub-sample of plots.¹⁹ The data are available publicly in the World Bank Microdata Library²⁰ and cover almost 20,000 plots across over 8,000 farms. We complement this data with data from a crop cut and survey experiment, the *Methodological Experiment on Measuring Maize Productivity, Soil Fertility and Variety* (MAPS) in Uganda. Table 1 summarizes our data sources.

The main unit of analysis for this paper is the plot with separate information for each crop cultivated on the plot. This high level of disaggregation is a feature of our data compared to other approaches that typically analyze losses at the farm level or higher.²¹ In Ethiopia, the analysis focuses on two of the major staple crops in the country, maize and sorghum, while in Malawi the paper analyzes maize as the country’s foremost crop cultivated. Furthermore, we also build a plot-level aggregate of total harvest value on the plot (i.e. across all crops cultivated for every plot of a farm) in local currency units in Malawi to verify the robustness of our analysis to the inclusion of all crop types. The analysis centers around the main agricultural season in each of the two countries, the *Meher* season in Ethiopia and the rainy season in Malawi. Where applicable, the paper expresses production quantities in total kilogram terms as well as in per hectare terms. The plot sizes used are measured by GPS if available or for small plots with rope and compass in the case of Ethiopia. To correct for outliers in self-reported harvest and input variables as well as our crop losses estimates, per hectare variables are winsorized at the 0.01 percentile.

Table 1: Data sources and sample sizes.

Country	Data Source	Sample Size	Survey Rounds
Ethiopia	Ethiopian Socioeconomic Survey (ESS)	13,570 plots	2013/14, 2015/16, 2018/19
Malawi	Integrated Household Panel Survey (IHPS)	5,998 plots	2010/11, 2013/14, 2016/17, 2019/20
Uganda	Methodological Experiment on Measuring Maize Productivity, Soil Fertility and Variety (MAPS2)	213 plots	2016

¹⁹ Crop cuts have been termed the “gold standard” to obtain an objective measure of crop yield (Carletto, Gourlay and Winters, 2015).

²⁰ <https://microdata.worldbank.org/index.php/catalog/lsms>

²¹ However, it is important to note that in case of urgent need of swift policy action in response to a sudden disaster, more aggregate approaches to measuring crop losses retain validity for the relative speed at which they can be conducted.

4. Analysis and Results

We divide our analysis into two key issues, the *quantification* of losses, i.e. determining how much was lost, and the *attribution* of losses to specific disasters. For each of these issues, the paper's analysis is guided by a set of concrete questions that are empirically assessed in the LSMS-ISA data.

Quantification

- Question 1: Can we separately measure planted area and plant productivity losses in survey data and form a robust, compound loss estimate?
- Question 2: How do farmer reports compare to crop losses based on expert assessments and can we establish an objective, reliable benchmark ('gold standard') for crop losses?
- Question 3: Can data from past harvests serve as a suitable benchmark for attainable harvest?
- Question 4: Can farmers' post-planting harvest expectations be used to construct an estimate of harvest losses?

Attribution

- Question 5: How uniform is farmers' self-reported shock identification of covariate shocks?
- Question 6: How localized are drought losses and how confidently can we attribute losses to adverse rainfall based on self-reports?

Table 2: Main questions and approaches for analysis.

Issue	Approach	Estimate of interest	Proxy	Availability
Quantification	Q1: Measuring different types of losses and building a compound estimate	Losses due to reductions in planted area ("planted area losses")	Approximately what percent of [PLOT] planted with [CROP] was harvested?	Malawi: IHPS19 Ethiopia: ESS13, ESS15, ESS18
		Losses due to productivity decline among harvested crops ("plant productivity losses")	What is the percentage of damage on [CROP]?	Ethiopia: ESS13, ESS15, ESS18
	Q2: Using losses assessed by an 'expert'	Consistency of self-reports and expert assessments	What percentage of the potential maize harvest was lost due to damage on [PLOT]?	Uganda: MAPS
	Q3: Exploiting the cross sectional or time dimension of survey data	Attainable yield based on past yield or unaffected districts	Past yield realizations or yields from similar but unaffected plots	
	Q4: Relying on farmers' self-reported harvest expectations	Attainable yield based on farmers' post-planting expectations	How much of [CROP] do you expect to harvest during the main rainy season?	Malawi: IHPS10, IHPS13, IHPS16, IHPS19 Ethiopia: ESS15, ESS18
Attribution	Q5: Relying on farmers' self reported shock identification	Consistency of self-reported shock exposure	Plot-level reports of losses cause; Household-level shock reports; Community reports	Malawi: IHPS10, IHPS13, IHPS16, IHPS19 Ethiopia: ESS13, ESS15, ESS18
	Q6: Index-based shock identification	Consistency of self-reports with geospatial rainfall data	Geospatial/sensor data on rainfall intensity	Malawi: IHPS10, IHPS13, IHPS16, IHPS19 Ethiopia: ESS13, ESS15, ESS18

4.1. Quantification: How high are losses?

Question 1: Can we separately measure planted area and plant productivity losses in survey data and form a robust, compound loss estimate?

In the stylized framework set out in Section 3, losses can manifest in two complementary ways capturing different types of losses. The first are losses due to reductions in the size of the harvested

area vis-à-vis what was planted (“planted area losses”).²² However, losses may also occur if harvest could occur on part of the plot but plants in the harvested areas have some sort of damage that makes them less productive (e.g. pests destroying some cobs or some plants not developing properly due to poor weather conditions). The latter is what this paper calls “plant productivity losses”.²³ Consequently, it is necessary to find survey-based measures for reductions in the area harvested²⁴ and reductions in productivity in harvested areas.²⁵

In relation to the general form expressed in Equation (1), we observe GPS measured plot area, h^c , and realized production Y^r . The two key unknowns to measure in the survey data are the reduction in the planted area and the reduction in productivity. We use a question on the percent share of the planted area that was harvested, $s_{i,j,t}$, to obtain the harvested area, h^r , and the realized yield in harvested areas, y^r . We then obtain the attainable yield, y^c , from y^r by accounting for plant productivity losses through a question on the percentage of damage on the crop, $d_{i,j,t}$. Therefore, in Ethiopia, the loss estimate built in this section will be a compound one consisting of a) planted area losses (the first term), and b) productivity (damage) losses (the second term).

$$L_{i,j,t} = \left[\left(\frac{Y_{i,j,t}^r}{ha_{i,j,t}^c \times s_{i,j,t}} \times \frac{1}{1-d_{i,j,t}} \right) \times (ha_{i,j,t}^c - ha_{i,j,t}^c \times s_{i,j,t}) \right] + \left[\left(\left(\frac{Y_{i,j,t}^r}{ha_{i,j,t}^c \times s_{i,j,t}} \times \frac{1}{1-d_{i,j,t}} \right) - \frac{Y_{i,j,t}^r}{ha_{i,j,t}^c \times s_{i,j,t}} \right) \times ha_{i,j,t}^c \times s_{i,j,t} \right] \quad (2)$$

We next take this approach to the data.²⁶

²² One example may be the partial flooding of a plot which destroys the harvest completely in one part but due to plot slopes etc. leaves the rest of the plot unaffected.

²³ For example, pests or a crop disease may completely wipe out the harvest in parts of the plot but only lead to partial losses in other parts of the plot that could still be harvested.

²⁴ In the LSMS-ISA data, farmers are asked for each crop on each plot at the post-harvest visit whether the planted area was greater than the area harvested and (if applicable) what the percent share of the planted area that was harvested was. A value of 0 percent indicates a full loss and values between 1 percent and 99 percent denote partial losses of the planted area. Farmers are also asked to choose up to two reasons for their loss from a list of shocks or specify another reason.

²⁵ This information is only available in the Ethiopia LSMS-ISA data. Farmers are asked for each crop on each plot whether there was any damage on the crop and what the percentage share of and reason for this damage was. In fact, this set of questions is asked *twice*, once at the post-planting visit and again at the post-harvest visit. Intuitively, one would assume for damage on the crop to only ever increase. However, we find that, depending on the survey wave, in up to 25 percent of cases damage reports at the post-planting visit exceed those at the post-harvest visit.

Furthermore, this is not due to trivial differences that could originate from rounding error or the like. We take note of this and trust post-harvest reports for the purpose of loss calculation as these should constitute the most updated farmer estimates of damage.

²⁶ Unless specified otherwise, the analysis excludes plots where the harvest was completely lost (the equivalent of an area loss of 100 percent on the plot). The main reasons for this is to avoid bias in our estimates stemming from a (by construction) perfect fit between the area loss measure and the realized harvest quantity. Furthermore, the distinction between area and plant productivity losses is not meaningful in this case and there is no way to infer the attainable yield according to Equation (2) in the lack of any realized harvest. To estimate total loss on fully destroyed plots, the attainable yield can, for instance, be inferred from the median attainable yield on similar or near plots with at most partial losses.

What do we see in the data? ²⁷

Comparing the area-based and productivity-based proxy in a histogram suggests that the productivity-based proxy is more sensitive to small losses. That is, farmers report having any plant productivity losses more frequently, but these losses tend to be smaller. Conversely, the planted-area-based proxy records fewer losses but these tend to be higher. Across years and crops, the incidence of plant productivity losses is much higher with about 1 in 2 plots recording some damage on the crop whereas only every fifth plot records reductions in the area harvested compared to what was planted.²⁸ Conditional on recording a loss though, planted-area-based losses tend to be larger.

Consistent with what has been found in other cases in the survey methodology literature, respondents tend to round numbers to convenient values divisible by ten or 25. A particularly conspicuous spike can be found at 50 percent likely reflecting the convenience of this value and easy interpretation. This pattern is consistent with the presence of considerable satisficing, the use of heuristics or “mental shortcuts” to reach a (in the respondent’s eyes) satisfactory rather than optimal answer (Krosnick and Presser, 2010).

Three factors mediate the degree of satisficing to be expected (Krosnick, Narayan and Smith, 1996): (i) the difficulty of the task; (ii) the farmer’s ability to complete the task; and (iii) the farmer’s motivation to provide an accurate estimate. Each of these concerns likely applies to our context. First, estimating average losses, particularly when differentiating between average reductions in productivity and the area completely lost, is cognitively challenging. Furthermore, percentages may not be a straightforward concept and lack an intuitive “benchmark” to correctly assess differences in magnitude along a percentage scale.²⁹ Secondly, the respondents in our sample are overwhelmingly low-educated with over two-thirds not having any school education and only a rough one-fifth having completed primary education in the case of Ethiopia. They also likely have little practice in performing this kind of task as the LSMS-ISA survey only has a two- to three-year frequency. Thirdly, the LSMS-ISA questionnaires are of substantial length, frequently requiring interviewing respondents over several hours and on multiple days. This may increase respondent fatigue and decrease the motivation to provide an accurate estimate. As a result, a percent difference of 10 percent or 20 percent, well within the conceivable margin of error for the amount of heaping we observe, may appear small to the farmer (and hence of satisfactory accuracy) but would lead to substantial measurement error (“weak satisficing”). Alternatively,

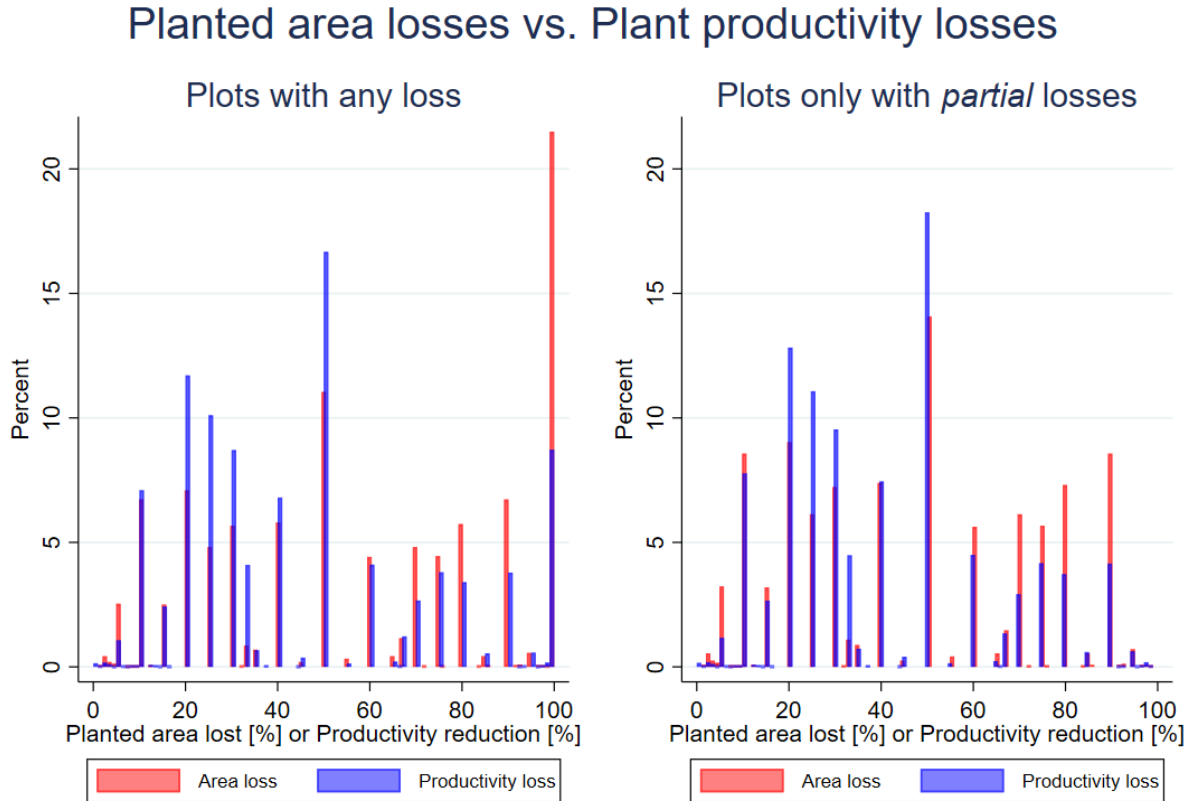
²⁷ Our exposition will focus on Ethiopia where we can measure both planted area losses and plant productivity reductions. In Malawi, we only have loss reports from one round (IHPS19) and also only for planted area losses which equates to a scenario where $d_{i,j,t} = 0$, that is, where realized yields in harvested areas are equal to attainable yields in harvested areas and the only losses occur due to complete destruction of parts of the plot. When limiting our analysis to this proxy though, results between Ethiopia and Malawi are largely comparable with some notable differences in the incidence of area losses (60 percent of plots in Malawi vs. 20 percent of plots across rounds in Ethiopia). This may reflect the planted area loss question compensating for a lack of a separate question asking for productivity reductions.

²⁸ To avoid a hard-to-read histogram due to a dominant spike at 0 percent area-based loss and 0 percent damage, we constrain the histograms depicted to those plots who record any loss (i.e. positive values on the horizontal axis).

²⁹ For a similar reason, the use of visual aids is common when eliciting probabilities or other relative numbers from subjects with low numeracy (Delavande, Giné and McKenzie, 2011).

farmers may choose to simply report a value that they believe will *appear* credible in the eyes of the enumerator (“strong satisficing”).

Figure 1: Histograms of planted area loss and plant productivity loss measures.

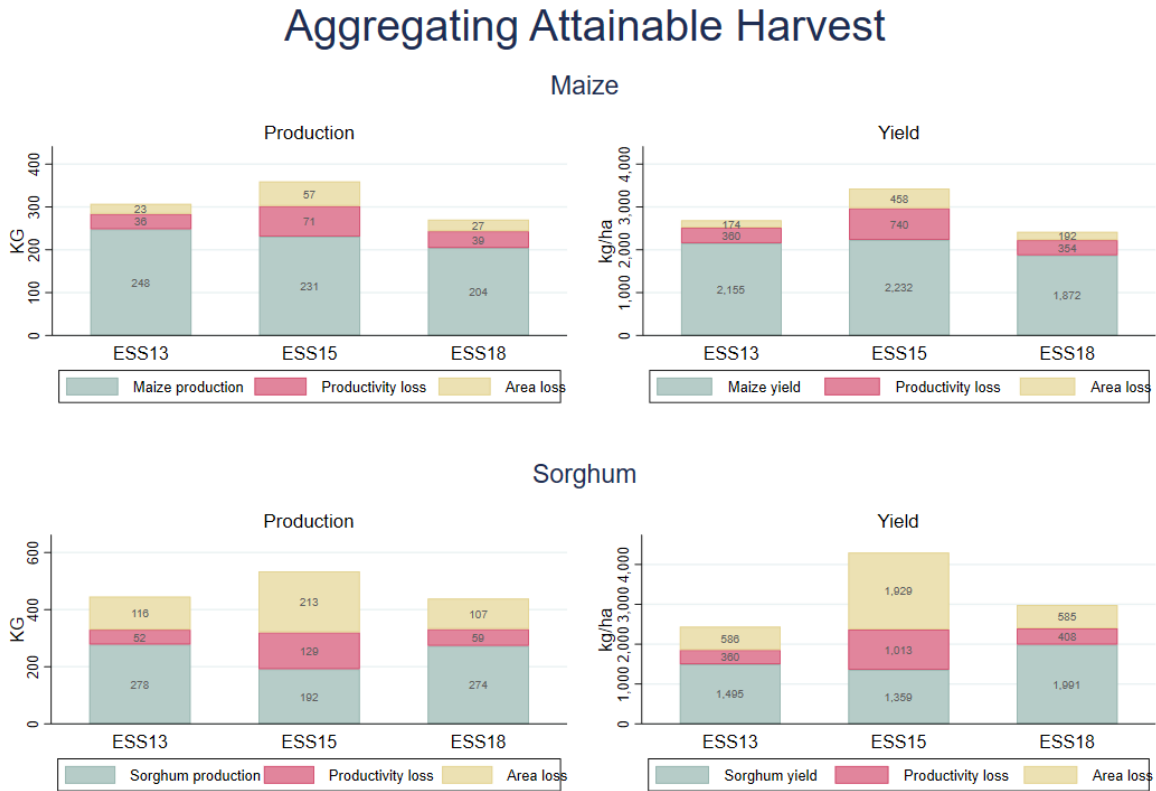


Excluding those w/o any losses

What does a compound loss estimate from these two proxies look like in the Ethiopia LSMS-ISA data? Figure 2 depicts the average realized harvest, plant productivity loss, and planted area loss for each round and crop.³⁰

³⁰ For production, area loss and crop damage correspond to the first and second term, respectively, in Equation (2) with realized production being Y^r . The entire stacked column corresponds to the average attainable production on plots without full losses. For yield, values are per hectare of the planted area h^c . Taken together, the entire column corresponds to the average attainable yield on maize (sorghum) plots without full losses.

Figure 2: Mean realized harvest and mean planted area and plant productivity loss by round and crop.



Note: Excluding full losses. Yield is per hectare of the planted area.

As an initial observation, accounting for plant productivity reductions crops in the area harvested naturally increases the loss estimate. The size of the overall (average) loss and its composition furthermore varies between years and crops. While plant productivity reductions account for the largest share of losses for maize, planted area losses are notably larger in size for sorghum. Furthermore, the 2015/16 survey round clearly sticks out with larger losses, particularly for sorghum. This is consistent with the literature reporting the worst drought in decades for this agricultural season (Philip *et al.*, 2018; Sohnesen, 2020).

If our proxies were accurate measures of production losses, perfectly complementary, and losses the only significant factor according to which realized yields would vary over the years, we would expect the average attainable harvest to be roughly comparable in size over the years. Instead, the average attainable yield, i.e. the height of the stacked bar, considerably varies across rounds and is particularly high in 2015/16. This should most likely not be taken as an indication that yields would have been exceptionally high if it had not been for the drought. Rather, it is evidence in favor of some overlap between our planted area and plant productivity loss measures and resulting overreporting of losses. This may be related to the salience of losses with small losses perhaps more likely to be overlooked or underreported in years without severe shocks but large losses at the same time overreported if shock salience is high as was likely the case in 2015.

Table 3: Loss-to-harvest ratio across survey rounds and crops in Ethiopia.

			Area Harvested %	Plant productivity reduction %	Kg	Kg/Ha
Maize	ESS13	Loss/Harvest Ratio	91.5	18.1	0.27	1.01
		Loss/Harvest Ratio (any loss)			0.28	0.96
	ESS15	Loss/Harvest Ratio	88.9	25.3	0.61	1.44
		Loss/Harvest Ratio (any loss)			0.58	1.2
	ESS18	Loss/Harvest Ratio	90.1	18.1	0.36	0.94
		Loss/Harvest Ratio (any loss)			0.34	1.01
Sorghum	ESS13	Loss/Harvest Ratio	88.1	19.8	0.66	0.66
		Loss/Harvest Ratio (any loss)			1.79	1.75
	ESS15	Loss/Harvest Ratio	79.4	40	2.08	2.51
		Loss/Harvest Ratio (any loss)			3.32	3.73
	ESS18	Loss/Harvest Ratio	88.4	20	0.63	1.61
		Loss/Harvest Ratio (any loss)			0.53	1.16

Note: The table reports the ratio of estimated losses to self-reported harvest on maize and sorghum plots for the whole sample and only for plots who have had any loss, by survey round.

Our graphical analysis is further complemented by Table 3. Generally, there is little variation between rounds in the average percent plant productivity reduction and percent area harvested. However, compared to the two other survey rounds, the Ethiopia Socioeconomic Survey 2015/16 (ESS15) clearly stands out as

- More plots had any plant productivity reduction (irrespective of crop)
- More sorghum plots lost some of the planted area
- The average size of losses for maize (plant productivity loss) and sorghum³¹ (plant productivity loss and planted area loss) was notably larger.

Conspicuously though, while our losses proxies tell a clear story consistent with the notion of a large drought in 2015/16, this is not supported by average self-reported or crop cut yields. This finding is not unique to our analysis and is extensively discussed in Sohnesen, (2020) who finds that while meteorological data and self-reported shock reports are consistent with a severe drought, this is not supported by vegetation indices and crop yields. In our case, higher than usual reports of area and productivity reductions combined with yet relatively normal realized production leads the loss-to-production ratio to be considerably larger in ESS15 with a few unrealistically large outliers in loss estimates (cases where the loss per hectare exceeds the highest realized yield on any plot in the sample for the respective year).

³¹ Higher losses for sorghum are not in line with the literature on crop drought resilience though and should thus be a reflection of sorghum being planted more often where drought (and thus losses) is/are more severe. This is confirmed by comparing GIS rainfall data to 14-year average rainfall - for maize and sorghum plots, respectively.

One possible reason for losses to be inflated would be if the two proxies were not complementary but overlapped instead. The accuracy of the loss estimate crucially hinges on farmers being able to clearly differentiate between planted area- and plant productivity losses. For example, the differentiation between planted area and plant productivity losses may not always be as clear cut as in theory and may result in farmers not treating the questions on the share of planted area harvested and the percentage of damage on the crop as independent, complementary questions. As a result, there may be some overlap between both questions which would lead to our loss estimate being inflated. Another potential source of error is heterogenous plant productivity reductions across the plot. The question for productivity reductions we are using relies on farmers' ability to accurately report the *average* plant productivity loss in harvested areas. However, productivity reductions may be very heterogenous affecting some plants more and others less, especially when the plot is large. Forming an accurate estimate of the average percentage share of productivity reductions is thus arguably difficult. As a result, farmers may struggle to report the average plant productivity reductions accurately and rather satisfice.

Table 4: Yield regression with area and plant productivity loss measures.

ETHIOPIA: Yield regressions with area- and productivity-based losses - Excl. Full Losses						
	(1)	(2)	(3)	(4)	(5)	(6)
	ESS13	ESS15	ESS18	ESS13	ESS15	ESS18
VARIABLES	Maize	Maize	Maize	Sorghum	Sorghum	Sorghum
Planted area harvested (%)	0.00567** (0.00238)	0.00359* (0.00213)	0.00813*** (0.00292)	0.00713 (0.00551)	0.00329 (0.00252)	0.00415* (0.00224)
Damage on crop (%)	-0.00862*** (0.00169)	-0.0114*** (0.00170)	-0.00708*** (0.00265)	-0.00838*** (0.00201)	-0.0168*** (0.00267)	-0.00541** (0.00266)
Constant	4.659*** (0.352)	4.640*** (0.284)	4.944*** (0.431)	4.776*** (0.604)	5.182*** (0.371)	5.406*** (0.313)
Observations	2,853	2,846	1,586	2,370	2,291	1,230
Adjusted R-squared	0.370	0.346	0.190	0.284	0.391	0.318
District FE	YES	YES	YES	YES	YES	YES
Controls	FULL	FULL	FULL	FULL	FULL	FULL

Note: Regression of log yield on loss proxies. Control variables: Inputs - log plot area, log total labor on the plot per hectare, hired labor dummy, log seed quantity per hectare, improved seed dummy, log inorganic fertilizer quantity per hectare, organic fertilizer dummy; Household Characteristics - dependency ratio, dummy for any HH member with at least primary education, agricultural asset index; Manager Characteristics - female dummy, age, dummy for at least primary education; Plot - intercropping dummy, soil fertility index; and a number of district fixed effects. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1"

While this paper cannot determine the degree of overlap and inaccuracy with certainty, it is possible to explore whether both proxies independently add explanatory power to a yield regression and carry the expected signs. This is confirmed in Table 4, however, the increase in explanatory power when controlling for productivity reductions on top of area reductions is relatively small. Furthermore, the actual loss estimate carries less predictive power than the percentage-based area and productivity reduction variables (not reported). Therefore, while the area and productivity measure seem somewhat complementary and serve as suitable proxies for

different types of losses, combining them does not seem to lead to an accurate estimate of crop losses.

What do we conclude and still need to know?

We find that proxying for losses on a percent scale captures ordinal differences in losses fairly well, i.e. their rough magnitude, however, they are inherently hard to handle and lack a tangible benchmark for farmers. This makes them prone to measurement error where, in the eyes of a farmer, even seemingly small inaccuracies in their reports of percent area and productivity reductions create considerable error in absolute terms. Furthermore, there is a risk of overlap between the two measures that is particularly high if shocks were more severe. At the same time, small losses may go undetected or may be underreported in years with low shock intensity. As a result, we observe an overestimation of losses in years with severe shocks. The accuracy of our measurement may be further compromised by very heterogeneous plant productivity reductions which would make it hard to report a single, averaged productivity reductions figure. The two measures thus serve as approximations of the magnitude of losses but giving them an interval-scale interpretation is less accurate.

Therefore, by measuring losses due to reductions in the planted area that was harvested and losses due to reductions in productivity, respectively, we gain complementary information on the size of crop losses. Together, the two measures also correctly identify years in which losses were conceivably particularly high. However, the questions are prone to satisficing by farmers and can cause large measurement error in absolute loss estimates. Consequently, the survey questions used may need improvements to be suitable for actually inferring the attainable harvest (and thus loss). Several (non-exclusive) options emerge:

- Improving the phrasing of both proxies, clearly differentiating planted area losses from productivity reductions
- Providing (visual) aid in handling percentages more accurately and consistently
- Directly asking for losses by merging both proxies into a single question eliciting the loss benchmarked against what farmers believe could have been their attainable harvest

Question 2: How do farmer reports compare to crop losses based on expert assessments and can we establish an objective, reliable benchmark ('gold standard') for crop losses?

Eliciting losses directly from farmers may be unbureaucratic and easily scalable, however, the previous analysis has indicated that it comes with a number of potential drawbacks. In order to better assess the accuracy of farmers' self-reports, there is a need to understand the way in which farmers form their estimates and compare them to a credible benchmark or 'gold standard'. For example, methodological research comparing self-reports to an objective measure uncovered systematic measurement error in farmers' reports of plot size and harvest quantity (Carletto, Savastano and Zezza, 2013; Carletto, Gourlay and Winters, 2015; Gourlay, Kilic and Lobell, 2019). Using data from the Uganda MAPS experiment, we can compare farmers' loss reports to

losses as assessed by the enumerator/crop cutter and gain insights into potential heuristics used by farmers to form their estimates.

In the Uganda MAPS experiment, two questions, one covering losses on the extensive margin and another one the intensive margin, are answered by both the enumerator (at the crop cut visit) and the farmer (at the post-harvest visit).


- Extensive margin losses: *Was there any damage to the maize crop prior to harvest?*
- Intensive margin losses: *What percentage of the potential maize harvest was lost due to damage on [PLOT]?*

We can therefore compare the answers farmers give to each of the questions to those of the enumerator (the “expert”) and verify their alignment. Good alignment may speak in favor of the accuracy of self-reports and increase our confidence in the ability of enumerator reports to serve as a ‘gold standard’.

What do we see in the data?

On average, we find evidence for systematic overreporting of losses by farmers compared to enumerators both on the extensive and intensive margin. Both farmers (95 percent of plots) and experts (89 percent), overwhelmingly report that there was at least some damage on the plot suggesting that crop losses in the MAPS sample are the norm. While agreement on the extensive margin is good with no statistically significant difference between farmer and expert reports, farmers report significantly higher losses when moving to the intensive margin, that is, in terms of the percentage of crop lost per plot. On average, enumerators/experts report losses to be at 21 percent compared to 34 percent when asking farmers (Figure 3).

Figure 3: Expert-assessed damage compared to farmer reports



Next, we run simple bivariate regressions of expert reports on farmers' reports to assess their correlation. Table 5 finds farmers' reports to be statistically significantly associated with enumerators' assessments. Agreement, however, is imperfect and the R^2 particularly low for whether there was any damage. The fit is somewhat better on the intensive margin, yet Figure 4 illustrates that we observe quite a few cases in the data where farmers report a high loss but enumerators barely any loss. This is surprising as such discrepancies lie outside any reasonable boundaries for the expectable margin of error when comparing the assessment of two different people (the farmer and the enumerator). A tentative explanation of this phenomenon may relate the discrepancy to difficulties in adequately handling the percentage scale (see discussion in Question 1), likely by farmers since the enumerator should have more routine in conducting such assessments. At the same time, large discrepancies in the assessment of losses may also drive the low observed fit on the extensive margin. In cases where enumerators do not report any loss but farmers do, farmer-reported losses are not marginal in size but rather cluster between 10 percent and 40 percent. This may point to a heuristic consistent with satisficing in which farmers simply report a number that appears reasonable if they are uncertain.

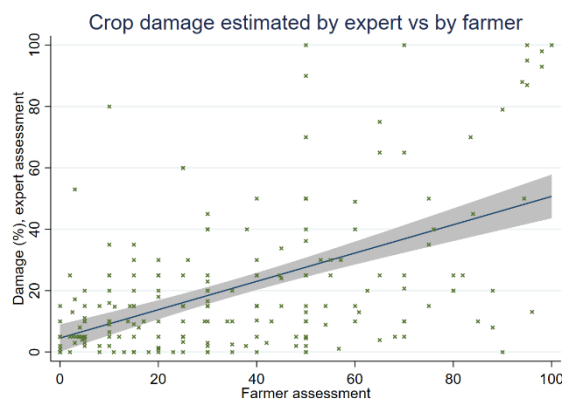
Table 5: Regression of expert-assessed on farmer-assessed damage.

Bivariate regressions of expert assessed damage on farmer assessment		
VARIABLES	(1)	(2)
	Expert assessment	
Damage on plot (Y/N), farmer assessment	0.270* (0.147)	
Farmer assessment		0.461*** (0.0690)
Constant	0.636*** (0.146)	4.544** (1.933)
Observations	213	213
Adjusted R-squared	0.032	0.276

Note: The table shows a bivariate regression (linear probability model) of expert assessments of whether there was any damage on the plot based on farmer assessments (Column 1) and expert assessments of the percentage amount of damage on the crop on farmer assessed values.

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Figure 4: Percent of damage on the plot, expert vs. farmer.



To further investigate what may explain the differences in the size of reported losses, we regress the log-ratio and absolute difference in farmers' and enumerators' reports on a number of potential drivers (Annex Table A 3). We do not find any significant factors relating to plot or plot manager characteristics. Notably, there does not seem to be a relationship between farmers' education and the agreement with expert assessments. However, over half of respondents in the MAPS sample have completed five or more years of education, making them better educated than farmers in the Ethiopia or Malawi sample.

On the other hand, it seems that when damage occurred due to drought (based on farmers' self-reports), farmers' overestimation is more pronounced. This may be due to drought being a "slow-onset disaster" whose salience depends greatly on its intensity and may be subject to differences in perception. What enumerators classify as a drought may thus differ from farmers' perceptions, especially since the former likely are basing their assessment on comparisons with other farms

surveyed. Importantly, earth observation data suggest that 2016, the year the MAPS data was collected, was a drought year (FAO, no date). The data may thus reflect a similar situation as in 2015 in Ethiopia where self-reported losses were conspicuously high compared to average yields and suggest that in drought years we observe an “overreaction” in farmers’ reports.

What do we conclude and still need to know?

While farmer reports and expert assessments show some overlap, they do not match very well and display some striking discrepancies that are not explainable by small differences that would be expected in the appraisal process. This may serve as further suggestive evidence that while eliciting losses on a percent scale gives a decent ordinal proxy in most cases, forming an exact point estimate is cognitively challenging, may lead to satisficing behavior, and is thus prone to error. This seems to be exacerbated in drought years where farmers’ reports are particularly high. At the same time, we cannot ascertain the accuracy of enumerator reports with confidence either as enumerators may err just as well when it comes to losses at the extensive and intensive margins. This makes it hard to proclaim enumerator assessments a credible ‘gold standard’. In order to better understand the accuracy in farmers’ self-reported losses, finding a dependable, “objective” measure of losses is thus still an open question for future methodological research.

Question 3: Can data from past harvests or unaffected plots serve as a suitable benchmark for attainable harvest?

The FAO methodology heavily relies on the ability to obtain an accurate estimate of the attainable yield. For this, it suggests drawing on past harvest data to calculate the multi-year trend (or average) of realized yields as an indication of in-season attainable yields (Conforti, Markova and Tochkov, 2020). Alternatively, using same-season realized yields from plots unaffected by any disaster, for example average yields from an agriculturally comparable district, would be a possibility (FAO, 2020).

Simplifying Equation (1) from the framework in Section 3 and using that

$$y_{i,j,t}^r = \frac{Y_{i,j,t}^r}{ha_{i,j,t}^r},$$

we can write

$$L_{i,j,t} = p_{i,g,t-1} \left(f \left(\sum_{z=1}^n y_{i,j,t-z}^r \right) ha_{i,j,t}^c - Y_{i,j,t}^r \right) + \gamma \quad (3)$$

where the counterfactual yield y^c is a function of past yield realizations or

$$L_{i,j,t} = p_{i,g,t-1} \left(f \left(\sum_{\substack{a=1 \\ a \neq j}}^n y_{i,a,t}^r \right) ha_{i,j,t}^c - Y_{i,j,t}^r \right) + \gamma \quad (4)$$

where y^c has been replaced by a function of the realized yields on other plots that share similar characteristics to the plot in question but did not incur any losses.³² Multiplying this by the planted area ha^c thus gives the attainable production. Losses then constitute a deviation of the realized harvest Y^r from what would have been expected based on past years or the performance of yields on plots in unaffected areas, respectively.

The remainder of this section looks at the data requirements of both approaches as well as their likely accuracy.

What do we see in the data?

We start by regarding the suitability of using past yield data to proxy for attainable yields from a data availability perspective. In order to ascertain a robust trend or average in yields, we require high frequency (e.g. annual) panel data reaching back several years. Micro-level panels of this frequency are exceedingly rare and, in our case, the LSMS-ISA panel data only cover recurring survey rounds every two to three years. While such high frequency data is currently unavailable at a sufficiently disaggregated level³³ it may become available in the future.³⁴

Another challenge in this regard relates to split-off households and those that move. The LSMS-ISA survey in both Ethiopia and Malawi tracks households that move to a new place whereas the latter also tracks households that split from the originally sampled household. For both groups of households, comparable past information on yields would be missing.³⁵

Irrespective of this, yields in the context of the LSMS-ISA data are measured at the plot level. Ideally, the panel data used would thus also be identified at the plot-crop level in order to hold differences in soil quality and other plot characteristics constant. However, plot boundaries and crop choice frequently vary across years and the LSMS-ISA data are thus household-level (and in the case of Malawi and Ethiopia parcel-level) panels but not necessarily plot-level panels (Carletto, Gourlay and Winters, 2015). As a result, it would likely be necessary to move to farm-level aggregate yields in order to reduce the idiosyncratic uncertainty associated with varying plot management decisions and growing conditions.

More fundamental than these concerns on data requirements is a conceptual issue with relying on past yield data. By definition, we only observe past yield *realizations* which will differ from attainable harvest in the presence of crop losses. In Ethiopia and Malawi, however, harvest losses are very common affecting, for instance, between 55 percent and 65 percent of maize plots across years in Malawi and between 41 and 56 percent of maize plots in Ethiopia. As such, any past

³² As a reminder, Y^r and ha^c are always observed in the LSMS-ISA data. They are the realized production and GPS measured area, respectively.

³³ Data sources such as FAOSTAT's yield database, for example, contain yearly data but cover the country-level which may be too coarse to accurately accommodate sub-national, sub sub-national, or even micro-level variation.

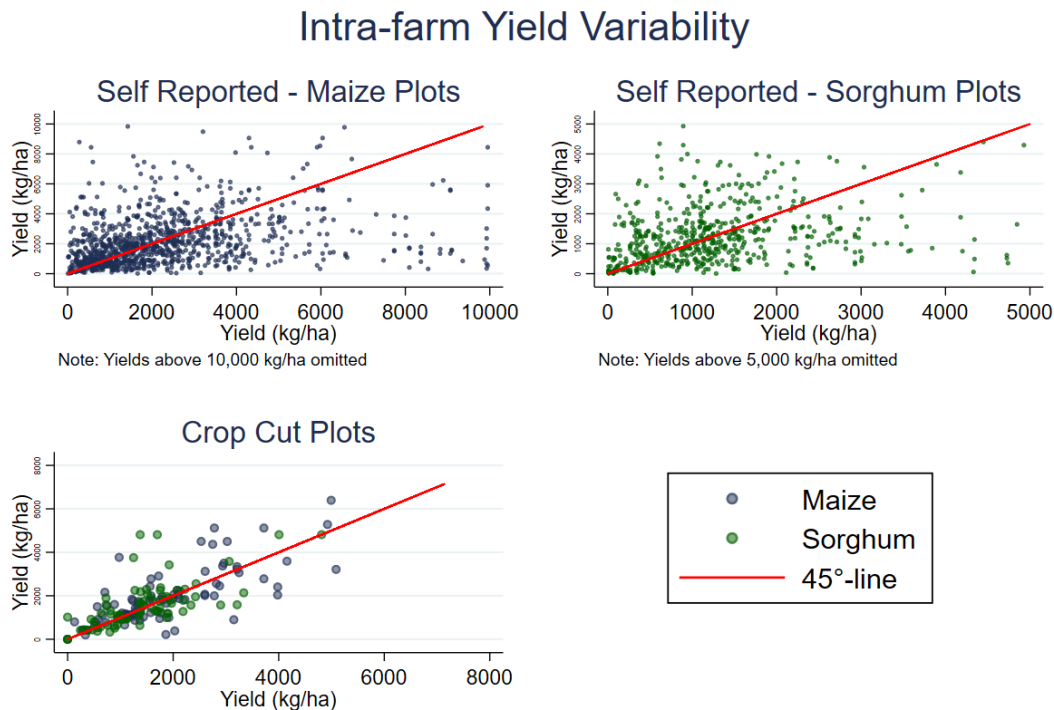
³⁴ Projects such as the 50x2030 initiative are planning on collecting agricultural statistics in recurring patterns including a core of agricultural data at yearly frequency (50x2030 Initiative, 2020).

³⁵ Between the first and third round of the Integrated Household Panel Survey (IHPS) in Malawi, for example, almost 40 percent of households moved to a different location.

harvest data is likely to provide a downward-biased estimate of the attainable harvest and carry over the effect of past shocks into the current-season losses estimate.³⁶

A possible remedy for this would be to use the maximum achieved yield in the past in the hopes that the same farmer (the same farm) has gotten at least close to the attainable yield before. However, there is no guarantee this has been the case and the method may simply pick up outliers due to measurement error. Furthermore, this would only be a suitable option if all other potentially confounding factors such as input intensities would be somewhat similar across rounds and thus the only remaining difference is constituted by the existence of exogenous shocks (or absence thereof). However, the literature finds yields to vary substantially even within the same farm over time with differences in crop management playing an important role on top of idiosyncratic shocks (Ronner *et al.*, 2018; Silva *et al.*, 2018; van Loon *et al.*, 2019).

Figure 5: Ethiopia: Intra-farm yield variability based on self-reports and crop cuts.



Alternatively, one could consider (average) realized yields in agriculturally comparable districts known to have been unaffected by large scale disasters as a proxy for attainable yields in disaster-affected districts. Similar to the use of past yield data, the accuracy of this method crucially depends on the degree to which factors other than disaster impact are constant between disaster-affected and the “control” plots. This may not be the case as the literature suggests the presence of large inter-farm differences in yields due to idiosyncratic shocks but notably also measurement error, varying input intensities, and soil quality (Gollin and Udry, 2021). This is furthermore

³⁶ Naturally, this is of particular concern when the time dimension of the panel data is small as is the case with most current agricultural household survey data.

highlighted in the data by comparing yields across different plots of the same farm. Comparing plots of the same farm should hold many of the aforementioned factors constant. However, as Figure 5 emphasizes, there is considerable heterogeneity in yields even within the same farm and even when only regarding pure-stand plots without losses or comparing crop cut yields on plots without losses.

What do we conclude and still need to know?

These considerations highlight the difficulties in finding a valid counterfactual based solely on realized yields, either exploiting the time or cross-sectional dimension of survey data. Using past harvest observations as a proxy for current season attainable harvest requires panel data of higher frequency than is common in current agricultural surveys and would not capture any households that have changed location or entered the panel. Fundamentally though, they would provide an estimate of attainable harvest that is downward biased to the degree that shocks have also occurred in previous survey rounds, a likely scenario based on our sample, and be prone to bias from inter-temporal differences in plot management decisions. Furthermore, the large inter-farm variability in yields due to factors other than shocks that is reported in the literature and heterogeneity in yields across plots of the same farm in our data cast doubt on the suitability of using realized yields from other farms to proxy for the attainable harvest.

However, using rich survey data as is provided by the LSMS-ISA, it may be possible to account for possible confounding factors and predict (attainable) yields using econometric or machine learning models based on past yield data, data from several similar but unaffected plots, and/or planting season crop management decisions. An exploration of this possibility is beyond the scope of this paper but may be a relevant avenue for future research.

Question 5: Can farmers' post-planting harvest expectations proxy for attainable harvest?

As previously discussed, the LSMS-ISA data contains detailed information on plot, input, and management characteristics as well as agricultural production down to the plot-crop level. This is an important feature as accounting for these differences acknowledges the fact that yields can vary greatly between different plots. For example, different field characteristics and management practices can create variation not only in the realized yield but also the attainable yield. Furthermore, taking Ethiopia as an example, between 35 percent and 44 percent of plots (depending on survey round and crop) are intercropped. This means that we may not have accurate information on the actual share of the plot that was planted with the crop and thus the attainable production (in absolute terms).³⁷

To potentially address these pitfalls, we can rely on farmers' ability to somewhat accurately estimate the attainable harvest for each crop on each plot at the post-planting visit. At this point,

³⁷ The LSMS-ISA data does contain a question that asks for the share of the plot that was planted with the crop for mixed stand plots. However, the estimate is rather coarse (less than a quarter, 50 percent, 75 percent, more than 75 percent).

the majority of crop management, plot preparation, and other input decisions should have been taken. Asking farmers for their expected harvest³⁸ may thus give an estimate that takes into account between-plot heterogeneity in the aforementioned factors. We could then define a loss as an “expectation gap”, the (positive) difference between expected and realized production.³⁹ Simplifying the general form of the loss equation in (1) (and dropping subscripts for simplicity) we can thus write

$$L_{i,j,t} = p_{i,g,t-1}(y^c ha^c - y^r ha^r) + \gamma \quad (5)$$

where the first term in parentheses is the expected production taken from farmers’ reports at the post-planting visit and the second term is realized production taken from their reports at the post-harvest visit.

Naturally, using farmers’ post-planting harvest expectations is not without its challenges and drawbacks either. Most importantly, farmers’ expectations are endogenous to past shock experience and future shock expectations as discussed in Section 2. For example, farmers may adjust expectations according to the amount of information (e.g. accurate weather forecasts) they have (Rosenzweig and Udry, 2019) or in light of recent shock exposure (Karlan *et al.*, 2014). Inevitably, this will contaminate the expected harvest as an unbiased estimate of the attainable harvest without shocks. Another fundamental question is whether expectations are consistent with actual behavior, i.e. whether they accurately relate to the plot management decisions that determine the attainable harvest. Furthermore, farmers’ expectations may also not capture any losses that already occurred during the planting season and before they were surveyed or suffer from downward bias as a result.⁴⁰ Lastly, the difference between expected and realized harvest may be too coarse an estimate of crop losses at the micro-level and for small losses.⁴¹

In assessing the suitability of proxying for the attainable harvest through farmers’ expectations, we will pursue a number of questions.

³⁸ In both countries (with small variations) farmers are asked: “How much of [CROP] do you expect to harvest during the [MAIN AG SEASON]?”

³⁹ Naturally, there is a possibility that realized harvest will exceed expectations. For the sake of this analysis, we have to interpret these instances as cases where no losses happened. This may not necessarily be true for small losses; however, it is conceivable that the proxy at least captures substantial losses.

⁴⁰ For instance, it is conceivable that a disaster, particularly one that hits suddenly such as a storm, may have occurred during the planting season. This may create a situation in which re-planting is necessary, but farmers manage on time for the growing season so that those losses are not reflected in farmers expectations (such “hidden” planting season losses are of course a potential source of bias in all approaches discussed in this paper). Alternatively, farmers may choose to not perform any post-disaster management on the plot resulting in their expected production being low. In this case, farmers’ expectations will lead to downward biased estimates of the attainable harvest (and thus disaster crop losses) since they are endogenous to planting season shocks. We drop those cases where farmers do not expect to harvest anything and do not end up harvesting anything, interpreting these cases as planting season full losses that the methodology in Question 4 cannot account for.

⁴¹ For example, crops that perform unexpectedly well can, by this definition, not record a loss if realized harvest outperformed expected harvest. However, one could easily imagine a crop performing very well in a given year and still recording some minor damage. Similarly, the unavoidable measurement error associated with using farmers’ self-reported expectations (in a unit of their choosing) and taking any positive deviations from realized yield as an estimate of the loss may make the measure too coarse to accurately capture small losses. Conversely, substantial losses as do matter for aggregate measurement of losses and policy should likely be captured by the proxy.

- Do expectations regularly exceed harvest realizations and does an overestimation coincide with the occurrence of a loss?
- Are expectations aligned with the actual management decisions (e.g. input intensities) made on the plot that determine the harvest that is attainable?
- Is the size of the (positive) expectation gap a reasonable estimate of the quantity lost?
- How are expectations formed?

What do we see in the data

Figure 6: Mean realized vs. expected harvest.

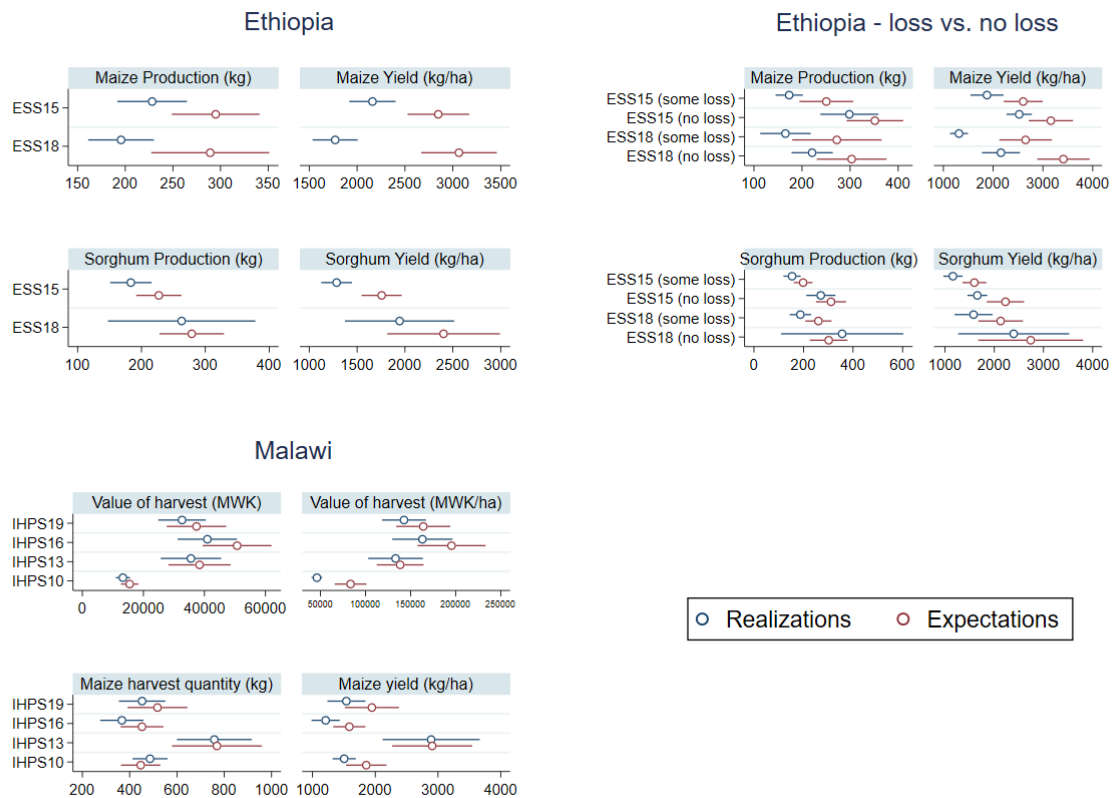


Figure 6 plots the mean expected and realized crop production and yield by country and survey round. There is indicative evidence for systematic over-expectation of harvest if both expectation and realization are based on farmers' self-reports. However, large confidence intervals mean that oftentimes we cannot reject the null hypothesis that the means are equal. Further, if we split the sample between plots with and without losses (based on self-reports), realized *and* expected yields are higher for those plots with no losses suggesting that part of the loss is already anticipated at the post-planting visit.⁴² There is also a tendency for expectation gaps, the difference between the

⁴² We can take further evidence for the adjustment of expectations based on damage that has already occurred at the post-planting visit by fitting a local polynomial of harvest expectations against self-reported crop damage at the

mean harvest expectation and mean realized harvest, to be larger where there was at least some loss. These observations provide suggestive evidence of an association between over-expectation and loss occurrence.

Taking an over-expectation as an indication of a loss, between 49 percent (ESS15, maize) and 57 percent (ESS18, maize) of plots record a loss. Notably, this also results in less variation between crops and there is no clear pattern of alignment between the loss incidence based on planted area and/or plant productivity losses and based on harvest expectations exceeding realizations (Table 6).⁴³

Table 6: Loss incidence across different proxies.

Percent of plots with crop losses based on..						
Country	Crop	Round	Area	Damage	Area or damage	Overexpectation
Ethiopia	Maize	ESS13	12.7%	40.9%	41.4%	
		ESS15	19.4%	55.7%	56.2%	48.7%
		ESS18	17.1%	43.3%	45.3%	56.5%
	Sorghum	ESS13	21.2%	45.9%	46.7%	
		ESS15	33.8%	74.5%	75.0%	49.8%
		ESS18	25.5%	53.3%	54.8%	53.3%
Malawi	Maize	IHPS10	56.5%			36.3%
		IHPS13	55.1%			44.6%
		IHPS16	65.3%			60.0%
		IHPS19	61.0%			51.8%

Note: Crop loss incidences (percent) according to different proxies, by survey round and crop: area harvested < area planted (area loss), damage on the crop (damage), any area or damage-related loss, or expected production < realized production (overexpectation). ESS = Ethiopia Socioeconomic Survey, IHPS = Integrated Household Panel Survey.

For post-planting expectations to provide a realistic proxy of the attainable harvest, they need to be aligned with the actual management decisions on the plot.⁴⁴ Table A 4 and Table A 5 ascertain this through two equivalent yield regressions for each survey round and crop: one using realized harvest as the outcome variable and one using harvest expectations. They generally confirm that the alignment of expectations with production factors, based on the model's R^2 , is generally equally good or better than for realized harvest. This increases confidence in farmers' expectations' alignment with their reported behavior. However, some notable exceptions in Ethiopia (ESS15 across both crops) and Malawi (IHPS10) exist.

post-planting visit (not reported). This, again, suggests that some of the loss already occurred before expectations were elicited which will likely bias them downward.

⁴³ For cross-country, longitudinal figures on crop and livestock losses in urban vs. rural areas see (Chaudhary, 2021).

⁴⁴ This is because farmers' input use and management practices place an effective upper boundary on the harvest that can be achieved. If farmers thus declare one thing as their harvest expectation but their input use is inconsistent with this, the relationship between post-planting expectation and realized harvest is distorted by factors other than shocks, rendering them an unsuitable proxy for the attainable harvest.

Table 7: Ethiopia - Regressions of self-reported harvest realizations on expected harvest and a range of controls.

Ethiopia - Yield

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	ESS15				ESS18			
	Maize yield (kg/ha)		Sorghum yield (kg/ha)		Maize yield (kg/ha)		Sorghum yield (kg/ha)	
Expected maize yield (kg/ha)	0.183*** (0.0416)	0.126*** (0.0349)			0.0997** (0.0460)	0.0652** (0.0330)		
Expected sorghum yield (kg/ha)			0.256*** (0.0387)	0.124*** (0.0324)			0.290** (0.113)	0.197* (0.107)
Planted area harvested (%)		10.30*** (2.943)		3.926** (1.732)		8.554*** (2.776)		10.69** (4.853)
Damage on crop (%)		-8.277*** (3.137)		-11.99*** (1.810)		-8.818*** (2.813)		-6.854 (5.473)
Constant	1,640*** (109.0)	1,587 (1,021)	837.0*** (82.59)	1,902*** (558.4)	1,465*** (102.1)	-447.6 (1,013)	1,246*** (270.3)	2,846 (3,272)
Observations	2,928	2,841	2,397	2,336	1,685	1,629	1,302	1,240
Adjusted R-squared	0.086	0.217	0.171	0.341	0.044	0.199	0.083	0.223
Controls	NO	YES	NO	YES	NO	YES	NO	YES

Note: Regression of self-reported realized yield on expected yield, planted area and plant productivity loss, and a number of plot-, manager-, and household controls as well as inputs and zone fixed effects. Controls: Inputs - Plot area, total labor p. ha, any hired labor dummy, seed quantity p. ha, improved variety dummy, inorganic fertilizer quantity p. ha, any organic fertilizer dummy; Household - dependency ratio, any HH member has primary education dummy, agricultural asset index; Manager - female dummy, age, primary education dummy; Plot - intercropped dummy, precipitation of wettest month, soil fertility index; Zone fixed effects. Clustered standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 7 and Table 8 explore the alignment of expectations with harvest realizations by regressing realized harvest on expected harvest, first in a bivariate regression and later including a full set of controls. Expectations prove a significant predictor of self-reported realized harvest with the association much stronger in Malawi. Furthermore, consistent with the notion of systematic over-expectation, coefficients are positive but less than one. Adding the planted area- and plant productivity loss measures as additional covariates to the model, we gain some but not much explanatory power with the coefficients carrying the expected signs.⁴⁵ This again indicates that the two measures (one in Malawi) are associated with losses but do not provide exact point estimates.

⁴⁵ We verify this result graphically by comparing the mean (positive) expectation gap to the mean estimated loss from Question 1 by country, crop, and survey round and by plotting the expectation gap against the two different losses measures (area and plant productivity losses in percent), respectively. Alignment is imperfect and the general pattern consistent with the finding of Sohnesen (2020) that farmers' shock (loss) reports and realized yields (and thus the expectation gap) tell two different stories.

They therefore do not explain much of the remaining variance when already controlling for the attainable harvest through expectations.⁴⁶

Table 8: Malawi - Regressions of self-reported harvest realizations on expected harvest and a range of controls.

Malawi - Maize Yield								
VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	IHPS19		IHPS16		IHPS13		IHPS10	
	Bivariate	Multiv.	Bivariate	Multiv.	Bivariate	Multiv.	Bivariate	Multiv.
Expected maize yield (kg/ha)	0.388*** (0.0768)	0.214** (0.0826)	0.536*** (0.0546)	0.398*** (0.0472)	0.687*** (0.0883)	0.315*** (0.0991)	0.136*** (0.0239)	0.0618*** (0.0196)
Planted maize area harvested (%)		6.385*** (1.642)						
Constant	784.3*** (115.0)	783.3 (1,684)	359.7*** (70.56)	-1,836 (1,344)	894.4*** (264.2)	10,025 (7,343)	1,252*** (86.41)	3,271*** (934.1)
Observations	568	544	493	490	442	442	842	842
Adjusted R-squared	0.373	0.509	0.385	0.448	0.468	0.556	0.090	0.333
Controls	NO	YES	NO	YES	NO	YES	NO	YES

Note: Regression of self-reported realized maize yield on expected maize yield, planted area loss (only IHPS19), and a number of plot-, manager-, and household controls as well as inputs and district fixed effects. Controls: Inputs - Plot area, total labor p. ha, any hired labor dummy, seed quantity p. ha, improved variety dummy (only IHPS19 and IHPS15), inorganic fertilizer quantity p. ha, any organic fertilizer dummy; Household - dependency ratio, head employed for wage dummy, any HH member owns non-ag enterprise dummy, any HH member has primary education dummy, agricultural asset index; Manager - female dummy, age, primary education dummy; Plot - intercropped dummy, precipitation of wettest month, soil fertility index; Zone fixed effects. Clustered standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

As a final exercise, Table 9 and Table 10 investigate the factors that underlie farmers' expectation building. Neither a set of respondent characteristics nor past shock exposure or other factors potentially associated with the process of forming an expected harvest estimate are significant predictors of the expectation gap across crops and countries.⁴⁷ We gain some explanatory power by constraining the sample to those plots where expectations exceeded realized harvest, i.e. those

⁴⁶ In Ethiopia, we also have data from crop cut harvests for a sub-sample of plots. This allows us to verify whether expectations are also strong predictors of what is commonly taken as the "gold standard" measure of realized harvest. We find that expectations do not seem to be a reliable predictor for crop cut harvest, and in particular crop cut yields. Furthermore, this does not seem to be due to bias only inherent in self-reported expectations as there is only a slightly stronger association even between self-reported harvest *realizations* and the crop cut measure. We thus conclude that expectations only provide a fair proxy of the attainable harvest when also using self-reported harvest realizations. The reason for this is likely that both share a similar underlying mental model and thus similar biases that cancel out when comparing them.

⁴⁷ The fact that we do not find a robust effect of past shock exposure, measured either by community reports for the past two years, on expectation building is somewhat surprising as there is strong evidence in the literature that past shock exposure affects risk attitudes and expectations about the likelihood of future shocks. This may be because the shock proxies used are not specific enough, i.e. they do not necessarily measure adverse experiences with past *agricultural shocks* and may not capture individual exposure.

where we hypothesized that some crop may have been lost. However, the adjusted R^2 remains low throughout. This is not unsurprising since there are likely complex behavioral processes underlying farmers' expectation building that are hard to model in this simple framework and without a (quasi)-experimental set-up. Notably, adding the percent-based planted area loss and plant productivity reduction estimates (not reported) adds a significant amount of predictive power to the model, however, this is driven by plots with full losses where the association is mechanical.

Table 9: Ethiopia - Drivers of over-expectation

Ethiopia	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Maize				Sorghum			
	ESS18		ESS15		ESS18		ESS15	
	Base	Overexp.	Base	Overexp.	Base	Overexp.	Base	Overexp.
Community reported drought in last 2y	0.351 (0.287)	0.504 (0.338)	-0.105 (0.127)	-0.0268 (0.151)	0.770*** (0.220)	0.421 (0.306)	0.0277 (0.124)	0.231 (0.140)
Community reported flood in last 2y	0.0318 (0.272)	0.0832 (0.268)	-0.0828 (0.168)	-0.331* (0.181)	0.362 (0.296)	0.395* (0.209)	-0.422* (0.242)	-0.382* (0.225)
Community reported crop disease in last 2y	0.0349 (0.182)	-0.181 (0.202)	-0.0892 (0.103)	-0.112 (0.113)	-0.292 (0.204)	-0.111 (0.252)	0.0540 (0.152)	0.266 (0.215)
log Plot area (ha)	-0.122** (0.0475)	-0.219*** (0.0512)	0.0356 (0.0292)	-0.0374 (0.0295)	-0.0198 (0.0377)	-0.106** (0.0501)	0.0369 (0.0450)	0.0531 (0.0470)
Intercropped	0.178 (0.142)	0.228 (0.163)	-0.0526 (0.0835)	0.0812 (0.0959)	-0.0490 (0.104)	-0.0437 (0.129)	-0.227** (0.109)	-0.188 (0.120)
Soil fertility index	-0.111 (0.0870)	-0.224** (0.108)	-0.0514 (0.0515)	-0.0325 (0.0557)	-0.00537 (0.117)	-0.128 (0.195)	-0.00517 (0.0781)	0.0126 (0.0906)
Improved seed	0.489*** (0.168)	0.558*** (0.199)	-0.144* (0.0863)	-0.183* (0.105)	0.366** (0.171)	-0.0551 (0.185)	0.644 (0.485)	0.870 (0.649)
Precipitation of Wettest Month (mm)	-0.00225** (0.00103)	-0.00289** (0.00135)	-0.00320*** (0.000954)	-0.00430*** (0.00120)	0.00192 (0.00360)	0.00409 (0.00639)	-0.000423 (0.00162)	-0.00209 (0.00210)
Dependency ratio	-0.130*** (0.0490)	-0.119** (0.0517)	0.0187 (0.0302)	-0.00715 (0.0422)	0.0216 (0.0409)	-0.0180 (0.0585)	0.0766* (0.0445)	0.0882 (0.0678)
Any HH member has primary edu	-0.123 (0.0995)	-0.164 (0.139)	0.0325 (0.0747)	-0.0300 (0.107)	-0.0504 (0.101)	-0.0477 (0.108)	0.122 (0.0928)	0.107 (0.126)
Ag Asset Index	0.205*** (0.0672)	0.149* (0.0817)	-0.127* (0.0664)	-0.186** (0.0875)	-0.0152 (0.0745)	-0.0494 (0.103)	-0.101 (0.0801)	-0.179* (0.103)
Female plot manager	-0.0700 (0.122)	-0.131 (0.144)	0.000363 (0.0830)	-0.121 (0.115)	-0.109 (0.0897)	-0.0539 (0.141)	0.0643 (0.148)	-0.0608 (0.195)
Plot manager age	-0.00692** (0.00313)	0.00210 (0.00378)	-0.00121 (0.00270)	0.00205 (0.00338)	0.00454 (0.00310)	0.00394 (0.00312)	-0.00163 (0.00298)	-0.00208 (0.00368)
Plot manager has primary edu	0.0691 (0.0992)	0.225* (0.131)	-0.0506 (0.0989)	-0.000532 (0.119)	0.122 (0.134)	0.125 (0.157)	-0.0871 (0.126)	0.0557 (0.157)
Constant	0.853** (0.395)	1.200** (0.475)	1.210*** (0.325)	2.087*** (0.407)	-0.609 (0.838)	-0.424 (1.417)	-0.0174 (0.443)	1.322** (0.556)
Observations	1,631	930	2,848	1,379	1,247	659	2,347	1,163
Adjusted R-squared	0.080	0.122	0.062	0.134	0.084	0.122	0.084	0.129
Region FE	YES	YES	YES	YES	YES	YES	YES	YES

Note: Regressions of the log of the ratio of expectations (+1) to realizations (+1) on a number of potential correlates. Clustered standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 10: Malawi – Drivers of over-expectation

Malawi (maize)	(1)		(2)		(3)		(4)		(5)		(6)		(7)		(8)	
	IHPS19		IHPS16		IHPS13		IHPS10		IHPS19		IHPS16		IHPS13		IHPS10	
	Base	Overexp.	Base	Overexp.	Base	Overexp.	Base	Overexp.	Base	Overexp.	Base	Overexp.	Base	Overexp.	Base	Overexp.
Community reported drought in last 2	-0.169 (0.298)	0.0694 (0.151)	-0.237 (0.384)	-0.472 (0.471)	-0.103 (0.154)	0.273 (0.188)	-0.211 (0.184)	-0.315** (0.143)								
Community reported flood in last 2y	-0.272 (0.265)	-0.133 (0.115)	1.515*** (0.426)	0.952* (0.514)	-0.136 (0.643)	1.058 (0.698)	-0.184 (0.262)	0.170 (0.179)								
log Plot area (ha)	0.363 (0.422)	0.0870 (0.582)	0.432 (0.375)	0.443 (0.649)	-0.0356 (0.267)	-0.316** (0.145)	-2.174*** (0.429)	-0.591 (0.398)								
Intercropped	-0.0330 (0.139)	-0.159 (0.193)	-0.0593 (0.233)	-0.435 (0.273)	0.105 (0.123)	-0.0713 (0.107)	0.241 (0.230)	0.437** (0.208)								
Soil fertility index	-0.101 (0.0996)	-0.165** (0.0640)	-0.185 (0.277)	-0.0698 (0.289)	-0.0719 (0.136)	0.00453 (0.158)	-0.0278 (0.103)	0.145** (0.0685)								
Precipitation of Wettest Month (mm)	-0.00365 (0.00651)	-0.000472 (0.00457)	0.000458 (0.0100)	0.00362 (0.0116)	0.00392 (0.00862)	-0.00251 (0.00565)	0.00436 (0.00561)	0.00668 (0.00538)								
Plot has improved maize variety	0.0298 (0.123)	-0.0248 (0.146)	0.0311 (0.205)	-0.123 (0.247)												
Dependency Ratio	0.0360 (0.118)	-0.000936 (0.104)	0.0601 (0.116)	0.0347 (0.143)	0.0231 (0.0756)	-0.0943 (0.0952)	0.0685 (0.0764)	0.0155 (0.0673)								
Any HH member has primary edu	0.265* (0.132)	0.204 (0.143)	0.180 (0.327)	0.441* (0.261)	0.111 (0.119)	-0.226 (0.175)	0.0150 (0.175)	-0.0541 (0.168)								
Ag Asset Index	-0.324 (0.330)	-0.443 (0.276)	-1.757** (0.670)	-1.686*** (0.622)	-0.0168 (0.251)	0.192 (0.259)	-1.134*** (0.338)	-0.748* (0.395)								
HH wealth index	-0.753* (0.443)	-0.378 (0.351)	0.0155 (1.055)	0.696 (1.012)	-0.495 (0.357)	-0.712* (0.363)	-0.641* (0.367)	-0.848* (0.471)								
HH head has wage employment	-0.0261 (0.172)	-0.0727 (0.192)	0.149 (0.416)	-0.241 (0.368)	0.255 (0.164)	0.361* (0.187)	-0.132 (0.150)	0.0977 (0.130)								
HH member owns non-ag enterprise	0.231 (0.176)	0.102 (0.125)	0.234 (0.241)	0.218 (0.263)	-0.156 (0.115)	-0.0367 (0.159)	-0.161 (0.175)	-0.187 (0.140)								
Female plot manager	0.103 (0.102)	-0.152 (0.116)	-0.0189 (0.226)	0.197 (0.314)	-0.0552 (0.120)	0.0582 (0.120)	-0.162 (0.129)	0.0647 (0.120)								
Plot manager age	-0.000786 (0.00430)	-0.00153 (0.00404)	0.00427 (0.00774)	0.00146 (0.00981)	0.000357 (0.00443)	0.00523 (0.00315)	0.00129 (0.00372)	0.00765 (0.00481)								
Manager has primary education	-0.110 (0.179)	-0.000140 (0.167)	-0.0524 (0.216)	-0.346 (0.290)	-0.169 (0.114)	-0.0710 (0.102)	-0.0150 (0.189)	0.241 (0.149)								
Constant	0.990 (1.570)	0.888 (1.196)	0.151 (2.905)	0.132 (3.234)	-1.001 (2.130)	1.086 (1.384)	-0.524 (1.471)	-1.195 (1.513)								
Observations	568	296	493	297	442	199	842	306								
Adjusted R-squared	0.029	0.078	0.088	0.242	-0.001	0.078	0.150	0.498								
District FE	YES	YES	YES	YES	YES	YES	YES	YES								

Note: Regressions of the log of the ratio of expectations (+1) to realizations (+1) on a number of potential correlates. Clustered standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

What do we conclude and still need to know?

From our analysis, farmers' self-reported harvest expectations generally prove fair, albeit likely endogenous, proxies of the attainable harvest. Expectations align relatively well with harvest realizations as long as both rely on the same underlying "mental model" (are both self-reports) and therefore likely share the same biases. We find limited evidence for an association between over-expectation and loss occurrence when comparing direct measures of losses (planted area and plant productivity losses) to the gap between expected and realized harvest. Consequently, there remains a sizeable "black box" of (behavioral) factors explaining farmers' expectation building and thus driving the expectation gap. Here, further research drawing on a (quasi-) experimental set-up is

needed to better understand the factors that underlie this mental model and, possibly on top of losses, explain the accuracy of expectations and their suitability to proxy for the attainable harvest.

We conclude that inferring losses from the difference between the expected and realized harvest is a coarse measure and may not be preferable over more fine-grained approaches asking farmers for planted area and plant productivity losses directly. However, expectations may be better suited in the case of full losses (when the methodology from Question 1 is not applicable) or in general when capturing large losses is the primary goal. This is because the expectation gap is likely not sensitive enough to adequately represent small losses but provides a decent approximation of the attainable harvest when small biases matter less, e.g. after severe disasters. In future methodological research, trying to elicit farmers' expectations even more accurately, for example by priming them to think about the attainable harvest without shocks, may furthermore decrease endogeneity in farmers' expectations as proxies of the attainable harvest. Lastly, expectations perhaps may not serve as the sole proxy for the attainable harvest but could be an important feature for econometric models that try to predict (attainable) yields.

4.2. Attribution: Linking losses to specific disasters

Measuring crop losses due to disasters for SDG 1.5.2 and Sendai indicator C-2 crucially hinges on the accurate attribution of losses to specific shocks. Most commonly, survey data has relied on self-reports as a fine-grained but subjective measure of shock exposure. At the same time, technological advancements have increased the accuracy of remote-sensed approaches over the last years, potentially alleviating concerns about measurement error at the micro-level and providing an objective measure. Importantly, self-reported data on shocks and index-based data have frequently been found to be at odds with no clear verdict in favor of any one approach (Meze-Hausken, 2004; Sohnesen, 2020). For instance, survey data may give a better representation of idiosyncratic (highly localized) shocks whereas basing shock occurrence on an index may capture covariate shocks, particularly meteorological disasters, more consistently across farms.

The most common disaster in our study region is drought with 82 percent of its impact in agriculture and accounting for the majority of disaster damages and losses in the sector (FAO, 2018, 2021). This is confirmed by the LSMS-ISA data where adverse rainfalls, either too much or too little rain, constitute the most frequently reported reasons for crop losses. They are therefore the focus of this paper. Our assumption is that these shocks are likely covariate, i.e. they should affect all plots with the same crop within a reasonably small geographic aggregate. Exploring several possible survey-based proxies, the analysis in this section thus aims to ascertain the degree of idiosyncrasy in self-reported data on shock exposure and discern variation due to differences in perception from variation due to actual loss incurrence.⁴⁸

⁴⁸ Farmers' plot-level shock reports are conditional on having experienced a crop loss due to the shock in question. To reduce the risk that we are measuring differences in disaster risk reduction technology between plots rather than the consistency of shock identification, we focus on rainfed plots that constitute the overwhelming majority of our sample.

Question 5: How uniform is farmers' self-reported shock identification of covariate shocks?

In order to confidently attribute losses to disasters based on self-reports, it is necessary to know how reliable farmers' shock identification is. If the main shocks in the data, adverse rainfalls, are really covariate, plots in vicinity should face the same shocks and self-reports should thus be relatively homogenous within small geographic aggregates.

In the context of attributing losses to disasters, this papers' primary interest is for farmers' shock identification in the face of a shock, i.e. the consistency with which self-reports identify a shock if a shock actually occurred. In the absence of a clear, objective measure of shock exposure, the analysis cannot compare self-reports to a "true value" but will rather study their internal consistency in "grey zone" cases where we cannot confidently rule out the possibility that a certain shock occurred. The analysis therefore excludes cases where a specific shock, for example a drought, was not identified on a single plot within the respective geographic unit of analysis (for example the village-level). This allows to focus on cases where it is conceivable that a shock occurred.

We study the consistency of farmers' shock reports at the farm- and enumeration area (EA-)level⁴⁹ based on several shock proxies. Specifically, four (three in Malawi) different questions in the LSMS-ISA data for self-reported shock identification, with some differences in shock definition and reporting periods, can be considered (Table 11).

Table 11: Shock identification questions in LSMS-ISA data.

Question	Level of ident.	Reporting period	Answer options
Q1: Why was the area harvested less than the area planted?	Plot-crop level	Current agr. season	Up to two reasons from a list including drought, irregular rains, insects, animals, floods (Malawi only)
Q2: What was the main cause of the damage on [CROP]?	Plot-crop level	Current agr. season	Main reason from a list including too little and too much rain
Q3: During the last 12 months, was your household affected by [SHOCK]?	Household level	Last 12 months	Yes/No response to a number of shocks including drought
Q4: Please describe important events that have taken place in this community since two years ago including any events that have occurred this year.	Community level (~ EA level)	This year and past two years	Up to four adverse events (if the same shock occurred multiple times, each time counts separately)

Q1 and Q2 are asked for each crop on each plot that recorded a planted area loss or plant productivity loss, respectively. They are thus the most suitable questions for this paper as they are

⁴⁹ For sampling purposes, 10-12 agricultural households (rural areas in Ethiopia) or 16 households (baseline sample in Malawi) form an enumeration area. In both countries, the typical EA is roughly equivalent to the village or community.

asked at the most disaggregated level and specifically in relation to crop losses in the agricultural season of interest. At the same time, they come with the caveat that they cannot discern idiosyncrasy that is due to differences in farmers' perception from idiosyncrasy in shock reports due to actual differences in highly localized loss occurrence. The latter matters if there are differences in shock exposure between plots, for example because of differences in soil quality (Iizumi and Wagai, 2019) or the adoption of improved varieties of the same crop (Lombardi *et al.*, 2019).

Furthermore, the LSMS-ISA data contains a question whether the household was negatively affected by a number of shocks, including drought, in the last 12 months (Q3). Lastly, the community module of the LSMS-ISA data features a question where a congregation of knowledgeable community members, roughly corresponding to the enumeration area, are asked about the four most significant adverse events in the last 3 years⁵⁰ (Q4).⁵¹

What do we see in the data?

The most common shocks in this study's region are rain-related (droughts, irregular/heavy rains, and floods) with some expected variation in incidence between the different proxies (Table A 8 and Table A 9). On an ordinal scale though, they broadly agree and capture similar variation between years and crops. Most notably, proxies consistently identify a large-scale drought for the 2015/16 harvesting season in Ethiopia as well as Malawi. Between the different proxies, the question on reasons for plant productivity reductions seems to pick up shocks most sensitively and leads to the highest incidences.⁵²

In the following, we study the consistency of shock identification at different levels of aggregation.⁵³ At the farm level, if drought damage (Q2) is recorded on any plot, this is usually also the case on the remaining plots with the same crop (Table A 10). Particularly in 2015 when drought was most severe in Ethiopia, about eight in ten households with more than one maize plot and nine in ten households with more than one sorghum plot record drought damage on all plots planted with the crop (conditional on at least one plot identifying some drought damage). In

⁵⁰ Ethiopia uses a different calendar that does not align with the Gregorian calendar. "Years" thus roughly correspond to a time period from September to September.

⁵¹ Household and community reports are what Sohnesen (2020) describes as "socio-economic" drought indicators, i.e. they may relate to a loss of consumption rather than necessarily an agricultural or meteorological drought. Compared to plot-level farmer reports, they have the drawback of not defining shocks in relation to a clear benchmark (e.g. a loss of production) which may cause farmers to define a drought differently and leaves room for endogeneity in responses (e.g. hardship is rationalized ex-post by reporting a drought) (Sohnesen, 2020).

⁵² In Ethiopia, it is also notable that loss incidences are higher for sorghum than maize. Using GIS rainfall data for the past 12 months and comparing this to the 14-year average, it seems like this pattern is explained by sorghum being more frequently planted in regions prone to adverse rainfalls.

⁵³ Since only the data in Ethiopia contains all four shock proxy questions (most importantly the damage question which seems most sensitive) it is the focus of our discussion. However, the main conclusions presented are robust across both countries and for the use of different shock proxies than crop damage where applicable. In particular, this should also alleviate concerns that the differences in farmers' shock identification we observe based on crop damage are because of a different prioritization of loss reasons across farms (farmers may only provide one reason for crop damage but two for area-based losses and the household-level question is asked for all households).

Malawi, where only information on planted area losses was elicited, drought identification is relatively homogenous within farms. Sample sizes are rather small though.⁵⁴

Table 12 confirms the importance of drought intensity for the uniformity of shock reports at the EA level through bivariate regressions. While there is a significant association between plot-level (Q1 or Q2) and community shock reports (Q4) across rounds, the explanatory power of this simple model is very low except, again, for ESS15 where drought was most intense.⁵⁵ Conversely, in Malawi, the association between community and plot-level reports is significant in all rounds but IHPS16 but the R^2 is low throughout.

Table 12: Regressions of community-reported drought exposure on plot-level drought reports.

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	ESS18	Ethiopia ESS15	ESS13	IHPS19	IHPS16	Malawi IHPS13	IHPS10
Drought damage	0.118** (0.0668)	0.572*** (0.0532)	0.237*** (0.0848)				
Drought area loss (maize)	0.209 (0.129)	0.0646 (0.0675)	0.108 (0.113)	0.381*** (0.144)	0.0364 (0.106)	0.185* (0.108)	0.405*** (0.0901)
Constant	0.0734*** (0.0251)	0.202*** (0.0353)	0.115*** (0.0316)	0.209*** (0.0605)	0.524*** (0.0872)	0.382*** (0.0728)	0.231*** (0.0559)
Observations	3,090	5,558	5,623	1,575	1,062	975	991
R-squared	0.044	0.349	0.069	0.039	0.001	0.009	0.160

Note: OLS regressions of community-reported drought exposure (last 2 calendar years) on plot-level crop loss reports due to drought. Ethiopia: Sorghum and maize plots. Malawi: Maize plots. Clustered standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A drought is a slow-onset disaster with no objective benchmark (or “start point”) for the community congregation. Furthermore, high drought salience (e.g. as observed on surrounding vegetation, based on conversations in the community, or the news) may prime farmers to attribute crop losses to drought. Similarly, the intensity of the shock may matter for the idiosyncrasy in loss occurrence across plots (e.g. soil with superior water-retention capacity may be able to buffer for little rain on some plots *to some extent*) and the clear attribution of crop losses to drought as opposed to other reasons. Therefore, idiosyncratic variation in self-reported drought exposure is

⁵⁴ As noted previously, shock identification of the plot-level proxies is conditional on there being some crop production loss. More precisely, we are thus comparing the consistency of loss occurrence due to drought between plots. This leaves the chance that disagreement between plots is because some plots may not record a loss in the first place since they are less exposed to drought (e.g. due to irrigation or superior soil quality). We assume that this distinction is less relevant after excluding irrigated plots from the analysis and particularly in years with severe drought but investigate this possibility in more detail in the following section.

⁵⁵ Importantly, community reports use yearly (in Ethiopia corresponding to a period between mid-September and mid-September) reporting which means time periods are imperfectly aligned with seasonal patterns in agriculture. We therefore combine community shock reports for the current and previous year. Our main conclusions are robust to the use of only current year or previous year shock reports and indeed correlations are even lower in these cases.

least in years with severe drought where marginal differences in perceptions, benchmarks, or plot characteristics matter least for shock identification.

An important caveat of the previous analysis is that it does not account for the fact that agreement may occur due to pure chance. In the LSMS-ISA data, farmers only face a binary “choice” between reporting or not reporting a shock (a loss due to a specific shock). This increases the chance for random agreement if farmers face some uncertainty when the choice is not obvious to them.⁵⁶ In the absence of a verifiable “true value”, there is thus a need for a *chance-corrected* agreement measure. We take this from the literature on “interrater reliability” (or “agreement coefficients”) which has one of its most common applications in clinical studies where the suitability of a rating scale to elicit consistent answers from different raters is explored. These agreement coefficients provide a measure of homogeneity (or consensus) in ratings of one or more “subjects” given by different “raters”.⁵⁷

This paper adapts this framework to the context at hand: Shock occurrence (the *subject*) is evaluated by farmers on each plot within an EA (the *raters*) and classified into shock or no shock (the *rating categories*). Various agreement coefficients exist in the literature. They differ in their definition of “chance” and thus the correction they apply to observed rates of agreement. This paper opts for Gwet’s AC₁ which features the highest flexibility regarding different numbers of raters, categories, subjects, and missing values while also being the most robust to a number of anomalies in other coefficients (Gwet, 2008, 2014). The standardized, chance-corrected agreement coefficients that are calculated for shock reports within an EA can lie in the range between -1 and 1. A value of zero means observing as much agreement in farmers’ shock reports within the same EA as would be expected by pure chance. A value of positive 1 (negative 1), conversely, would imply perfect agreement (disagreement). The interpretation of all other values in the interval between -1 and 1 is somewhat arbitrary, however, a widely cited reference for interpretation is given by the Landis and Koch (1977) benchmark scale. Following the recommendation in Gwet (2014), we do not assign the coefficients obtained deterministically to the intervals on the Landis and Koch (LK) scale but acknowledge the probabilistic nature of the point estimates. We therefore rather assign coefficients to the highest interval on the LK scale for which the cumulative probability that the value falls into the interval exceeds 95 percent. Taking this framework to the data, we analyze the homogeneity in farmers’ loss identification due to specific shocks within the same EA (Table 13, Table A 11).⁵⁸

We calculate agreement coefficients for three different cases of shocks.

1. Whether there were any maize productivity reductions on the plot
2. Whether there were any drought productivity reductions on the plot
3. Whether there were any plant productivity reductions due to too much rain on the plot

⁵⁶ In other terms, raw agreement does not necessarily reflect *intrinsic* agreement.

⁵⁷ One practical example may be studying the degree of agreement between several doctors in their diagnosis of a sample of patients.

⁵⁸ The analysis focuses on maize plots which constitute the largest share of the sample and shock reports based on damage on the crop which showed the highest agreement and greatest sensitivity to shocks thus far.

Furthermore, we also calculate a fourth agreement coefficient value that takes drought and rain damage as separate observations (separate “subjects”) on each plot and calculates a coefficient of chance-corrected agreement within an EA for this larger sample of observations. Throughout, we obtain a single, chance-corrected agreement coefficient based on comparing shock reports on plots within the same EA, i.e. we only compare those plots where it is reasonable to assume they face the same weather shocks. As before, we constrain our sample to those EAs where at least one plot recorded a certain shock.⁵⁹

Table 13: Ethiopia - Chance-adjusted agreement coefficients for self-reported shock exposure in EAs where at least one plot recorded the respective shock.

Round	Shock	Gwet's		N	N	Obs.	Exp.	LB	UB
		AC	SE	Raters	Subjects	Agreement	Agreement		
ESS18	Any	0.359	0.045	1326	131	0.657	0.465	0.2	0.4
	Drought	0.389	0.080	546	56	0.685	0.485	0.2	0.4
	Rains	0.527	0.117	327	24	0.725	0.417	0.2	0.4
	All shocks	0.396	0.072	797	80	0.698	0.500	0.2	0.4
ESS15	Any	0.623	0.030	2641	201	0.773	0.398	0.4	0.6
	Drought	0.655	0.040	1732	143	0.785	0.378	0.4	0.6
	Rains	0.688	0.067	492	30	0.781	0.299	0.4	0.6
	All shocks	0.604	0.041	2051	173	0.784	0.455	0.4	0.6
ESS13	Any	0.427	0.034	2415	196	0.697	0.472	0.2	0.4
	Drought	0.458	0.061	877	85	0.719	0.481	0.2	0.4
	Rains	0.512	0.093	547	35	0.698	0.381	0.2	0.4
	All shocks	0.425	0.060	1186	120	0.712	0.500	0.2	0.4

Note: Interrater reliability coefficients (Gwet's AC) for plots within the same enumeration area. Classification into benchmark interval is probabilistic, not deterministic; LB = lower boundary, UB = upper boundary; Landis and Koch (1977) benchmark scale for agreement coefficients: < 0 poor; 0-20 slight; 0.21 - 0.40 fair; 0.41 - 0.60 moderate; 0.61 - 0.80 substantial; 0.81- 1.0 almost perfect.

The coefficients obtained mostly fall in the range between 0.2 and 0.4 suggesting only limited agreement between farmers. Agreement tends to be lowest for whether there were any plant productivity reductions on the plot which is expected since idiosyncratic shocks may play a role here. On the other hand, rain shocks (either too much or too little rain) should be covariate within the EA which is reflected in slightly higher agreement. Coefficients are again highest in the round with the highest shock incidence, i.e. ESS15 lending further support to the hypothesis that the salience and intensity of the disaster plays a role for the consistency with which farmers record losses and attribute them to the same shock. Since we do not observe the “true value” of whether

⁵⁹ If we do not impose this constraint and compare shock reports for the full sample, we find high rates of agreement. These, however, are driven by low shock incidences in most rounds which, naturally, lead farmers to straightforwardly rule out the occurrence of a shock, e.g. a drought. From this, we can conclude that farmers relatively consistently rule out the occurrence of a shock, however, this does not imply how trustworthy farmers' reports are *if they claim to have incurred, for example, damage due to drought*. Arguably, the latter is more relevant for the issue of attributing losses to disasters and we therefore only compare a sample in which there seems to be at least some chance that a shock actually occurred (see also the previous discussion in this section).

a shock occurred, this variation in farmers' shock identification, particularly when drought was not severe, suggests that there is either non-negligible idiosyncrasy in shock identification or disaster crop losses itself are highly localized.⁶⁰

What do we conclude and still need to know?

There are several options to attribute losses to shocks based on self-reports in the LSMS-ISA data, the most disaggregated of which are given by farmers' reports of loss reasons at the plot-crop level. The most common shocks in the data are rainfall-related and should thus be covariate at small geographic aggregates. In general, the different shock proxies point in similar directions but vary in the shock incidence they suggest and are imperfectly correlated. The most sensitive measure of drought comes from a question on crop damage cause.

Agreement in self-reported shock exposure is limited between plots of the same EA. While in most cases farmers consistently rule out the occurrence of a shock, the primary interest of this section is in the consistency of farmer reports *in the face of a shock*. We therefore focus on cases where a shock, for example a drought, may have occurred based on farmers reports and find the consistency of shock identification between plots to be generally rather low. However, agreement depends on shock intensity which may alleviate idiosyncrasies due to different *perceptions* (a source of variation we would like to eliminate) and due to differences in actual loss *occurrence* (variation we would like to retain).

As this section cannot conclusively discern between the two sources of idiosyncratic variation, we conclude that it is difficult to confidently attribute losses to specific disasters based on self-reports alone. As a next step, it is thus necessary to further disentangle variation in self-reported shock identification due to differences in perception between farmers (a source of measurement error) from the variation that is due to highly localized shock impact at the sub-EA level.

Question 6: How localized are drought losses and how confidently can we attribute losses to adverse rainfall based on self-reports?

The previous analysis found notable idiosyncrasy in farmers' self-reported shock identification but was unable to discern what part of this variation comes from differences in perceptions between farmers and what part is due to actual differences in loss occurrence at the micro-level. This section compares self-reported shock exposure data at the plot level to geospatial rainfall data and a set of observable plot characteristics. We hypothesize that any differences in loss occurrence within the same EA do not stem from differences in the occurrence or intensity of rainfall shocks but differences in their actual impact on agricultural production on different plots. For example, a number of sustainable land management (SLM) practices and plot characteristics, on top of irrigation, can reduce the susceptibility of a plot to disaster crop losses (Nkonya *et al.*, 2018; Iizumi and Wagai, 2019; Lombardi *et al.*, 2019; McCarthy *et al.*, 2021).

⁶⁰ We reach a similar conclusion using shock reports based on area-losses or household-level reports and in Malawi.

By virtue of the rich LSMS-ISA data and in combination with geospatial rainfall data, we can control for many of these factors and identify the amount of variation in self-reported shock (loss) reports that is due differences in *observable* factors (such as rainfall intensity or SLM practices). Conversely, we assume that the remaining variation should then be attributable to *unobservable* factors such as idiosyncratic perceptions of farmers. This allows us to (i) identify the amount of variation in farmers' reports that should be retained as it reflects differences in shock exposure and contrast this with (ii) undesirable variation that should ideally be eliminated.

What do we see in the data?

To obtain an objective drought measure, we exploit geospatial, yearly total rainfall data for the year in which the agricultural season in question started provided by the National Oceanic and Atmospheric Administration Climate Prediction Center's (NOAA CPC) Africa Rainfall Climatology 2 (ARC2) dataset at 0.1° (~10km) resolution. As drought is a slow onset disaster and effects on crop losses should be greatest during the planting and growing season, this provides a basic, local measure of drought intensity.

Table 14: Drought incidence for self reported exposure and a GIS rainfall data-based drought indicator.

Percent of plots with drought exposure based on..				
Country	Crop	Round	Self-reports	GIS rainfall
Ethiopia	Maize	ESS13	12.5%	0%
		ESS15	39.3%	39.1%
		ESS18	15.4%	33.3%
	Sorghum	ESS13	18.7%	0%
		ESS15	64.1%	64.7%
		ESS18	29.0%	37.0%
Malawi	Maize	IHPS10	28.3%	53.8%
		IHPS13	5.2%	69.4%
		IHPS16	21.7%	19.8%
		IHPS19	3.7%	0.6%
	Any	IHPS10	33.8%	49.9%
		IHPS13	8.1%	68.9%
		IHPS16	32.3%	24.9%
		IHPS19	7.1%	0.5%

Note: Percent of plots with drought exposure based on farmers' self reports (area or productivity losses due to drought) and based on total annual rainfall below the long-term average (since 2001). ESS = Ethiopia Socioeconomic Survey, IHPS = Integrated Household Panel Survey

Comparing self-reported drought exposure, taken from reports of a planted area or plant productivity loss due to drought, to a drought indicator taken from 12-month total rainfall intensity being below the long-term average (since 2001), differences in incidence do not seem systematic (Table 14). Across years, farmers report either a notably higher or a notably lower incidence of drought exposure than the geospatial rainfall drought dummy. One notable exception is the 2015/16 drought in Ethiopia where incidences between both measures are strikingly aligned

lending further evidence in favor of the reliability of farmers' self-reports in the face of a severe drought.

Table 15: Ethiopia: Regression of self-reported drought exposure on measured rainfall intensity and controls.

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
	ESS13 Maize	ESS13 Sorghum	ESS15 Maize	ESS15 Sorghum	ESS18 Maize	ESS18 Sorghum
12-month total rainfall (mm)	-0.000247*** (6.22e-05)	-0.000608*** (0.000151)	-0.000672*** (6.78e-05)	-0.000744*** (8.03e-05)	-0.000103 (9.95e-05)	-0.000147 (0.000123)
Soil fertility index	-0.00704 (0.0136)	-0.00305 (0.0199)	-0.0350 (0.0267)	-0.0175 (0.0299)	-0.0131 (0.0291)	0.0508 (0.0550)
Plot Potential Wetness Index	-0.00352 (0.00328)	0.00108 (0.00616)	-0.00820 (0.00739)	-0.00776 (0.00754)	0.00138 (0.00432)	0.00158 (0.00581)
Plot Slope (percent)	0.000656 (0.00101)	8.42e-05 (0.000938)	-0.000754 (0.00211)	-0.000852 (0.00197)	-0.00567*** (0.00217)	-0.00652** (0.00259)
Intercropped	-0.0226 (0.0166)	-0.0780** (0.0381)	-0.0349 (0.0343)	0.0435 (0.0389)	0.0563 (0.0351)	0.00169 (0.0340)
Erosion control	-0.00801 (0.0194)	-0.0558 (0.0452)	-0.0118 (0.0419)	-0.00165 (0.0438)	0.0692* (0.0395)	-0.0303 (0.0439)
Any organic fertilizer	0.00545 (0.0175)	0.0600 (0.0514)	0.0395 (0.0272)	0.0374 (0.0382)	-0.0207 (0.0292)	0.0266 (0.0441)
Any inorganic fertilizer on plot	-0.0739** (0.0296)	0.0319 (0.0416)	-0.0716 (0.0504)	0.0977* (0.0580)	-0.00839 (0.0295)	0.0271 (0.0583)
Improved seed	0.0397 (0.0356)	0.179 (0.161)	0.0814 (0.0499)	0.0511 (0.0486)	0.00775 (0.0437)	-0.0621 (0.0773)
Constant	0.706*** (0.149)	1.245*** (0.276)	1.516*** (0.160)	1.501*** (0.166)	0.192 (0.137)	0.264 (0.183)
Observations	2,908	2,342	2,931	2,344	1,605	1,210
Adjusted R-squared	0.293	0.220	0.383	0.498	0.115	0.384
Region FE	YES	YES	YES	YES	YES	YES

Linear Probability Model regressing self-reported drought exposure on 12-month rainfall intensity and a number of controls that are potentially relevant for drought exposure. Clustered standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 15 and Table A 12 show regressions of self-reported crop damage due to drought on 12-month total rainfall intensity and a number of controls hypothesized to be relevant for plot-level drought exposure (Nkonya *et al.*, 2018; Iizumi and Wagai, 2019; Lombardi *et al.*, 2019; McCarthy *et al.*, 2021). The findings suggest that self-reports are significantly associated with low measured rainfall in the year in which the crop was planted and grown. Particularly in the 2015/16 Ethiopian Meher season, which featured widely reported drought (Philip *et al.*, 2018; Sohnesen, 2020), observed rainfall explains a considerable share of the variation in farmers' self-reports based on the adjusted R². However, there is no robust set of plot-level predictors of self-reported drought exposure. This, in turn, leaves a significant share of the variation unexplained by observable

factors.⁶¹ This suggests that a non-negligible share of the idiosyncrasy in farmers' self-reports are due to differences in perception rather than actual exposure to drought and casts doubt on the reliability of self-reported shock identification in order to attribute losses to shocks in years without a severe drought.

What do we conclude and still need to know?

Farmers' self-reported drought exposure features considerable idiosyncrasy that is unexplainable by observable factors such as rainfall intensity or SLM practices. Therefore, differences in perception between farmers, rather than differences in actual exposure to drought, seem to drive the considerable variation in drought reports within the same EA that we observed earlier. This is somewhat attenuated in years with severe drought where local rainfall accounts for a substantial share of the variation in self-reported drought exposure.

Therefore, for the attribution of losses to disasters and with rainfall data of sufficient resolution available, the "basis risk" of overlooking some truly idiosyncratic drought crop losses should be less severe than the measurement error incurred by relying on self-reports alone. One direction for future research could thus be the use of weather stations at a high density, e.g. the community level, or microsensors at even higher resolution to accurately contextualize self-reports and calibrate accurate measures of local shock exposure.

5. Discussion and Takeaways

Every year, disasters threaten millions of lives and livelihoods depending on agriculture in LMICs. What is more, this burden will likely increase in the coming decades due to the effects of climate change (Agnolucci *et al.*, 2020). In order to inform robust disaster risk reduction policies, accurate measurement of crop losses is of paramount importance and an important component of SDG 1.5.2 and Sendai indicator C-2.

In recent years, advances in the availability of and methodological innovations in farm surveys have made them rich and reliable resources of primary data in agriculture (Carletto and Gourlay, 2019). In the coming years, the *50x2030 initiative to close the agricultural data gap* will furthermore provide an unprecedented opportunity to measure crop losses at a highly disaggregated level, in large samples across 50 LMICs, and over time. Survey data thus have the potential to complement existing approaches to measure disaster crop losses that come with demanding data requirements in the context of LMICs by providing data that is

- (i) high in spatial resolution and density
- (ii) scalable to collection and trialed and tested for the specifications of smallholder agriculture

⁶¹ Our results are qualitatively robust to the use of a drought dummy (based on rainfall below the 14-year average) instead of absolute rainfall and for the use of a logit model instead of a linear probability model. Furthermore, they are robust to the use of a dummy whether there is any self-reported drought based on damage on the crop *or* area loss.

(iii) effortless to integrate with other survey information, e.g. on the behavioral components of crop management in the face of disasters and crop losses.

As such, survey data may facilitate the future development of well-calibrated, accurate, and scalable hybrid models using sensor, remote-sensed, and survey data to measure disaster crop losses in LMICs.

To tap into this potential, however, it is essential to understand the feasibility of measuring crop losses through survey data and develop a robust methodology. This paper pursues the question of operationalizing the recent methodology developed by Conforti, Markova and Tochkov (2020) for micro-level management through existing and future survey data. Specifically, it first analyzes the suitability of existing survey data to ascertain the magnitude of losses and secondly the ability to attribute the observed losses to specific disasters.

The papers' main findings are threefold. First, directly asking for different types of losses based on farmers' self-reports provides relatively consistent estimates on an ordinal scale. However, assessing losses is a cognitively complex task and invites satisficing in farmers' responses. As a result, forming accurate point-estimates of crop losses is impaired by concerns of numeracy and question format. Improvements in the phrasing and approach of asking for crop losses may allow for more precise answers.

Secondly, we find that post-planting harvest expectations often seem fair estimates of the attainable harvest and may thus allow to proxy for losses indirectly, perhaps most so at the extreme when losses are large.

Thirdly, farmers' self-reported shock identification exhibits considerable variation even within small geographic aggregates and for shocks that should be covariate in nature. To a large extent, the variation in self-reported drought exposure is not explainable by observable factors such as rainfall intensity or sustainable land management (SLM) practices and seems thus driven by differences in farmers' idiosyncratic perceptions. This is particularly conceivable for slow-onset disasters such as drought that lack a clear-cut benchmark or starting point. Importantly though, the intensity of the shock matters, with severe, salient shocks leading to more consistent shock identification between plots in the same area.

Moving forward, we propose the following.

For the quantification of losses

Changes in the wording of the direct proxy questions for crop losses and perhaps the use of additional aids may reduce the cognitive complexity of the task and elicit answers more accurately. Alternatively, the risk of misreporting may be reduced by combining existing proxies into a single, direct question benchmarked against farmers' self-reported attainable harvest. A third option could make use of farmers' post-planting expectation information as a measure of the attainable harvest. While expectations perhaps may not serve as the sole proxy for the attainable harvest, they could be an important feature for econometric models that try to predict (attainable) yields. This approach may be particularly suited to capture losses on fully destroyed plots or otherwise large losses. Furthermore, the endogeneity in expectations as a proxy of the attainable harvest may be reduced

by priming farmers to think about the attainable harvest in the absence of shock. Testing these different options should be the subject of a future survey experiment.

For the attribution of losses to disasters

In our study region, rainfall-related shocks present the most common reasons for crop losses. While these shocks are covariant, they lack a clear-cut benchmark that would make relying on farmers' self-reports alone a dependable strategy to attribute losses to disasters. While for large-scale drought, it may be feasible to rely on self-reports, more fine-grained measurement in years with less severe drought would benefit from a high density of weather stations or the distribution of micro-sensors to accurately capture differences in drought exposure.

Lastly, there is ample scope for future research. So far, we can mostly only study the internal consistency of different methods of measuring crop loss in survey data with no benchmark that would allow for an appraisal of their accuracy. An important element to ascertain the suitability of different options to estimate crop losses through survey data will thus be establishing a credible "gold standard" against which different methods can be compared. Furthermore, farmers' expectation-building still leaves a large (behavioral) black box. Better understanding the factors that underlie this mental model would provide crucial insights into the feasibility of using farmers' expectations to proxy for the attainable harvest. Similarly, future research may look into the possibility of drawing on a large set of possible predictors, including farmers' expectations, planted area and plant productivity loss proxies, inputs, and plot characteristics to predict yields. Finally, one pivotal question this paper leaves for future research is the representativeness of losses elicited from a sample. In order to build accurate national and sub-national aggregates of losses and deliver evidence for disaster risk reduction policies in general and SDG 1.5.2 and Sendai indicator C-2 in particular, it will be central to develop a sound methodology to representatively aggregate losses measured through surveys. An important validation exercise will then be the comparison between post-disaster needs assessments (PDNAs) and existing aggregate estimates of disaster crop losses and the estimates formed from micro survey data. In this realm, it will also be essential to keep investing in the collection of accurate ground data on disaster incidences and impacts at different scales, explore synergies between different data sources, and link this information to preventive disaster risk reduction interventions.

6. References

- 50x2030 Initiative (2020) *A Guide to the 50x2030 Data Collection Approach: Questionnaire Design*. Available at: <https://bit.ly/3tL6upb> (Accessed: 25 April 2021).
- Adamopoulos, T. and Restuccia, D. (2018) *Geography and Agricultural Productivity: Cross-Country Evidence from Micro Plot-Level Data*. w24532. Cambridge, MA: National Bureau of Economic Research, p. w24532. doi:10.3386/w24532.
- Agnolucci, P. *et al.* (2020) 'Impacts of rising temperatures and farm management practices on global yields of 18 crops', *Nature Food*, 1(9), pp. 562–571. doi:10.1038/s43016-020-00148-x.
- Amare, M. *et al.* (2018) 'Rainfall shocks and agricultural productivity: Implication for rural household consumption', *Agricultural Systems*, 166, pp. 79–89. doi:10.1016/j.agsy.2018.07.014.
- Arthi, V. *et al.* (2018) 'Not your average job: Measuring farm labor in Tanzania', *Journal of Development Economics*, 130, pp. 160–172. doi:10.1016/j.jdeveco.2017.10.005.
- Assefa, B.T. *et al.* (2020) 'Unravelling the variability and causes of smallholder maize yield gaps in Ethiopia', *Food Security*, 12(1), pp. 83–103. doi:10.1007/s12571-019-00981-4.
- Asseng, S. *et al.* (2013) 'Uncertainty in simulating wheat yields under climate change', *Nature Climate Change*, 3(9), pp. 827–832. doi:10.1038/nclimate1916.
- Barnett, B.J., Barrett, C.B. and Skees, J.R. (2008) 'Poverty Traps and Index-Based Risk Transfer Products', *Special Section (pp. 2045-2102)*. *The Volatility of Overseas Aid*, 36(10), pp. 1766–1785. doi:10.1016/j.worlddev.2007.10.016.
- Beza, E. *et al.* (2017) 'Review of yield gap explaining factors and opportunities for alternative data collection approaches', *European Journal of Agronomy*, 82, pp. 206–222. doi:10.1016/j.eja.2016.06.016.
- Brown, P. *et al.* (2018) 'Natural disasters, social protection, and risk perceptions', *World Development*, 104, pp. 310–325. doi:10.1016/j.worlddev.2017.12.002.
- Burke, M. and Lobell, D.B. (2017) 'Satellite-based assessment of yield variation and its determinants in smallholder African systems', *Proceedings of the National Academy of Sciences*, 114(9), pp. 2189–2194. doi:10.1073/pnas.1616919114.
- van Bussel, L.G.J. *et al.* (2015) 'From field to atlas: Upscaling of location-specific yield gap estimates', *Field Crops Research*, 177, pp. 98–108. doi:10.1016/j.fcr.2015.03.005.
- Carletto, C. and Gourlay, S. (2019) 'A thing of the past? Household surveys in a rapidly evolving (agricultural) data landscape: Insights from the LSMS-ISA', *Agricultural Economics*, 50(S1), pp. 51–62. doi:10.1111/agec.12532.
- Carletto, C., Gourlay, S. and Winters, P. (2015) 'From Guesstimates to GPStimates: Land Area Measurement and Implications for Agricultural Analysis', *Journal of African Economies*, 24(5), pp. 593–628. doi:10.1093/jae/ejv011.
- Carletto, C., Savastano, S. and Zezza, A. (2013) 'Fact or artifact: The impact of measurement errors on the farm size–productivity relationship', *Journal of Development Economics*, 103, pp. 254–261. doi:10.1016/j.jdeveco.2013.03.004.

- Cassar, A., Healy, A. and von Kessler, C. (2017) 'Trust, Risk, and Time Preferences After a Natural Disaster: Experimental Evidence from Thailand', *World Development*, 94, pp. 90–105. doi:10.1016/j.worlddev.2016.12.042.
- Challinor, A.J. *et al.* (2014) 'A meta-analysis of crop yield under climate change and adaptation', *Nature Climate Change*, 4(4), pp. 287–291. doi:10.1038/nclimate2153.
- Chaudhary, N. (2021) *Weather- and disease-related shocks in agriculture using data from the Rural Livelihoods Information System (RuLIS)*. RuLIS Brief. Rome: Food and Agriculture Organization. Available at: <https://www.fao.org/3/cb5593en/cb5593en.pdf> (Accessed: 17 December 2021).
- Cole, S. *et al.* (2013) 'Barriers to Household Risk Management: Evidence from India', *American Economic Journal: Applied Economics*, 5(1), pp. 104–35. doi:10.1257/app.5.1.104.
- Conforti, P., Markova, M. and Tochkov, D. (2020) *FAO's methodology for damage and loss assessment in agriculture*. FAO Statistics Working Paper 19–17. Rome: Food and Agriculture Organization. doi:10.4060/ca6990en.
- Delavande, A., Giné, X. and McKenzie, D. (2011) 'Eliciting probabilistic expectations with visual aids in developing countries: how sensitive are answers to variations in elicitation design?', *Journal of Applied Econometrics*, 26(3), pp. 479–497. doi:10.1002/jae.1233.
- Drechsler, M. and Soer, W. (2016) *Early Warning, Early Action: The Use of Predictive Tools in Drought Response through Ethiopia's Productive Safety Net Programme*. Policy Research Working Papers. Washington, DC: World Bank. doi:10.1596/1813-9450-7716.
- European Union, UN Development Group and World Bank (2013) *Post Disaster Needs Assessment: Agriculture, Livestock, Fisheries & Forestry*. PDNA Guidelines Volume B. Available at: <https://bit.ly/3tyFBEJ>.
- FAO (2018) *The impact of disasters and crises on agriculture and food security 2017*. Rome: Food and Agriculture Organization of the United Nations.
- FAO (2020) *Introduction to FAO's Damage and Loss Assessment Methodology, FAO elearning Academy*. Available at: <https://elearning.fao.org/course/view.php?id=608> (Accessed: 12 April 2021).
- FAO *et al.* (2020) *The State of Food Security and Nutrition in the World 2020*. Rome: FAO, IFAD, UNICEF, WFP and WHO. doi:10.4060/ca9692en.
- FAO (2021) *The impact of disasters and crises on agriculture and food security: 2021*. Rome: FAO. doi:10.4060/cb3673en.
- FAO (no date) *Drought Intensity: Uganda, Earth Observation*. Available at: <http://www.fao.org/giews/earthobservation/country/index.jsp?lang=en&type=11111&code=UGA>.
- Freudenreich, H. and Kebede, S. (2019) 'Experience of shocks and expectation formation – Evidence from smallholder farmers in Kenya', in. *168th Seminar, February 6-7, 2019, Uppsala, Sweden*, European Association of Agricultural Economists. doi:10.22004/ag.econ.289587.

- Giné, X., Townsend, R. and Vickery, J. (2007) 'Statistical Analysis of Rainfall Insurance Payouts in Southern India', *American Journal of Agricultural Economics*, 89(5), pp. 1248–1254. doi:10.1111/j.1467-8276.2007.01092.x.
- Giné, X., Townsend, R. and Vickery, J. (2008) 'Patterns of Rainfall Insurance Participation in Rural India', *The World Bank Economic Review*, 22(3), pp. 539–566. doi:10.1093/wber/lhn015.
- Giné, X., Townsend, R.M. and Vickery, J. (2009) *Forecasting when it matters: Evidence from Semi-Arid India*. Working Paper. Available at: <https://economics.yale.edu/sites/default/files/files/Workshops-Seminars/Development/gine-100405.pdf> (Accessed: 16 April 2021).
- Gloede, O., Menkhoff, L. and Waibel, H. (2015) 'Shocks, Individual Risk Attitude, and Vulnerability to Poverty among Rural Households in Thailand and Vietnam', *Vulnerability to Poverty in South-East Asia: Drivers, Measurement, Responses, and Policy Issues*, 71, pp. 54–78. doi:10.1016/j.worlddev.2013.11.005.
- Gollin, D., Lagakos, D. and Waugh, M.E. (2014) 'Agricultural Productivity Differences across Countries', *American Economic Review*, 104(5), pp. 165–170. doi:10.1257/aer.104.5.165.
- Gollin, D. and Udry, C. (2021) 'Heterogeneity, Measurement Error, and Misallocation: Evidence from African Agriculture', *Journal of Political Economy*, 129(1), pp. 1–80. doi:10.1086/711369.
- Gourlay, S., Kilic, T. and Lobell, D.B. (2019) 'A new spin on an old debate: Errors in farmer-reported production and their implications for inverse scale - Productivity relationship in Uganda', *Journal of Development Economics*, 141, p. 102376. doi:10.1016/j.jdeveco.2019.102376.
- Grassini, P. *et al.* (2015) 'How good is good enough? Data requirements for reliable crop yield simulations and yield-gap analysis', *Field Crops Research*, 177, pp. 49–63. doi:10.1016/j.fcr.2015.03.004.
- Gwet, K.L. (2008) 'Computing inter-rater reliability and its variance in the presence of high agreement', *British Journal of Mathematical and Statistical Psychology*, 61(1), pp. 29–48. doi:10.1348/000711006X126600.
- Gwet, K.L. (2014) *Handbook of inter-rater reliability: the definitive guide to measuring the extent of agreement among raters ; [a handbook for researchers, practitioners, teachers & students]*. 4. ed. Gaithersburg, MD: Advanced Analytics, LLC.
- Hallegatte, S. and Rozenberg, J. (2017) 'Climate change through a poverty lens', *Nature Climate Change*, 7(4), pp. 250–256. doi:10.1038/nclimate3253.
- Hanaoka, C., Shigeoka, H. and Watanabe, Y. (2018) 'Do Risk Preferences Change? Evidence from the Great East Japan Earthquake', *American Economic Journal: Applied Economics*, 10(2), pp. 298–330. doi:10.1257/app.20170048.
- Hazell, P.B.R. and Hess, U. (2010) 'Drought insurance for agricultural development and food security in dryland areas', *Food Security*, 2(4), pp. 395–405. doi:10.1007/s12571-010-0087-y.
- Hill, R.V. and Porter, C. (2017) 'Vulnerability to Drought and Food Price Shocks: Evidence from Ethiopia', *World Development*, 96, pp. 65–77. doi:10.1016/j.worlddev.2017.02.025.

- lizumi, T. and Wagai, R. (2019) 'Leveraging drought risk reduction for sustainable food, soil and climate via soil organic carbon sequestration', *Scientific Reports*, 9(1), p. 19744. doi:10.1038/s41598-019-55835-y.
- Imran, M., Zurita-Milla, R. and Stein, A. (2013) 'Modeling Crop Yield in West-African Rainfed Agriculture Using Global and Local Spatial Regression', *Agronomy Journal*, 105(4), pp. 1177–1188. doi:10.2134/agronj2012.0370.
- van Ittersum, M.K. *et al.* (2013) 'Yield gap analysis with local to global relevance—A review', *Field Crops Research*, 143, pp. 4–17. doi:10.1016/j.fcr.2012.09.009.
- Jain, M. *et al.* (2016) 'Mapping Smallholder Wheat Yields and Sowing Dates Using Micro-Satellite Data', *Remote Sensing*, 8(10), p. 860. doi:10.3390/rs8100860.
- Jensen, N.D., Barrett, C.B. and Mude, A.G. (2016) 'Index Insurance Quality and Basis Risk: Evidence from Northern Kenya', *American Journal of Agricultural Economics*, 98(5), pp. 1450–1469. doi:10.1093/ajae/aaw046.
- Jensen, N.D., Mude, A.G. and Barrett, C.B. (2018) 'How basis risk and spatiotemporal adverse selection influence demand for index insurance: Evidence from northern Kenya', *Food Policy*, 74, pp. 172–198. doi:10.1016/j.foodpol.2018.01.002.
- Karlan, D. *et al.* (2014) 'Agricultural Decisions after Relaxing Credit and Risk Constraints*', *The Quarterly Journal of Economics*, 129(2), pp. 597–652. doi:10.1093/qje/qju002.
- Krosnick, J.A., Narayan, S. and Smith, W.R. (1996) 'Satisficing in surveys: Initial evidence', *New Directions for Evaluation*, 1996(70), pp. 29–44. doi:10.1002/ev.1033.
- Krosnick, J.A. and Presser, S. (2010) 'Question and Questionnaire Design', in Marsden, P.V. and Wright, J.D. (eds) *Handbook of survey research*. Second edition. Bingley, UK: Emerald.
- Landis, J.R. and Koch, G.G. (1977) 'The Measurement of Observer Agreement for Categorical Data', *Biometrics*, 33(1), p. 159. doi:10.2307/2529310.
- Lichand, G. and Mani, A. (2020) 'Cognitive Droughts', *University of Zurich, Department of Economics, Working Paper No. 341* [Preprint]. doi:10.2139/ssrn.3540149.
- Liu, B. *et al.* (2016) 'Similar estimates of temperature impacts on global wheat yield by three independent methods', *Nature Climate Change*, 6(12), pp. 1130–1136. doi:10.1038/nclimate3115.
- Lobell, D.B., Cassman, K.G. and Field, C.B. (2009) 'Crop Yield Gaps: Their Importance, Magnitudes, and Causes', *Annual Review of Environment and Resources*, 34(1), pp. 179–204. doi:10.1146/annurev.enviro.041008.093740.
- Lobell, D.B., Schlenker, W. and Costa-Roberts, J. (2011) 'Climate Trends and Global Crop Production Since 1980', *Science*, 333(6042), pp. 616–620. doi:10.1126/science.1204531.
- Loïc, V. *et al.* (2018) 'Yield gap analysis extended to marketable grain reveals the profitability of organic lentil-spring wheat intercrops', *Agronomy for Sustainable Development*, 38(4), p. 39. doi:10.1007/s13593-018-0515-5.

- Lombardi, N. *et al.* (2019) *Disaster risk reduction at farm level: multiple benefits, no regrets : results from cost-benefit analyses conducted in a multi-country study, 2016-2018*. Available at: <http://www.fao.org/3/ca4429en/CA4429EN.pdf> (Accessed: 29 April 2021).
- van Loon, M.P. *et al.* (2019) 'Can yield variability be explained? Integrated assessment of maize yield gaps across smallholders in Ghana', *Field Crops Research*, 236, pp. 132–144. doi:10.1016/j.fcr.2019.03.022.
- Mani, A. *et al.* (2013) 'Poverty Impedes Cognitive Function', *Science*, 341(6149), pp. 976–980. doi:10.1126/science.1238041.
- Markhvida, M. *et al.* (2020) 'Quantification of disaster impacts through household well-being losses', *Nature Sustainability*, 3(7), pp. 538–547. doi:10.1038/s41893-020-0508-7.
- McCarthy, N. *et al.* (2021) 'Droughts and floods in Malawi: impacts on crop production and the performance of sustainable land management practices under weather extremes', *Environment and Development Economics*, pp. 1–18. doi:10.1017/S1355770X20000455.
- Meze-Hausken, E. (2004) 'Contrasting climate variability and meteorological drought with perceived drought and climate change in northern Ethiopia', *Climate Research*, 27(1), pp. 19–31.
- Miranda, M.J. and Farrin, K. (2012) 'Index Insurance for Developing Countries', *Applied Economic Perspectives and Policy*, 34(3), pp. 391–427.
- Mueller, N.D. *et al.* (2012) 'Closing yield gaps through nutrient and water management', *Nature*, 490(7419), pp. 254–257. doi:10.1038/nature11420.
- Nkonya, E. *et al.* (2018) 'Climate Risk Management through Sustainable Land and Water Management in Sub-Saharan Africa', in Lipper, L. *et al.* (eds) *Climate Smart Agriculture : Building Resilience to Climate Change*. Cham: Springer International Publishing, pp. 445–476. doi:10.1007/978-3-319-61194-5_19.
- Our World In Data (no date) *Depth of the Food Deficit (Kilocalories per person per day)*. Available at: <https://ourworldindata.org/grapher/depth-of-the-food-deficit?tab=chart> (Accessed: 12 April 2021).
- Philip, S. *et al.* (2018) 'Attribution Analysis of the Ethiopian Drought of 2015', *Journal of Climate*, 31(6), pp. 2465–2486. doi:10.1175/JCLI-D-17-0274.1.
- Ray, D.K. *et al.* (2015) 'Climate variation explains a third of global crop yield variability', *Nature Communications*, 6(1), p. 5989. doi:10.1038/ncomms6989.
- Restuccia, D. and Rogerson, R. (2013) 'Misallocation and productivity', *Review of Economic Dynamics*, 16(1), pp. 1–10. doi:10.1016/j.red.2012.11.003.
- Ronner, E. *et al.* (2018) 'Farmers' use and adaptation of improved climbing bean production practices in the highlands of Uganda', *Agriculture, Ecosystems & Environment*, 261, pp. 186–200. doi:10.1016/j.agee.2017.09.004.
- Rosenzweig, C. *et al.* (2014) 'Assessing agricultural risks of climate change in the 21st century in a global gridded crop model intercomparison', *Proceedings of the National Academy of Sciences*, 111(9), pp. 3268–3273. doi:10.1073/pnas.1222463110.

- Rosenzweig, M.R. and Udry, C.R. (2019) *Assessing the Benefits of Long-Run Weather Forecasting for the Rural Poor: Farmer Investments and Worker Migration in a Dynamic Equilibrium Model*. National Bureau of Economic Research.
- Said, F., Afzal, U. and Turner, G. (2015) 'Risk taking and risk learning after a rare event: Evidence from a field experiment in Pakistan', *Economic Experiments in Developing Countries*, 118, pp. 167–183. doi:10.1016/j.jebo.2015.03.001.
- Schilbach, F., Schofield, H. and Mullainathan, S. (2016) 'The Psychological Lives of the Poor', *American Economic Review*, 106(5), pp. 435–40. doi:10.1257/aer.p20161101.
- Shah, A.K., Mullainathan, S. and Shafir, E. (2012) 'Some Consequences of Having Too Little', *Science*, 338(6107), pp. 682–685. doi:10.1126/science.1222426.
- Silva, J.V. *et al.* (2018) 'Intensification of rice-based farming systems in Central Luzon, Philippines: Constraints at field, farm and regional levels', *Agricultural Systems*, 165, pp. 55–70. doi:10.1016/j.agsy.2018.05.008.
- Sohnesen, T.P. (2020) 'Two Sides to Same Drought: Measurement and Impact of Ethiopia's 2015 Historical Drought', *Economics of Disasters and Climate Change*, 4(1), pp. 83–101. doi:10.1007/s41885-019-00048-w.
- Tadesse, M.A., Shiferaw, B.A. and Erenstein, O. (2015) 'Weather index insurance for managing drought risk in smallholder agriculture: lessons and policy implications for sub-Saharan Africa', *Agricultural and Food Economics*, 3(1), p. 26. doi:10.1186/s40100-015-0044-3.
- Tversky, A. and Kahneman, D. (1973) 'Availability: A heuristic for judging frequency and probability', *Cognitive Psychology*, 5(2), pp. 207–232. doi:10.1016/0010-0285(73)90033-9.
- UN ECLAC (2003) *Handbook for Estimating the Socio-economic and Environmental Effects of Disasters*. Economic Commission for Latin America and the Caribbean.
- Waldman, K.B. *et al.* (2020) 'Agricultural decision making and climate uncertainty in developing countries', *Environmental Research Letters*, 15(11), p. 113004. doi:10.1088/1748-9326/abb909.
- van Wart, J. *et al.* (2013) 'Estimating crop yield potential at regional to national scales', *Field Crops Research*, 143, pp. 34–43. doi:10.1016/j.fcr.2012.11.018.
- Webber, H. *et al.* (2018) 'Diverging importance of drought stress for maize and winter wheat in Europe', *Nature Communications*, 9(1), p. 4249. doi:10.1038/s41467-018-06525-2.
- Wollburg, P., Tiberti, M. and Zezza, A. (2021) 'Recall length and measurement error in agricultural surveys', *Food Policy*, 100, p. 102003. doi:10.1016/j.foodpol.2020.102003.
- World Bank (no date) *World Development Indicators*. Available at: <https://data.worldbank.org/country/XM> (Accessed: 11 April 2021).
- Zhao, C. *et al.* (2017) 'Plausible rice yield losses under future climate warming', *Nature Plants*, 3(1), p. 16202. doi:10.1038/nplants.2016.202.

7. Appendix

Table A 1: Loss-to-harvest ratio using only planted area losses.

		Loss-to-harvest Ratio			
<i>Ethiopia</i>	Crop	Loss / Production	Loss p. ha / Yield	Loss / Production (any loss)	Loss p. ha / Yield (any loss)
ESS13	Maize	0.06	0.06	1.00	1.01
	Sorghum	0.21	0.14	1.69	0.99
ESS15	Maize	0.14	0.11	1.31	0.91
	Sorghum	0.33	0.36	1.58	1.80
ESS18	Maize	0.10	0.09	0.91	1.02
	Sorghum	0.22	0.16	1.53	0.96
Malawi					
IHPS19	Maize	0.79	0.74	0.71	0.70
	Harvest Value	1.57	1.43	1.36	1.26

Table A 2: Yield regressions with different loss proxies.

MALAWI: Yield regressions with planted area losses proxies -- excl. full losses

VARIABLES	(1) log Value/Ha	(2) log Maize Yield	(3) log Value/Ha	(4) log Maize Yield
Planted area harvested (%)	0.0243*** (0.00286)			
Planted maize area harvested (%)		0.00965*** (0.00139)		
log Area-based loss (MWK/ha)			-0.0903*** (0.0131)	
log Area-based loss (kg/ha)				-0.0342*** (0.0116)
Constant	9.898*** (0.559)	3.890*** (0.301)	12.14*** (0.665)	4.531*** (0.287)
Observations	2,087	1,507	2,087	1,507
Adjusted R-squared	0.263	0.450	0.232	0.407
District FE	YES	YES	YES	YES
Controls	YES	YES	YES	YES

Regression of log harvest value per hectare and log maize yield on loss proxies and a range of household, plot, and plot manager characteristics as well as inputs. Clustered standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table A 3: Correlates of discrepancy between farmers' self-assessed crop damage and enumerators' assessments.

Correlates of discrepancy		
VARIABLES	(1) ln(SR/CC)	(2) SR - CC
Farmer reports drought damage	0.515*** (0.186)	11.52*** (2.914)
Farmer reports termite damage	-0.467*** (0.178)	0.228 (2.779)
Farmer reports damage from crop disease	-0.115 (0.232)	-10.29*** (3.629)
GPS-Based Plot Area in Hectares	-0.338 (0.529)	3.532 (8.273)
HH head is female	-0.155 (0.245)	-3.623 (3.831)
Plot manager is respondent	-0.312 (0.277)	3.712 (4.336)
Plot manager is female	-0.0530 (0.205)	2.384 (3.213)
Age of plot manager (years)	0.00861 (0.00639)	-0.181* (0.0999)
Plot manager education (years)	0.00194 (0.0201)	-0.0282 (0.315)
Constant	0.862** (0.414)	22.64*** (6.472)
Observations	202	202
Adjusted R-squared	0.075	0.119

Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A 4: Ethiopia - Regression of realized harvest and expected harvest on inputs, household, and plot manager characteristics.

ETHIOPIA: Yield regressions: Realized Harvest vs. Expected Harvest

VARIABLES	ln(Realized Yield)				ln(Expected Yield)			
	ESS18		ESS15		ESS18		ESS15	
	Maize (1)	Sorghum (2)	Maize (3)	Sorghum (4)	Maize (5)	Sorghum (6)	Maize (7)	Sorghum (8)
log Plot area (ha)	0.179** (0.0754)	-0.0931 (0.0657)	-0.184*** (0.0487)	-0.249*** (0.0695)	-0.0954 (0.0623)	-0.0872 (0.0652)	-0.192*** (0.0384)	-0.197*** (0.0574)
log Tot. labor days/ha	0.114** (0.0459)	0.0990** (0.0426)	0.0212 (0.0345)	0.0821 (0.0507)	0.0967* (0.0499)	0.161*** (0.0492)	0.0536* (0.0289)	0.0815* (0.0432)
Any labor hired	-0.236 (0.186)	-0.105 (0.173)	0.0858 (0.142)	0.0395 (0.127)	-0.0242 (0.110)	-0.111 (0.102)	0.159 (0.102)	0.0204 (0.110)
log Seeds (kg/ha)	0.311*** (0.0766)	0.300*** (0.0739)	0.280*** (0.0556)	0.211*** (0.0740)	0.376*** (0.0498)	0.400*** (0.0726)	0.366*** (0.0489)	0.410*** (0.0655)
Improved seed	0.0631 (0.196)	-0.0992 (0.239)	0.180 (0.114)	-0.772 (0.482)	0.463*** (0.114)	0.328 (0.222)	0.213** (0.0946)	-0.173 (0.377)
log Inorganic fertilizer (kg/ha)	0.0498 (0.0330)	0.0284 (0.0299)	0.0772*** (0.0201)	0.108*** (0.0287)	0.0493** (0.0226)	-0.00618 (0.0289)	0.0489*** (0.0184)	0.0723*** (0.0242)
Any organic fertilizer	0.125 (0.130)	0.0818 (0.167)	0.0808 (0.0900)	-0.0515 (0.186)	0.0515 (0.110)	0.00349 (0.106)	0.115* (0.0646)	0.107 (0.117)
Dependency ratio	0.110** (0.0518)	-0.0839 (0.0598)	0.0701* (0.0366)	-0.117** (0.0563)	-0.0881* (0.0451)	-0.0390 (0.0360)	0.0201 (0.0303)	0.00877 (0.0491)
Any HH member has primary edu	0.0444 (0.124)	0.178 (0.115)	0.0625 (0.0961)	-0.169 (0.103)	-0.0438 (0.0934)	0.0955 (0.0901)	0.0840 (0.0724)	-0.0336 (0.0940)
Ag Asset Index	-0.151 (0.0944)	0.0819 (0.0942)	0.167** (0.0790)	0.188* (0.0962)	0.130** (0.0577)	-0.0107 (0.0731)	0.0247 (0.0484)	0.0788 (0.0690)
Female plot manager	0.319** (0.151)	0.0366 (0.143)	-0.0487 (0.111)	-0.132 (0.209)	0.0473 (0.0802)	-0.110 (0.110)	-0.0148 (0.0732)	-0.0523 (0.142)
Plot manager age	0.00153 (0.00371)	-0.0138*** (0.00411)	0.00470 (0.00314)	0.00699** (0.00321)	-0.000558 (0.00290)	-0.00433* (0.00257)	0.00102 (0.00235)	0.00586* (0.00328)

Plot manager has primary edu	0.0581 (0.114)	-0.363** (0.167)	-0.00541 (0.0873)	0.0442 (0.121)	0.0488 (0.102)	-0.0537 (0.110)	-0.0207 (0.0875)	0.0463 (0.115)
Intercropped	-0.652*** (0.155)	-0.288** (0.142)	-0.263*** (0.0940)	-0.146 (0.155)	-0.276*** (0.0777)	-0.365*** (0.122)	-0.250*** (0.0879)	-0.347*** (0.111)
Soil fertility index	0.311** (0.148)	-0.0517 (0.306)	0.182** (0.0730)	0.131 (0.115)	-0.0515 (0.122)	-0.245 (0.175)	0.0858 (0.0615)	0.00963 (0.0778)
Constant	5.088*** (0.458)	5.963*** (0.431)	4.269*** (0.326)	4.919*** (0.360)	4.598*** (0.402)	5.359*** (0.351)	4.441*** (0.243)	4.408*** (0.298)
Observations	1,678	1,281	2,971	2,444	1,630	1,240	2,845	2,345
Adjusted R-squared	0.287	0.259	0.309	0.268	0.304	0.349	0.234	0.183
District FE	YES	YES	YES	YES	YES	YES	YES	YES

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A 5: Malawi - Regression of realized harvest and expected harvest on inputs, household, and plot manager characteristics.

MALAWI - Yield regressions: Realized Maize Harvest vs. Expected Maize Harvest

	ln(Maize Yield)				ln(Expected Maize Yield)			
	IHPS19 (1)	IHPS15 (2)	IHPS13 (3)	IHPS10 (4)	IHPS19 (5)	IHPS15 (6)	IHPS13 (7)	IHPS10 (8)
log Plot area (ha)	-0.707* (0.386)	-0.828 (0.554)	-0.967** (0.385)	-0.998*** (0.277)	-0.0961 (0.413)	-1.007** (0.467)	-0.715*** (0.257)	-2.053*** (0.432)
log Tot. labor days/ha	0.164*** (0.0509)	0.379 (0.264)	0.302** (0.123)	0.210*** (0.0510)	0.123 (0.0769)	0.124 (0.115)	0.334*** (0.0909)	0.00704 (0.0765)
Any labor hired	0.689*** (0.171)	0.433 (0.312)	0.317** (0.125)	0.191* (0.0970)	0.285** (0.140)	0.434*** (0.126)	0.437*** (0.139)	-0.137 (0.123)
log Maize seeds (kg/ha)	0.369*** (0.0564)	-0.154 (0.237)	0.142** (0.0593)	-0.0368 (0.0465)	0.567*** (0.0910)	0.372* (0.202)	0.297*** (0.0443)	0.495*** (0.0484)
Plot has improved maize variety		0.113	0.0669		0.226**	0.110		

	(0.0999)	(0.221)			(0.102)	(0.157)		
log Inorganic fertilizer (kg/ha)	0.104***	0.108*	0.106***	0.0786***	0.162***	0.0410	0.0775***	0.0347
	(0.0317)	(0.0585)	(0.0257)	(0.0210)	(0.0273)	(0.0377)	(0.0209)	(0.0273)
Any organic fertilizer	0.227**	-0.171	0.0803	0.126	0.0619	-0.0537	0.0482	-0.170
	(0.105)	(0.324)	(0.179)	(0.0977)	(0.115)	(0.178)	(0.198)	(0.167)
Dependency Ratio	0.0296	0.0360	0.0805	-0.000781	0.0234	0.0542	0.0750*	0.0228
	(0.0871)	(0.128)	(0.0517)	(0.0483)	(0.108)	(0.104)	(0.0403)	(0.0786)
HH head has wage employment	0.127	0.199	-0.160	-0.0425	0.0104	0.0820	0.0937	-0.176
	(0.210)	(0.390)	(0.200)	(0.0963)	(0.189)	(0.273)	(0.156)	(0.152)
Any HH member owns non-ag enterprise	-0.0686	-0.0639	0.119	0.120	0.111	0.0902	-0.0371	-0.0972
	(0.126)	(0.266)	(0.154)	(0.0744)	(0.133)	(0.267)	(0.128)	(0.168)
Ag Asset Index	0.374	2.084***	0.523**	0.816***	0.114	-0.0790	0.227	-0.101
	(0.384)	(0.543)	(0.210)	(0.235)	(0.291)	(0.330)	(0.192)	(0.249)
Female plot manager	0.0227	-0.280	-0.142	-0.0638	0.126	-0.505*	-0.163	-0.290***
	(0.117)	(0.199)	(0.121)	(0.0820)	(0.124)	(0.268)	(0.120)	(0.101)
Plot manager age	-3.23e-06	-0.000150	-0.00217	-0.00455	-0.000845	0.00658	-0.00182	-0.00284
	(0.00362)	(0.00800)	(0.00333)	(0.00302)	(0.00520)	(0.00614)	(0.00381)	(0.00293)
Edu years of plot manager	0.0309*	-0.0407	0.0310**	0.0225*	0.0250	0.0140	0.0210	0.000883
	(0.0182)	(0.0403)	(0.0128)	(0.0113)	(0.0184)	(0.0191)	(0.0128)	(0.0180)
Intercropped	0.0815	0.168	-0.199	-0.275**	0.0226	0.0670	-0.0886	0.177
	(0.116)	(0.290)	(0.159)	(0.127)	(0.149)	(0.152)	(0.165)	(0.178)
Soil fertility index	0.0146	-0.0879	0.0251	0.0944*	-0.0859	-0.263	-0.0171	0.0330
	(0.0875)	(0.306)	(0.164)	(0.0567)	(0.0753)	(0.177)	(0.0889)	(0.0913)
Constant	4.190***	4.444**	4.868***	6.035***	3.440***	4.373***	4.148***	5.789***
	(0.339)	(1.676)	(0.738)	(0.439)	(0.743)	(0.694)	(0.570)	(0.598)
Observations	547	490	441	839	547	490	441	839
Adjusted R-squared	0.354	0.246	0.340	0.471	0.421	0.383	0.421	0.391

Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A 6: Malawi - Regressions of self-reported harvest realizations on expected harvest and a range of controls.

Malawi

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Harvest value (MWK)	IHPS19		IHPS16		IHPS13		IHPS10	
Expected harvest value (MWK)	0.514*** (0.129)	0.402*** (0.0927)	0.523*** (0.102)	0.336*** (0.109)	0.725*** (0.102)	0.437*** (0.110)	0.354 (0.225)	0.257 (0.193)
Planted area harvested (%)		181.2*** (38.06)						
Constant	13,434*** (4,447)	4,466 (24,474)	14,445*** (4,991)	-16,564 (28,693)	7,743** (3,442)	66,365 (51,522)	7,808** (2,970)	16,974 (16,379)
Observations	2,114	2,112	1,379	1,379	1,247	1,247	1,258	1,258
Adjusted R-squared	0.541	0.646	0.368	0.469	0.510	0.593	0.144	0.327
Controls	NO	YES	NO	YES	NO	YES	NO	YES
Panel B: Harvest value (MWK/ha)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Expected harvest value (MWK/ha)	0.551*** (0.0343)	0.425*** (0.0345)	0.530*** (0.0513)	0.409*** (0.0501)	0.678*** (0.111)	0.379*** (0.0820)	0.0936*** (0.0271)	0.0486** (0.0210)
Planted area harvested (%)		877.4*** (139.8)						
Constant	52,167*** (6,216)	147,796 (94,256)	59,532*** (9,307)	-20,166 (101,789)	39,442*** (11,478)	378,749 (241,858)	38,228*** (2,751)	80,479** (34,676)
Observations	2,114	2,112	1,379	1,379	1,247	1,247	1,258	1,258
Adjusted R-squared	0.392	0.445	0.294	0.350	0.364	0.489	0.085	0.298
Controls	NO	YES	NO	YES	NO	YES	NO	YES
Panel C: Maize production (kg)	(17)	(18)	(19)	(20)	(21)	(22)	(23)	(24)
Expected maize harvest (kg)	0.783*** (0.0488)	0.545*** (0.0730)	0.654*** (0.0528)	0.440*** (0.0765)	0.580*** (0.105)	0.112 (0.114)	0.129*** (0.0282)	0.0276 (0.0209)
Planted maize area harvested (%)		1.828*** (0.510)						
Constant	47.25** (19.89)	-278.1 (525.7)	71.41*** (24.72)	-482.3 (583.6)	312.0*** (81.05)	2,710* (1,418)	428.3*** (35.87)	1,023*** (356.9)
Observations	568	544	493	490	442	442	842	842
Adjusted R-squared	0.666	0.718	0.460	0.557	0.249	0.386	0.026	0.398
Controls	NO	YES	NO	YES	NO	YES	NO	YES

Note: Regression of self-reported realized harvest value (per hectare) (Panels A + B) or production (Panel C) on expected harvest value (per hectare) (Panels A + B) or production (Panel C), planted area loss (only IHPS19), and a number of plot, manager, and household controls as well as inputs and district fixed effects. Controls: Inputs - Plot area, total labor (p. ha), any hired labor dummy, seed quantity (p. ha), improved variety dummy (only IHPS19 and IHPS15), inorganic fertilizer quantity (p. ha), any organic fertilizer dummy; Household - dependency ratio, head employed for wage dummy, any HH member owns non-ag enterprise dummy, any HH member has primary education dummy, agricultural asset index; Manager - female dummy, age, primary education dummy; Plot - intercropped dummy, precipitation of wettest month, soil fertility index, cash crop dummy; Zone fixed effects. Clustered standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A 7: Regressions of self-reported harvest realizations on expected harvest and a range of controls.

Ethiopia - Production

VARIABLES	(1)	(2)	ESS15		(5)	(6)	ESS18	
	SR maize Production (kg)	SR maize Production (kg)	SR sorghum Production (kg)	SR sorghum Production (kg)	SR maize Production (kg)	SR maize Production (kg)	SR sorghum Production (kg)	SR sorghum Production (kg)
Expected maize production (kg)	0.285** (0.115)	0.176** (0.0817)			0.395*** (0.0800)	0.179** (0.0755)		
Expected sorghum production (kg)			0.253*** (0.0924)	0.113* (0.0602)			0.478*** (0.137)	0.279*** (0.104)
Planted area harvested (%)		0.618* (0.361)		0.328 (0.225)		0.882*** (0.337)		0.508 (0.353)
Damage on crop (%)		-1.072** (0.478)		-1.914*** (0.332)		-0.581* (0.339)		-0.401 (0.423)
Constant	144.1*** (28.29)	177.3 (119.2)	126.0*** (19.27)	196.3 (124.0)	81.48*** (18.71)	-183.6** (86.99)	129.8*** (48.08)	374.6* (208.9)
Observations	2,928	2,841	2,397	2,336	1,685	1,629	1,302	1,240
Adjusted R-squared	0.181	0.359	0.109	0.379	0.303	0.544	0.053	0.528
Controls	NO	YES	NO	YES	NO	YES	NO	YES

Note: Regression of realized production on expected production, planted area and plant productivity loss, and a number of plot-, manager-, and household controls as well as inputs and zone fixed effects. Controls: Inputs - Plot area, total labor, any hired labor dummy, seed quantity, improved variety dummy, inorganic fertilizer quantity, any organic fertilizer dummy; Household - dependency ratio, any HH member has primary education dummy, agricultural asset index; Manager - female dummy, age, primary education dummy; Plot - intercropped dummy, precipitation of wettest month, soil fertility index; Zone fixed effects. Clustered standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table A 8: Ethiopia - Shock prevalence across different proxies.

		Planted area loss					
	Round	Any	Drought	Irregular rains	Insects	Animals	Other
Maize	ESS13	0.13	0.04	0.04	0.01	0.03	0.03
	ESS15	0.19	0.11	0.11	0.01	0.03	0.02
	ESS18	0.17	0.02	0.05	0.03	0.03	0.04
Sorghum	ESS13	0.21	0.03	0.13	0.08	0.04	0.02
	ESS15	0.34	0.21	0.20	0.01	0.03	0.01
	ESS18	0.25	0.02	0.16	0.01	0.04	0.02
		Plant productivity loss					
	Round	Any	Drought	Too much rain	Insects		Other
Maize	ESS13	0.41	0.13	0.03	0.05		0.19

	ESS15	0.56	0.39	0.02	0.02		0.12
	ESS18	0.43	0.16	0.05	0.08		0.15
Sorghum	ESS13	0.46	0.18	0.01	0.12		0.14
	ESS15	0.74	0.65	0.02	0.01		0.07
	ESS18	0.53	0.29	0.06	0.05		0.13
HH reported							
	Round	Any	Drought	Heavy rains	Flood	Landslides	Other crop damage
HH lvl.	ESS13	0.14	0.08	0.01	0.02	0.00	0.05
	ESS15	0.37	0.30	0.03	0.01	0.01	0.10
	ESS18	0.15	0.09	0.02	0.03	0.00	0.04
Community reported							
	Round	Any (last 2y)	Drought (this year)	Drought (last year)	Flood (this year)	Flood (last year)	
EA lvl.	ESS13	0.27	0.02	0.13	0.03	0.14	
	ESS15	0.50	0.23	0.23	0.01	0.06	
	ESS18	0.17	0.03	0.20	0.04	0.10	

Table A 9: Malawi - Shock prevalence across different proxies.

Planted area loss						
	Round	Any	Drought	Irregular rains	Flood	Other
Maize	IHPS10	0.57	0.36	0.13		0.03
	IHPS13	0.55	0.07	0.14	0.15	0.05
	IHPS16	0.65	0.28	0.34	0.01	0.02
	IHPS19	0.61	0.05	0.15	0.25	0.04
HH reported						
	Round	Drought	Flood			
HH lvl.	IHPS10	0.48	0.05			
	IHPS13	0.63	0.19			
	IHPS16	0.89	0.16			
	IHPS19	0.68	0.38			
Community reported						
	Round	Any (last 2y)	Drought (this year)	Drought (last year)	Flood (this year)	Flood (last year)
EA lvl.	IHPS10	0.33	0.01	0.26	0.01	0.05
	IHPS13	0.41	0.12	0.21	0.09	0.04
	IHPS16	0.63	0.18	0.42	0.04	0.19
	IHPS19	0.49	0.06	0.13	0.20	0.26

Table A 10: Intra-farm consistency in drought exposure reports.

Panel A: Ethiopia				Panel B: Malawi			
Round	Crop	N	All Drought Damage	Round	Crop	N	All Drought Area loss
ESS13	Maize	92	64%	IHPS10	Maize	91	90%
	Sorghum	150	68%	IHPS13	Maize	22	82%
ESS15	Maize	333	77%	IHPS16	Maize	73	67%
	Sorghum	405	89%	IHPS19	Maize	30	30%
ESS18	Maize	103	50%	<i>Note: Share of farms with with more than one maize plot and at least one plot with planted area loss due to drought that have drought planted area loss on every maize plot.</i>			
	Sorghum	124	70%				

Note: Share of farms with with more than one maize/sorghum plot and at least one plot with drought damage that have damage on every maize/sorghum plot.

Table A 11: Malawi - Chance-adjusted agreement coefficients for self-reported shock exposure in EAs where at least one plot recorded the respective shock.

Round	Shock	Gwet's AC	SE	N Raters	N Subjects	Obs. Agreement	Exp. Agreement	LB	UB
IHPS19	Drought	0.579	0.126	643	26	0.696	0.277	0.2	0.4
	Flood	0.206	0.085	1057	51	0.579	0.471	0	0.2
	Irregular Rain	0.412	0.083	1340	58	0.641	0.389	0.2	0.4
	All shocks	0.370	0.059	1504	135	0.629	0.411	0.2	0.4
IHPS16	Drought	0.192	0.114	804	56	0.595	0.499	0	0.2
	Flood	0.809	0.108	90	5	0.881	0.374	0.4	0.6
	Irregular Rain	0.237	0.090	1000	66	0.618	0.500	0	0.2
	All shocks	0.234	0.069	1039	127	0.617	0.500	0	0.2
IHPS13	Drought	0.501	0.097	401	25	0.662	0.324	0.2	0.4
	Flood	0.326	0.090	501	31	0.633	0.457	0	0.2
	Irregular Rain	0.432	0.080	727	47	0.648	0.380	0.2	0.4
	All shocks	0.416	0.048	881	103	0.647	0.396	0.2	0.4
IHPS10	Drought	0.204	0.080	643	48	0.600	0.497	0	0.2
	Irregular Rain	0.326	0.112	529	39	0.613	0.425	0	0.2
	All shocks	0.224	0.076	803	87	0.605	0.491	0	0.2

Note: Interrater reliability coefficients (Gwet's AC) for plots within the same enumeration area. Classification into benchmark interval is probabilistic, not deterministic; LB = lower boundary, UB = upper boundary; Landis and Koch (1977) benchmark scale for agreement coefficients: < 0 poor; 0-20 slight; 0.21 - 0.40 fair; 0.41 - 0.60 moderate; 0.61 - 0.80 substantial; 0.81- 1.0 almost perfect.

Table A 12: Malawi: Regression of self-reported drought exposure on measured rainfall intensity and controls.

VARIABLES	(1) IHPS13	(2) IHPS16	(3) IHPS19
12-month total rainfall (mm)	-0.000582*** (0.000217)	-0.000305** (0.000128)	-3.88e-05 (0.000241)
Soil fertility index	-0.0269 (0.0239)	0.0110 (0.0304)	0.0142* (0.00834)
Plot Potential Wetness Index	0.00874* (0.00512)	0.0123* (0.00631)	0.00105 (0.00201)
Plot Slope (percent)	0.000432 (0.00243)	-0.00328 (0.00382)	-0.000594 (0.000946)
Intercropped	0.0523 (0.0390)	0.114*** (0.0411)	0.0321*** (0.0122)
Plot has any erosion control structure	-0.0226 (0.0240)	0.0461 (0.0335)	-0.0115 (0.0189)
Any organic fertilizer	0.0147 (0.0256)	0.0340 (0.0239)	0.0346* (0.0187)
Inorganic fertilizer was applied on plot	-0.0613* (0.0350)	0.0105 (0.0322)	0.0117 (0.0162)
Plot has improved maize variety		0.0130 (0.0221)	-0.000186 (0.0119)
Constant	0.639*** (0.225)	0.158 (0.172)	0.0379 (0.301)
Observations	925	1,004	1,490
Adjusted R-squared	0.052	0.370	0.111
District FE	YES	YES	YES

Linear Probability Model regressing self-reported drought exposure on 12-month rainfall intensity and a number of controls that are potentially relevant for drought exposure. Clustered standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1