Potential for Application of a Probabilistic Catastrophe Risk Modelling Framework to Poverty Outcomes

General form Vulnerability Functions Relating
Household Poverty Outcomes to Hazard Intensity
in Ethiopia

Catherine Porter

Emily White

WORLD BANK GROUP

Finance and Markets Global Practice Group

June 2016

Abstract

This paper analyzes the potential to combine catastrophe risk modelling (CAT risk modeling) with economic analysis of vulnerability to poverty using the example of drought hazard impacts on the welfare of rural households in Ethiopia. The aim is to determine the potential for applying a derived set of damage (vulnerability) functions based on realized shocks and household expenditure/consumption outcomes, onto a forward-looking view of drought risk. The paper outlines the CAT risk modeling framework and the role of the vulnerability module, which describes the response of an affected exposure to a given hazard intensity. The need to explicitly account for different household characteristics that determine vulnerability within our model is considered, analogous to how a CAT risk model would differentiate damage functions for buildings by different

classes of construction. Results for a regression model are presented, estimating ex-post drought impacts on consumption for heterogeneous household types (e.g. with cattle, safety-net access, illness). Next, the validity/generalizability of the derived functions are assessed, to infer applicability of the derived relationships within a CAT risk modelling framework. In particular, the analysis focuses on external validity: whether the relationships established in the dataset can be used for forecasting outside of the sample used for analysis. The model is stress-tested using statistical methods of resampling. This involves randomly splitting the data into “training” and “testing” datasets. The tests show consistency of results across the datasets. Finally, future plans are outlined with regard to developing a fuller catastrophe risk model to combine with the consumption results.

This paper is a product of the Disaster Risk Financing and Insurance Program (DRFIP), a partnership of the World Bank's Finance and Markets Global Practice Group and the Global Facility for Disaster Reduction and Recovery, with funding from the UK Department for International Development. 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://econ.worldbank.org>. The authors may be contacted at Catherine.porter@hw.ac.uk.

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.

Potential for application of a probabilistic catastrophe risk modelling framework to poverty outcomes:

General form vulnerability functions relating household poverty outcomes to hazard intensity in Ethiopia

Catherine Porter¹ and Emily White²

JEL Classifications: *C18; C21; I31; G2.*

Keywords: disaster risk finance; catastrophe risk modelling; poverty; vulnerability, safety net.

This paper is a product of the Disaster Risk Financing and Insurance Program (DRFIP), a partnership of the World Bank's Finance and Markets Global Practice Group and the Global Facility for Disaster Reduction and Recovery, with funding from the UK Department for International Development. 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://econ.worldbank.org>. The author may be contacted at Catherine.porter@hw.ac.uk.

This paper was prepared with statistical analysis in R by Qhelile Nyathi and data preparation by Douglas Scott. We thank Emmanuel Skoufias, Ruth Hill, and Robert Muir-Wood for helpful comments.

¹ Heriot-Watt University, Edinburgh, Scotland.

² Consultant, Finance and Markets Global Practice, The World Bank Group, Washington DC, USA.

Extended Abstract

This paper brings together two strands of research that have thus far been developed independently: catastrophe risk modelling, and economic analysis of vulnerability to poverty. The analysis focuses on a specific example to fix ideas: the impact of drought hazard on the welfare of rural households in Ethiopia. The aim is to determine the potential for applying a derived set of damage (vulnerability) functions based on realized shocks and household expenditure/consumption outcomes, onto a forward-looking view of drought risk. The paper outlines the contribution that combining the two analyses can bring, shows preliminary results for the validation of a regression model estimating ex-post drought impacts, and outlines future plans with regard to catastrophe risk modeling.

The catastrophe risk modeling framework (CAT risk modeling) and the role of the vulnerability module, which describes the response of an affected exposure (household, building, crop) to a given hazard intensity, is first outlined. CAT risk models have been mainly used for modelling financial losses, and the paper describes some of the challenges in applying the CAT risk modelling framework, and specifically the vulnerability module, to create a view of the impact of hazards on household poverty. The need to explicitly account for different household characteristics that determine vulnerability within our model, is considered, analogous to how a CAT risk model would differentiate damage functions by different classes of construction. The importance of this differentiation in external validity of any derived relationships is acknowledged in the development of the analysis methodology. In the applied micro-economics literature, considerable progress has been made in estimating the ex-post impact of realized shocks on welfare outcomes, notably household consumption and associated poverty. Hill and Porter (2014) extended such a backward looking model with simulations from the historical distribution of drought, to quantify vulnerability, defined as high probability of falling into poverty in future periods. The analysis herein then attempts to develop the model of drought and consumption further, in particular to extend the heterogeneous impacts on different household types. The validity/generalizability of the derived functions are then assessed, to infer applicability of the derived relationships within a CAT risk modelling framework.

For the applied regression analysis representative cross-section data for Rural Ethiopia in 2005 and 2011 is used, as also used by Hill and Porter (2014) for vulnerability analysis (Household Income Consumption and Expenditure Survey, HICES, and Welfare Monitoring Survey, WMS). The data comprise household characteristics, consumption outcomes (our variable of interest for poverty analysis) and measures of realized shocks such as household member illness, food prices, crop damage, unemployment. The analysis also incorporates the drought data as compiled by the world food program (WFP, LEAP data). This drought measure shows the expected crop losses at community level based on water adequacy specific to the crops grown in that community. A regression model is fitted to the pooled data, and several different specifications are compared. The main goal is to assess the stability of the drought impact across years, and also across “types” of household (analogous to catastrophe risk modelling e.g. of hurricane impacts on different building structures). The results show significantly different impact of the drought for households with assets, access to the PSNP, and financial access: these characteristics are seen to be mitigating the impact of the drought. However, households that suffered other crop damage experienced a heavier impact of the drought.

An attempt is then undertaken to “stress-test” the model by using statistical methods of resampling. This involves randomly splitting the data into “training” and “testing” datasets. The training data is used to fit the model, and then the goodness of fit is tested on the holdout sample or testing

dataset. The methodology attempts several cuts of the data, randomly selected, separately between 2005 and 2011, and also non-random holdout of each of four agro-climatic regions of Ethiopia. Overall the model performs consistently across different 'cuts' of the data when the sample is randomly divided into training and testing data. The outcome variable (root-mean-squared error of the model) is not very different across either the models or the choice of training and testing dataset. The analysis also tests the support of non-linearity in the drought impact, and we find that a quadratic model (a squared term in the regression) best fits the data. However, there is data paucity at higher levels of the shock (low probability but high impact events), so we cannot be certain that the model would accurately predict drought impact for e.g. 70% crop failure and above.

Based on these results, the paper concludes that the analysis demonstrates some level of external and internal validity of the derived relationships. However, several caveats remain. An attempt was also made to test the model on 2012 data from a different source (Ethiopian Rural Socioeconomic Survey), which has fewer observations, and a slightly different questionnaire, leading to difficulties in comparability of the main welfare measure. Further, the econometric question of identification is never completely solved when using cross-section data, so a follow-up robustness check with panel data would be extremely useful.

With these caveats in mind, the paper concludes that the derived relationships could form the basis of a 'vulnerability module' within a CAT risk model, and outlines how the model developed could be then applied to a CAT risk modeling framework. The impact of drought at various levels of severity on poverty is shown, and a suggestion is made that after some further robustness checks a full CAT risk model could be built. This would entail development and application of a stochastic model of rainfall that could be applied into the drought (LEAP) framework to produce values for actual evapotranspiration in the index calculation. The model would need to be compatible with the time and geographic resolution of the derivation of the vulnerability/damage functions in order for these to be applied. The final outcome would then link the damage functions with poverty outcomes in financial terms, for example the total poverty gap and associated fiscal burden.

1. Introduction

Probabilistic catastrophe risk models, used extensively in the international insurance and reinsurance markets, develop a view of risk beyond the historical occurrence of catastrophes. This is done through the generation of thousands of synthetic stochastic events whose characteristics, evolution and pathways are calibrated based on historical event occurrence and a physical knowledge of the potential of the system that generates them. This framework is powerful as it allows for changes in exposed population and assets over time, and considers an extensive range of possible event scenarios well beyond the historical record (see Appendix). This is of particular value when evaluating low recurrence frequency catastrophe events, which by nature have a sparse historical record.

Probabilistic catastrophe risk modelling frameworks have yet to be applied to poverty outcomes at the household level³. If these frameworks could provide perspectives on the relationship between the local severity of hazards arising from natural catastrophes and indicators of poverty/welfare, practitioners would have tools to assess the impacts of shocks under a forward-looking view of potential catastrophe occurrence, with a view to providing assistance or insurance (Muir-Wood, 2014, Anttila –Hughes and Sharma, 2014).

This paper takes the premise that the principal challenge in the application of probabilistic catastrophe risk modelling frameworks to estimate poverty outcomes at the household level is the development of general form relationships between hazard occurrence and indicators of welfare outcome (Anttila-Hughes and Sharma 2014 for further discussion). These relationships would comprise the ‘vulnerability module’ in a catastrophe risk modelling framework as outlined in the Appendix: a model of the relationship between the modelled hazard occurrence and the impact on the exposure. We note here that much work in economics that has been done to show that the ex-post impact of shocks on consumption is a “reduced form” approach which evaluates the impact after households have used every strategy available to them to mitigate the shock (diversification, asset sales etc – see below for further discussion, as well as Porter, 2012). If such strategies include opting for low-risk, low-return activities (Dercon, 1996), then risk also carries an ex-ante cost, which cannot be measured with this particular approach. However, this cost can be difficult to quantify and is very data intensive. It is also less compatible with a perspective of quantifying realized financial losses to be compensated ex-post, if for example index insurance were offered to households.

Poverty, for the purpose of policy analysis, is most often conceptualized as a level of consumption or expenditure that is below some pre-specified level that is considered the minimum for an acceptable standard of living (e.g. based on a basket of food to reach a calorific minimum plus other basic expenditures). Other outcomes such as ownership of household assets, or human capital have also been used in this context. Much poverty analysis up to the turn of the Millennium had been in a static context, but the World Development Report 2000/01 highlighted issues of risk and a hitherto neglected stochastic component of poverty, noting that poor people have low incomes, but also higher risks, and fewer resources to cope with shocks, should they hit.

³ The potential to use probabilistic catastrophe risk models outside of the context of insurance, has been recognized in recent years. This has resulted in developments such as the Pacific Risk Information System³ and the CAPRA³ Program which both apply the probabilistic catastrophe risk modelling framework developed by the insurance markets for disaster risk management. A similar framework has also been applied to estimate food security needs through the Africa RiskView platform. See <http://pcrafi.sopac.org/about/>; <http://www.ecapra.org/>; <http://www.rockefellerfoundation.org/uploads/files/fa08d48b-08ef-4fc7-8991-4872f6e929b0-africa.pdf>

Vulnerability has variously been conceptualized in the social sciences/development studies as a high level of insecurity. In economics, the term has been mainly used to broaden the concept of poverty to include dynamics, and an element of risk – vulnerability to poverty conceptualized as a high probability that at some future point, consumption will fall below the poverty threshold (see Hill and Porter, 2014 for a discussion). The main constraint in this literature to accurately estimating vulnerability to poverty has been an imperfect ability to model future states of the world, with a tendency to rely on the cross-sectional distribution of wellbeing and shocks (and their correlation) to model the probability of any given household falling below the poverty line, though there have been a number of microeconomic studies linking specific stochastic events such as hurricanes, droughts or floods *ex-post* (Anttila –Hughes and Sharma, 2014).

The vulnerability module within a catastrophe risk model contains damage functions⁴ that represent a mean response to a given hazard intensity, with response typically given as a damage ratio (the level of damage expressed as a percentage of total potential damage⁵) (see Jain, V, 2010). The relationships are continuous functions, increasing monotonically with hazard intensity – their shape denoting the form of response of the impacted asset to the shock. Some models accounting for the uncertainty in the damage ratio, by describing a probability distribution around the mean damage ratio (as shown in figure 1.1) for a given level of hazard intensity.

Figure 1.1: Generic damage functions for earthquake demonstrating the response of three structures with different resistances to hazard intensity

Source: CoreLogic Eqecat.

Household welfare outcomes after a catastrophe are influenced by many complex, often interacting, factors beyond the direct damage to household physical assets. Such factors include different adaptive behaviours that lead to different outcomes for a given level of physical asset damage. For example, Hurricane Katrina provides an interesting case study of how the speed of recovery can vary substantially between different cultural groups due to social structures and networks (Leong et al. 2007). This paper asks whether it is possible to reduce these complex relationships into functions that can be used to model poverty outcomes under a forward-looking view of catastrophe occurrence i.e. the feasibility of the development of a ‘vulnerability module’ that relates a specified hazard occurrence to a household-level welfare outcome.

In order to evaluate the feasibility of developing such a module, the work presented herein comprises two components:

⁴ Also referred to as fragility or vulnerability curves or functions

⁵ For buildings, the damage ratio is the cost of repair as a proportion of the total replacement value

- 1) **A regression model to derive quantitative relationships between a selected drought hazard measure and household poverty outcome** for rural households in Ethiopia. This is a survey-weighted regression model combining historical household data with historical data on drought hazard, which effectively constitutes our ‘vulnerability module’;
- 2) **Testing of the derived ‘vulnerability module’ to evaluate its robustness**, and therefore the validity of its future application onto a forward-looking probabilistic view of drought occurrence generated from a catastrophe risk modelling framework. This evaluation is conducted through interpretation of the regression results and the application of Statistical Learning Methods (drawing from James et al (2013)) as described in Section 4.

Our evaluation of validity of the resulting damage functions will be centred around concepts of internal and external validity as described by Anttila-Hughes and Sharma (2014). Anttila-Hughes and Sharma emphasize that ensuring that estimates of damage functions are both internally and externally valid is the major econometric challenge in the development of general form damage functions from historical data on disaster impacts. In the context examined, internal validity is considered as the extent to which impacts statistically associated with disaster occurrence can actually be causally linked to the disaster occurrence, i.e., that the estimates are econometrically “well-identified.” External validity is considered as the extent to which relationships estimated in the context of one disaster event can be generalized to other contexts and locations. More detail on these concepts is outlined in section 4.

The framework for damage function development described above, allows for the application of the derived functions out of the context in which they were derived. This is a fundamental feature of catastrophe risk models, as their purpose is to provide a forward-looking view of risk beyond the historical events used in their development.

2. A methodology for deriving vulnerability relationships for the impact of drought on poverty in Ethiopia

For our chosen area of focus – the impact of drought hazard on the welfare of rural households in Ethiopia – we aim to determine the validity of applying a derived set of damage functions onto a forward-looking view of drought risk. The methodology follows a purely empirical approach using a survey-weighted regression model combining historical household survey data with historical data on drought hazard. Within catastrophe risk modelling, there are two principal methods used to derive damage functions; empirical and analytical derivation. For example, content for the Global Earthquake Model⁶ demonstrates statistical methods for empirical derivation of damage functions from historical damage and loss data, and also analytical methods based on numerical models of response of different structures/structural components (for further information, see D’Ayala et al, 2014 and Rossetto et al, 2014). The analogy for drought damage to crops, for example, would be the use of mechanistic agro-meteorological models based on a process approach versus statistical methods using historical crop loss and drought hazard data.

This model will also attempt to separate out observed relationships for different vulnerability-determining characteristics, analogous to how catastrophe risk models incorporate distinct damage functions for different classes of exposed asset; for example, by structural ‘class’ for buildings or by crop type. For example, a set of primary building characteristics that determine seismic response (such as those shown in table 2.1 from the Global Earthquake Model project) will form the basis of a

⁶ <http://www.nexus.globalquakemodel.org/gem-vulnerability/posts/guidelines-for-empirical-vulnerability-assessment>

set of distinct damage functions. These functions may then be modified using empirical and analytical methods to account for the presence of secondary characteristics that increase or decrease the base vulnerability of the class. Such secondary characteristics include roof pitch for wind damage, foundation type for earthquake shaking hazard and minimum floor height for flood damage (see (World Bank, 2013) for a list of secondary characteristics modelled by AIR for PCRAFI).

Table 2.1: Global Earthquake Model examples of parameter types characterizing building seismic response

Type of Parameter	Examples
Mechanical Characteristics	Strength of the material of the lateral load-resisting system
Dimension Characteristics	Total height / Storey height Number of storeys Plan dimensions - Bay length
Structural Detailing	Tie spacing at the column Reinforcement ratio at the column Hardening ratio of steel
Geometric Configuration	Perimeter Frame Building - Space Frame Building Rigid Roof / Deformable Roof Column orientation

Source: Global Earthquake Model from D'Ayala et al, 2014.

The ability to differentiate the response of a building (or other asset) to hazard intensity according to vulnerability-determining characteristics is an important facilitator of the 'external validity' of damage functions. For example, damage functions developed for buildings in one region can be modified for use in another, if sufficient detail around differences in construction types and quality is available. This out-of-context application is more challenging when considering the impact of hazards on poverty, rather than physical damage, outcomes.

The pathway of impact of disasters on household welfare outcomes is complex and involves many indirect mechanisms. In economics, much research has focused on the ex-post welfare impacts of shocks, using either a "reduced-form" or "structural approach" (Chetty, 2008). The reduced form approach is taken in this paper. Practically, this means that the shock impact on the final outcome is estimated without a structural model of behavioural change. The structural model approach places (as the name suggests) more structure on preferences and behavior of households. In particular, this approach allows the researcher to understand the potential ex-ante effect of behavioural responses, that may cause lower returns (such as adopting low risk, low return crops or activities), and also represent a more complete picture of the welfare loss associated with risk and shocks (Dercon 1996). However, structural models require some quite restrictive assumptions. Further, data scarcity poses significant challenges here, as the complexity of causal mechanisms demands substantial data to be sure of the validity of the apparent relationships observed in the data (Chetty, 2008).

We further acknowledge in our examination of rainfall and poverty that extreme drought is not the only cause of famine for example, the work of Sen (1981) shows that famine is also a political process, and has its root causes in the failure of entitlements. Devereux (2009) shows that people are often pushed into hunger when rainfall may fail only moderately, but if asset prices fall considerably, or food prices increase rapidly, then coping strategies may break down.

With this discussion in mind, we pursue our more positivist approach and obtain the reduced form estimates, following directly from Hill and Porter (2014) though noting that by not estimating ex-ante risk impacts, we are potentially underestimating the full impact of the shock, rather, showing

the actual financial losses that would be experienced by households under different scenarios, conditional on households adopting all other risk management strategies at their disposal (which may have already reduced their welfare in advance of the shock).

The challenges noted above are evident in existing research to elaborate quantitatively the relationships between hazard occurrence and household-level poverty outcomes, with backward looking data. A few examples for further reading are listed below:

- The derivation of quantitative relationships between typhoon wind speed and infant mortality from historical typhoon data in the Philippines (Anttila-Hughes and Hsiang 2013);
- Quantitative examination of the impact of rainfall variability on household consumption in Ethiopia (Hill and Tsehaye 2014, Porter 2012);
- Quantitative examination of the impact of both climate variability and disaster occurrence on the annual probability of permanent household migration in Indonesia (Bohra-Mishra et al. 2014);
- Use of panel datasets to examine the impact of climate variability on allocation time among child labour activities as well as participation in education and labour activities (Colmer 2013).

a. Selection of hazard measure

The measure of drought chosen for this study is an index of crop yield shortfall, taken from the World Food Programme's LEAP (Livelihoods Early Assessment and Protection) software. The index has been developed for the two main cropping seasons in Ethiopia, the *Meher* (main) and *Belg* (minor) seasons.

The yield shortfall calculation uses time-variable meteorological recordings combined with data tables on soil and crop characteristics to calculate yield reductions relative to the expected production under non-limiting water conditions. The yield shortfall calculation is a seasonal value and can be calculated for both *Belg* and *Meher* cropping seasons. It is available at the Woreda⁷ level through creation of a composite index for the relevant crop basket.

The methodology for the calculation is described in Box 2.1 below. If we consider the methodology in the context of a probabilistic catastrophe risk modelling framework, the key 'hazard' input into the process is the rainfall data used in the calculation of Actual Evapotranspiration. It is the high levels of variability around rainfall that are determining the occurrence or otherwise of meteorological drought (translated into crop yield loss), and this is where the large covariate shocks arise that probabilistic catastrophe models seek to capture.

Recorded climate data on temperature, humidity, windspeed and solar radiation are also used in the methodology to determine the evaporating power of the atmosphere at a specific location and time of the year. However, whilst simulating trends in these parameters over the long term could be an interesting exercise on the impact of climate change, the low variability of these parameters over time compared to rainfall make them less interesting candidates for probabilistic simulation in the context of drought.

The LEAP software makes available a combination of satellite and weather station datasets from which historical precipitation values can be taken or estimated. The primary precipitation dataset comes from the National Meteorological Agency of the Republic of Ethiopia (NMA). NMA provides rainfall (and other climate data) on a daily and dekadal⁸ basis for its network of weather stations.

⁷ *Woreda* is an administrative unit in Ethiopia, equivalent to a county in the United States.

⁸ Ten days accumulated

These point estimates are converted to gridded estimates at a spatial resolution of 0.1 x 0.1 degrees (approximately 10km x 10km) for the purposes of calculation. The software also provides an option to use NMA station data modified with satellite data on cloud cover to improve the grid of estimated precipitation⁹.

Box 2.1: WFP LEAP calculation of yield shortfall

LEAP calculates yield shortfall by combining a model for water balance (the FAO's Water Requirement Satisfaction Index (WRSI)) with a model describing crop yield response to water stress (Doorenbos and Kassam, 1979):

$$100 - ((1 - (1 - A/B) * K_y) * 100)$$

Where:

- A is the **Actual Evapotranspiration**;
- B is the **Total Water Requirement without water stress**; and
- K_y is a crop specific factor – “**Yield Response Factor**” - for growing seasons or stages of growing seasons, derived empirically from actual measured crop yield responses to water under good growing conditions¹⁰.

The WRSI gives the ratio of **Actual Evapotranspiration** to **Total Water Requirement** (i.e. A/B in the above) for a season.

- The **Actual Evapotranspiration** represents the actual amount of water withdrawn from the soil water reservoir and is calculated indirectly using rainfall data within a model of soil water balance. It requires inputs of time-variable rainfall data and soil and crop-specific data (soil water holding capacity and crop-specific maximum allowable water depletion);
- The **Total Water Requirement** is calculated as (Potential Evapotranspiration x Crop Coefficient). It requires empirically-derived crop coefficients to relate the reference evapotranspiration to the relevant crop, and time-variable climate data inputs of;
 - o Solar radiation (sunshine);
 - o Air temperature;
 - o Humidity; and
 - o Wind speed.

The reference evapotranspiration expresses the evaporating power of the atmosphere at a specific location and time of the year. Both average and recorded values are suitable for use in the methodology described above, with recorded (actual) values preferred.

Sources: (Hoefsloot & Calmanti, 2012), (Frère, M, and Popov, G, 1979), (Frère, M, and Popov, G, 1986), (Abraha, M, 2013), (Doorenbos, J. & Kassam, A, 1979).

The LEAP protocol for assessing crop yield loss estimates only meteorological drought arising from rainfall variability. This has been criticized as limiting the extent to which the LEAP crop loss figures capture experience on the ground, as they do not capture the many other factors (such as pests) impacting yields¹¹. It is worth reiterating here that the purpose of this paper is not to derive relationships between household poverty and crop yield losses. Rather, the purpose is to establish the validity of relationships derived between a selected poverty indicator and a measure of natural hazard that could be modelled within established probabilistic catastrophe risk modelling frameworks. We are interested in the covariate shocks that can arise from rainfall variability, as this

⁹ <http://www.hoefsloot.com/Downloads/The%20SEDI%20interpolation%20method%20.pdf>

¹⁰ Doorenbos and Kassam, 1979

¹¹ As demonstrated by a 2007 ground-truthing exercise described in: FSCB/WFP Workshop Livelihood Early Assessment and Protection (LEAP): its potential application, benefits and limits, NOTE FOR THE RECORD. 21st January 2008.

is where probabilistic catastrophe risk models can add value (see introduction and Appendix). Measures of meteorological drought therefore serve our objectives. Since we are using crop yield loss indices based on rainfall variability and tables of stationary data on crop and soil characteristics¹², our regression analysis described in Section 3 asks the question:

“What is the relationship between our selected poverty indicator (reduction in household consumption) and crop yield losses arising from rainfall variability?”

The household survey data forming the basis of the regression methodology contains self-reported crop damage arising from sources other than rainfall variability. This field is controlled for in the regression methodology in order to extract as direct as possible a relationship between household consumption and meteorological drought.

Using a rainfall-based index also confers the advantage of objectivity in the measure, removing challenges such as reporting bias that are present when working with reported crop yield statistics (Verma et al, 1988, Fermont and Benson, 2011). It is also more plausibly exogenous than a self-reported measure. However, we acknowledge the limitations of working with a modelled estimate rather than a direct measure of yield impact.

The relative importance of the Belg and Meher seasons for each Woreda was accounted for by taking a weighted-index of the two seasonal values for crop yield shortfall. The regional summary is shown in table 2.1 below:

Table 2.1: Weighting of cropping seasons, by region

	Cropping Season	
	Meher	Belg
	Prod. Share	Prod. Share
Tigray	0.99	0.01
Afar	0.85	0.15
Amhara	0.99	0.01
Oromia	0.95	0.05
Somali	0.89	0.11
B.G.	0.99	0.01
S.N.N.P.R	0.81	0.19
Gambella	0.75	0.25
Harari	0.97	0.03
Addis Ababa	n/a	n/a
Dire Dawa	0.98	0.02
Total	0.95	0.05

Source: Central Statistics Agency documentation, see annex for details.

The key questions to be answered in terms of the household regression model have been agreed as:

- To what extent can impacts statistically associated with a drought hazard measure be causally attributed to the drought (internal validity/robustness)?

¹² See Box 1 for data inputs, which include additional non-stationary, but low-variability climate data

- To what extent can the regression results be applied to predict outcomes from hazard in other contexts (external validity/robustness)?

b. Selection of household characteristics for regression analysis & Regression Model definition

The model specified is: $y=D(h,s)$ where y our outcome of interest is the log of consumption per adult equivalent, h is the community level annual crop loss as predicted by LEAP (defined above), and s are other household and community characteristics (including other shocks experienced by the household). We note also that the relationship between h and y will be attenuated by certain household and community characteristics (s), and the regression model will therefore seek to draw out these attenuating impacts separately to increase the external validity/robustness of the model.

The regression model will therefore output a relationship between consumption and crop loss, specific to different household 'categories'. Households will be separated into categories according to the characteristics known to attenuate the impact of drought on consumption; each distinct set of characteristics receiving a distinct relationship between crop loss and consumption. The relationship in all cases being defined by a selected functional form and coefficient output from the regression model.

Regression specification

The base specification is based on initial work by Hill and Porter (2014) that derived a general model of consumption for Ethiopian households using all areas, rural and urban, and focused on the impact of drought, food prices, and other idiosyncratic shocks on ln consumption per adult. We make several modifications due to the new aims of the report, mainly concerned with achieving precision on the relationship between drought and consumption for differing household characteristics.

The **dependent variable** is the natural logarithm of consumption per adult equivalent at household level, as used by Hill and Porter (2014). The drought-crop-loss variable or **hazard** is the LEAP estimated crop losses at woreda level (see above for definitions).

Household characteristics included in the model are: HH head gender, age and education, and a dummy for household head not in agriculture; HH assets including cattle, sheep, chickens, land, good roof, toilet; Idiosyncratic shocks including crop-loss, animal illness or death, hh member illness or death, food price shocks; other characteristics including financial capital and household composition. Community characteristics include: agro-climatic zone, region, distance to town, market access.

As noted above, the model also seeks to capture what in risk-modelling are termed attenuating factors, in econometrics as heterogeneous impacts, through interaction terms included in the model (e.g. LEAP*varname). The statistical learning methods are used to perform model selection for the appropriate specification of the relationship between drought-crop-loss and consumption, including potential non-linear specification (higher powers of LEAP e.g. LEAP-squared, cubed etc) and selected interaction terms. The shortlist of interacting variables includes characteristics of the household/community that may affect household ability to cope with shock, these include: ability of head to access coping strategies (head education, head not in agriculture, hh doesn't own cattle, dependency ratio); Other shocks that may already be stretching household ability to cope (crop-damage, livestock shock, illness); Access to institutional coping strategies (distance to town/market, access to financial capital, public safety net).

Table 2.2: Household characteristics and proxy variables

Household characteristics	Interacting variables
Ability of head to access coping strategies	Head education level, sector of occupation, gender
Household composition that allows labour response	Dependency ratio, ratio of able-bodied
Household assets that mitigate shock	Cattle, other livestock, jewelry
Other shocks that compromise ability to mitigate shock	Illness, livestock disease, crop damage from pests
Access to institutional coping strategies	Distance to market, access to financial products (insurance, credit), public safety net access.

Finally, agro-climatic zone is discussed in a separate section below.

c. Internal validity

We keep with the notation of Antilla-Hughes and Sharma (2014, AHS) who outline succinctly the issues in establishing internal validity when estimating hazard relationships: in our case the internal validity is concerned with the establishment of a causal (or, at least, externally replicable) relationship between crop loss and consumption. In the absence of a randomised experiment assigning some households to better community-level crop yields than others, we must assume implicitly that (in some specific sense) the crop loss in a given year is exogenous to (uncorrelated with) unobservable household and community characteristics that affect household consumption. We acknowledge that there are some issues around this. For example, if crop yields are highly correlated over time this may lead to lower average consumption in the community, and therefore a bias towards households who are less intrinsically productive living in high crop loss communities. Without a panel dataset we are not able to deal with this and to the extent that other unobservable household and community characteristics that impact consumption are not included in the regression, and *are correlated with the level of the drought shock in a given year*, we may be capturing biased estimates of the relationship between drought and consumption. We can include average rainfall at the community level to help control for this, or alternatively, a fixed-effect at the zonal level (there are 68 zones in Ethiopia). The datasets used in the analysis do allow a rich set of controls in which we aim to capture as many characteristics as is feasible, without over-fitting the model to the dataset at hand. With this in mind, we also use the statistical learning procedures to validate the model.

The question of internal validity may be more weakly stated as being a “stable” observable relationship between drought-crop-loss and consumption over several specifications (and external validity over several datasets), and answerable by looking at whether the point estimates and associated p-values established through the regression work, and statistical learning through the bootstrapping method described later. By giving a level of significance around the relationship between consumption and crop loss (defined by the coefficients output from the regression, and their associated functional form), the derived p-values give us information on the strength of fit of the modelled relationship. The bootstrapping method generates a mean estimate of the relationship and a related standard error (see below for more detailed methodology) which allows us to examine the stability of the coefficients that have been derived.

However, it must be noted that from an econometric perspective, the absence of random application of the shock to households (e.g. if at all there may be a bias for households more susceptible to the impacts of shocks locating themselves in hazard-prone areas) limits the extent to which internal robustness/validity can be examined using cross-sectional data. The purpose of the research sidesteps this issue somewhat, as the question of interest is not primarily about causality *per se* but rather, an externally replicable association. That is, whether it is possible to capture a

relationship between shock and impact for other households of the same type (type being defined according the list of vulnerability-determining household characteristics we identify as attenuating factors in the drought-consumption relationship) that can be replicated in other contexts/ using other datasets. As noted below, one of the datasets used in the analysis has now been updated to a 2-period panel, which could potentially be an informative extension to the analysis.

d. External validity

External validity is concerned with whether results of any individual piece of analysis are generalizable to other contexts. Whilst in scientific contexts, it may be possible to establish complete external validity for certain experiments, in the context of social science, the question is rather *how* generalizable are the results to other contexts? In our case, the model is to be estimated on rural Ethiopian (nationally representative) data, so as long as we believe the results to be internally valid, then we could say that they are representative for the population in that particular time and place (Ethiopia, 2011 for example). However, it is more difficult to establish whether the relationships we find between rainfall and consumption in 2011 are generalizable to other years with e.g. more extreme drought.

It is helpful to define three specific sub-topics within overall external validity that are relevant in our context. These are referred to as EV1) out-of-sample shock estimates EV2) time-stability of estimated relationships and EV3) concerns around over-fitting to the training data. We define this in more detail below.¹³

EV1) Antilla-Hughes and Sharma (2014, hereafter, AHS) note that the main threat to external validity in the context of specifying the damage function $D(\cdot)$ as outlined above, lies in mis-specifying the function for values of the hazard (shock) that were unavailable for the analysis- in particular, extreme values that occur only very infrequently (e.g. a once-in-a-generation severe drought, typhoon, tsunami etc).

For the initial model estimation Hill and Porter (2014) combined the years 2004-5 and 2010-11 in order to allow the broadest (artificially) cross-sectional range of values for our shock. In some pre-analysis, we conducted a careful analysis of the range of data (see LEAP document). The issue in the context of our dataset is the limited number of observations available at some of the high levels of crop loss (e.g. greater than 50 percent crop loss).

Figure 2.1 shows the distribution of the crop loss data for 2011 and 2005 combined, and it can be seen that the distribution is left-skewed. Where we identify data paucity that challenges the strength of fit of the modelled relationship at this extreme end of the crop loss spectrum, we need to highlight a caveat regarding the sparse data above 60% crop loss. Table 2.3 outlines in more detail the number of observations in each “bin” of data at 5% intervals. 2005 is somewhat worse year than 2011, with a higher mean crop loss, and 300 observations above 50%. However in 2011, just under 100 observations lie between 50 and 60%, and none at all are above 60%. The distribution can then be said to be somewhat different between years, and we therefore examine the difference between the pooled dataset that has a broader distribution, with the use of one round of data as the training data, and another round as the testing data, in order to test the stability of the drought-consumption relationship over different drought distributions.

¹³ EV1) is discussed as overlapping within the scope of the internal validity checking.

Figure 2.1 Frequency and 5% bin of Drought-Crop-Loss Data (LEAP)¹⁴

Table 2.3: Distribution of crop –loss data in 2005 and 2011 survey data

% of crop loss from LEAP data

	0	0 to < 5	5 to < 10	10 to < 15	15 to < 20	20 to < 25	25 to < 30	30 to < 35	35 to < 40	40 to < 45	45 to < 50
2005 survey	1434	3467	816	548	621	790	373	420	575	122	393
2011 survey	1955	4311	934	653	576	564	690	579	154	114	140
2005 & 2011	3389	7778	1750	1201	1197	1354	1063	999	729	236	533

	50 to < 55	55 to < 60	60 to < 65	65 to < 70	70 to < 75	75 to < 80	80 to < 85	85 to < 90	90 to < 95	95 to =100	Total
2005 survey	72	24	1	1	23	0	56	70	59	0	8431
2011 survey	43	49	0	0	0	0	0	0	0	0	8807
2005 & 2011	115	73	1	1	23	0	56	70	59	0	17238

EV2) The second issue is over what length of time relationships established can be considered as valid. For the purpose of CAT risk modelling, it is helpful to consider relationships to be stable over a five-year horizon. To establish this, we attempt to use a dataset from a different survey year (2012, as opposed to 2005 and 2010-11) as validation data. We do not anticipate the external validity of the derived relationships to be applicable over very broad timeframes (i.e. 20 years or mre), due to structural changes in the economy that lead to changes in lifestyle/behavior and potential exposure to shocks. We discuss the testing dataset for 2012 in a later section of the paper.

Also related to temporal aspects of shock exposure, AHS also raise the concern of recurrence times and macroeconomic (or second order) effects. An example of recurrence times is that, in one crisis households with livestock may sell something in order to protect their consumption, but if another crisis hits, the impact of the second shock is likely to be higher. We are not likely able to address this latter point with the data available, and so must simply take the ex-post distribution as being standard.

¹⁴ NB: Distributions use 2004 LEAP (Meher and Belg) for 2005 survey, 2010 LEAP (Meher and Belg) for 2011 survey (interviews Jan - June 2011) 2009 Meher and 2010 Belg LEAP for 2011 survey (interviews July-Dec 2010) .

EV3) Roe and Just (2009) define the concept of *ecological validity*, which is useful in this context: “a study has ecological validity to the extent that the context in which subjects cast decisions is similar to the context of interest.” (p1267) In the context of this broader external validity, we may ask “can the consumption model be applied to other countries than Ethiopia?” A slightly narrower and more tractable initial question can be attempted: “is the specified model valid across Ethiopia?” We do conclude however that it is unlikely to be possible to derive a valid model for pastoralists using the crop loss parameters, given their reliance on grazing rather than crop yields. In one of the specifications, the regression model will also explicitly treat agro-climatic region as an attenuating factor, testing whether the relationship between crop loss and consumption will not hold across regions. This is consistent with the findings of Hsiang and Narita (2012), which suggest that the efficacy of adaptive behaviour to weather shocks varies with average climate conditions.

Box 2.2: Impact Amplification

The limited number of survey years available for the regression poses a number of challenges. One such challenge is the impossibility of capturing impacts arising from the severity of the seasonal drought on a national or regional scale, due to the very small number of data points available to describe this across the two examined years.

In the context of catastrophe risk models, these impacts are captured as **post-event loss amplification**, or **demand surge**¹⁵. For the purposes of the exercise presented herein, the equivalent would be impact amplification.

Claims data from large catastrophes in developed insurance markets shows that the cost of a given level of damage increases when that damage is situated within a catastrophe of such severity as to have a significant impact at the national or regional scale. In the context of building damage, modelling firm RMS (Souch, C, 2010) describes the drivers of this loss amplification from the perspective of insurers as being:

- Economic demand surge leading to inflated costs. This arises from an imbalance in supply and demand for key materials and labour required for recovery and reconstruction due to the extent of the disaster footprint.
- Deterioration vulnerability leading to damage escalation. This arises from delays in addressing the impacts from the disaster.
- Claims inflation leading to increased payouts from insurers. This arises from a relaxation in insurers’ claims assessment procedures due to the scale of the task.
- Coverage expansion leading to increased payouts from insurers. This can arise from political pressure for insurers to honour claims that would not have qualified for payouts under the terms of the insurance contracts in place at the time of the event.

These impacts are tied to the overall scale of the event, rather than the local intensity used to calculate damage in the vulnerability methodology laid out in this paper, and as such, require data from a variety of catastrophes of different magnitudes in order to derive factors for amplifying losses/impacts. Although we are not looking through the lens of insurance, a number of aspects of post-event loss amplification as described above are relevant to our analysis, and as such are disclosed here as limitations.

For example, in the aftermath of a large drought, impacts on food prices (and other commodities for which demand has increased), can be expected. This can impact household consumption. Changes in the labour market should also be considered, as a large drought event may push a large number of households to turn to alternative employment as a coping mechanism. This could disrupt the balance of supply and demand for certain types of labour, impacting the effectiveness of this option as a coping mechanism. A large drought may also prompt a political response that influences household consumption in a way not seen in the historical dataset used to derive damage functions.

These non-local factors are complex and would be difficult to capture, even with a more comprehensive dataset. However, the fact that the impact on a household is influenced by the scale of a disaster, as well as by its local intensity, is indisputable, and not accounted for in the methodology presented in this paper.

¹⁵ In some models, demand surge is taken as a sub-component of post-event loss amplification

A final note of caution around the estimation results is that there is most likely measurement error in our variables that could potentially lead to downward bias in the estimates. This is not dealt with in the current draft.

We now present the data and regression results, followed by the statistical analysis and robustness testing.

3. Presentation of results of the derivation of vulnerability relationships for the impact of drought on poverty in Ethiopia

This section discusses the results of the vulnerability relationships estimated. We first present the descriptive statistics, and then the tables of results. The next section discusses the statistical learning methods attempting to validate the regression results.

a. Data

The data used by Hill and Porter (2014) to establish vulnerability to poverty relationships is used as the main dataset for the initial regression results. We use the data collected in the 2005 and 2011 rounds of the nationally representative Household Income and Consumption Expenditure and Welfare Monitoring Surveys (HICES/WMS). These contain information on just under 25,000 households in each year. The key information recorded in the HICE used to calculate vulnerability is expenditure on food and other items. The WMS records household assets and characteristics as well as a fairly detailed module on self-reported adverse events (referred to as shocks throughout). In both years they were administered by Ethiopia's Central Statistics Agency (CSA).

The advantage of using the HICES-WMS for vulnerability analysis is that they are relatively large, nationally representative, comparable across years and allow measures of vulnerability to be estimated at the household level. This allows us to look at the relative importance of geographic and household factors in determining vulnerability, and it also allows us to examine how vulnerability varies across certain groups of households.

As noted above, we exclude pastoral areas of Ethiopia from the preliminary analysis as we believe the relationship between rainfall and consumption to be structurally different to that of other regions. The descriptive statistics for the two datasets are shown in table 3.1.

Table 3.1: Descriptive statistics

Variable	2005	2011
Ln adult monthly expend	7.27 (0.50)	7.28 (0.50)
LEAP crop loss	16.33 (18.53)	11.58 (13.25)
Femalehead	0.23 (0.42)	0.23 (0.42)
Age hhead	43.24 (15.73)	44.43 (15.74)
Head not agri	0.16 (0.37)	0.14 (0.35)
Head school	0.25 (0.43)	0.30 (0.46)
Lotsplots	0.57 (0.49)	0.53 (0.50)
Cattle	0.66 (0.47)	0.67 (0.47)

Sheep	0.49 (0.50)	0.51 (0.50)
Chicken	0.60 (0.49)	0.59 (0.49)
Financial access	0.25 (0.43)	0.50 (0.50)
Distance to town	326.83 (229.93)	378.64 (285.74)
Dependency ratio	0.49 (0.24)	0.50 (0.23)
Death shock	0.08 (0.27)	0.02 (0.13)
Illness shock	0.23 (0.42)	0.08 (0.27)
Cropdamage shock	0.10 (0.30)	0.03 (0.17)
Livestock shock	0.09 (0.29)	0.05 (0.21)
Jobloss shock	0.01 (0.09)	0.00 (0.04)
Price shock	0.02 (0.14)	0.18 (0.38)
Psnp beneficiary	0.00 (0.00)	0.15 (0.36)
Highlands-drought	0.39 (0.49)	0.37 (0.48)
Highlands reliable	0.38 (0.48)	0.34 (0.47)
Lowlands reliable	0.03 (0.18)	0.11 (0.3)1
Lowlands onset	0.19 (0.40)	0.18 (0.39)
Good roof	0.22 (0.42)	0.34 (0.47)
Electricity	0.00 (0.00)	0.06 (0.24)
Improved loo	0.21 (0.41)	0.59 (0.49)
Household size	4.91 (2.28)	5.00 (2.20)
Number obs	8431	8807

b. Regression results

In this section we discuss the result for the vulnerability relationships tested. We modelled the dependent variable log of consumption per adult equivalent in the household, and included (the same list of) household and community characteristics in each model. We then estimated a baseline model with simply the LEAP % loss in linear form.

We tested for a non-linear relationship of LEAP but as results below show, the linear model is the preferred specification. We noted above that in line with the disaster modelling, we explicitly incorporate characteristics of households that may mitigate or exacerbate the impact of the drought shock. The analogy can be drawn with table 2.2 above which shows characteristics of buildings with different wall structures. Table 3.2 below shows the list of mediating characteristics that were included in the regression model. In economics, this would be termed as heterogeneity in the

drought impact, and in the regression model, interaction terms are included between the shock and each characteristic.

Table 3.2: Household characteristics that affect drought shock impact

Characteristic	Expected to mitigate/exacerbate?
HOUSEHOLD RESOURCES TO MITIGATE	
Cattle owner (yes/no)	Mitigate
Not in agriculture (yes no)	Mitigate
Female headed(yes/no)	Exacerbate
Head schooling (yes/no)	Mitigate
Dependency ratio	?
INSTITUTIONAL RESOURCES	
PSNP (safety-net) access (yes/no)	Mitigate
Financial access	Mitigate
Market access (distance to town)	Mitigate
SHOCKS THAT COMPROMISE COPING	
Illness shock	Exacerbate
Crop damage (non-drought)	Exacerbate
Livestock illness/death	Exacerbate

It can be seen in table 3.3 below that there are some household characteristics that are significantly correlated, though the correlation coefficient is never above 0.3. Female-headed is correlated with not being in agriculture and having no schooling, as well as being more likely to have access to the PSNP but less access to financial services. Surprisingly, PSNP and financial services are positively correlated.

Table 3.3: Correlation between household characteristics

	Cattle	Notagri	femalehead	hschool	psnpb	finaccess	disttown07	dependency	illness	cropdam	livshock
Cattle	1										
notagri	0.2168*	1									
femalehead	0.1670*	0.3325*	1								
hschool	-0.011	-0.0164	-0.2297*	1							
psnpb	0.0471*	0.0271*	0.0636*	-0.0558*	1						
finaccess	-0.0245*	0.0254*	-0.0508*	0.1071*	0.0505*	1					
disttown07	0.0871*	-0.0182	0.0191	0.0063	-0.0359	-0.0627*	1				
dependency	-0.0742*	0.0199*	0.0558*	-0.0476*	0.0545*	-0.0330*	-0.0292*	1			
illness	0.0271*	0.0370*	0.0186	-0.0282*	-0.0653	-0.0746*	-0.0019	0.0092	1		
cropdam	-0.0001	-0.0222*	-0.0179	-0.0016	-0.0471*	-0.0781*	0.0319*	0.0223*	0.1518*	1	
livshock	-0.0480*	-0.0374*	-0.0287*	-0.008	-0.0294*	-0.0732*	0.0815*	0.0317*	0.2078*	0.1794*	1

Notes: See descriptive statistics for full variable names. * significant at 1%

Therefore our four specifications include the baseline (no interactions); **model 1**, a parsimonious model including only female-headed household, head schooling and access to the social safety net (PSNP). **Model 2** includes further interaction terms of distance to market and the dependency ratio. **Model 3** includes a full set of interaction terms, including the financial access indicator, and the other shocks that may compromise household ability to cope with drought. We do note the caveat here that this latter set of variables are self-reported and therefore are potentially endogenous, especially if having experienced a drought in any way affects the response to the questions asked, which we discuss further below.

Summary of results:

Table 3.4 below shows the results. We present the coefficient on the LEAP variable (divided by 10, to ease interpretation). As the dependent variable is in log form, we can read the coefficients as “a 10% worsening of the LEAP crop loss leads to an X% reduction in consumption per adult”. For example, the baseline result shows that for every 10% worsening of LEAP, consumption falls on average by 1.5%. We computed bootstrapped standard errors across all of the statistical learning components (see next section) and therefore also present a 95% confidence interval for the point estimate based on the bootstrapping results. The LEAP drought variable is significant in all specifications. In the following columns, interaction terms are included, so the point estimate is interpreted as the drought impact for households that are not defined by any of the characteristics included. So, for column 2, we include female headed, psnp, and schooling. The coefficient of 2.0% represents the impact of a 10% increase in the LEAP drought on consumption of male headed households with no schooling and no access to the PSNP. There is a zero coefficient on schooling and female headed, so we do not see any difference in the impact for those groups. However PSNP access mitigates the drought impact by 0.5% = therefore households with PSNP access would experience a 1.5% decrease in consumption (rather than a 2% decrease with no access).

Reading across the columns we see no differential impact in any of the specifications for female headed households, educated household heads, head not in agriculture, distance to town, dependency ratio.

We do find significantly different impact for households with access to the PSNP, with cattle, and financial access: these characteristics are seen to be mitigating the impact of the drought. However, households that suffered other crop damage experienced a heavier impact of the drought.

Noting above the potential concern that financial access and shocks such as crop damage are self-reported.

Table 3.4 Main regression results

	Baseline	Specparsl	Spec2	Spec4	Spec3
(Intercept)	7.794*** (0.018)	7.800*** (0.018)	7.815*** (0.020)	7.821*** (0.020)	7.815*** (0.020)
Drought shock	-0.015*** (0.002)	-0.020*** (0.003)	-0.028*** (0.006)	-0.033*** (0.007)	-0.058*** (0.007)
boot.se	0.0025	0.0028	0.0066	0.0067	0.0078
boot.ci	(-0.020, -0.0104)	(-0.0242, -0.0130)	(-0.0408, -0.0148)	(-0.0463, -0.0202)	(-0.0728, -0.0425)
Drought* Head school		0.000 (0.001)	0.000 (0.001)	0.000 (0.001)	0.000 (0.001)
Drought*Femalehead		0.000 (0.000)			
Drought* PSNP		0.005*** (0.001)	0.005*** (0.001)	0.005*** (0.001)	0.005*** (0.001)
Drought*non-agri hh			0.000 (0.001)	0.000 (0.001)	-0.000 (0.001)
Drought*dist. market			-0.000 (0.000)	0.000 (0.000)	0.000*** (0.000)
Drought* dependency			0.001 (0.001)	0.001 (0.001)	0.001 (0.001)
Drought* anycattle			0.014** (0.004)	0.014** (0.004)	0.018*** (0.004)
Drought* Fin. access				0.001** (0.000)	0.001** (0.000)
Drought* Illness				0.000 (0.001)	0.000 (0.001)
Drought* Cropdamage				-0.001* (0.001)	-0.001* (0.001)
Drought* livestock shock				-0.000 (0.001)	-0.000 (0.001)

D* highlands reliable					0.005*** (0.001)
D* lowlands reliable					-0.004*** (0.001)
D* lowlands enset					0.004*** (0.001)
R-squared	0.245	0.246	0.247	0.247	0.251
Adj. R-squared	0.244	0.245	0.245	0.246	0.249
AIC	20047.21	20022.04	20016.28	20010.07	19937.65
BIC	20271.92	20270.00	20287.49	20312.27	20263.10
Num. obs.	17134	17134	17134	17134	17134
boot.cv.cropleap	0.249	0.248	0.248	0.248	0.248
boot.cv	0.170	0.171	0.170	0.170	0.169
RMSEtest05	0.0232	0.0281	0.0338	0.0272	0.0291
RMSEtest11	0.0438	0.0431	0.0427	0.0429	0.0490
RMSEtestHRE	0.0514	0.0546	0.0595	0.0602	0.0602
RMSEtestHDR	0.0302	0.0297	0.0519	0.0544	0.0544
RMSEtestLOE	0.0042	0.0060	0.0055	0.0055	0.0055
RMSEtestLOR	0.0602	0.0080	0.0034	0.0013	0.0013

Notes: s.e.=standard error cv=cross validation. R2=r-squared; AIC=Aikike information Criterion; BIC= ; boot=bootstrap; RMSE=root mean squared error. Sample size = 17238. Full results in appendix.

To explore whether the impact of drought differs across regions we created dummy variables for the four agro-ecological zones of Ethiopia, and included these in the final regression model. Unreliable rainfall has historically underpinned much of the discussion on vulnerability in Ethiopia given the widespread predominance of livelihoods that are dependent on rainfed production systems. This characterization of vulnerability has resulted in a widespread understanding of a geographic footprint of vulnerability. Until recently the Government of Ethiopia framed rural policy discussions around “three Ethiopias”: drought-prone highlands, moisture-reliable highlands and pastoral lowlands. This classification was recently been expanded to a concept of “five Ethiopias” according to agricultural productivity and agroecological conditions (EDRI 2009). The five areas are drought prone highlands, moisture-reliable cereals areas, moisture-reliable enset areas, humid moisture-reliable lowlands and pastoral areas.

Pastoral areas are excluded from the analysis. In the final column of table 3.4 we show interaction terms with the regions. All three interactions are significant. With highlands drought as the baseline category, we find that the highlands reliable suffer significantly less due to drought. In fact, a test of whether the drought impact is equal to zero for the highlands reliable region cannot be rejected. The impact is also less for enset growing region, though the test of zero impact is rejected at 5%. Surprisingly, lowlands reliable is the region with the highest impact, 3% higher per 10% leap loss than the highlands drought region. This evidence does show that the relationship between drought and crop loss is different in each of the agro-ecological zones, and we further check the models in section 4 below when we implement statistical testing.

Non-linear impacts of drought

We estimated the same model and performed statistical testing on a quadratic and cubic form to test for nonlinearity of the relationship between drought and consumption, which is plausible. The full set of results are presented in Annex 1, Table A2. The cubic model results are presented in Table A3. Adding these higher powers did not change the coefficients on the interaction terms, so we present here in figure 3.1 a graph showing the simulated shape of the curve using the squared and cubic models for ease of interpretation.

Figure 3.1

The cubic model appears to have a second turning point around 70% crop loss- which is around the point at which we lose support for the data in 2011, so we may not have enough values of the data to create a plausible estimate for any further nonlinearity than a squared term.

4. A methodology for testing the predictive power/robustness of the derived vulnerability relationships for drought and poverty in Ethiopia

The proposed methodology to testing the predictive power of the vulnerability relationships is to use Statistical Learning Methods (re-sampling and cross-validation). This section draws on James et al (2013). Statistical learning is used for:

- Assessing model accuracy,
- Checking the performance of different functional forms of $D(.)$
- Choosing the model with the lowest “test” MSE (the best prediction) instead of just the lowest “training” MSE (the best fit on the currently used historic data).

(James et al 2013 “Introduction to Statistical Learning”)

a. Concept of training data and test data

We discuss the datasets used below. However, we outline both a method that can be applied using one (existing, clean) dataset, and another that can be applied to a second dataset. We propose that the two methods answer somewhat different questions and something can be learned from comparing the results. This methodology has not been much used in Economics, however, Todd and Wolpin (2007) use the concept of “holdout data” in a study of test scores for children in the US, in order to compare competing models of cognitive achievement accumulation (e.g. deciding whether to include lagged test scores in the model). The authors choose the model that performs best in terms of RMSE. Recently, the World Bank has applied the method of training and testing datasets in the validation of poverty scorecard methodology (Diamond et al, 2015).

The resampling method works in general by examining the fit of the model when we apply the results from the regression work on a training dataset (where the model is initially fitted) to a testing dataset (ideally a separate dataset or subset of a dataset that has never been used to fit the model).

In our case, we can take the initial dataset, and partition into a (randomly selected) training dataset that is used to fit candidate models. Then, the rest of the data are used as if they were a new dataset, as test data, in order to validate external validity.

b. Model selection: Cross-validation approach

K-fold Cross validation can be used to confirm what one might call pseudo-external validity by explicitly dividing the data into training and testing datasets: the testing dataset is a randomly selected subset of the original dataset (with n/k observations). The procedure is repeated by allowing each of the k folds of the dataset to be excluded from model fitting and treated as if they were external datasets.

A number of functional forms, with associated coefficients derived from the regression model will be compared using this methodology. A limited number of regression models will be compared using this method. We limit the number of possible models because of the aim of the study, which is to examine the effect of drought on consumption. Hence, we consider community characteristics as control variables and household characteristics as potentially interacting variables, to be dropped if found insignificant.

The k -fold cross validation begins by randomly allocating all n observations of the data into k -parts (folds or groups) of approximately equal sizes. The first fold is treated as a “testing” set and withheld while the model is fitted to the remaining $k-1$ folds of the data (the “training” data). The observations in the first fold are then fitted to the estimated model and the mean squared error (MSE) MSE calculated, or MSE1. The procedure is repeated, each time using a different fold as the validation dataset (and $k-1$ folds as the training dataset). The testing MSE is defined as the average of MSE1, MSE2, ...MSE k .

$$CV_k = \frac{1}{k} \sum_{i=1}^k MSE_i \quad (5.1)$$

This methodology will give us information on the strength of fit of the models with respect to each other, and in absolute terms, using the testing MSE, in particular we examine the relative improvement in the MSE relative to the complexity of the model.

We will also use the Akaike Information Criteria (AIC), which penalizes a model for including “unnecessary” variables or complexity, to complement the k -fold cross validation method of model selection.

c. Stability of coefficients within the model

Our discussion thus far of model testing comprises concerns around the issues of “fit”. However it is worth discussing where the question of ‘predictive power’ of the derived relationship between drought and consumption may require testing of the stability of the proposed relationship. Given that the end-purpose of the modelling is to “bolt” the model on to simulations of possible drought outcomes, we need to ensure that the model does produce a stable prediction of the key relationship between crop loss and log consumption.

With bootstrapping, M distinct data sets with n observations are drawn (with replacement), by repeatedly sampling observations from the original data set, the same observation can occur more than once in the bootstrap data set. The model is fitted to each of these M datasets. The coefficient of interest is computed M times, and we can compute the standard error of these M bootstrap estimates.

We will then compare the bootstrap estimates of the drought parameter across the candidate models. The model with the least bootstrap variance would be selected. Thus, bootstrapping can be considered as a test of internal validity. The relevant internal validity question can be phrased as: given a model that has been estimated using a nationally representative sample, if we resampled many times, would we estimate a stable relationship?

Note here that either method (bootstrapping or k-fold) can be used to validate the parameters of an estimated model.

Outputs from the analysis

The shortlisted functions of $D(.)$ in the conceptual section are candidate models 1-4 outlined above. We proposed that the control variables should be selected on the basis of economic theory/empirical findings/common sense. The interactions should also be selected on that basis, however the number of interactions, and the functional form is to be decided using the statistical analysis. The next step in the exercise is to test the performance of the candidate models in predicting consumption in different contexts (or out of sample predictions) using different values of X , than those based on which the parameters for the $D(.)$ functions were estimated (i.e. minimizing or reducing “test” MSE).

How to select candidate functions of $D(.)$?

Recall that $D(.)$ comprises the hazard I and household characteristics X . There are two issues:

- a) Specifying the functional form for $D(h)$ whilst holding X constant and
- b) Which characteristics enter the vector X ?

For a) we investigate vulnerability functions specific to household “types” as noted in the description of the regression model above. We also investigate ii) whether the vulnerability function for all types should be non-linear by including squared and cubed terms in the model.

For b) we use economic theory to choose the characteristics of households and communities that enter the function. These will be constrained by the availability of data in the household survey datasets.

d. How significant is the variation of the derived functional forms across geographical regions within Ethiopia?

The first round of testing discussed above involved random division into k parts. However, also split groups according to characteristics that vary across the dataset, but that we have not explicitly identified as determinant of vulnerability, to test *external validity, or generalizability across ecological contexts*. We use non-random sampling for the k-Fold Cross Validation method, separating the k groups according to geographical area (e.g. Woreda) but controlling for differences expected from changes in agro-climatic conditions. This will tell us if, when we separate out households according to what we have determined to be the key vulnerability characteristics, the relationships hold across geographies.

However, we have reasons to believe that the relationship between rainfall, crop losses and consumption differs considerably in pastoral areas as discussed above so we exclude the pastoral regions from the analysis.

e. External Validity using a second dataset

To establish external validity of any statistical results across time and contexts beyond reasonable doubt, it would be necessary to conduct identical analysis many times (replication) and then derive bounds for the relationships. This is not possible in the timeframe of the initial analysis, but has potential for a future exercise.

Up to now, we have examined statistical learning techniques using the original dataset to create a “pseudo-external validity”. To further explore the concept of external validity, we attempted replicating the analysis of the initial vulnerability function on a second dataset. However, the consumption methodology was incomparable so we briefly review the method here.

To answer this question, we would first employ model selection techniques to identify the model which gives the most similar results to the first dataset. In this case, there is no resampling involved, we use the second data set as the validation/testing dataset.

The first step involves fitting a regression model using the original/training dataset. Thereafter the second dataset or testing dataset is fit to the model and the MSE is calculated.

Given that the second dataset consists of n observations, $\{(y_1, x_1), (y_2, x_2), \dots (y_n, x_n)\}$, we would find the predicted values, $\hat{f}(x_i)$ predicted by the regression model obtained using the training dataset and we calculate the MSE;

$$MSE = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{f}(x_i))^2 \quad (5.2)$$

We would then select the model for which the test MSE is the smallest.

To measure the accuracy of the model or the extent to which it gives the most similar results to the first dataset we would then calculate the bootstrap estimates, this time resampling from the second dataset.

5. Presentation of results from the statistical analysis with interpretation

As noted in section 4, statistical learning can give information on how informative a model is when applied to a holdout sample. As this was an exploratory exercise, rather than presenting one methodology and set of results, the exercise begins with the q-fold replication. One subset of the data is taken as the “training dataset” and then we assess the model fit in terms of the RMSE on the holdout sample or “testing dataset”. There are thus several different cuts of the data that have been applied to assess the information needed.

We began with the combined 2005/11 dataset, and conducted q-fold (10 rep) cross-validation on the shortlisted 4 models with a linear specification for the drought variable. To be specific, this meant that for the pooled dataset, we fit the model onto 9 tenths of the data, and assess the fit on the one remaining tenth. The exercise is repeated 10 times. As discussed in section 4, the models are assessed based on the average mean squared error (MSE) over all ten repetitions as defined in equation (5.1). The results are shown in table 6.1 below. There is little difference between the models, with model 4 having the lowest average MSE by a small margin.

Table 5.1: Testing the models with 2005 and 2011 as training and “holdout” data

	Baseline	Model 1	Model 2	Model 3	Model 4
Original R^2	0.245	0.246	0.247	0.247	0.251
Original Adj. R^2	0.244	0.245	0.245	0.246	0.249
Original AIC	20047.21	20022.04	20072.26	20065.60	20001.73
Original BIC	20271.92	20270.00	20335.72	20360.06	20319.43
Num. obs.	17134	17134	17134	17134	17134
boot.cv	0.170	0.171	0.170	0.170	0.169
RMSEtest05	0.0232	0.0281	0.0338	0.0272	0.0291
RMSEtest11	0.0438	0.0431	0.0427	0.0429	0.0490

Notes: cv=cross-validation, R^2 =r-squared; AIC=Aikike information Criterion; BIC= ; boot=bootstrap; RMSE=root mean squared error. HRE=Highlands, Reliable Region; HDR=Highland, drought-prone Region; LOE=Lowlands Enset growing region; LOR=Lowlands reliable region. 05,11 refer to the datasets collected in 2005, 2011 respectively (and defined above).

We then repeat the exercise, but this time considering the whole dataset for 2005/2011 as the training/testing dataset respectively. In this way one might imagine going back in time to 2005, predicting consumption with that dataset for 2011, and assessing the results. As a though exercise we also repeat the process but with the training and testing roles reversed. When 2005 is the training dataset, Model 2 is the best predictor (lowest MSE). However when 2011 is the training dataset, the baseline model performs best.

Comparing how useful the two datasets are as “training” datasets. For all the models the MSE is smaller when 2005 is the testing dataset (& 2011 is training) than when 2011 is the testing dataset (& 2005 is the training dataset). This means that the 2005 dataset fits the model predicted by the 2011 dataset better than the 2011 dataset fits the model predicted by the 2005 dataset. In other words, the model fitted on the 2011 dataset predicts the 2005 data more accurately than the model from 2005 predicts the 2011 data. This is a somewhat unexpected result, given that the 2005 dataset has more variation in the rainfall/drought variable. However, it must also be remembered that the PSNP safety net was introduced in mid-2005, just after the survey was completed, which may have changed the structural relationship between drought and consumption in ways that are not captured perfectly by the regression model.

In all cases there is actually fairly small difference in the fit of the model, which suggests that the relationship between drought and consumption is actually fairly homogenous, and stable.

Table 5.2: Testing the models with 2005 and 2011 as training and “holdout” data – squared model

	Baseline	Model 1	Model 2	Model 3	Model 4
R^2	0.245	0.246	0.244	0.245	0.248
Adj. R^2	0.244	0.245	0.243	0.243	0.246
AIC	20044.22	20020.87	20068.66	20065.60	20003.63
BIC	20276.69	20276.58	20339.87	20360.06	20329.08
Num. obs.	17134	17134	17134	17134	17134
boot.cv	0.170	0.171	0.170	0.170	0.169
RMSEtest05	0.0297	0.0386	0.0338	0.0344	0.0348
RMSEtest11	0.0412	0.0408	0.0422	0.0424	0.0453

Notes: cv=cross validation. R^2 =r-squared; AIC=Aikike information Criterion; BIC= ; boot=bootstrap; RMSE=root mean squared error.

The results using the squared models appear very similar to those of the linear model. The cross-validation results are identical to three decimal places. Using 2005 as training data and 2011 as testing, the parsimonious model has the lowest RMSE, and for the opposite specification, the baseline model has lowest RMSE. Comparing the overall results for the linear and quadratic models, it appears that the linear model has lower RMSE for all specifications when 2005 is the testing dataset. But when the 2011 data is the testing dataset, the quadratic model performs best. The

lowest RMSE overall for 2005 as testing, is the baseline, linear model. For 2011 testing the best performing model is the parsimonious model.

Regional Results

Table 5.3: Results of excluding regions

	Baseline	Model 1	Model 2	Model3
RMSEtestHRE	0.0514	0.0546	0.0595	0.0602
RMSEtestHDR	0.0302	0.0297	0.0519	0.0544
RMSEtestLOE	0.0042	0.0060	0.0055	0.0055
RMSEtestLOR	0.0602	0.0080	0.0034	0.0013

Notes: cv=cross validation. R2=r-squared; AIC=Aikike information Criterion; BIC= ; boot=bootstrap; RMSE=root mean squared error. HRE=Highlands, Reliable Region; HDR=Highland, drought-prone Region; LOE=Lowlands Enset growing region; LOR=Lowlands reliable region.

For the regional validation we follow the example of Todd and Wolpin (2007), by non-randomly holding out one region at a time to use as a testing dataset. Using the Highlands non-drought as the testing region (row 1), the baseline performs best. Using the highlands drought as the testing region, the parsimonious model (1) is the best performer. Using the Enset lowlands as testing region, the baseline performs best. Using the lowlands-reliable as testing, the full model performs best.

For the validation using an external dataset we attempted to use the 2012 Ethiopian Rural Sociological Survey (ERSS). This survey was also implemented by the Central Statistical Agency (CSA) of Ethiopia. However the consumption (and some other) indicators were deemed to be incomparable for the purpose of the exercise. The survey was conducted over a full year period, revisiting households three times during the year. See Central Statistical Agency and The World Bank (2013) for further details. The main three differences pertinent to this study are first, the much smaller sample size (3500 observations). Second, that the consumption module has slightly different recall periods for food consumption (7 days as opposed to 3 days in HICES) and non-food consumption (a less detailed short-term component, with 3 and 12 month recall). Finally, the data are not nationally representative.

Regarding the validation results overall, the full model with all interactions and a quadratic crop loss function is the preferred specification. We use this model to illustrate below for a simulation of CAT risk modelling outcomes, using the pooled 2005-2011 model as the training data.

6. Application of the derived vulnerability relationships within a probabilistic catastrophe risk modelling framework

The overall purpose of the exercise presented in this paper is to determine whether valid general form vulnerability functions for poverty outcomes can be applied within a probabilistic catastrophe risk modelling framework. Sections 3 to 5 above describe the outcome of the exercise. The purpose of this section is to demonstrate how the derived functions would be applied in practice, and to propose future work on the basis of the findings of the paper.

a. Outcomes for consumption simulation exercise

In the absence of a probabilistic hazard model for rainfall variability in Ethiopia at the resolution required, and compatible with the LEAP protocol, we have produced illustrative examples based on the results from the regression model.

Figure 6.1

Figure 6.1 shows the simulated crop loss with some examples of heterogeneous types (using the linear model for illustration): The baseline (in dark blue) shows crop losses up to 100% and subsequent impact on consumption (e.g. at 80% crop loss, consumption would fall by just under 30%). The most effective mitigation is that of cattle owners whose consumption is approximately half as affected as those without cattle.

For the simulated impacts on consumption overall, each household's characteristics are aggregated, and the net impact of the drought should be calculated (e.g. households with PSNP and cattle will be impacted even less than those with only one of the two, but if they experience another idiosyncratic shock, the impacts would be less mitigated).

Finally, if the policy interest is **poverty impacts of drought**, then this should be incorporated into the module. E.g. a 20% drop in consumption will push households already below the poverty line into deeper poverty, those well above the line may not fall into poverty but those whose consumption is less than 20% above the poverty line will fall below the line.

The components for the vulnerability module that would be needed are as follows:

- A model of geo-referenced exposure, as assets or population at risk (exposure module);

2011 household survey data taken as the exposure dataset. This contained a breakdown of households with relevant vulnerability-determining characteristics (as defined in the regression model) aggregated to the Woreda level.

- A model of the frequency, severity and location of possible hazard occurrence (hazard module);

A full stochastic meteorological model (as described in Box 1.1) for rainfall variability compatible with the LEAP protocol inputs (as described in Box 3.1) was not available at the required resolution for this exercise. However, we illustrate using the "worst year scenario" where each community (woreda) experiences the worst year of crop losses.

- A model of the relationship between the modelled hazard occurrence and the impact on the exposure (vulnerability module).

The results of the regression model can then be used as ‘vulnerability functions’ and applied according to the coefficients in table A2 in the annex, which includes non-stochastic components (e.g. household size, assets, occupation) the quadratic specification for the crop loss, and the full set of interaction terms.

An illustration is shown below – which shows simulated headcount poverty for each level of crop loss as predicted by the model. The full vulnerability module would thus combine these impacts with a more fully developed risk model and the final result would allow policymakers to understand the likely poverty burden in future time periods.

Figure 6.2

b. Conclusion and extension of the exercise

We have explored the possibility of combining a regression-based model of shock impact on consumption with a catastrophe risk (CAT risk) model for the purposes of producing a forward-looking instrument for policy. The relationship between consumption and a drought measure comprised of crop losses based on water adequacy has been calculated. This relationship has then been stress tested using two datasets, and validation techniques.

The results show that the impact of drought is significant across all models examined; with the baseline result showing that for every 10% worsening of the LEAP drought variable, consumption falls on average by 1.5%, and the other models showing typically around a 2% fall in consumption per 10% drought worsening. The results also show an apparent mitigating impact on this relationship from certain community and household characteristics; for example, access to a safety net (PSNP) mitigates the drought impact on consumption by 0.5%. However, whilst the results show consistency, we consider there still to be some caveats around the modeling of this relationship, given that the two years of data available do not show the worst rainfall experienced in Ethiopia, and therefore the relationship may not be valid, e.g. because coping strategies break down at more extreme levels of drought.

As a potential extension of the exercise presented herein, the authors recommend that a full probabilistic catastrophe risk model be used to replace the example approach applied above for the hazard module. This would entail development and application of a stochastic model of rainfall that could be applied into the LEAP framework to produce values for actual evapotranspiration in the index calculation. The model would need to be compatible with the time and geographic resolution of the derivation of the vulnerability/damage functions (see Appendix) in order for these to be applied. Sensitivity analyses could be applied within the hazard modelling to consider potential outcomes in the longer term under climate change scenarios. For example, increases in the rates of occurrence of extreme rainfall variability could be used to look beyond the near term view. Similarly, projections of population increase and composition change could be applied to the exposure dataset to demonstrate different future outcomes.

Data paucity is a restricting factor in the methodology and analysis applied, and also in the interpretation of results. It would also be interesting to see how the drought-consumption relationship, and the attenuation of this by household and community characteristics, differs for other geographies with sufficient data to support this type of analysis.

7. References

- Abraha, M, (2013). Assessment of Drought Early Warning in Ethiopia, a comparison of WRSI by surface energy balance and water balance. PhD Thesis. University of Twente, The Netherlands.
- Anttila-Hughes and Hsiang (2011) Destruction, disinvestment, and death: economic and human losses following environmental disaster. Working Paper.
- Anttila -Hughes and Sharma (2014) Linking Risk Models to Microeconomic Indicators, World Bank publications, 2014.
- Bohra-Mishra, Oppenheimer and Hsiang (2014), Nonlinear permanent migration response to climatic variations but minimal response to disasters, PNAS, 2014.
- Cavallo and Noy (2011) Natural Disasters and the Economy — A Survey, International Review of Environmental and Resource Economics, 2011.
- Central Statistical Agency and The World Bank (2013) Living Standards Measurement Study Integrated Surveys on Agriculture (LSMS-ISA): Ethiopia Rural Socioeconomic Survey (ERSS) Basic information Document. Addis Ababa, September 2013.
- Chetty, R (2008). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. No. w14399. National Bureau of Economic Research,
- Colmer, J (2013) Climate Variability, Child Labour and Schooling: Evidence on the Intensive and Extensive Margin. Working paper.
- D'Ayala D., Meslem A., Vamvastikos D., Porter, K., Rossetto T., Crowley H., Silva V. (2014) Guidelines for Analytical Vulnerability Assessment of low-mid-rise Buildings – Methodology, Vulnerability Global Component project.
- Dercon, S. "Risk, crop choice, and savings: Evidence from Tanzania." Economic development and cultural change 44.3 (1996): 485-513.
- Diamond, P, et al (2015) Estimating Poverty Rates in Target Populations: An Assessment of the Simple Poverty Scorecard and Alternative Approaches. World Bank Group.
- Doorenbos, J. & Kassam, A, (1979). Yield response to water. FAO Irrigation and Drainage Paper No. 33. Rome, FAO
- FEMA (2002). Rapid visual screening of buildings for potential seismic hazards. FEMA, 2002. Retrieved from: http://www.fema.gov/media-library-data/20130726-1646-20490-8071/fema_154.pdf.
- Fermont, A, and Benson, T, (2011). Estimating Yield of Food Crops Grown by Smallholder Farmers: A Review in the Uganda Context. International Food Policy Research Institute, IFPRI Discussion Paper 01097, June 2011
- Hill, Ruth and C. Porter (2014). The Changing Nature of Vulnerability In Ethiopia. In Ethiopia Poverty Assessment 2014. World Bank Group. 2015. Washington, DC.
<https://openknowledge.worldbank.org/handle/10986/21323>

Hsiang and Jina (2014) The Causal Effect of Environmental Catastrophe on Long-Run Economic Growth: Evidence From 6,700 Cyclones, NBER Working Paper.

Jain, V (2010). Anatomy of a Damage Function, AIR Worldwide online publications. Retrieved from: www.air-worldwide.com, February 2015.

James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani. An introduction to statistical learning. New York: Springer, 2013.

Leong, Airriess, Li, Chen, and Keith (2007) Resilient History and the Rebuilding of a Community: The Vietnamese American Community in New Orleans East, Journal of American History, 2007.

Muir-Wood (2014) The use of probabilistic catastrophe risk models to assess the impact of disasters on poverty, World Bank publications, 2014.

Porter (2012) Shocks, Consumption and Income Diversification in Rural Ethiopia, the Journal of Development Studies, 2012.

RMSI (2009). Economic Vulnerability and Risk Assessment for Malawi. World Bank publications, 2009.

Rossetto T., Ioannou I., Grant D.N., Maqsood T. (2014) Guidelines for empirical vulnerability assessment, GEM Technical Report 2014-X, GEM Foundation, Pavia

Smith M, and Steduto, P, Yield response to water: the original FAO water production function Food and Agricultural Organisation Irrigation and Drainage Paper No. 66. Rome, FAO

Souch, C, (RMS) (2010). Post Catastrophe Loss Amplification. International Underwriting Association of London (IUA), 2010.

Todd, P E., and K. I. Wolpin. "The production of cognitive achievement in children: Home, school, and racial test score gaps." *Journal of Human capital* 1.1 (2007): 91-136.

Verma, V, Marchant, T, and Scott, C, (1988). Evaluation of Crop-Cut Methods and Farmer Reports for Estimating Crop Production: Results of a Methodological Study in Five African Countries. London, Longacre Agricultural Development Centre Ltd

World Bank (2013). Pacific Catastrophe Risk Assessment and Financing Initiative – Catastrophe Risk Assessment Methodology, World Bank publications, 2013.

World Bank (2014) Ethiopia Poverty Assessment, World Bank Working Paper.

World Food Programme (2008). FSCB/WFP Workshop Livelihood Early Assessment and Protection (LEAP): its potential application, benefits and limits, NOTE FOR THE RECORD. 21st January 2008.

APPENDIX

Table A1: Simple linear models:

	Baseline	Model 1	Model 2	Model 3	Model 4
(Intercept)	7.794*** (0.018)	7.800*** (0.018)	7.769*** (0.019)	7.775*** (0.019)	7.765*** (0.019)
Drought: proportion lost crops	-0.015*** (0.002)	-0.020*** (0.003)	-0.020** (0.006)	-0.025*** (0.007)	-0.048*** (0.007)
boot.se	0.0025	0.0030	0.0065	0.0137	0.0162
boot.ci	(-0.020, -0.0104)	(-0.0256, -0.0139)	(-0.0322, -0.0069)	(0.0260, 0.0798)	(-0.0181, 0.0455)
Dr_hschool		0.004 (0.005)	0.004 (0.005)	0.003 (0.005)	0.001 (0.005)
Dr_femalehead		0.002 (0.005)			
Dr_psnpb		0.053*** (0.010)	0.053*** (0.010)	0.050*** (0.010)	0.053*** (0.010)
Dr_newcattle			-0.007* (0.003)	-0.007* (0.003)	-0.005 (0.003)
Dr_notag			0.007 (0.006)	0.006 (0.006)	0.006 (0.006)
Dr_disttown07			-0.000 (0.000)	0.000 (0.000)	0.000*** (0.000)
Dr_dependency			0.006 (0.009)	0.007 (0.009)	0.008 (0.009)
Dr_finaccess				0.014** (0.005)	0.012* (0.005)
Dr_illness				0.006 (0.006)	0.006 (0.006)
Dr_cropldam				-0.015* (0.007)	-0.015* (0.007)
Dr_livshock				-0.003 (0.007)	-0.001 (0.007)
Dr_highlands_reliable					0.047*** (0.006)
Dr_lowlands_reliable					-0.029** (0.011)
Dr_lowlands_enstet					0.033** (0.010)
R ²	0.245	0.246	0.247	0.247	0.251
Adj. R ²	0.244	0.245	0.245	0.246	0.249
AIC	20047.21	20022.04	20072.26	20065.60	20001.73
BIC	20271.92	20270.00	20335.72	20360.06	20319.43
Num. obs.	17134	17134	17134	17134	17134
boot.cv.cropleap	0.249	0.248	0.248	0.248	0.248
boot.cv	0.170	0.171	0.170	0.170	0.169
RMSEtest05	0.0232	0.0281	0.0338	0.0272	0.0291
RMSEtest11	0.0438	0.0431	0.0427	0.0429	0.0490
RMSEtestHRE	0.0514	0.0546	0.0595	0.0602	0.0602
RMSEtestHDR	0.0302	0.0297	0.0519	0.0544	0.0544
RMSEtestLOE	0.0042	0.0060	0.0055	0.0055	0.0055
RMSEtestLOR	0.0602	0.0080	0.0034	0.0013	0.0013
RMSEtest12	0.6444	0.6413	0.6414	0.6420	0.6241
05RMSEtest12	0.6318	0.6324	0.6346	0.6346	0.6088
11RMSEtest12	0.6493	0.6494	0.6520	0.6475	0.6356

*** p < 0.001, ** p < 0.01, * p < 0.05

Table A2: Quadratic regression models

	Baseline	Model 1	Model 2	Model 3	Model 4
(Intercept)	7.787*** (0.018)	7.794*** (0.018)	7.767*** (0.019)	7.774*** (0.019)	7.765*** (0.019)
Drought: proportion lost crops	-0.005 (0.005)	-0.011* (0.006)	-0.012 (0.007)	-0.018* (0.007)	-0.050*** (0.009)
boot.se	0.0052	0.0057	0.0131	0.0077	0.0095
boot.ci	(-0.0156, 0.0051)	(-0.0225, 0.0001)	(0.0345, 0.0858)	(-0.0339, -0.0038)	(-0.0690, -0.0314)
cropleap ²	-0.0182* (0.081)	-0.0146 (0.082)	-0.0246* (0.104)	-0.0195 (0.105)	0.040 (0.123)
boot.se	0.0094	0.0118	0.3703	0.0095	0.0090
boot.ci	(-0.0327, 0.0042)	(-0.0420, 0.0044)	(1.816, 3.269)	(-0.0355, 0.0016)	(-0.0349, 0.0006)
Dr hschool		0.003 (0.005)	0.003 (0.005)	0.002 (0.005)	0.001 (0.005)
Dr femalehead		0.003 (0.005)			
Dr psnpb		0.052*** (0.010)	0.052*** (0.010)	0.049*** (0.010)	0.053*** (0.010)
Dr newcattle			-0.007* (0.003)	-0.007* (0.003)	-0.005 (0.003)
Dr notag			0.007 (0.006)	0.006 (0.006)	0.006 (0.006)
Dr disttown07			0.000 (0.000)	0.000 (0.000)	0.000** (0.000)
Dr dependency			0.006 (0.009)	0.007 (0.009)	0.008 (0.009)
Dr finaccess				0.013** (0.005)	0.012* (0.005)
Dr illness				0.005 (0.006)	0.006 (0.006)
Dr cropldam				-0.014* (0.007)	-0.015* (0.007)
Dr livshock				-0.003 (0.007)	-0.001 (0.007)
Dr highlands reliable					0.048*** (0.007)
Dr lowlands reliable					-0.028* (0.012)
Dr lowlands onset					0.034** (0.011)
R ²	0.245	0.246	0.244	0.245	0.248
Adj. R ²	0.244	0.245	0.243	0.243	0.246
AIC	20044.22	20020.87	20068.66	20065.60	20003.63
BIC	20276.69	20276.58	20339.87	20360.06	20329.08
Num. obs.	17134	17134	17134	17134	17134
boot.cv.cropleap	0.249	0.248	0.248	0.248	0.248
boot.cv	0.170	0.171	0.170	0.170	0.169
RMSEtest05	0.0297	0.0386	0.0338	0.0344	0.0348
RMSEtest11	0.0412	0.0408	0.0422	0.0424	0.0453
RMSEtestHRE	0.0504	0.0553	0.0583	0.0598	0.0598
RMSEtestHDR	0.0884	0.0880	0.1036	0.1067	0.1067
RMSEtestLOE	0.0067	0.0081	0.0057	0.0057	0.0057
RMSEtestLOR	0.0126	0.0120	0.0144	0.0167	0.0167
RMSEtest12	0.6451	0.6419	0.6426	0.6429	0.6222
05RMSEtest12	0.6326	0.6331	0.6345	0.6345	0.6017
11RMSEtest12	0.6539	0.6560	0.6520	0.6525	0.6419

*** p < 0.001, ** p < 0.01, * p < 0.05

Table A3: Cubic regression models

	Baseline	Specparsl	Spec2	Spec4	Spec3
(Intercept)	7.772*** (0.018)	7.777*** (0.018)	7.746*** (0.019)	7.752*** (0.019)	7.751*** (0.019)
Drought shock	0.052*** (0.011)	0.056*** (0.011)	0.060*** (0.012)	0.053*** (0.013)	0.013 (0.016)
boot.se	0.0113	0.0111	0.0130	0.0136	0.0167
boot.ci	(0.0299, 0.0742)	(0.0341, 0.0777)	(0.0351, 0.0860)	(0.0265, 0.0798)	(-0.0202, 0.0451)
cropleap ²	-0.025*** (0.004)	-0.028*** (0.004)	-0.030*** (0.004)	-0.030*** (0.004)	-0.021*** (0.005)
boot.se	0.0041	0.0040	0.0042	0.0043	0.0048
boot.ci	(-0.0326, -0.0167)	(-0.0363, -0.0205)	(-0.0386, -0.0221)	(-0.0377, -0.0211)	(-0.0300, -0.0113)
cropleap ³	0.002*** (0.000)	0.002*** (0.000)	0.003*** (0.000)	0.003*** (0.000)	0.002*** (0.000)
boot.se	0.0004	0.004	0.0004	0.004	0.0004
boot.ci	(0.0013, 0.0027)	(0.0017, 0.0030)	(0.0018, 0.0033)	(0.0018, 0.0033)	(0.0010, 0.0026)
Dr hschool		0.002 (0.005)	0.002 (0.005)	0.001 (0.005)	0.001 (0.005)
Dr femalehead		0.003 (0.005)			
Dr psnpb		0.063*** (0.010)	0.063*** (0.010)	0.060*** (0.010)	0.060*** (0.010)
Dr newcattle			-0.007* (0.003)	-0.007 (0.003)	-0.006 (0.003)
Dr notag			0.009 (0.006)	0.008 (0.006)	0.007 (0.006)
Dr disttown07			-0.000 (0.000)	-0.000 (0.000)	0.000 (0.000)
Dr dependency			0.009 (0.009)	0.010 (0.009)	0.009 (0.009)
Dr finaccess				0.012* (0.005)	0.012* (0.005)
Dr illness				0.004 (0.006)	0.005 (0.006)
Dr cropdam				-0.015* (0.007)	-0.016* (0.007)
Dr livshock				-0.002 (0.007)	-0.001 (0.007)
Dr highlands reliable					0.036*** (0.007)
Dr lowlands reliable					-0.031** (0.012)
Dr lowlands onset					0.017 (0.011)
R ²	0.247	0.249	0.247	0.247	0.249
Adj. R ²	0.245	0.247	0.245	0.246	0.247
AIC	20011.03	19975.44	20020.43	20016.47	19982.96
BIC	20251.24	20238.90	20299.39	20326.42	20316.16
Num. obs.	17134	17134	17134	17134	17134
boot.cv.cropleap	0.249	0.248	0.248	0.248	0.248
boot.cv	0.170	0.170	0.170	0.169	0.169
RMSEtest05	0.0081	0.00004	0.0171	0.0159	0.0097
RMSEtest11	0.0425	0.0421	0.0421	0.0422	0.0445
RMSEtestHRE	0.0442	0.0477	0.0484	0.0501	0.0501
RMSEtestHDR	0.0026	0.0029	0.0361	0.0358	0.0358
RMSEtestLOE	0.0017	0.0030	0.0021	0.0020	0.0020
RMSEtestLOR	0.0022	0.00001	0.0120	0.0144	0.0144
RMSEtest12	0.6429	0.6430	0.6423	0.6426	0.6273
05RMSEtest12	0.6329	0.6334	0.6335	0.6336	0.6062
11RMSEtest12	0.6369	0.6386	0.6272	0.6280	0.6225

*** p < 0.001, ** p < 0.01, * p < 0.05

Table A4: Full regression results, linear model

	(1)	(2)	(3)	(4)
	Ln consumption per adult (1996 prices) b/se	Ln consumption per adult (1996 prices) b/se	Ln consumption per adult (1996 prices) b/se	Ln consumption per adult (1996 prices) b/se
Drought: proportion lost crops	-0.002*** (0.00)	-0.003*** (0.00)	-0.003*** (0.00)	-0.006*** (0.00)
Female headed hh	-0.019 (0.01)	-0.016 (0.01)	-0.016 (0.01)	-0.016 (0.01)
Age of household head	-0.001*** (0.00)	-0.001*** (0.00)	-0.001*** (0.00)	-0.001*** (0.00)
HH head not in agriculture	0.058*** (0.01)	0.054*** (0.01)	0.055*** (0.01)	0.058*** (0.01)
HH head completed primary or secondary	0.098*** (0.01)	0.098*** (0.01)	0.101*** (0.01)	0.101*** (0.01)
HH owns more than 3 plots	0.013 (0.01)	0.013 (0.01)	0.013 (0.01)	0.013 (0.01)
Household owns cattle	-0.056*** (0.01)	-0.076*** (0.01)	-0.076*** (0.01)	-0.081*** (0.01)
Household owns sheep or goats	0.017* (0.01)	0.017* (0.01)	0.018* (0.01)	0.016* (0.01)
Household owns chickens	0.003 (0.01)	0.002 (0.01)	0.003 (0.01)	0.003 (0.01)
Access to financial coping (gift/loan/bank)	0.062*** (0.01)	0.062*** (0.01)	0.045*** (0.01)	0.045*** (0.01)
Distance to town	-0.000*** (0.00)	-0.000*** (0.00)	-0.000*** (0.00)	-0.000*** (0.00)
HH dependency ratio	-0.104*** (0.02)	-0.118*** (0.02)	-0.119*** (0.02)	-0.119*** (0.02)
Household suffered death of member	-0.015 (0.02)	-0.015 (0.02)	-0.015 (0.02)	-0.015 (0.02)
Household suffered illness of member	0.017 (0.01)	0.017 (0.01)	0.010 (0.01)	0.009 (0.01)
Household suffered from crop damage	-0.091*** (0.01)	-0.091*** (0.01)	-0.065*** (0.02)	-0.068*** (0.02)
Household suffered from livestock shock	-0.002 (0.01)	-0.001 (0.01)	0.000 (0.02)	0.002 (0.02)
Household suffered job loss of member	-0.038 (0.05)	-0.039 (0.05)	-0.040 (0.05)	-0.031 (0.05)
Household suffered from price shock	0.014 (0.01)	0.014 (0.01)	0.013 (0.01)	0.016 (0.01)
HH received income from PSNP 2010	-0.178*** (0.02)	-0.177*** (0.02)	-0.174*** (0.02)	-0.184*** (0.02)
Highlands_drought	0.068*** (0.01)	0.067*** (0.01)		
Lowlands_enstet	0.035*** (0.01)	0.034** (0.01)	-0.033** (0.01)	-0.088*** (0.01)
Lowlands_reliable	0.012 (0.02)	0.013 (0.02)	-0.051** (0.02)	-0.032 (0.02)
Year 2011	-0.037*** (0.01)	-0.037*** (0.01)	-0.038*** (0.01)	-0.042*** (0.01)
Household size	-0.099*** (0.00)	-0.099*** (0.00)	-0.099*** (0.00)	-0.099*** (0.00)
household uses flush toilet/pit latrine	0.029*** (0.01)	0.029*** (0.01)	0.030*** (0.01)	0.036*** (0.01)
HH has electricity	0.133*** (0.02)	0.134*** (0.02)	0.131*** (0.02)	0.122*** (0.02)
HH has corrugated iron roof	0.146*** (0.01)	0.145*** (0.01)	0.145*** (0.01)	0.144*** (0.01)
femalehead* cropleap	0.000 (0.00)			
hschool* cropleap	0.000 (0.00)	0.000 (0.00)	0.000 (0.00)	0.000 (0.00)
psnpb* cropleap	0.005*** (0.00)	0.005*** (0.00)	0.005*** (0.00)	0.005*** (0.00)
cattle* cropleap		0.001** (0.00)	0.001** (0.00)	0.002*** (0.00)
notag* cropleap		0.000 (0.00)	0.000 (0.00)	-0.000 (0.00)
disttown07* cropleap		-0.000 (0.00)	0.000 (0.00)	0.000*** (0.00)
dependency* cropleap		0.001 (0.00)	0.001 (0.00)	0.001 (0.00)
highlands_reliable			-0.067*** (0.01)	-0.131*** (0.01)
finaccess* cropleap			0.001** (0.00)	0.001** (0.00)
illness* cropleap			0.000	0.000

croptam* cropleap			(0.00)	(0.00)
			-0.001*	-0.001*
			(0.00)	(0.00)
livshock* cropleap			-0.000	-0.000
			(0.00)	(0.00)
highlands_reliable* cropleap				0.005***
				(0.00)
lowlands_enet* cropleap				0.004***
				(0.00)
lowlands_reliable* cropleap				-0.004***
				(0.00)
Constant	7.800***	7.815***	7.888***	7.946***
	(0.02)	(0.02)	(0.02)	(0.02)
r2	0.246	0.247	0.247	0.251

Table A5: Full regression results, quadratic model

	(1) Ln consumption per adult (1996 prices) b/se	(2) Ln consumption per adult (1996 prices) b/se	(3) Ln consumption per adult (1996 prices) b/se	(4) Ln consumption per adult (1996 prices) b/se
Drought: proportion lost crops	-0.001* (0.00)	-0.002** (0.00)	-0.003*** (0.00)	-0.006*** (0.00)
Squared LEAP crop-loss	-0.000 (0.00)	-0.000* (0.00)	-0.000 (0.00)	0.000 (0.00)
Female headed hh	-0.020 (0.01)	-0.016 (0.01)	-0.016 (0.01)	-0.016 (0.01)
Age of household head	-0.001*** (0.00)	-0.001*** (0.00)	-0.001*** (0.00)	-0.001*** (0.00)
HH head not in agriculture	0.058*** (0.01)	0.053*** (0.01)	0.055*** (0.01)	0.059*** (0.01)
HH head completed primary or secondary	0.100*** (0.01)	0.101*** (0.01)	0.102*** (0.01)	0.101*** (0.01)
HH owns more than 3 plots	0.012 (0.01)	0.013 (0.01)	0.013 (0.01)	0.013 (0.01)
Household owns cattle	-0.056*** (0.01)	-0.075*** (0.01)	-0.075*** (0.01)	-0.082*** (0.01)
Household owns sheep or goats	0.017* (0.01)	0.017* (0.01)	0.017* (0.01)	0.016* (0.01)
Household owns chickens	0.003 (0.01)	0.003 (0.01)	0.003 (0.01)	0.003 (0.01)
Access to financial coping (gift/loan/bank)	0.062*** (0.01)	0.062*** (0.01)	0.046*** (0.01)	0.045*** (0.01)
Distance to town	-0.000*** (0.00)	-0.000*** (0.00)	-0.000*** (0.00)	-0.000*** (0.00)
HH dependency ratio	-0.104*** (0.02)	-0.119*** (0.02)	-0.120*** (0.02)	-0.119*** (0.02)
Household suffered death of member	-0.015 (0.02)	-0.015 (0.02)	-0.015 (0.02)	-0.015 (0.02)
Household suffered illness of member	0.017 (0.01)	0.017 (0.01)	0.010 (0.01)	0.009 (0.01)
Household suffered from crop damage	-0.090*** (0.01)	-0.092*** (0.01)	-0.067*** (0.02)	-0.067*** (0.02)
Household suffered from livestock shock	-0.002 (0.01)	-0.002 (0.01)	0.000 (0.02)	0.002 (0.02)
Household suffered job loss of member	-0.037 (0.05)	-0.038 (0.05)	-0.039 (0.05)	-0.031 (0.05)
Household suffered from price shock	0.014 (0.01)	0.013 (0.01)	0.012 (0.01)	0.016 (0.01)
HH received income from PSNP 2010	-0.178*** (0.02)	-0.178*** (0.02)	-0.175*** (0.02)	-0.184*** (0.02)
highlands_drought	0.064*** (0.01)	0.065*** (0.01)		
lowlands_enet	0.034** (0.01)	0.033** (0.01)	-0.032** (0.01)	-0.091*** (0.02)
lowlands_reliable	0.008 (0.02)	0.017 (0.02)	-0.047** (0.02)	-0.036 (0.02)
Year 2011	-0.038*** (0.01)	-0.038*** (0.01)	-0.038*** (0.01)	-0.042*** (0.01)
Household size	-0.099*** (0.00)	-0.099*** (0.00)	-0.099*** (0.00)	-0.099*** (0.00)
household uses flush toilet/pit latrine	0.030*** (0.01)	0.030*** (0.01)	0.030*** (0.01)	0.036*** (0.01)
HH has electricity	0.133*** (0.02)	0.134*** (0.02)	0.131*** (0.02)	0.122*** (0.02)
HH has corrugated iron roof	0.146*** (0.01)	0.145*** (0.01)	0.145*** (0.01)	0.144*** (0.01)
femalehead* croleap	0.000 (0.00)			
hschool* croleap	0.000 (0.00)	0.000 (0.00)	0.000 (0.00)	0.000 (0.00)
psnpb* croleap	0.005*** (0.00)	0.005*** (0.00)	0.005*** (0.00)	0.005*** (0.00)
cattle* croleap		0.001** (0.00)	0.001** (0.00)	0.002*** (0.00)
notag* croleap		0.000 (0.00)	0.000 (0.00)	-0.000 (0.00)
disttown07* croleap		0.000 (0.00)	0.000 (0.00)	0.000** (0.00)
dependency* croleap		0.001 (0.00)	0.001 (0.00)	0.001 (0.00)
highlands_reliable			-0.066*** (0.01)	-0.134*** (0.01)
finaccess* croleap			0.001** (0.00)	0.001** (0.00)

illness* cropleap			0.000 (0.00)	0.000 (0.00)
cropleap* cropleap			-0.001 (0.00)	-0.001* (0.00)
livshock* cropleap			-0.000 (0.00)	-0.000 (0.00)
highlands_reliable* cropleap				0.005*** (0.00)
lowlands_enstet* cropleap				0.004*** (0.00)
lowlands_reliable* cropleap				-0.003** (0.00)
Constant	7.794*** (0.02)	7.813*** (0.02)	7.885*** (0.02)	7.949*** (0.02)
r2	0.246	0.247	0.248	0.251

Data source for calculation of Meher/Belg season yields

CSA Ethiopia: Agricultural Sample Survey 2004/05-2010/11 Private Peasant holdings, Meher and Belg seasons

Introduction to probabilistic catastrophe risk modelling frameworks

Probabilistic catastrophe risk models combine a view of possible hazard occurrence and associated probabilities with a view of assets/population/households exposed to the hazard occurrence ('exposure'), and a view of the damageability of the 'exposure' given different levels of hazard. A probabilistic curve of impact versus probability or return period¹⁶ is the output of the process (see figure A1.1).

Figure A1.1) Modular probabilistic catastrophe risk modelling framework

Source: authors, using images listed at the bottom of this section.

¹⁶ A return period is the average recurrence period at which a given level of impact or higher is expected.

Although informed by data on past disaster occurrence, the modelling framework allows for a forward-looking view of risk, considering the potential occurrence of hazard beyond that presented in the historical record (although bounded by the potential of the physical system to produce catastrophe events). Thus catastrophe risk models typically contain tens of thousands of events for physical systems where the numbers of events recorded are orders of magnitude fewer (see figure A1.2). This approach is fundamental to catastrophe risk modelling, where the low recurrence period of events means that even a comprehensive historical record will not capture all the future possibilities of event occurrence, particularly for extreme events.

Figure A1.2a) Historical storm tracks (Pacific basin) Figure A1.2b) Selected stochastic Cat 5 storm tracks (Pacific basin)

Source: AIR Worldwide from (World Bank, 2013)

An example of how this modelling framework works in practice is described below:

Modelling process for a building exposed to earthquake

In a probabilistic earthquake model, possible financial loss to an insurance company underwriting a policy for a particular building would be determined as follows:

- 1) The hazard component would supply a view of spectral acceleration of the building arising from each of some tens of thousands of different modelled stochastic events taking into account factors such as the size of the earthquake, its location relative to the building, and local conditions such as soil type/potential for liquefaction that are reflected in the local resulting hazard value;
- 2) The resulting damage experienced by the building would next be determined based on the level of hazard (spectral acceleration) and the physical building characteristics that determine vulnerability, such as building material, number of stories and year of build;
- 3) The level of damage of the building is converted into a total financial loss based on factors such as replacement value, and the policy conditions such as deductible and limit are applied to give the insurance company's perspective.

Sources: as cited, with images from: UNEP GRID (population density for exposure, historical cyclone tracks for stochastic events); USGS (Shakemap of Tohoku earthquake for local intensity); CoreLogic EQECAT (vulnerability curves for damage estimation); PCRAFI (probabilistic impact curve for earthquake and tropical cyclone).