Doctoral School of Economics – DSE

Sapienza University of Rome

PHD DISSERTATION

PHD IN ECONOMICS AND FINANCE

Curriculum: Development Economics

30th Edition

ESSAYS ON WEATHER, CLIMATE AND DEVELOPMENT

Candidate

Marco Letta

Supervised by

Prof. Pierluigi Montalbano

Prof. Richard S.J. Tol

Coordinated by Prof. Rita Laura D'Ecclesia

Academic Year 2017 - 2018

Acknowledgments

This work would not have been possible without the advice and support of many people.

First and foremost, I want to expressly thank my supervisors, Professor Pierluigi Montalbano and Professor Richard S.J. Tol, from whose teachings I will benefit for my entire career.
Pierluigi is the reason why I started a PhD in the first place. It was his excellent supervision during my Bachelor's and Master's dissertations, together with his intense passion for research, to spark my interest in economics. His precious support, encouragement and tireless guidance have been, and still are, invaluable.
I thank Richard because he believed in me without even knowing me, and gave me a chance to work with him as a Visiting PhD Student at the University of Sussex. He taught me to distinguish objectivity from hypocrisy and to ground my opinions in evidence. It is and will always be a pleasure and a privilege to work with him.

I am grateful to Professor D'Ecclesia, to the boards of the PhD in Economics and Finance and of La Sapienza PhD School of Economics, as well as to the Department of Economics of the University of Sussex.

I am also grateful to L.Alan Winters, Melissa Dell, Kalle Hirvonen, Veronica Notaro, the FAO RAP and RIGA Teams, the World Bank LSMS Team and seminar participants at Sapienza University of Rome and University of Sussex.

Very special thanks go to a host of relatives and friends. There is no need for what would be a very long list: each one of you already knows my affection and gratitude. Here I just want to say that all of you have been extremely important to me, in different ways and places and at different times. See you down the road.

Finally, this thesis is dedicated to the four fundamental pillars of my life: Anna, Bruno, Paolo and Veronica. Your love is the source of my happiness, and I owe you everything I am and will ever be.

Summary*

This thesis examines the relationship between weather shocks, climate change and development. The research constituting this work thus lies at the intersection between environmental and development economics.

Although it is by now generally acknowledged that it is the poor who will suffer most from the negative impacts of climate change, the literature on the relationship between climate change and development is still far from reaching definitive conclusions.

In particular, from a macro point of view, the debate on the potential impacts of global warming on the growth rate of the economy, and not just on the level of output, is still unsettled. At the same time, research on the links between climate change and micro welfare dynamics, especially in developing countries, is scarce.

Drawing from both the new climate-economy literature and the development literature on household growth and poverty traps, we address, at multiple scales of analysis, the research question of whether there is a causal relationship between weather shocks and economic growth.

We specifically focus our attention on temperature shocks, because the ultimate aim is to provide policy guidelines on the future impacts of climate change. Although we cannot capture in our analysis long-run phenomena such as intensification effects and adaptation, we employ robust empirical methods and panel data so to infer from short-run elasticities and inform thinking about the impacts of long-run, permanent changes in climate.

The main contribution of this thesis lies in its integrated, multi-level framework, able to provide empirical evidence and insights on the crucial links, still not well understood, between weather shocks, climate change, poverty and economic growth. More broadly, this work contributes to a deeper understanding of the much-debated historical relationship between climate and development.

This thesis consists of three self-contained essays.

Essay 1 investigates the relationship between weather shocks and the growth rate of total factor productivity (TFP), the key driver of long-run development, at the macro level. We look at TFP growth to shed light on whether climate change will affect the growth rate of the economy, as recently hypothesized in influential theoretical studies. If true, this would imply a radical upward revision of

* I hereby declare that this thesis has not been and will not be submitted in whole or in part to another university for the award of any other degree. Material included in Essays 1 and 2 has been incorporated in two working papers co-authored with my advisors. However, I hereby state that the bulk of the original research presented in this thesis, including all the empirical applications in each of the following essays, is my own work.

current impact estimates. Using a sample of 60 countries covering the time span 1960-2006, we show that temperature shocks affected TFP growth significantly and negatively only in poor countries. This finding provides direct evidence of the dynamic effects of temperature shocks in poor countries. If such causal relationship between temperature, poverty and TFP growth will persist in the long-run, climate change in poor countries will affect not only the level of output, but also its growth rate through the TFP channel, thus further increasing concerns over inequality of future impacts.

Essay 2 is a shift of perspective from the macro to the micro point of view. Specifically, the aim is to understand if the pattern of inequality of impacts observed at the macro level also holds within-country, other than between-country. Here, using LSMS-ISA World Bank household panel data, we explore the short-run elasticities between weather shocks and consumption growth in rural Tanzania during the period 2008-2013. The core results are: i) temperature shocks slowed the convergence process among households and ii) the existence of critical poverty thresholds above which households are immune to temperature shocks. Furthermore, we also provide empirical evidence on the transmission channels (labour productivity and crop yields) responsible for the heterogeneity of impacts.

Essay 3 is partially an extension of this approach, to make up for the short-run nature of the analysis in Essay 2. Research on micro growth dynamics in developing countries, in fact, is hampered by the lack of long household panels. Thanks to the creation of synthetic panels, we obtain a longitudinal panel of cohorts for rural Tanzania from 2000 to 2013. Using an *ad hoc* measure of household resilience to food insecurity developed by the Food and Agriculture Organization (FAO), we provide evidence that, despite the convergence process among households in rural Tanzania also holds in the long-run, temperature shocks still have a diverging effect for the least resilient households. We then show the existence of resilience thresholds that entail a bifurcation of impacts from temperature shocks on food consumption growth, which is particularly meaningful for adaptation policies in developing countries.

On the whole, the core result of this work is a sharp and remarkable pattern of heterogeneity of impacts: temperature shocks only affect poorest households and countries.

This finding, although consistent with most previous studies, goes beyond them by showing how the causal relationship between temperature, poverty and economic growth is deeper, more persistent and more extensive than previously thought, and may even be ubiquitous. This points to the paramount role played by development in dampening the effects of weather shocks on human welfare dynamics. Extrapolating from weather to climate, but also acknowledging the issue of external validity in doing so, our research suggests that climate change will cause, first and foremost, a fractal increase in inequality, both between- and within-country: worse-off countries and households will be disproportionately affected by the negative impacts of global warming.

Such a conclusion points to the importance of poverty reduction as a complementary strategy to greenhouse gas emission reduction and as a fundamental element of climate policy.

Essay 1

Weather, climate and total factor productivity

Abstract

Recently it has been hypothesized that climate change will affect total factor productivity growth. Given the importance of TFP for long-run economic growth, if true this would entail a substantial upward revision of current impact estimates. Using macro TFP data from a recently developed dataset in Penn World Tables, we test this hypothesis by directly examining the nature of the relationship between annual temperature shocks and TFP growth rates in the period 1960-2006. The results show a negative relationship only exists in poor countries, where a 1°C annual increase in temperature decreases TFP growth rates by about 1.1-1.5 percentage points, compared to an impact indistinguishable from zero in rich countries. Extrapolating from weather to climate, the possibility of dynamic effects of climate change in poor countries increases concerns over the distributional issues of future impacts and, from a policy perspective, restates the case for complementarity between climate policy and poverty reduction.

Introduction

Since the path-breaking work of Nordhaus (1991), economists have argued in favour of a modest carbon tax. Although frequently challenged in favour of more stringent climate policy, estimates of the social cost of carbon have not increased over the years (Tol, 2015). Three independent author teams (Moore & Diaz, 2015; Dietz & Stern 2015; Moyer, Woolley, Matteson, Glotter & Weisbach, 2014) have recently hypothesized that, should climate change negatively affect total factor productivity, then the estimate of the Pigou tax increases drastically. In this essay, we present econometric evidence of the impact of weather and climate on total factor productivity growth. While not disputing the sign of the hypothesized effect, we show the effect size is small.

Most impact studies of climate change have taken the form of comparative statics impact estimates. These studies show that climate change would have a modest negative impact of human welfare, i.e., a few percent over a century (Tol, 2015), but they have been criticized because they could not fully capture the potential damage by future climate change (Pindyck, 2012 & 2013; Stern, 2013; Weitzman, 2009 & 2011).

Besides static impacts on welfare, there are also dynamic ones: climate change affects the growth rate of the economy (Fankhauser & Tol, 2005; Hallegatte, 2005). The distinction between static, or "level" effects, and dynamic, or "growth" effects of climate change on economic activity is of first order importance in terms of the magnitude of future impacts. While the so-called *level* effects are temporary and intrinsically reversible, *growth* effects compound over time and permanently reduce output. An impact of hot temperatures on a given year's agricultural yields would represent a *level* effect, while an impact on investments or institutions would affect the economy's ability to *grow*, altering its future path. Fankhauser and Tol (2005) argue that climate change may affect labour supply, capital depreciation and productivity (rather than productivity growth). They find that, if these effects are negative, economic growth would be suppressed. The resulting welfare loss would be similar in size to the estimates of the static welfare losses.

Since the onset of growth economics and the pioneering Solow model (Solow, 1956) TFP has been considered a key element to explain long-run development. TFP, as is widely known, represents a combination of labour and capital productivity, which accounts for increase in total output not due to labour or capital inputs, and traditionally has been seen as a rough measure of technological progress. Recently, a number of theoretical studies have hypothesized a future impact of global warming on

TFP growth (Stern, 2013; Moore & Diaz, 2015; Dietz & Stern 2015; Moyer et al., 2014). Given the preeminent importance of TFP for long-run economic growth, if climate change will really harm TFP growth rates, this would entail a radical revision of impact estimates.

Dietz and Stern (2015) change the workings of DICE[1], one of the most used Integrated Assessment Models (IAMs), to allow climate impacts to affect TFP growth[2]. They find a much stronger case for stringent emission abatement.

Similarly, Moyer et al. (2014) argue that the IAMs used by the US federal Interagency Working Group (IWG)[3] on the Social Cost of Carbon may not capture the full range of consequences of climate change, and contest the fact that "(IAMs) implicitly assume that society will grow far wealthier in the future even if temperatures increase by amounts that many scientists believe may cause substantial hardships". Consequently they change DICE and allow climate impacts to directly affect TFP growth, finding, consistently with Dietz and Stern (2015) large effects on future growth and a much higher value of the Social Cost of Carbon (SCC) than the IWG one[4].

However, these works do not provide any empirical evidence for this claim and the consequent simulations (Tol, 2015). In fact, while these calibrated models are very sensitive to assumptions about the impact of climate change on TFP growth, the assumptions are just that: they are not grounded in observations. The current essay estimates the impact of weather variability and climate change on total factor productivity growth.

There is a large and growing body of empirical literature which focuses on the relationship between climate and economic activity. Jared Diamond (Diamond, 1999) revived the spirit of Ellsworth Huntington (Huntington, 1922), arguing that geography and climate are the fundamental drivers of economic development. Olsson and Hibbs (2005) provide empirical support. Gallup, Sachs, and Mellinger (1999) argue that geography and climate are important, but that their impact can be modified by technology. In sharp contrast to this environmental determinism, (Acemoglu, Johnson, & Robinson, 2000; Easterly & Levine, 2003; Rodrik, Subramanian, & Trebbi, 2004) argue for institutional determinism and find that, in a direct statistical contest, institutional variables have

[1] DICE stands for "Dynamic Integrated Climate-Economy Model", and is a computer-based Integrated Assessment Model developed by Professor William Nordhaus of Yale University. See Nordhaus (2008).

[2] Further changes to the DICE framework they undertake are allowing for convexity of the damage function (Weitzman, 2010) and for high values of the climate sensitivity parameter (Weitzman, 2009 & 2011).

[3] DICE (Nordhaus, 2008), FUND (Anthoff, Tol, & Yohe, 2009) and PAGE (Hope, 2006).

[4] Also, they notice how impacts on growth would contribute to settle the debate on the discount rate sparked after the publication of the Stern Review (Stern, 2007). See also Nordhaus (2007), Stern (2013) and Tol et al. (2006).

predictive power but climate and geography variables do not. The institutional view has been challenged by Alsan (2014) and Andersen, Dalgaard and Selaya (2016). Alsan (2014) shows that the tse-tse fly is a major factor in the underdevelopment in Sub-Saharan Africa. Andersen, Dalgaard and Selaya (2016) show that UV radiation (but not climate) plays a role in explaining the pattern of development across the world.

These cross-section analyses of the climate-income relationship suffer from a range of endogeneity and confounders problems. A literature has emerged that uses robust panel studies that try to isolate the effect of temperature or other meteorological variables on economic activity and growth.[5] A comprehensive review is carried out in Dell, Jones, and Olken (2014).

As far as climate change is concerned, though, this literature is problematic for a number of reasons. First, as emphasized by Tol (2015), weather impacts are assumed to be informative about climate impacts; put differently, short-term elasticities are used to assess long-term effects. Second, since the Industrial Revolution global temperature has risen of almost 1°C (IPCC, 2013) while increases in temperature during the 21st century will very likely be of 2°C or more (IPCC, 2013) which means these studies extrapolate far beyond historical experience. Third, it is by no means guaranteed that historical relationships will continue to hold in the future as technologies and institutions evolve. However, while external validity is debatable, there are techniques, as for example long differences, that can alleviate these concerns. Thus, these *caveats* notwithstanding, recently panel methods have been employed to disentangle level effects from growth effects.

For example, in a global sample from 1950-2003, Dell, Jones, and Olken (2012) find temperature shocks have significant negative effects on GDP growth of poor countries, but not of rich ones. Interestingly, using weather lags and long differences, they find evidence for persistence of impacts, which suggests temperature shocks are only slowly absorbed by the economy and have long-lasting effects in poor countries, leading them to conclude that temperature also affects the growth rate of GDP in poor countries, other, or rather, than output level. Bansal and Ochoa (2011) do not exploit country-specific temperature shocks, but global average temperature shocks, and find tropical countries are the most vulnerable and that on average a 1°C global increase reduces growth by 0.9%. A study on windstorms by Hsiang and Jina (2014) for 28 Caribbean countries over the 1970-2006 period shows similar results. Burke, Hsiang, and Miguel (2015), studying 166 countries between 1960 and 2010, find that productivity peaks at about 13°C and declines non-linearly thereafter, leading

---

[5] As they explain: "panel data exploit the exogeneity of cross-time weather variation, allowing for causative identification".

them to predict impacts much larger than previously estimated.
These studies focus on the recent past, which saw only limited climate change. This could, on the one hand, lead one to speculate that these impacts could be exacerbated by further increases or non-linear effects which lie outside historical experience and, on the other, that weather impacts must be interpreted with caution given both the difference between a 1°C shock in a given year and place and a permanent 1°C global increase, and the fact that in the long-run adaptation may take place and substantially mitigate negative impacts. It is the controversial but ultimately difficult to solve "intensification vs adaptation" debate over which of these two long-term effects will eventually outweigh the other (Dell, Jones & Olken, 2014).

A first consequence of this new wave of empirical studies on climate and growth has been to induce practitioners to use these new estimates to derive empirically-based projections and implement them in IAMs to see how these respond to the relaxation of assumptions about exogenous economic growth. Moore and Diaz (2015) show that if DICE is modified and calibrated on Dell, Jones and Olken (2012), the predicted impacts go up, and so the consequent SCC, compared to the baseline scenario in which climate change does not affect growth. Lemoine and Kapnick (2015) convert estimates of past economic costs of regional warming into projections of the economic costs of future global warming. They do recognize, though, that this is mostly relevant only for relatively small changes in climate.

Using TFP data from the most recent version of the Penn World Table, we use a panel dataset for 60 countries, covering the period 1960 – 2006, to test the hypothesis of a causal relationship between temperature shocks and annual TFP growth rates. What emerges from our analysis is that temperature shocks affect annual TFP growth rates only in poor countries. Of course, this conclusion is subjected to *caveats* and must be interpreted with caution. Nonetheless, it basically confirms the results of Dell, Jones and Olken (2012) and rejects the conclusions of Burke, Hsiang, and Miguel (2015). We also show that the assumptions of Dietz and Stern (2015), Moore and Diaz (2015) and Moyer et al. (2014) have no empirical grounding.

The contributions of this essay are the following: first, it provides a useful empirical test for the plausibility of the recent hypothesis of an impact of climate change on TFP growth. Second, to our knowledge this is the first study to examine the macro relationship between temperature shocks and TFP growth. Third, unlike other previous works on temperature and economic growth, this analysis can provide direct, and not just indirect, evidence on the persistence of weather impacts on economic activity in the medium or long-run, since it focuses on TFP, and not GDP, growth rate.

Fourth, we show the main reason behind the impact on TFP growth in poor countries is labour productivity, thus linking the existing macro literature with the recent micro studies on the relationship between temperature and human physiology.

The outline of the rest of this essay is as follows. Section 1 provides a theoretical background on the potential TFP-climate change relationship. Section 2 presents data and descriptive statistics. Section 3 describes the identification strategy. Section 4 presents empirical results. Section 5 performs robustness checks. Section 6 discusses the implications of the results with regard to climate change. Section 7 sums up, illustrates some *caveats* and concludes.

Section 1
Background on the TFP impact channel

We follow Dietz and Stern (2015) to show how climate change could affect technological progress. Consider the standard DICE model: a Ramsey-Cass-Koopmans growth model with an added climate externality and emission abatement costs:

$$Y_t = (1 - \Omega_t^Y)(1 - \Lambda_t)[A_t N_t^{1-\alpha} K_t^{\alpha}] \tag{1}$$

where A_t and N_t are specified exogenously, K_t evolves according to the standard equation:

$$K_{t+1} = K_t (1-\delta) + sY_t \tag{2}$$

Λ_t are emission abatement costs and Ω_t^Y is a quadratic damage function of the change in global temperature relative to the global mean in 1900[6]:

$$\Omega_t^Y = 1 - \frac{1}{1+\pi_1 \triangle T_t + \pi_2 \triangle T_t^2} \tag{3}$$

Equation (1) represents the impact function in case of only level effects: in this model, a portion of output in each time period is simply "thrown away" due to the impacts of climate change Ω_t^Y .

[6] The damage function is usually calibrated *ad hoc* on the basis of impact studies of climate change. The quadratic form has been criticized because it does not allow for convexity of damages [*sic*] (Stern, 2013; Weitzman, 2010).

In this framework, climate impacts affect long-run economic growth as climate change reduces current output, and hence savings and investment, which in turn reduce future capital and future output. The savings rate may also be affected, as the returns to investment fall. Both effects have been shown to be quantitatively small (Fankhauser & Tol, 2005; Moyer et al., 2014).

If, instead, climate change also affects TFP, things change substantially. Specifically, TFP is endogenous and grows according to the following law of motion:

$$A_{t+1} = (1 - \Omega_t^{A})(1 - \delta_t^{A}) A_t + \alpha(I_t) \tag{4}$$

where δ_t^{A} is the net depreciation rate for productivity, $\alpha(I_t)$ is a "spillover function" that converts the flow of capital investment in each period into a flow of capital externalities, and Ω_t^{A} are the impacts of climate change on TFP, while the remaining share of damages still affects output level.
Damages are then partitioned between output and TFP:

$$\Omega_t^{A} = f^{A} \cdot \Omega_t \tag{5}$$

$$\Omega_t^{Y} = 1 - \frac{(1-\Omega_t)}{(1-\Omega_t^{A})} \tag{6}$$

where f^{A} is the fraction of impacts of climate change that harms TFP growth.

The effects of this modification depend on the share of impacts directly affecting TFP, but even a small share leads to a radically different consumption growth path: Dietz and Stern (2015) assume that $f^{A} = 0.05$ and find that consumption per capita in year 2205 is reduced from more than 15 times the 2005 level to 11.4 times higher.
Moyer et al. (2014) explore the consequences of different values of f^{A} between 1% and 100%. They show that $f^{A} = 0.05$ leads to a 70% drop in consumption per capita in 2300 relative to the no climate change case.
Similar qualitative results are obtained by Moore and Diaz (2015) when they alter the DICE model to let climate change affect TFP growth on the basis of parameters calibrated on the estimates of Dell, Jones and Olken (2012). As Dietz and Stern (2015) sum up: "in this formulation some part of the instantaneous impacts of climate change falls on TFP, permanently reducing future output possibilities".

Section 2
Data and Descriptive Statistics

A. Data

Data for this study are taken from a range of different sources.

TFP Data

Data on total factor productivity of countries come from the most recent version of the Penn World Table, PWT 8.1 (PWT 8.1, 2016). In particular, in our study we use $RTFP^{NA}$ data[7]. $RTFP^{NA}$, where the prefix R stays for "real", is a country-specific index of TFP where in the benchmark year, 2005, $RTFP^{NA}$ is 1 for all countries. $RTFP^{NA}$ can be used to study within-country productivity growth over time. In our specifications, we use the natural logarithm of the $RTFP^{NA}$ index. This means that the 2005 benchmark value is 0 for all countries in the logarithmic specification. We calculate annual $RTFP^{NA}$ growth rates by first-differencing, and check for stationarity[8]. Henceforth, from now on, "TFP growth rate" it is intended as the annual growth rate of the natural logarithm of the $RTFP^{NA}$ index as taken from PWT 8.1. For further information on the $RTFP^{NA}$ index and data, see Appendix (1).

Temperature and Precipitation Data

These data are taken from the *Terrestrial Air Temperature and Precipitation: 1900 – 2006 Monthly Time Series* (Matsuura & Willmott, 2007), from the University of Delaware (UDEL), as aggregated to the country-year level by (Dell, Jones & Olken, 2012), using population weights, where the weights are constructed from 1990 population data at 30 arc second resolution from the *Global Rural Urban Mapping Project* (Balk et al., 2004). Importantly, given temperature levels are trend-stationary, in order to exclude potentially spurious results and ensure stationary residuals in our regressions, we transform data by first-differencing and check for stationarity. We do the same with precipitation data.

GDP Data

We use per capita GDP data to distinguish between impacts in rich and poor countries. These data come from the Maddison Project ('Maddison Project', 2016.).

[7] Note that this series has only recently become available. Previous studies of the impact of climate change on economic growth, reviewed above, therefore did not have access to these data.
[8] For the panel unit root tests for annual TFP growth, temperature change and precipitation change, see Appendix (2), Table A.1 – A.6.

B. Descriptive Statistics

The main dataset is composed of 60 countries[9] and covers the period 1960 – 2006. Figure 1 is a scatterplot of TFP and temperature levels in 2006, and the linear prediction. As can be seen, there is a negative correlation between the two. This correlation is not a causal relationship, but could be due to confounding factors such as institutions. There is no reverse causality.

Table 1 provides some descriptive statistics for the main variables. There is a huge variation both in the annual growth rates of TFP, with an average of about 5% annual increase but a minimum and a maximum that are respectively -56% and 27%, and in terms of temperature changes as well, where the mean annual change in temperature is very small but the extremes are between 2°C and 3°C. Finally, precipitation exhibits even greater variability.

Section 3
Empirical Strategy

We use a fixed-effect panel as the estimation method to isolate the impact of weather shocks on the growth rate of total factor productivity[10]. Our identification strategy is straightforward and follows Dell, Jones, and Olken (2012). The baseline specification of our model is the following:

$$TFP_{it} = \alpha + \beta \Delta Temp_{it} + \gamma \Delta Pre_{it} + \mu_i + \theta_{rt} + \varepsilon_{it} \tag{7}$$

Where TFP_{it} represents the annual growth rate of TFP, and $\Delta Temp_{it}$ is annual temperature change. ΔPre_{it} represents annual change in precipitation levels, which is used only as a control variable following the recommendation in Auffhammer, Hsiang, Schlenker and Sobel (2013). By excluding precipitation we would run the risk of omitted variable bias. Furthermore, in order to investigate for heterogeneous effects of temperature shocks, we follow Dell, Jones and Olken (2014) and interact the vector of temperature changes with dummies that capture the heterogeneity of interest, in particular dummies for being a "poor" or a "hot" country.

As for the other elements in the equation, μ_i are country fixed effects, θ_{rt} are region x time fixed effects, where this interaction allows for differentiated trends in different regions, as recommended by Dell, Jones and Olken (2014), in order to isolate idiosyncratic local shocks[11]. Finally, ε_{it} are error

[9] The choice of the countries has been made on the basis of data availability. For the list of countries, see Appendix (5).
[10] For the appropriateness of the FE approach compared to a random effects (RE) specification, see Appendix (2), Table A.7.
[11] For the list of regions, see Appendix (6).

terms adjusted for clustering at the country level.

There is no reason to be concerned about reverse causality. Confounding variables are a minor worry. TFP is constructed rather than observed. If weather variations would cause mismeasurement in the size of the labour force or the capital stock, then we would wrongly attribute this to TFP. We are not aware of a way to test this for our data.

TFP is total factor productivity. By construction, when measured at a national, annual resolution, TFP is a mix of a wide range of factors. Changes in TFP can be due to technological change, the standard but flawed interpretation. Changes in TFP can also be due to managerial or behavioural change, changes in the structure of the economy or company entry and exit within sectors, changes in regulation or taxation, changes in the provision of public goods, changes in market power, or changes in international trade. The results below show that temperature variations affect TFP growth, but our data do not allow us to precisely identify the channel through which TFP is affected. That said, our approach is a step forward compared to previous studies which looked at economic growth, an even more convoluted measure.

Section 4
Empirical Results

Table 2 reports the results for the baseline specification of equation (7). Column (1) only includes annual changes in temperature and precipitation levels. A first inspection shows that the coefficient for the annual change in temperatures, ΔTemp, is negative and significant at the 5% level, suggesting that a 1°C annual increase in temperature would lower TFP growth rates of countries by 0.49%. Column (2), however, reveals that adding an interaction between temperature change and a dummy for being poor – with "poor" being defined as having a below median GDP per capita in the initial year of our panel, 1960 – substantially changes the picture: this interaction in fact is negative and strongly significant, while the coefficient for temperature changes is now negative but statistically insignificant, which suggests the negative effects of temperature on TFP growth rates are concentrated in poor countries.

This is confirmed by looking at the net impact of temperature change in poor countries, at the bottom of Column (2), which suggests a 1°C annual increase in temperature in poor countries would decrease TFP growth rates by about 1.5 percentage points, with a significance at the 1% level.

This finding is somewhat weakened when we add an interaction between temperature changes and a dummy for being hot, with “hot” being defined as having an above median average temperature in the 1960s. The results are shown in Column (3): the coefficient of the Poor x ΔTemp interaction is now -1.2 %, and significant at 5%, while the “hot” interaction turns out to be insignificant, and so its net effect. Importantly, the total effect of temperature in poor countries is also diminished both in terms of magnitude and significance[12]. The fact that the negative effect of temperature changes in poor countries is somewhat weakened could be explained in two different ways: the first is that the negative effect of temperature on TFP growth rates comes not only through being poor, but also, partially, through being hot, and the second is that the definitions of “hot” and “poor” overlap to a good extent and thus the inclusion of an “hot” interaction partially offsets the results for poor countries. The distinction matters a great deal when it comes to conclusions with regard to future climate change: it is a completely different picture whether the negative effects of temperature shocks appear only in poor countries or also, even if slightly, in hot countries regardless whether rich or poor.

In order to shed light on the issue, in Column (4) we use an alternative definition of poor, with “poor” being now defined as having a below median GDP per capita, where median GDP per capita is now calculated over the whole 1960 – 2006 period and not just in 1960 as above. The “poor” interaction is again strongly significant, with the coefficient of Poor_2 x ΔTemp again very similar, with a value of -1.43 percentage points, the “hot” interaction again negative but statistically insignificant (and so its net total impact), and the total impact in poor countries again significant at the 1% level. Therefore, this variation suggests that only TFP growth rates of poor countries are affected by temperature shocks.

Finally, to enhance confidence in this finding, in Column (5) and (6) we consider a different definition of “hot” country, with the dummy for hot that has value 1 for countries with an average temperature in the 1960s above the 75% percentile, and repeat our specifications. The results, while confirming the negative impact of temperature shocks on the TFP growth rate of poor countries, also show that there is a negative and 5% significant impact of temperature shocks in hot countries, with a net effect of about -1 percentage point on the annual TFP growth. In other words, even though the negative effect of annual temperature comes through being poor, there also seems to be weak evidence of an

[12] Incidentally, it is also worth remarking how precipitation change has a negative and significant effect, but this control variable has proved to be very sensitive to specifications throughout the entire empirical analysis and its results should therefore be interpreted with caution and are no further discussed here.

impact in hot countries. Given the importance of this distinction, in Section V we investigate more closely the relationship between temperature shocks and TFP growth, by performing a variety of robustness checks.

Section 5
Robustness Checks

Thirteen robustness checks are performed: the repetition of the baseline specification for a different dataset, comprising 68 countries and covering the period 1970–2006; the repetition of the main specification in both datasets using different weather data; the repetition of the regressions without the precipitation variable; regressions including interactions with poor and hot dummies for precipitation as well; a specification including an interaction between temperature shocks and a dummy for being rich; an investigation of the poor subsample of our dataset; a specification using a joint interaction term for countries which are both poor and hot; two alternative specifications which include respectively an interaction with GDP per capita and one with a measure of institutional quality; a repetition of the main specification in which we use labour productivity growth as the dependent variable in place of TFP growth; regressions on changes in the number of persons employed and capital stock; the use of Driscoll-Kraay (1998) standard errors in place of clustered standard errors for the baseline analysis in both samples.

A. Different sample

We run the same regressions using a different sample of the same dataset, changing the composition of countries and the time period. In particular, we add 8 countries to the main sample: Bulgaria, Hungary, Kuwait, Panama, Paraguay, Poland, Qatar and Saudi Arabia. Some of these countries are hot and rich, increasing the statistical power to distinguish between heat and affluence. The new time period is 1970 – 2006. Table 3 provides some descriptive statistics for the new dataset, Table 4 the results for the main specification.

As for the impact in poor countries, the results are very similar: the previous findings are confirmed in terms of magnitude, sign and significance, and if anything reinforced. This is probably due to the fact that some of the added countries, such as the Arab oil states, are very rich, very hot and with high TFP growth (although concentrated in one sector). The robustness check conducted on Sample B reinforces the main thesis of this work: a negative causal relationship between annual TFP growth rates and temperature shocks only exists in poor countries, while the TFP growth rates of rich countries, regardless whether they are hot or cold, do not appear to suffer from temperature changes.

In other words, the impacts of temperature on total factor productivity are conditional on the level of GDP per capita.

B. Different weather data

Since both TFP and weather data are notably affected by measurement errors, to partially alleviate these concerns we perform exactly the same analysis, in both samples, but using another weather dataset, the *CRUCY Version 3.23* by the Climatic Research Unit (CRU) of the University of East Anglia (CRU, 2016) . Furthermore, this dataset uses a different weighting scheme with respect to our main source of weather data: the CRU data are aggregated at the country levels using area weights, rather than population weights as in the first case, which means that the aggregated data now represent the average weather experienced by a place, as opposed to the average weather experienced by a person (Dell, Jones & Olken, 2014). This is not a trivial difference: in countries like United States, Australia, Canada, China, large and scarcely populated areas will dominate the national average temperature when using area weights. This double difference, both of source and aggregation method, takes to weather data that are quite different from those used in our main specification[13], and thus we reckon this constitutes a useful and reliable check for the robustness of our findings[14]. Table 5 replicates the specification of Table 2 for the main dataset using the CRU data.

The results are remarkably consistent with those emerged from the baseline analysis: the negative effect of temperature shocks on TFP growth rates only comes through being poor, not through being hot, and there is no such causal relationship in rich countries. This consistency is further confirmed when repeating the same exercise but using Sample B. The table for this check can be found in Appendix (3), Table A.9: results are similar.

C. Excluding precipitation

Even though we highlighted how the risk of omitted variable bias is a concrete one when dealing with weather variables, one could also make the case against including precipitation in our specification. Therefore, in Table 6 we show what happens if we exclude precipitation from the regressions: the coefficients for temperature, and their significance, remain almost unchanged.

[13] See the Appendix (3), Table A.8 for descriptive statistics of the CRU weather variables.

D. *Heterogeneity with respect to precipitation impacts*

Dell, Jones & Olken (2012) also include control variables which interact precipitation with "poor" and "hot" dummies. We do not do this in our main specification, because we are primarily interested in temperature. However, one could suspect that adding interactions for precipitation as well could change the results. To make sure this is not the case, and also to investigate whether there is heterogeneity of precipitation impacts with regard to poor or hot countries, in Table 7 we add these precipitation interactions in Table 7, but with do not find any significant evidence supporting this heterogeneity. In addition, the impacts from temperature shocks in poor countries are not diminished, but actually slightly larger.

E. *Using temperature and precipitation levels*

In our main specification, we regress annual TFP growth rate and annual temperature and precipitation changes. Dell, Jones & Olken (2012), instead, regress their dependent variable, GDP growth, on temperature and precipitation levels. We argue above that we choose to using annual changes for our weather regressors because of the trend-stationary nature of temperature data.
However, to ensure our results are not driven by the choice of first-differencing weather data, in Table 8 we regress TFP growth on temperature and precipitation levels, and check for the heterogeneity of temperature impacts as in our main specification reported in Table 2.
Our core findings are not altered by the use of levels instead of first differences.

F. *Non-linearity of temperature impacts*

One might object that we fail to account for non-linearity of temperature impacts.
Indeed, this is what Burke, Miguel and Hsiang (2015) results seem to suggest: "with most poor countries on the downward slope of the response function [between the temperature level and the GDP growth rate] but rich countries distributed almost symmetrically around the optimum, a linear regression for the effect of temperature would recover a steep negative in poor countries but ambiguous (and closer to zero) slope for rich countries"[15].
To take this possibility into account, in Table 9 we present a different specification in which we include the square of annual temperature change as an additional independent variable in the regressions, and we also interact it with the "poor" and "hot" dummies like we do for annual temperature change in the main specification.

---

[15] Supplementary Information, page 20.

As Columns (1) – (6) show, the square of annual temperature change is almost always insignificant, even in poor countries. The only effect its inclusion seems to have on the estimates is to make all the cumulative impacts slightly bigger.

Anyway, what the results suggest is the we can rule out the possibility of meaningful non-linear effects of temperature shocks, and that a linear function is the best approximation of the TFP-temperature relationship for this dataset, in line with Dell, Jones & Olken (2012).

G. Exploring the "rich" interaction

We first check whether or not only in poor countries TFP growth is affected by temperature by inspecting its complement. We therefore run exactly the same specification of Table 2, but substitute the "poor" interaction with an interaction between annual temperature changes and a dummy for being rich, with "rich" being defined as having an above average GDP per capita in 1960. Additionally, we also include the alternative definition of "rich", Rich_2, defined as having an above average GDP per capita, and interact it with temperature shocks.

The results are shown in Table 10. Column (1) shows results for the baseline specification which only includes annual temperature and precipitation shocks and the "rich" interaction. Although at a first inspection the coefficient for ∆Temp and Rich*∆Temp being both strongly significant, but of opposite sign, their linear combination at the bottom of Column (2) makes clear that the total effect of temperature on the TFP growth rate of rich countries is small and statistically insignificant. When we add the "hot" interaction in Column (2), the total effect of shocks in rich countries is again very small and insignificant. We repeat the same exercise in Columns (3) and (4), using the alternative definition of "hot", with analogous results. Finally, in Column (5) and (6), we run two specifications with the different definition of "hot" as having above 75% percentile average temperature in the 1960s. Once again, the net effect in rich countries is again close to zero and insignificant.

H. Investigating the subsample of poor countries

In Table 11 we run a specification using only the subsample of poor countries, "poor" defined as having a below median GDP per capita in 1960. The coefficient for ∆Temp is negative and significant, predicting a -1.8 percentage point decrease in the TFP growth rate for a 1°C increase. This confirms again the negative causal relationship in poor countries, which is shown graphically in Figure 2.

I. Joint interactions with poor and hot dummies

Finally, we run two specifications in which we add in the regressions a double interaction term, namely between temperature changes, a dummy for being poor and a dummy for being hot, and we repeat these for both our definitions of poor and hot countries. Table 12 shows the results. With the joint interaction included, temperature shocks significantly affect TFP growth not only in poor countries but also in hot countries. The effect is larger in poor countries than in hot countries, but the difference is not significant. The joint effect is similar in size as above.

J. Interactions with GDP / per capita and Polity2

Additionally, we investigate two specifications which could affect the interpretation and validity of our findings. First, we run a specification in which we substitute the "poor" interaction with an interaction between temperature shocks and GDP per capita. The previous definitions of poor, in fact, are all based on a fixed classification between who is rich and who is poor. This is fine for estimation, but not for simulation. In almost fifty years countries that were poor in the beginning grew out of poverty, with the notable examples of South Korea, Malaysia and China. We would hope for other countries to follow their lead in the next fifty years. Interacting annual temperature changes with GDP per capita can overcome this, and provide evidence on whether the negative impact of temperature shocks on the growth rate of TFP gets smaller or disappears as countries grow richer.

As Column (1) in Table 13 shows, this is the case. The interaction with GDP per capita is positive and significant at the 1% level: solving the first derivative with respect to $\triangle$Temp, and re-transforming the natural logarithm of GDP in dollars, suggests that the marginal effect of a 1°C annual increase becomes zero when income is approximately \$34,400 per person per year for countries classified as "hot"[16], approximately \$14,900 per person per year for countries not classified as "hot"[17], and approximately \$25,600 per person per year for the sample as a whole[18] (see Figures 3, 4 and 5 for a graphical representation of the marginal effects, at different GDP per capita levels, for the three cases).

This indicates that, even though the estimates are inevitably imprecise, and the GDP level where the marginal effect of $\triangle$Temp turns zero depends on the initial temperature level, development always means reduced vulnerability and, ultimately, immunity from the impact of temperature shocks on TFP growth rate.

---

[16] In natural logarithm: 10.447 (SE = 1.234).
[17] In natural logarithm: 9.609 (SE = 0.351).
[18] In natural logarithm: 10.150 (SE = 0.283).

The second alternative specification includes an interaction between temperature changes and a measure of institutional quality, Polity2 ('Polity IV Project', 2014). We added this interaction because it could be the case that negative impacts come not through being a poor country, but through poor institutions, i.e. through low institutional quality. In the context of the well-known debate on the determinants of long-run development (Acemoglu, Jonson & Robinson, 2000; Diamond, 1999; Easterly & Levine, 2003; Gallup, Sachs, & Mellinger, 1999), the institution hypothesis is one of the two main currents (the other being the geography hypothesis). Institutions are considered by many (Acemoglu, Johnson & Robinson, 2000; Acemoglu, Johnson & Robinson, 2001; Easterly and Levine, 2003; Rodrik, Subramanian & Trebbi, 2004) as the fundamental cause of economic growth in the long-run. This specification thus constitutes a way of testing once again the relationship between climate, institutions and development.

We use Polity2 as a measure of institutions. Polity2 ranges from -10 to 10 and combines the democracy and autocracy scores from the Polity IV dataset. In order to investigate whether or not the impact of temperature appears also, or exclusively, through the institutional channel, we interact it with annual temperature changes and add this interaction to the baseline specification with the "poor" interaction.

Column (2) in Table 13 shows our finding is not altered: the negative impact of temperature still appears through being poor, and the coefficient for the total effect in poor countries is analogous both in significance and magnitude to the previous ones. There is some weak evidence that the interaction between temperature shocks and Polity2 has a positive effect on the TFP growth rate, but this is not enough to justify a rethinking of our main conclusion.

K. Labour productivity growth as the dependent variable

We find a negative effect of weather shocks on total factor productivity growth, but only in poor countries. This is probably due to the fact that poor countries have a much larger share of their GDP in the agricultural sector, much more outdoor work and lower adaptive capacity, which suggests that one of the channels could be an impact on (outdoor) labour productivity.

Labour productivity is one of the components of total factor productivity. We use labour productivity growth in place of TFP growth as an alternative dependent variable for two reasons: first, it represents an additional and useful to check the robustness for our core findings; and second, it could provide insights on the channels through which temperature affects TFP growth and on the reasons why this is only the case for poor countries. Hence, we repeat our basic specification, replacing annual TFP growth with annual labour productivity growth, where labour productivity is defined as annual output per person employed. Data on labour productivity have been obtained by Penn World Tables, PWT 8.1 (PWT 8.1, 2016), by dividing real GDP at constant national prices by the annual number of persons employed.

Table 14 shows the results for the baseline sample, Table A.10 for the alternative sample: the impact of temperature shocks on labour productivity growth is negative and significant only in poor countries, and the coefficients are remarkably consistent and very similar in magnitude and significance to those of the TFP regressions, which suggests, as discussed in further detail in Section 6, that this is indeed a key channel responsible for the temperature-TFP relationship in poor countries. This has also been shown in studies of microdata (Cachon, Gallino, & Olivares, 2012; Heal & Park, 2015; Niemelä, Hannula, Rautio, Reijula, & Railio, 2002; Sudarshan & Tewari, 2013).

L. Labour force and capital stock

Dell, Jones and Olken (2012) studied the impact of temperature variations on the growth rate of per capita income. Their results are qualitatively similar to ours: unusually hot years negatively affect growth, but only in poor countries. We investigate the growth rate of total factor productivity, and hypothesize that this explains Dell, Jones and Olken's results. However, their result could also be explained, at least partly, by changes in the labour force or capital stock.

Table 15 shows the results for regressions of the annual growth rate of the number of persons employed[19] and the annual growth rate of real capital stock on temperature and precipitation change. The explanatory variables are both statistically insignificant in the main specification, and only the total effect of temperature change on the growth rate of the capital stock in poor countries is positive and weakly significant at the 10% level. In other words, Dell, Jones and Olken's temperature impact on income growth is due to the effect of temperature on total factor productivity growth, perhaps dampened by an effect of temperature on capital deepening.

M. Regressions with Driscoll-Kraay standard errors

[19] Data on the size of the labour force were incomplete in PWT 8.1.

Countries are not independent from each other. In the specifications above, we do not check or correct for spatial autocorrelation. As Dell, Jones & Olken (2014) notice, in the weather-economy literature this is usually accomplished by making use of Conley (1999) standard errors which allow correlation to decay smoothly with distance. However, the use of Conley (1999) standard errors would make little sense in our sample, given that the choice of common distance cut-off points would be equally applied to countries as different in geographical size as China and Trinidad and Tobago. Hence, we opted for the use of Driscoll and Kraay (1998) standard errors, which are robust to cross-sectional / spatial and temporal dependence.

Table 16 reports the results of the baseline FE regressions for the main sample, Table A.11 for the alternative sample. The significance of the coefficients is slightly diminished in some of the specifications, but the overall picture is that our core findings are not altered when taking into account the possibility of spatial dependence between countries.

Section 6
Implications of climate change

What do these results mean for future climate change? The temperature in poor countries in the almost half century of our sample saw an increase of approximately 0.6°C, or on average 0.012°C per year. There were positive and negative shocks to the annual temperature but the positive shocks were, on average, 0.012°C larger. This means that, on average, negative shocks to the annual TFP growth rate were 0.012°C/year * 1.762%/°C (cf. Table 7) = 0.021% (SE = 0.006%) per year larger than positive shocks. The 21st century could see an additional global warming of 0.3-4.8°C[20] relative to the period 1986-2005 (IPCC 2013). If past relationships will continue to hold, and excluding both intensification and adaptation, annual TFP growth rate could be reduced by 0.005-0.085% per year. Over a hundred years, total factor productivity would be 0.5-8.2% below where it would be without climate change. This is an upper bound, as we estimated the short-run semi-elasticity rather than the long-run one. This extrapolation is not immune to concerns about external validity.

In the worst case scenario of a further 4.8°C warming, annual TFP growth in poor countries would be lowered by about 0.085% during this century. This is not trivial, considering that it would be an

[20] Given that the standard deviation for annual temperature change is 0.56°C (cf. Table 1), interannual variability is quite large relative to the projected trend, so while this extrapolation should be interpreted with the usual caution, its implications should not be *a priori* dismissed.

additional dynamic effect to be added to the current impact estimates, but it is much smaller than hypothesized and simulated in recent literature. In the simulation using DICE 2010 run by Dietz and Stern (2015), and in particular in their endogenous TFP model with standard assumptions about the damage function and climate sensitivity, annual *global* TFP growth rate is reduced by about 0.20 percentage points, for the period 2005-2205 and with a temperature increase of 5.7°C above pre-industrial levels. Using our estimates and their scenario, we find a value of 1.762*(4.9/200) = 0.04%[21], roughly five times lower and, importantly, *only* for poor countries.

Similarly, Moyer et al. (2014) alter the growth path of TFP in DICE, allowing for a reduction in the annual *global* growth rate by more than 0.20%, over a 300-year period and under a predicted temperature increase of 5.9°C above pre-industrial. Under these conditions, we would predict an annual decrease by 0.03%, but again *only* for poor countries.

In Moore and Diaz (2015), who endogenize TFP in a two-region (rich and poor) version of DICE 2013R, using parameters calibrated on the empirical findings of Dell, Jones & Olken (2012), the decrease in annual TFP growth rate in poor countries is approximately 0.52%, over the period 2015-2105, with a temperature increase over the century of about 3°C. Conversely, our derived calculations for this simulation point to a reduction in the annual growth rate of TFP in poor countries by about 0.06%, almost an order of magnitude lower that their projection.

Unlike the papers above, we stress that once a certain income per capita threshold is reached, these negative impacts would disappear altogether. Our estimates point to an upper threshold of $34,400 income per capita (for hot countries), a value which, according to global projections, will be largely surpassed during this century.

These results further increase concerns over distributional issues of future impacts. As Tol (2015) shows, it is widely accepted that poor countries will be the ones who will suffer the most from climate change impacts. This work confirms and reinforces this view. Additionally, as explained in Inklaar and Timmer (2013), Keller (2004), Griffith, Redding, and Van Reenen (2004), TFP growth as a determinant of long-run economic growth is more important in poor countries than in rich ones.

Finally, given that, as noted by Gillingham et al. (2015): "uncertainty in the growth of productivity

[21] In the DICE model, temperature in 2005 is already 0.83 °C above pre-industrial.

(or output per capita) is known to be a critical parameter in determining all elements of climate change", all this calls for complementarity between climate policy and poverty reduction (Schelling, 1992).

Section 7
Discussion and conclusion

We test the recently advanced hypothesis that climate change harms TFP growth by looking at the past relationship between TFP growth rates and temperature shocks. We find a negative relationship only in poor countries. The relationship is robust to alternative samples, alternative data, alternative specifications, and to spatial autocorrelation. There is some evidence that temperature shocks may have a negative effect in hot countries too. The estimated temperature effect on TFP growth probably explains the effect on economic growth found in previous papers, and is probably explained by temperature effects on labour productivity. While statistically significant, our upper bound estimate suggest that climate change would reduce TFP growth by less than 0.1%.

The findings of this work confirm the results of Dell, Jones and Olken (2012), who also found a statistically significant but modestly sized relationship between temperature levels and economic growth only in poor countries, and that showed using lags and long differences a persistence of weather impacts in the medium run which is likely to mean the presence of growth effects other, or rather, than level output effects. Our results contradict the conclusions of Burke, Miguel and Hsiang (2015), who found large impacts of temperature on productivity.

This work represents an advancement compared to previous literature because, using the first differences of TFP and temperature levels, not only it alleviates the issue of non-stationarity in panel analysis which may tend to produce spurious results, but also directly addresses the issue of potential long-run growth effects, since its main dependent variable is notably one of the main drivers of long-run economic growth (Solow, 1956). In this different perspective, an impact on annual TFP growth is already, *per se* a long-term impact. There is no need to use first differences, since in this scenario temperature shocks affects economic activity not through Equation (1), but directly through Equation (4). Conversely, Dell, Jones and Olken (2012) focused on GDP growth as the dependent variable of interest, and thus could not explicitly test for the presence of growth effects.

However, a number of limits and *caveats* for this work also need to be made clear. First: sample size and data quality. Both our samples only include less than 70 countries (60 and 68, respectively).

Although together they account for a large share of world GDP and population, sample size is indeed reduced. As for data quality, TFP data represent the so-called *Solow residual*, and in fact this is the way they are calculated in PWT 8.1 (Feenstra, Inklaar, & Timmer, 2013 & 2015; Inklaar & Timmer, 2013). Therefore, the estimates are potentially affected by measurement error and a whole host of errors in the specification and the estimation of the production function used to derive TFP. Unfortunately, to the best of our knowledge there is no availability of other TFP datasets at the country level covering such a long timespan. Weather data as well notably suffer from measurement error and different data quality in different countries. However, the issue of measurement error is at least partially alleviated here since the results appear to be robust to sample choices, to different specifications of key explanatory variables, and to different weather data with different aggregation methods.

Second, as already mentioned in the introduction, external validity with respect to future climate change. Again, weather variations are *not* climate variations: the first are random shorter-run temporal variations, the second are averages over several decades (Dell, Jones & Olken, 2014). In other words climate, as emphasized by Auffhammer et al. (2013), is a long average of weather at a given location. It is thus key to always keep in mind that a 1°C shock in a given year and place is not equivalent to a permanent 1°C global increase, and that projections like the simple extrapolation with regard to global warming we performed above typically suffer from this drawback. In other words, we only estimated the short-run semi-elasticity, whereas we need to know the long-run semi-elasticity.

Third, future climate change, especially if pronounced as it is projected in some extreme emission scenarios (IPCC, 2013) may well entail consequences and effects which lie outside historical experience. Substantial sea level rise, a thermohaline circulation slowdown, the release of methane from melting permafrost are all potential intensifying effects which are indeed not captured by this analysis, based on a period in which there was only limited climatic variability and limited warming. Such an intensification of impacts may well change the picture we depicted, both quantitatively and qualitatively.

Fourth, every forecast or projection based on this study implies the assumption that past historical relationship will continue to hold in the future. As argued in Dell, Jones and Olken (2014) and Tol (2015), this could indeed not be the case, either due to intensification of negative impacts or to adaptation through development in the long run.

Fifth, total factor productivity is an aggregate measure, and changes in total factor productivity are due to a variety of changes in underlying economic phenomena. With our data it is impossible to open this black box, but future research should attempt this using micro-data and natural experiments.

The central finding of this work is that TFP growth rates of poor countries are affected by temperature shocks in recent past. Once again, poverty means vulnerability. However, this causal relationship between temperature, poverty and productivity growth is subjected to *caveats* and should be interpreted with caution. What this analysis suggests is the fact that weather shocks affect economic growth through the TFP channel only when coupled with poverty, not that climate change will harm future economic growth by affecting technological progress, as hypothesized in literature. Hence, given the preeminent importance of TFP growth for long-run development, and under the assumption that weather impacts have at least some external validity with regard to climate change, the main conclusions that stem from this study are an increase of concerns over the inequality of future impacts, a policy guideline which considers poverty reduction as a crucial and paramount element of climate policy and, at the research level, a call for further studies on the potential dynamic effects of future climate change.

References

Acemoglu, Daron, Simon Johnson, and James A. Robinson. 2000. The Colonial Origins of Comparative Development: An Empirical Investigation. National bureau of economic research. http://www.nber.org/papers/w7771.

Alsan, M. (2014). The effect of the tsetse fly on African development. *The American Economic Review*, *105*(1), 382–410.

Andersen, T. B., Dalgaard, C.-J., & Selaya, P. (2016). Climate and the Emergence of Global Income Differences. *The Review of Economic Studies*, rdw006.

Anthoff, D., Tol, R. S., & Yohe, G. W. (2009). Risk aversion, time preference, and the social cost of carbon. *Environmental Research Letters*, *4*(2), 024002.

Auffhammer, M., Hsiang, S. M., Schlenker, W., & Sobel, A. (2013). Using weather data and climate model output in economic analyses of climate change. *Review of Environmental Economics and Policy*, ret016.

Bansal, R., & Ochoa, M. (2011). *Temperature, aggregate risk, and expected returns*. National Bureau of Economic Research. Retrieved from http://www.nber.org/papers/w17575

Burke, M., Hsiang, S. M., & Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature* (527): 235-239.

Cachon, G., Gallino, S., & Olivares, M. (2012). Severe weather and automobile assembly productivity. *Columbia Business School Research Paper*, (12/37).

Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of Econometrics*, *92*(1), 1–45.

CRU (2016), Climatic Research Unit, University of East Anglia. Retrieved 9 July 2016, from: http://www.cru.uea.ac.uk/data

Dell, M., Jones, B. F., & Olken, B. A. (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics*, *4*(3), 66–95.

Dell, M., Jones, B. F., & Olken, B. A. (2014). What do we learn from the weather? The new climate–economy literature. *Journal of Economic Literature*, *52*(3), 740–798.

Diamond, J. (1999). *Guns, germs, and steel: The fates of human societies*. WW Norton & Company.

Dietz, S., & Stern, N. (2015). Endogenous Growth, Convexity of Damage and Climate Risk: How Nordhaus' Framework Supports Deep Cuts in Carbon Emissions. *The Economic Journal*, *125*(583), 574–620.

Driscoll, J. C., & Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics*, *80*(4), 549–560.

Easterly, W., & Levine, R. (2003). Tropics, germs, and crops: how endowments influence economic development. *Journal of Monetary Economics*, *50*(1), 3–39.

Fankhauser, S., & Tol, R. S. (2005). On climate change and economic growth. *Resource and Energy Economics*, *27*(1), 1–17.

Feenstra, R. C., Inklaar, R., & Timmer, M. (2013). PWT 8.0–a user guide. *Groningen Growth and Development Centre, University of Groningen*.

Feenstra, R. C., Inklaar, R., & Timmer, M. P. (2015). The next generation of the Penn World Table. *The American Economic Review*, *105*(10), 3150–3182.

Gallup, J. L., Sachs, J. D., & Mellinger, A. D. (1999). Geography and economic development. *International Regional Science Review*, *22*(2), 179–232.

Gay-Garcia, C., Estrada, F., & Sánchez, A. (2009). Global and hemispheric temperatures revisited. *Climatic Change*, *94*(3–4), 333–349. https://doi.org/10.1007/s10584-008-9524-8

Gillingham, K., Nordhaus, W. D., Anthoff, D., Blanford, G., Bosetti, V., Christensen, P., … Sztorc, P. (2015). *Modeling Uncertainty in Climate Change: A Multi-Model Comparison* (Working Paper No. 21637). National Bureau of Economic Research. Retrieved from http://www.nber.org/papers/w21637

Griffith, R., Redding, S., & Van Reenen, J. (2004). Mapping the two faces of R&D: Productivity growth in a panel of OECD industries. *Review of Economics and Statistics*, *86*(4), 883–895.

Hallegatte, S. (2005). The long time scales of the climate–economy feedback and the climatic cost of growth. *Environmental Modeling & Assessment*, *10*(4), 277–289.

Harris, R. D., & Tzavalis, E. (1999). Inference for unit roots in dynamic panels where the time dimension is fixed. *Journal of Econometrics*, *91*(2), 201–226.

Heal, G., & Park, J. (2015). *Goldilocks economies? temperature stress and the direct impacts of climate change*. National Bureau of Economic Research. Retrieved from http://www.nber.org/papers/w21119

Hope, C. (2006). The marginal impact of CO2 from PAGE2002: An integrated assessment model incorporating the IPCC's five reasons for concern. *Integrated Assessment*, *6*(1). Retrieved from http://78.47.223.121:8080/index.php/iaj/article/viewArticle/227

Hsiang, S. M., & Jina, A. S. (2014). *The causal effect of environmental catastrophe on long-run economic growth: Evidence from 6,700 cyclones*. National Bureau of Economic Research. Retrieved from http://www.nber.org/papers/w20352

Huntington, E. (1922). *Civilization and climate*. Yale University Press.

Im, K. S., Pesaran, M. H., & Shin, Y. (2003). Testing for unit roots in heterogeneous panels. *Journal of Econometrics*, *115*(1), 53–74.

Inklaar, R., & Timmer, M. (2013). Capital, Labor and TFP in PWT8. 0. *University of Groningen (Unpublished)*. Retrieved from http://www.piketty.pse.ens.fr/files/InklaarTimmer13.pdf

Kaufmann, R. K., Kauppi, H., & Stock, J. H. (2010). Does temperature contain a stochastic trend? Evaluating conflicting statistical results. *Climatic Change*, *101*(3–4), 395–405. https://doi.org/10.1007/s10584-009-9711-2

Keller, W. (2004). International technology diffusion. *Journal of Economic Literature*, *42*(3), 752–782.

Lemoine, D., & Kapnick, S. (2015). A top-down approach to projecting market impacts of climate change. *Nature Climate Change* 6(5):514-519.

'Maddison Project' (2016). Retrieved 9 July 2016, from: http://www.ggdc.net/maddison/maddison-project/home.html

Moore, F. C., & Diaz, D. B. (2015). Temperature impacts on economic growth warrant stringent mitigation policy. *Nature Climate Change*, *5*(2), 127–131.

Moyer, E. J., Woolley, M. D., Glotter, M., & Weisbach, D. A. (2014). Climate impacts on economic growth as drivers of uncertainty in the social cost of carbon. *Journal of Legal Studies*, *43*(2), 401–425.

Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica: Journal of the Econometric Society*, 69–85.

Niemelä, R., Hannula, M., Rautio, S., Reijula, K., & Railio, J. (2002). The effect of air temperature on labour productivity in call centres—a case study. *Energy and Buildings*, *34*(8), 759–764.

Nordhaus, W. D. (1991). To slow or not to slow: the economics of the greenhouse effect. *The Economic Journal*, *101*(407), 920–937.

Nordhaus, W. D. (2007). A review of the Stern review on the economics of climate change. *Journal of Economic Literature*, *45*(3), 686–702.

Nordhaus, W. D. (2008). A question of balance: economic modeling of global warming. *New Haven, CT: Yale University Press. Retrieved*, *1*(10), 2011.

Olsson, O., & Hibbs, D. A. (2005). Biogeography and long-run economic development. *European Economic Review*, *49*(4), 909–938.

Pindyck, R. S. (2012). Uncertain outcomes and climate change policy. *Journal of Environmental Economics and Management*, *63*(3), 289–303.

Pindyck, R. S. (2013). Climate change policy: What do the models tell us? *Journal of Economic Literature*, *51*(3), 860–872.

'Polity IV Project' (2014): Retrieved 9 July 2016, from: http://www.systemicpeace.org/polity/polity4.htm

PWT 8.1 (2016) | PWT Releases | University of Groningen.. Retrieved 9 July 2016, from: http://www.rug.nl/research/ggdc/data/pwt/pwt-8.1

Rodrik, D., Subramanian, A., & Trebbi, F. (2004). Institutions rule: the primacy of institutions over geography and integration in economic development. *Journal of Economic Growth*, *9*(2), 131–165.

Schelling, T. C. (1992). Some economics of global warming. *The American Economic Review*, *82*(1), 1–14.

Solow, R. M. (1956). A contribution to the theory of economic growth. *The Quarterly Journal of Economics*, 65–94.

Stern, N. (2013). The structure of economic modeling of the potential impacts of climate change: grafting gross underestimation of risk onto already narrow science models. *Journal of Economic Literature*, *51*(3), 838–859.

Stern, N. H., Peters, S., Bakhshi, V., Bowen, A., Cameron, C., Catovsky, S., … others. (2006). *Stern Review: The economics of climate change* (Vol. 30). Cambridge University Press Cambridge.

Stocker, T. F., Qin, D., Plattner, G. K., Tignor, M., Allen, S. K., Boschung, J., … Midgley, P. M. (2013). *IPCC, 2013: Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, 1535 pp*. Cambridge Univ. Press, Cambridge, UK, and New York.

Sudarshan, A., & Tewari, M. (2013). *The economic impacts of temperature on industrial productivity: Evidence from indian manufacturing*. Working Paper. Retrieved from http://icrier.org/pdf/working_paper_278.pdf

Tol, R. S. J. (2015). *Economic impacts of climate change* (Working Paper Series No. 7515). Department of Economics, University of Sussex. Retrieved from https://ideas.repec.org/p/sus/susewp/7515.html

Tol, R. S., Yohe, G. W., & others. (2006). A review of the Stern Review. *World Economics-Henley on Thames*, *7*(4), 233.

Weitzman, M. L. (2009). On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics*, *91*(1), 1–19.

Weitzman, M. L. (2010). What Is The‘ Damages Function’ For Global Warming—And What Difference Might It Make? *Climate Change Economics*, *1*(01), 57–69.

Weitzman, M. L. (2011). Fat-tailed uncertainty in the economics of catastrophic climate change. *Review of Environmental Economics and Policy*, *5*(2), 275–292.

Table 1

Descriptive Statistics

	Mean	Var	sd	Min	Max	Obs
TFP growth rate	0.481	15.49	3.935	-56.05	26.76	2760
△Temp	0.0121	0.318	0.564	-2.952	2.442	2760
△Pre	-0.0142	5.942	2.438	-35.40	37.64	2760
GDP_percap	8.480	1.022	1.011	6.084	10.35	2820

TFP growth rate is the annual percentage change and expressed in natural logarithm.
Temperature change is annual and expressed in degree Celsius.
Precipitation change is annual and expressed in units of 100 mm per year.
GDP per capita is in natural logarithm of 1990 international Geary - Khamis dollars.

Table 2

Relationship between annual TFP growth rates and temperature changes

Dependent variable: annual TFP growth rate	(1)	(2)	(3)	(4)	(5)	(6)
△Temp	-0.485**	-0.029	0.057	0.098	0.008	0.051
	(0.216)	(0.136)	(0.143)	(0.123)	(0.134)	(0.120)
△Pre	-0.033	-0.042*	-0.047**	-0.048**	-0.049**	-0.051**
	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)
Poor x △Temp		-1.493***	-1.195**		-1.315***	
		(0.404)	(0.468)		(0.437)	
Hot x △Temp			-0.684	-0.612		
			(0.452)	(0.429)		
Poor_2 x △Temp				-1.425***		-1.513***
				(0.420)		(0.410)
Hot_2 x △Temp					-1.048**	-0.979**
					(0.484)	(0.481)
_cons	1.416***	1.338***	1.280***	1.271***	1.284***	1.273***
	(0.327)	(0.331)	(0.322)	(0.324)	(0.318)	(0.319)
N	2760	2760	2760	2760	2760	2760
R^2	0.208	0.215	0.216	0.217	0.216	0.218
adj. R^2	0.121	0.128	0.129	0.131	0.130	0.131
AIC	14749.211	14727.177	14725.510	14720.275	14723.930	14718.702
Total effect in poor countries		-1.523***	-1.139**	-1.327***	-1.307***	-1.462***
		(0.406)	(0.515)	(0.456)	(0.453)	(0.420)
Total effect in hot countries			-0.627	-0.515	-1.040**	-0.928*
			(0.402)	(0.388)	(0.473)	(0.477)

Notes:

All specifications include country FE and Region x Time FE.
Poor is a dummy with value 1 for countries with below median GDP per capita in 1960.
Hot is a dummy with value 1 for countries with above median average temperature in the 1960s.
Poor_2 is a dummy with value 1 for countries with below median GDP per capita.
Hot_2 is a dummy with value 1 for countries with average temperature in the 1960s above the 75%.
Temperature change is annual and expressed in degree Celsius.
Precipitation change is annual and expressed in units of 100 mm per year.
Standard errors are in parentheses and are clustered at the country level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3

Descriptive Statistics – Sample B

	Mean	Var	Sd	Min	Max	Obs
TFP growth rate	0.0475	18.07	4.251	-57.82	37.10	2448
△Temp	0.0220	0.327	0.572	-2.952	2.442	2448
△Pre	-0.0198	5.890	2.427	-35.40	37.64	2448
GDP_percap	8.619	0.945	0.972	6.084	10.67	2516

TFP growth rate is the annual percentage change and expressed in natural logarithm.
Temperature change is annual and expressed in degree Celsius.
Precipitation change is annual and expressed in units of 100 mm per year.
GDP per capita is in natural logarithm of 1990 international Geary - Khamis dollars.

Table 4

Relationship between annual TFP growth rates and temperature changes – Sample B

Dependent variable: Annual TFP growth rate	(1)	(2)	(3)	(4)	(5)	(6)
△Temp	-0.345*	0.053	0.087	0.071	0.077	0.075
	(0.193)	(0.111)	(0.114)	(0.112)	(0.113)	(0.113)
△Pre	-0.033	-0.041**	-0.043**	-0.046**	-0.043**	-0.046**
	(0.020)	(0.018)	(0.018)	(0.018)	(0.019)	(0.019)
Poor x △Temp		-1.200***	-1.125***		-1.198***	
		(0.318)	(0.328)		(0.318)	
Hot x △Temp			-0.230	-0.093		
			(0.334)	(0.323)		
Poor_2 x △Temp				-1.308***		-1.337***
				(0.314)		(0.314)
Hot_2 x △Temp					-0.244	-0.196
					(0.316)	(0.319)
_cons	1.362***	1.295***	1.282***	1.283***	1.290***	1.284***
	(0.307)	(0.304)	(0.302)	(0.303)	(0.302)	(0.303)
N	2448	2448	2448	2448	2448	2448
R^2	0.198	0.202	0.202	0.203	0.202	0.203
adj. R^2	0.121	0.126	0.126	0.127	0.126	0.127
AIC	13476.708	13464.522	13466.184	13463.170	13466.210	13463.021
Total effect in poor countries		-1.147***	-1.037***	-1.237***	-1.121***	-1.262***
		(0.351)	(0.370)	(0.354)	(0.355)	(0.354)
Total effect in hot countries			-0.143	-0.022	-0.166	-0.120
			(0.318)	(0.314)	(0.309)	(0.313)

Notes:

All specifications include country FE and Region x Time FE.
Poor is a dummy with value 1 for countries with below median GDP per capita in 1970.
Hot is a dummy with value 1 for countries with above median average temperature in the 1970s.
Poor_2 is a dummy with value 1 for countries with below median GDP per capita.
Hot_2 is a dummy with value 1 for countries with above median average temperature.
Temperature change is annual and expressed in degree Celsius.
Precipitation change is annual and expressed in units of 100 mm per year.
Standard errors are in parentheses and are clustered at the country level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5

Relationship between annual TFP growth rates and temperature changes – CRU Data

Dependent variable: Annual TFP growth rates	(1)	(2)	(3)	(4)	(5)	(6)
△Temp	-0.288	0.100	0.145	0.136	0.104	0.124
	(0.198)	(0.108)	(0.114)	(0.111)	(0.113)	(0.108)
△Pre	-0.034	-0.043	-0.046	-0.046	-0.043	-0.045
	(0.034)	(0.034)	(0.033)	(0.033)	(0.034)	(0.034)
Poor x △Temp		-1.411***	-1.297***		-1.400***	
		(0.396)	(0.435)		(0.428)	
Hot x △Temp			-0.337	-0.113		
			(0.383)	(0.402)		
Poor_2 X △Temp				-1.553***		-1.599***
				(0.446)		(0.416)
Hot_2 x △Temp					-0.076	-0.015
					(0.406)	(0.404)
_cons	1.409***	1.352***	1.333***	1.331***	1.350***	1.337***
	(0.329)	(0.330)	(0.334)	(0.335)	(0.328)	(0.330)
N	2760	2760	2760	2760	2760	2760
R^2	0.206	0.212	0.212	0.213	0.212	0.213
adj. R^2	0.119	0.125	0.125	0.126	0.125	0.126
AIC	14755.501	14737.214	14738.333	14734.320	14739.186	14734.410
Total effect in poor countries		-1.311***	-1.152**	-1.416***	-1.296***	-1.475***
		(0.401)	(0.470)	(0.478)	(0.449)	(0.434)
Total effect in hot countries			-0.192	0.023	0.028	0.109
			(0.343)	(0.371)	(0.381)	(0.385)

Notes:

All specifications include country FE and Region x Time FE.
Poor is a dummy with value 1 for countries with below median GDP per capita in 1960.
Hot is a dummy with value 1 for countries with above median average temperature in the 1960s.
Poor_2 is a dummy with value 1 for countries with below median GDP per capita.
Hot_2 is a dummy with value 1 for countries with above median average temperature.
Temperature change is annual and expressed in degree Celsius.
Precipitation change is annual and expressed in units of 100 mm per year.
Standard errors are in parentheses and are clustered at the country level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6

Omitting precipitation

Dependent variable: Annual TFP growth rate	(1)	(2)	(3)	(4)	(5)	(6)
ΔTemp	-0.480**	-0.031	0.049	0.089	0.003	0.045
	(0.215)	(0.135)	(0.143)	(0.123)	(0.133)	(0.119)
Poor x ΔTemp		-1.464***	-1.183**		-1.295***	
		(0.401)	(0.465)		(0.433)	
Hot x ΔTemp			-0.639	-0.568		
			(0.450)	(0.428)		
Poor_2 x ΔTemp				-1.407***		-1.488***
				(0.417)		(0.406)
Hot_2 x ΔTemp					-0.969**	-0.899*
					(0.477)	(0.475)
_cons	1.400***	1.320***	1.264***	1.254***	1.267***	1.255***
	(0.326)	(0.329)	(0.321)	(0.322)	(0.317)	(0.318)
N	2760	2760	2760	2760	2760	2760
R^2	0.207	0.214	0.215	0.216	0.215	0.217
adj. R^2	0.121	0.128	0.129	0.130	0.129	0.131
AIC	14748.538	14727.353	14726.134	14721.090	14724.826	14719.796
Total effect in poor countries		-1.496***	-1.134**	-1.318***	-1.292***	-1.444***
		(0.402)	(0.511)	(0.454)	(0.449)	(0.416)
Total effect in hot countries			-0.590	-0.480	-0.966**	-0.854*
			(0.399)	(0.386)	(0.468)	(0.472)

Notes:

All specifications include country FE and Region x Time FE.
Poor is a dummy with value 1 for countries with below median GDP per capita in 1960.
Hot is a dummy with value 1 for countries with above median average temperature in the 1960s.
Poor_2 is a dummy with value 1 for countries with below median GDP per capita.
Hot_2 is a dummy with value 1 for countries with above median average temperature.
Temperature change is annual and expressed in degree Celsius.
Standard errors are in parentheses and are clustered at the country level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7

Including precipitation interactions

Dependent variable: Annual TFP growth rate	(1)	(2)	(3)	(4)	(5)	(6)
ΔTemp	-0.485**	-0.031	0.055	0.096	0.007	0.050
	(0.216)	(0.136)	(0.143)	(0.123)	(0.134)	(0.120)
Poor x ΔTemp		-1.506***	-1.199**		-1.335***	
		(0.404)	(0.463)		(0.437)	
Hot x ΔTemp			-0.690	-0.618		
			(0.452)	(0.429)		
Poor_2 x ΔTemp				-1.427***		-1.534***
				(0.413)		(0.410)
Hot_2 x ΔTemp					-1.035**	-0.958*
					(0.483)	(0.481)
ΔPre	-0.033	-0.025	-0.002	-0.012	-0.042	-0.047
	(0.023)	(0.024)	(0.045)	(0.045)	(0.036)	(0.036)
Poor x ΔPre		-0.036	-0.030		-0.045	
		(0.046)	(0.044)		(0.049)	
Hot x ΔPre			-0.038	-0.030		
			(0.056)	(0.059)		
Poor _2 x ΔPre				-0.027		-0.044
				(0.050)		(0.052)
Hot_2 x ΔPre					0.029	0.033
					(0.049)	(0.049)
_cons	1.416***	1.333***	1.269***	1.262***	1.281***	1.272***
	(0.327)	(0.331)	(0.323)	(0.324)	(0.319)	(0.320)
N	2760	2760	2760	2760	2760	2760
R^2	0.208	0.215	0.216	0.217	0.216	0.218
adj. R^2	0.121	0.128	0.129	0.130	0.129	0.131
AIC	14749.211	14728.774	14724.777	14719.715	14723.171	14717.971
Total temperature effect in poor countries		-1.537***	-1.144**	-1.331***	-1.328***	-1.484***
		(0.405)	(0.509)	(0.450)	(0.453)	(0.419)
Total temperature effect in hot countries			-0.636	-0.523	-1.028**	-0.908*
			(0.402)	(0.387)	(0.473)	(0.480)
Total precipitation effect in poor countries		-0.061	-0.032	-0.039	-0.087*	-0.091*
		(0.040)	(0.056)	(0.066)	(0.047)	(0.052)
Total precipitation effect in hot countries			-0.040	-0.042	-0.013	-0.014
			(0.029)	(0.031)	(0.040)	(0.040)

Notes:
All specifications include country FE and Region x Time FE. Poor is a dummy with value 1 for countries with below median GDP per capita in 1960. Hot is a dummy with value 1 for countries with above median average temperature in the 1960s. Poor_2 is a dummy with value 1 for countries with below median GDP per capita. Hot_2 is a dummy with value 1 for countries with above median average temperature. Temperature change is annual and expressed in degree Celsius. Precipitation change is annual and expressed is in units of 100 mm per year. Standard errors are in parentheses and are clustered at the country level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8

Using temperature and precipitation levels

Dependent variable: Annual TFP growth rate	(1)	(2)	(3)	(4)	(5)	(6)
Temp	-0.254	0.284	0.270	0.286*	0.283*	0.304**
	(0.212)	(0.171)	(0.169)	(0.147)	(0.166)	(0.147)
Pre	0.015	-0.004	-0.003	-0.005	-0.004	-0.006
	(0.046)	(0.045)	(0.045)	(0.045)	(0.044)	(0.044)
Poor x Temp		-1.529***	-1.570***		-1.534***	
		(0.377)	(0.438)		(0.379)	
Hot x Temp			0.081	0.108		
			(0.402)	(0.385)		
Poor_2 x Temp				-1.686***		-1.644***
				(0.419)		(0.371)
Hot_2 x Temp					0.030	0.062
					(0.648)	(0.650)
_cons	5.638	13.008***	12.741***	13.485***	12.880**	13.577**
	(3.691)	(4.271)	(4.533)	(4.597)	(5.648)	(5.676)
N	2760	2760	2760	2760	2760	2760
R^2	0.205	0.211	0.211	0.212	0.211	0.212
adj. R^2	0.118	0.125	0.124	0.125	0.124	0.125
AIC	14757.567	14738.273	14740.227	14737.313	14740.269	14737.379
Total effect in poor countries		-1.245***	-1.300***	-1.399***	-1.251***	-1.339***
		(0.359)	(0.440)	(0.421)	(0.349)	(0.351)
Total effect in hot countries			0.351	0.395	0.313	0.367
			(0.406)	(0.390)	(0.694)	(0.692)

Notes:
All specifications include country FE and Region x Time FE.
Poor is a dummy with value 1 for countries with below median GDP per capita in 1960.
Hot is a dummy with value 1 for countries with above median average temperature in the 1960s.
Poor_2 is a dummy with value 1 for countries with below median GDP per capita.
Hot_2 is a dummy with value 1 for countries with above median average temperature.
Temperature is annual and expressed in degree Celsius.
Precipitation is annual and expressed is in units of 100 mm per year.
Standard errors are in parentheses and are clustered at the country level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 9

Checking for non-linearity of temperature impacts

Dependent variable: Annual TFP growth rate	(1)	(2)	(3)	(4)	(5)	(6)
ΔTemp	-0.490**	-0.037	0.050	0.092	0.001	0.045
	(0.215)	(0.134)	(0.140)	(0.120)	(0.132)	(0.117)
(ΔTemp)^2	-0.110	-0.143*	-0.151*	-0.135**	-0.151*	-0.131*
	(0.101)	(0.077)	(0.076)	(0.067)	(0.076)	(0.066)
ΔPre	-0.033	-0.042*	-0.046**	-0.048**	-0.049**	-0.051**
	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)
Poor x ΔTemp		-1.484***	-1.181**		-1.306***	
		(0.403)	(0.451)		(0.434)	
Poor x (ΔTemp)^2		0.182	0.207		0.175	
		(0.343)	(0.323)		(0.340)	
Hot x ΔTemp			-0.694	-0.611		
			(0.447)	(0.425)		
Hot x (ΔTemp)^2			0.017	0.090		
			(0.430)	(0.440)		
Poor_2 x ΔTemp				-1.421***		-1.507***
				(0.407)		(0.407)
Poor_2 x (ΔTemp)^2				0.094		0.088
				(0.346)		(0.353)
Hot_2 x ΔTemp					-1.051**	-0.980**
					(0.479)	(0.473)
Hot_2 x (ΔTemp)^2					0.081	0.127
					(0.621)	(0.629)
_cons	1.439***	1.355***	1.294***	1.284***	1.298***	1.287***
	(0.324)	(0.325)	(0.307)	(0.307)	(0.307)	(0.307)
N	2760	2760	2760	2760	2760	2760
R^2	0.208	0.215	0.216	0.217	0.216	0.218
adj. R^2	0.121	0.128	0.128	0.130	0.129	0.130
AIC	14750.611	14728.335	14724.544	14719.561	14722.997	14717.998
Total effect in poor countries		-1.523***	-1.132**	-1.332***	-1.307***	-1.464***
		(0.406)	(0.497)	(0.443)	(0.452)	(0.420)
Total effect in hot countries			-0.648	-0.522	-1.048**	-0.929**
			(0.400)	(0.387)	(0.465)	(0.464)

Notes: All specifications include country FE and Region x Time FE. Poor is a dummy with value 1 for countries with below median GDP per capita in 1960. Hot is a dummy with value 1 for countries with above median average temperature in the 1960s. Poor_2 is a dummy with value 1 for countries with below median GDP per capita. Hot_2 is a dummy with value 1 for countries with above median average temperature. Temperature change is annual and expressed in degree Celsius. Precipitation change is annual and expressed is in units of 100 mm per year. Standard errors are in parentheses and are clustered at the country level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10

Relationship between annual TFP growth rates and temperature changes in rich countries

Dependent variable: Annual TFP growth rate	(1)	(2)	(3)	(4)	(5)	(6)
ΔTemp	-1.523***	-1.139**	-1.667***	-1.327***	-1.307***	-1.462***
	(0.406)	(0.515)	(0.369)	(0.456)	(0.453)	(0.420)
ΔPre	-0.042*	-0.047**	-0.045*	-0.048**	-0.049**	-0.051**
	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)	(0.023)
Rich x ΔTemp	1.493***	1.195**			1.315***	
	(0.404)	(0.468)			(0.437)	
Hot x ΔTemp		-0.684		-0.612		
		(0.452)		(0.429)		
Rich_2 x ΔTemp			1.683***	1.425***		1.513***
			(0.375)	(0.420)		(0.410)
Hot_2 x ΔTemp					-1.048**	-0.979**
					(0.484)	(0.481)
_cons	1.338***	1.280***	1.323***	1.271***	1.284***	1.273***
	(0.331)	(0.322)	(0.333)	(0.324)	(0.318)	(0.319)
N	2760	2760	2760	2760	2760	2760
R^2	0.215	0.216	0.216	0.217	0.216	0.218
adj. R^2	0.128	0.129	0.130	0.131	0.130	0.131
AIC	14727.177	14725.510	14721.291	14720.275	14723.930	14718.702
Total effect in rich countries	-0.029	0.057	0.016	0.098	0.008	0.051
	(0.136)	(0.143)	(0.125)	(0.123)	(0.134)	(0.120)

Notes:

All specifications include country FE and Region x Time FE.
Rich is a dummy with value 1 for countries with above median GDP per capita in 1960.
Hot is a dummy with value 1 for countries with above median average temperature in the 1960s.
Rich_2 is a dummy with value 1 for countries with above median GDP per capita.
Hot_2 is a dummy with value 1 for countries with above median average temperature.
Temperature change is annual and expressed in degree Celsius.
Precipitation change is annual and expressed in units of 100 mm per year.
Standard errors are in parentheses and are clustered at the country level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 11

Annual TFP growth rates and temperature changes in poor countries

	Annual TFP growth rate
△Temp	-1.762***
	(0.459)
△Pre	-0.067
	(0.044)
_cons	1.745***
	(0.378)
N	1380
R^2	0.226
adj. R^2	0.073
AIC	7863.372

Notes
The specification includes country FE and Region x Time FE.
Temperature change is annual and expressed in degree Celsius.
Precipitation change is annual and expressed in units of 100 mm per year.
Standard errors are in parentheses and are clustered at the country level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 12

Specification with a double interaction term

Dependent variable: annual TFP growth rate	(1)	(2)
∆Temp	0.126	0.066
	(0.131)	(0.122)
∆Pre	-0.047**	-0.051**
	(0.023)	(0.024)
Poor x ∆Temp	-1.576**	
	(0.765)	
Hot x ∆Temp	-1.170***	
	(0.252)	
Poor x Hot x ∆Temp	0.919	
	(0.896)	
Poor_2 x ∆Temp		-1.570***
		(0.455)
Hot_2 x ∆Temp		-1.412***
		(0.427)
Poor_2 x Hot_2 x ∆Temp		0.571
		(0.814)
_cons	1.274***	1.258***
	(0.316)	(0.320)
N	2760	2760
R^2	0.216	0.218
adj. R^2	0.129	0.131
AIC	14723.706	14718.412
Total effect in hot and poor countries	-1.701***	-2.344***
	(0.449)	(0.515)

Notes:
All specifications include country FE and Region x Time FE.
Poor is a dummy with value 1 for countries with below median GDP per capita in 1960.
Hot is a dummy with value 1 for countries with above median average temperature in the 1960s.
Poor_2 is a dummy with value 1 for countries with below median GDP per capita.
Hot_2 is a dummy with value 1 for countries with above median average temperature.
Temperature change is annual and expressed in degree Celsius.
Precipitation change is annual and expressed is in units of 100 mm per year.
Standard errors are in parentheses and are clustered at the country level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 13

Specifications with GDP per capita and Polity2

Dependent variable: annual TFP growth rate	(1)	(2)
ΔTemp	-5.065**	-0.518
	(1.921)	(0.318)
ΔPre	-0.046**	-0.050**
	(0.022)	(0.022)
Poor x ΔTemp		-0.823**
		(0.328)
Polity2 x ΔTemp		0.062*
		(0.032)
Polity2		-0.012
		(0.034)
GDP_percap	-0.216	
	(0.689)	
GDP x ΔTemp	0.532***	
	(0.195)	
Hot x ΔTemp	-0.675	
	(0.454)	
_cons	3.158	1.441***
	(6.143)	(0.404)
N	2760	2705
R^2	0.214	0.224
adj. R^2	0.127	0.136
AIC	14355.504	6698.196
Total effect in poor countries		-1.342***
		(0.336)

Notes
All specifications include country FE and Region x Time FE. Poor is a dummy with value 1 for countries with below median GDP per capita in 1960. Hot is a dummy with value 1 for countries with above median average temperature in the 1960s. Temperature change is annual and expressed in degree Celsius. Precipitation change is annual and expressed in units of 100 mm per year. Standard errors are in parentheses and are clustered at the country level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 14

Relationship between annual labour productivity growth rates and temperature changes

Dependent Variable: Annual labour productivity growth	(1)	(2)	(3)	(4)	(5)	(6)
△Temp	-0.543**	-0.068	0.037	0.095	-0.027	0.037
	(0.227)	(0.159)	(0.164)	(0.131)	(0.155)	(0.131)
△Pre	-0.037	-0.047*	-0.052**	-0.054**	-0.054**	-0.056**
	(0.025)	(0.024)	(0.024)	(0.024)	(0.024)	(0.024)
Poor x △Temp		-1.559***	-1.197**		-1.366***	
		(0.412)	(0.481)		(0.448)	
Hot x △Temp			-0.831*	-0.717		
			(0.466)	(0.434)		
Poor_2 x △Temp				-1.522***		-1.644***
				(0.421)		(0.411)
Hot_2 x △Temp					-1.133**	-1.036**
					(0.503)	(0.498)
_cons	2.535***	2.454***	2.383***	2.374***	2.395***	2.381***
	(0.515)	(0.518)	(0.511)	(0.514)	(0.507)	(0.508)
N	2760	2760	2760	2760	2760	2760
R^2	0.211	0.218	0.219	0.221	0.219	0.221
adj. R^2	0.125	0.132	0.133	0.135	0.133	0.135
AIC	15259.399	15239.634	15237.136	15230.769	15236.537	15229.933
Total effect in poor countries		-1.627***	-1.160**	-1.427***	-1.394***	-1.607***
		(0.405)	(0.526)	(0.455)	(0.457)	(0.414)
Total effect in hot countries			-0.794*	-0.621	-1.161**	-0.999*
			(0.416)	(0.395)	(0.501)	(0.503)

Notes:

All specifications include country FE and Region x Time FE.
Poor is a dummy with value 1 for countries with below median GDP per capita in 1960.
Hot is a dummy with value 1 for countries with above median average temperature in the 1960s.
Poor_2 is a dummy with value 1 for countries with below median GDP per capita.
Hot_2 is a dummy with value 1 for countries with above median average temperature.
Temperature change is annual and expressed in degree Celsius.
Precipitation change is annual and expressed is in units of 100 mm per year.
Standard errors are in parentheses and are clustered at the country level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 15

Regressions with number of persons employed and capital stock as dependent variables

	(1)	(2)	(3)	(4)
	Growth rate in the number of persons employed	Growth rate in the number of persons employed	Capital stock growth rate	Capital stock growth rate
△Temp	0.068	0.016	3.610	0.663
	(0.079)	(0.096)	(2.335)	(0.560)
Poor x △Temp		0.172		9.661*
		(0.152)		(5.465)
△Pre	0.004	0.005	-0.174	-0.114
	(0.014)	(0.014)	(0.322)	(0.338)
_cons	2.533***	2.542***	4.106	4.608
	(0.439)	(0.437)	(4.347)	(4.338)
N	2760	2760	2760	2760
R^2	0.171	0.172	0.129	0.132
adj. R^2	0.080	0.081	0.034	0.037
AIC	10968.569	10969.314	27815.168	27808.350
Total effect in poor countries		0.188		10.323*
		(0.126)		(5.866)

Notes:
All specifications include country FE and Region x Time FE.
Poor is a dummy with value 1 for countries with below median GDP per capita in 1960.
Temperature change is annual and expressed in degree Celsius.
Precipitation change is annual and expressed is in units of 100 mm per year.
Standard errors are in parentheses and are clustered at the country level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 16

Baseline specification with Driscoll-Kraay standard errors

Dependent variable: annual TFP growth rate	(1)	(2)	(3)	(4)	(5)	(6)
△Temp	-0.485**	-0.029	0.057	0.098	0.008	0.051
	(0.200)	(0.130)	(0.164)	(0.164)	(0.134)	(0.142)
△Pre	-0.033	-0.042*	-0.047*	-0.048*	-0.049*	-0.051*
	(0.026)	(0.025)	(0.025)	(0.025)	(0.025)	(0.025)
Poor x △Temp		-1.493***	-1.195**		-1.315***	
		(0.432)	(0.566)		(0.448)	
Hot x △Temp			-0.684	-0.612		
			(0.487)	(0.470)		
Poor_2 x △Temp				-1.425**		-1.513***
				(0.565)		(0.443)
Hot_2 x △Temp					-1.048**	-0.979**
					(0.457)	(0.431)
_cons	0.184	0.402*	0.429**	0.480**	0.470**	0.519**
	(0.166)	(0.212)	(0.200)	(0.204)	(0.204)	(0.208)
N	2760	2760	2760	2760	2760	2760
Within R^2	0.208	0.215	0.216	0.217	0.216	0.218
Total effect in poor countries		-1.523***	-1.139*	-1.327**	-1.307***	-1.462***
		(0.462)	(0.660)	(0.644)	(0.487)	(0.472)
Total effect in hot countries			-0.627	-0.515	-1.040**	-0.928**
			(0.383)	(0.390)	(0.412)	(0.383)

Notes:
All specifications include country FE and Region x Time FE.
Poor is a dummy with value 1 for countries with below median GDP per capita in 1960.
Hot is a dummy with value 1 for countries with above median average temperature in the 1960s.
Poor_2 is a dummy with value 1 for countries with below median GDP per capita.
Hot_2 is a dummy with value 1 for countries with above median average temperature.
Temperature change is annual and expressed in degree Celsius.
Precipitation change is annual and expressed is in units of 100 mm per year.
Driscoll-Kraay standard errors are in parentheses, and allow up to two lags of autocorrelation.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 1

TFP levels and average temperatures in 2006

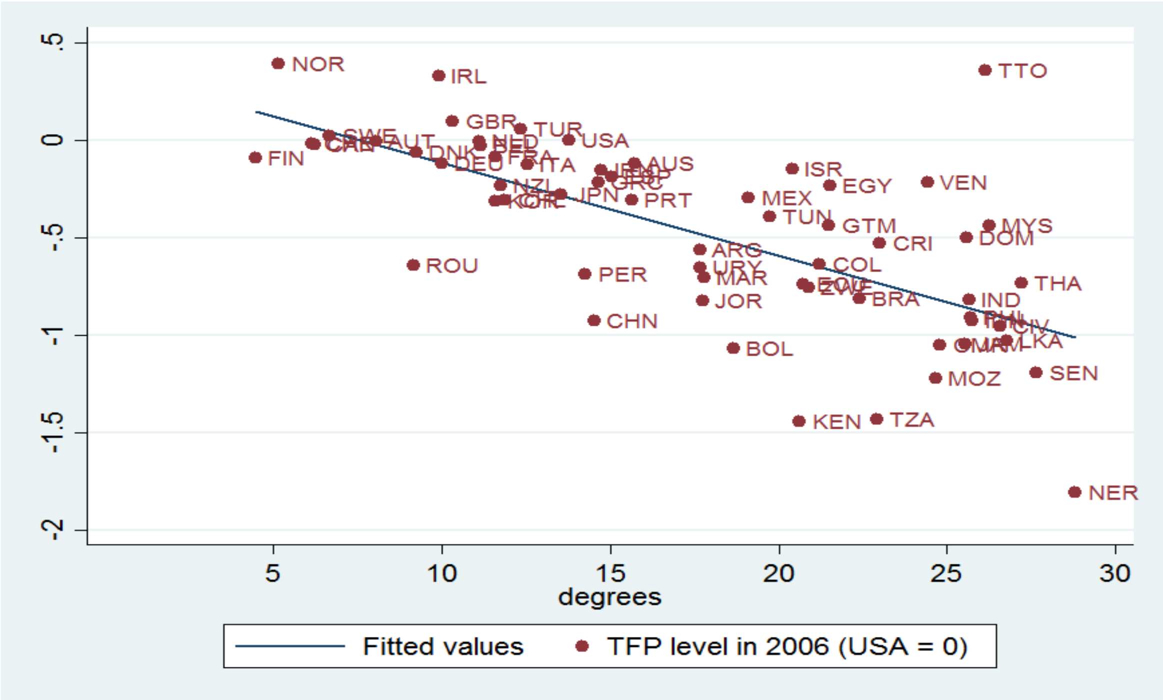

Figure 2

TFP growth rates and temperature shocks in poor countries

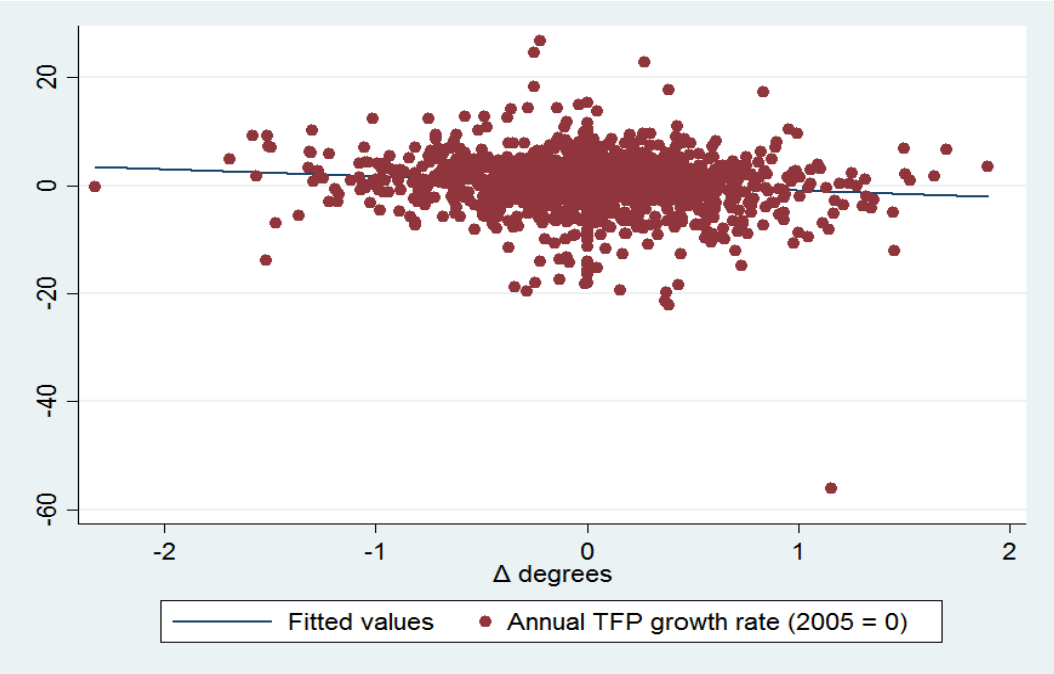

Figure 3

Marginal effect of ΔTemp at different GDP per capita levels – hot = 1

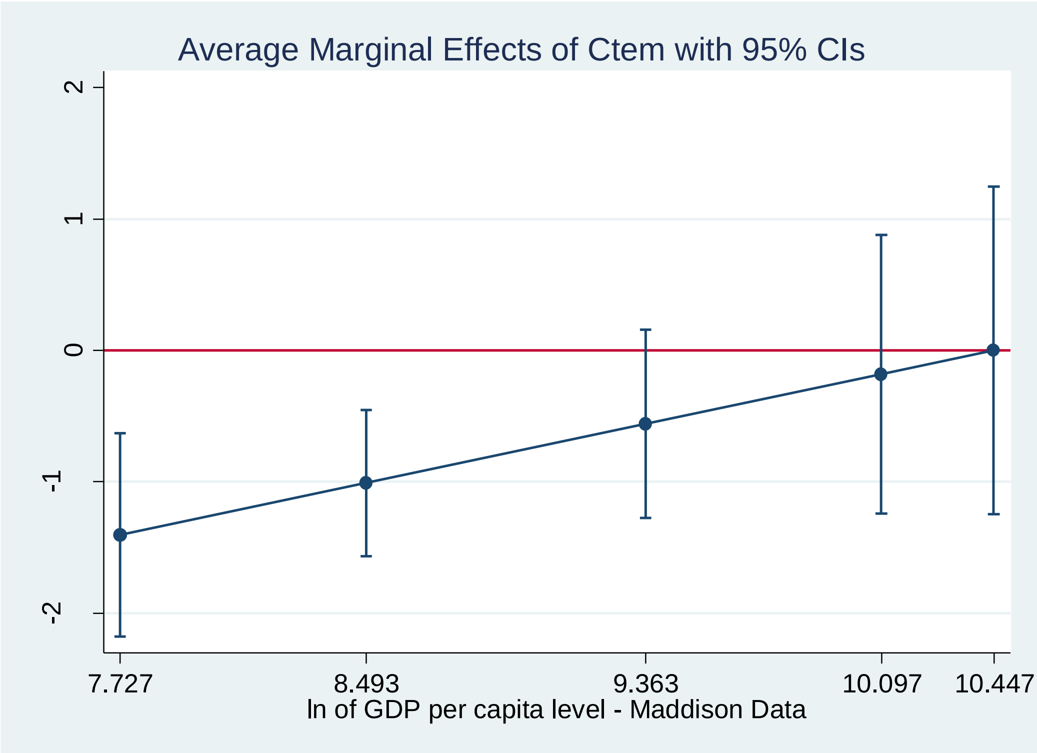

Figure 4

Marginal effect of ΔTemp at different GDP per capita levels – hot = 0

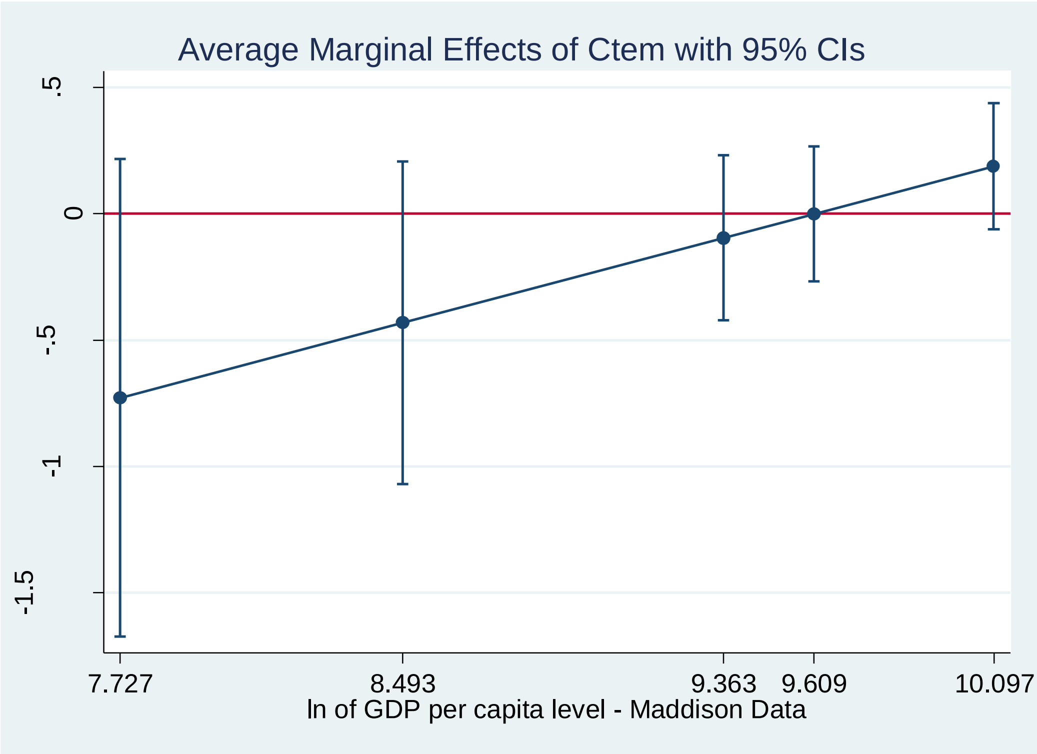

Figure 5

Marginal effect of ΔTemp at different GDP per capita levels – whole sample

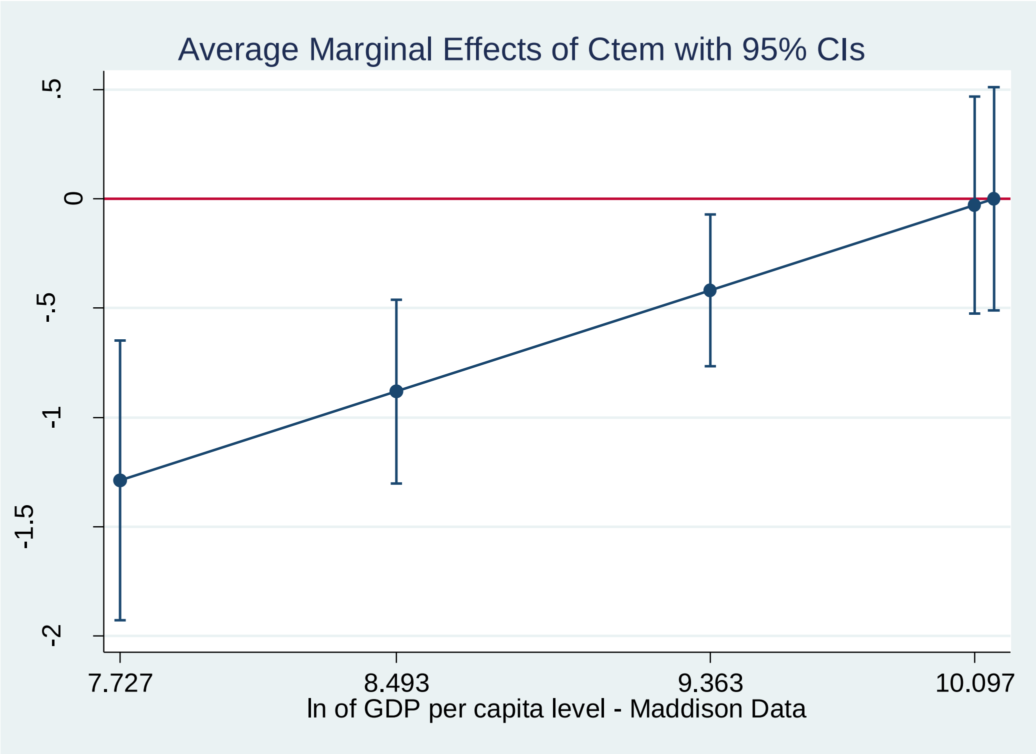

Appendices

(1) Construction of the RTFP^{NA} Index in PWT 8.1

Since version 8.0, the Penn World Tables include data on TFP at the country level (Feenstra, Inklaar & Timmer, 2013). In particular, there are two measures of TFP in PWT 8.1. The first one is CTFP, where the prefix C stays for "current year": this is a measure of TFP levels of countries in a given year compared to the US, whose TFP levels are 1 in each year. It is thus a measure of relative TFP levels which allows for comparisons among countries (Feenstra, Inklaar & Timmer, 2015), and can be seen as an index of technological catch-up or as the distance from the technological frontier (represented by the US).

The other, and the one used in this study is RTFP^{NA}. This index is calculated using the growth rate of real GDP from national accounts data, in conjunction with the growth rates of capital stock at constant national prices and of the labour force (Feenstra, Inklaar & Timmer, 2013). This is a standard process in TFP estimation since, as a residual representing a combination of labour and capital productivity, TFP is obviously dependent on the estimates of the other components. As discussed above, RTFP^{NA} is normalized to 1 in 2005 for all countries, and since we use natural logarithms in our specification, the normalized value for 2005 is 0.

More specifically, Inklaar and Timmer (2013) describe how the productivity measurement starts from the following general production function:

$$Y = Af(K, L) = AK^{\alpha}(Ehc^{1-\alpha}) \tag{A.1}$$

where, in the second equality, labour input is defined as the product of the number of workers in the economy *E* times their average human capital *hc* and α is the output elasticity of capital.

A second-order approximation to the production function f is represented by the Törnqvist quantity index of factor inputs Q^T, which can be used to compare inputs between *t-1* and *t* for a given country as follows:

$$\ln Q_{t,t-1}^{T} = \frac{1}{2}(\alpha_t + \alpha_{t-1})\ln\frac{K_t}{K_{t-1}} + \left[1 - \frac{1}{2}(\alpha_t + \alpha_{t-1})\right]\ln\frac{L_t}{L_{t-1}} \tag{A.2}$$

In order to implement this equation, they make the assumption that the output elasticity of capital is approximated by the country's share of GDP that is not earned by labour. They assume a common labour share neither across countries nor over time, i.e., the input index in equation (A.2) is the more flexible Törnqvist index rather than the more common Cobb-Douglas function.

Finally, growth of productivity over time is given by:

$$RTFP^{NA}_{t,t-1} = \frac{RGDP^{NA}_{t}}{RGDP^{NA}_{t-1}} / Q^{T}_{t,t-1} \tag{A.3}$$

where $RGDP^{NA}$ stands for real GDP at constant national prices.

For further information with regard to the construction of the RTFP^{NA} index, see Feenstra, Inklaar and Timmer (2013), Feenstra, Inklaar and Timmer (2015) and Inklaar and Timmer (2013).

(2) Statistical tests

A. Panel unit root tests

In order to check that our main variables are stationary, we performed panel unit root tests for annual TFP growth, annual temperature change and annual precipitation change. In particular, we used two unit root tests which are both[1] fit when N > T, as it is the case in our sample: the Im, Pesaran, and Shin (2003) test and the Harris and Tzavalis (1999) test. The results, reported in Tables A.1-A.6, confirm that the tested variables are stationary.

Table A.1

Im-Pesaran-Shin unit-root test for annual TFP growth

Ho: All panels contain unit roots	Number of panels = 60
Ha: Some panels are stationary	Number of periods = 46
AR parameter: Panel-specific	Asymptotics: T,N →Infinity sequentially
Panel means: Included	
Time trend: Included	Cross-sectional means removed
ADF regressions: No lags included	

			Fixed-N exact critical values		
	Statistic	p-value	1%	5%	10%
t-bar	-5.8532		-2.360	-2.310	-2.280
t-tilde-bar	-4.3796				
Z-t-tilde-bar	-27.9582	0.0000			

[1] The Im-Pesaran-Shit (2003) test is fit when N > T if a time trend is included.

Table A.2

Im-Pesaran-Shin unit-root test for ∆Temp

Ho: All panels contain unit roots	Number of panels = 60
Ha: Some panels are stationary	Number of periods = 46
AR parameter: Panel-specific	Asymptotics: T,N→Infinity
Panel means: Included	sequentially
Time trend: Included	Cross-sectional means removed
ADF regressions: No lags included	

	Statistic	p-value	Fixed-N exact critical values 1%	5%	10%
t-bar	-9.7663		-2.360	-2.310	-2.280
t-tilde-bar	-5.4829				
Z-t-tilde-bar	-38.5654	0.0000			

Table A.3

Im-Pesaran-Shin unit-root test for ∆Pre

Ho: All panels contain unit roots	Number of panels = 60
Ha: Some panels are stationary	Number of periods = 46
AR parameter: Panel-specific	Asymptotics: T,N →Infinity sequentially
Panel means: Included	
Time trend: Included	Cross-sectional means removed
ADF regressions: No lags included	

			Fixed-N exact critical values		
	Statistic	p-value	1%	5%	10%
t-bar	-10.2704		-2.360	-2.310	-2.280
t-tilde-bar	-5.5661				
Z-t-tilde-bar	-39.3644	0.0000			

Table A.4

Harris-Tzavalis unit-root test for annual TFP growth

Ho: Panels contain unit roots	Number of panels = 60
Ha: Panels are stationary	Number of periods = 46
AR parameter: Common	Asymptotics: N → Infinity,
Panel means: Included	T fixed
Time trend: Included	Cross-sectional means removed

	Statistic	z	p-value
rho	0.1823	-50.3642	0.0000

Table A.5

Harris-Tzavalis unit-root test for ∆Temp

Ho: Panels contain unit roots	Number of panels = 60
Ha: Panels are stationary	Number of periods = 46
AR parameter: Common	Asymptotics: N $\rightarrow$ Infinity,
Panel means: Included	T fixed
Time trend: Included	Cross-sectional means removed

	Statistic	z	p-value
rho	-0.3900	-93.9377	0.0000

Table A.6

Harris-Tzavalis unit-root test for ∆Pre

Ho: Panels contain unit roots	Number of panels = 60
Ha: Panels are stationary	Number of periods = 46
AR parameter: Common	Asymptotics: N → Infinity, T fixed
Panel means: Included	
Time trend: Included	Cross-sectional means removed

	Statistic	z	p-value
rho	-0.4215	-96.3329	0.0000

B. FE vs RE

To test the appropriateness of a fixed effects - panel approach rather than a random effects (RE) specification, we performed a test using the approach suggested by Mundlak (1978). The traditional Hausman test, in fact, is not recommended when time fixed effects are included in the regressions, and is based on the assumption of homoskedasticity, which is very unlikely to hold in our sample. The Mundlak test, in contrast, allows for heteroskedastic errors and serial intracorrelation.

Essentially, we performed a RE regression including panel-level means of our time-varying variables – in the specification we used, temperature change, precipitation change and the interaction between and the poor dummy – and then tested for the joint significance of the coefficients for the means time varying variables. The results, reported in Table A.7, are strongly in favour of a FE approach.

Table A.7

Mundlak test – Random Effects GLS regression with added panel-level means

Dependent variable: Annual TFP growth rate	(1)
△Temp	-0.029
	(0.136)
Poor x △Temp	-1.493***
	(0.405)
△Pre	-0.042*
	(0.023)
Mean_△Temp	-3.160
	(6.413)
Mean_ Poor x △Temp	29.436***
	(10.984)
Mean_△Pre	2.456**
	(1.138)
_cons	2.575***
	(0.723)
N	2760
R^2	0.226

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Test on the joint significance of the panel-level means for the time varying variables:

(1) Mean_△Temp = 0
(2) Mean_△Pre = 0
(3) Mean_ Poor x △Temp

chi2(3) = 9.40
Prob > chi2 = 0.0245

(3) Additional results

Table A.8

CRU Weather Data – Descriptive Statistics

	Mean	Var	sd	Min	Max	Obs
△Temp_2	0.0147	0.303	0.550	-3	2.700	2760
△Pre_2	0.00164	4.877	2.208	-16.36	16.61	2760

Temperature change is annual and expressed in degree Celsius.
Precipitation change is annual and expressed in units of 100 mm per year.

Table A.9

Relationship between annual TFP growth rates and temperature changes - Sample B & CRU Data

Dependent variable: Annual TFP growth rate	(1)	(2)	(3)	(4)	(5)	(6)
△Temp_2	-0.155	0.141	0.139	0.118	0.112	0.104
	(0.158)	(0.101)	(0.106)	(0.105)	(0.103)	(0.104)
△Pre_2	-0.043	-0.053*	-0.052*	-0.053*	-0.050*	-0.051*
	(0.029)	(0.027)	(0.028)	(0.027)	(0.028)	(0.028)
Poor x △Temp_2		-1.001***	-1.005***		-1.010***	
		(0.279)	(0.279)		(0.278)	
Hot x △Temp_2			0.014	0.109		
			(0.321)	(0.315)		
Poor_2 x △Temp_2				-1.126***		-1.108***
				(0.273)		(0.278)
Hot_2 x △Temp_2					0.322	0.368
					(0.312)	(0.311)
_cons	1.344***	1.310***	1.310***	1.313***	1.311***	1.310***
	(0.304)	(0.304)	(0.305)	(0.305)	(0.305)	(0.306)
N	2448	2448	2448	2448	2448	2448
R^2	0.197	0.200	0.200	0.200	0.200	0.200
adj. R^2	0.121	0.123	0.123	0.124	0.123	0.124
AIC	13479.345	13472.425	13474.424	13472.862	13473.928	13472.283
Total effect in poor countries		-0.860***	-0.866***	-1.007***	-0.898***	-1.004***
		(0.301)	(0.306)	(0.300)	(0.304)	(0.307)
Total effect in hot countries			0.153	0.228	0.434	0.472
			(0.302)	(0.299)	(0.305)	(0.304)

Notes:

All specifications include country FE and Region x Time FE.
Poor is a dummy with value 1 for countries with below median GDP per capita in 1970.
Hot is a dummy with value 1 for countries with above median average temperature in the 1970s.
Poor_2 is a dummy with value 1 for countries with below median GDP per capita.
Hot_2 is a dummy with value 1 for countries with above median average temperature.
Temperature change is annual and expressed in degree Celsius.
Precipitation change is annual and expressed in units of 100 mm per year.
Standard errors are in parentheses and are clustered at the country level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.10

Relationship between annual labour productivity growth rates and temperature changes – Sample B

Dependent variable: annual labour productivity growth rate	(1)	(2)	(3)	(4)	(5)	(6)
ΔTemp	-0.382*	0.115	0.106	0.064	0.104	0.076
	(0.215)	(0.133)	(0.134)	(0.135)	(0.131)	(0.136)
ΔPre	-0.042*	-0.053**	-0.053***	-0.055***	-0.052**	-0.055**
	(0.021)	(0.020)	(0.020)	(0.020)	(0.021)	(0.021)
Poor x ΔTemp		-1.498***	-1.519***		-1.499***	
		(0.377)	(0.409)		(0.377)	
Hot x ΔTemp			0.064	0.195		
			(0.429)	(0.417)		
Poor_2 x ΔTemp				-1.655***		-1.588***
				(0.403)		(0.380)
Hot_2 x ΔTemp					0.117	0.173
					(0.398)	(0.398)
_cons	2.512***	2.428***	2.431***	2.434***	2.430***	2.427***
	(0.480)	(0.475)	(0.475)	(0.477)	(0.474)	(0.476)
N	2448	2448	2448	2448	2448	2448
R^2	0.181	0.186	0.186	0.187	0.186	0.187
adj. R^2	0.103	0.109	0.109	0.109	0.109	0.109
AIC	14034.284	14018.674	14020.654	14018.717	14020.617	14018.779
Total effect in poor countries		-1.383***	-1.413***	-1.591***	-1.396***	-1.512***
		(0.399)	(0.444)	(0.425)	(0.408)	(0.407)
Total effect in hot countries			0.169	0.259	0.221	0.249
			(0.414)	(0.415)	(0.403)	(0.406)

Notes:

All specifications include country FE and Region x Time FE.
Poor is a dummy with value 1 for countries with below median GDP per capita in 1970.
Hot is a dummy with value 1 for countries with above median average temperature in the 1970s.
Poor_2 is a dummy with value 1 for countries with below median GDP per capita.
Hot_2 is a dummy with value 1 for countries with above median average temperature.
Temperature change is annual and expressed in degree Celsius.
Precipitation change is annual and expressed is in units of 100 mm per year.
Standard errors are in parentheses and are clustered at the country level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.11

Baseline specification with Driscoll-Kraay standard errors – Sample B

Dependent variable: annual TFP growth rate	(1)	(2)	(3)	(4)	(5)	(6)
∆Temp	-0.345***	0.053	0.087	0.071	0.077	0.075
	(0.128)	(0.172)	(0.150)	(0.138)	(0.154)	(0.141)
∆Pre	-0.033	-0.041*	-0.043**	-0.046**	-0.043**	-0.046**
	(0.020)	(0.022)	(0.021)	(0.021)	(0.020)	(0.022)
Poor x ∆Temp		-1.200***	-1.125**		-1.198***	
		(0.354)	(0.479)		(0.360)	
Hot x ∆Temp			-0.230	-0.093		
			(0.522)	(0.564)		
Poor_2 x ∆Temp				-1.308**		-1.337***
				(0.568)		(0.392)
Hot_2 x ∆Temp					-0.244	-0.196
					(0.724)	(0.710)
_cons	-0.543***	0.158	0.127	0.747***	-0.037	0.796***
	(0.084)	(0.120)	(0.135)	(0.137)	(0.064)	(0.115)
N	2448	2448	2448	2448	2448	2448
Within R^2	0.198	0.202	0.202	0.203	0.202	0.203
Total effect in poor countries		-1.147***	-1.037**	-1.237**	-1.121***	-1.262***
		(0.281)	(0.451)	(0.546)	(0.335)	(0.401)
Total effect in hot countries			-0.143	-0.022	-0.166	-0.120
			(0.552)	(0.589)	(0.825)	(0.765)

Notes:

All specifications include country FE and Region x Time FE.
Poor is a dummy with value 1 for countries with below median GDP per capita in 1970.
Hot is a dummy with value 1 for countries with above median average temperature in the 1970s.
Poor_2 is a dummy with value 1 for countries with below median GDP per capita.
Hot_2 is a dummy with value 1 for countries with above median average temperature.
Temperature change is annual and expressed in degree Celsius.
Precipitation change is annual and expressed is in units of 100 mm per year.
Driscoll-Kraay standard errors are in parentheses, and allow up to two lags of autocorrelation.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

(4) List of countries in the sample

Main Dataset:

Argentina

Australia

Austria

Belgium

Bolivia

Brazil

Cameroon

Canada

Chile

China

Colombia

Costa Rica

Denmark

Dominican Republic

Ecuador

Egypt

Finland

France

Germany

Greece

Guatemala

India

Indonesia

Iran

Ireland

Israel

Italy

Ivory Coast

Jamaica

Japan

Jordan

Kenya

Malaysia
Mexico
Morocco
Mozambique
Netherlands
New Zealand
Niger
Norway
Peru
Philippines
Portugal
Romania
Senegal
South Korea
Spain
Sri Lanka
Sweden
Switzerland
Tanzania
Thailand
Trinidad & Tobago
Tunisia
Turkey
United Kingdom
United States
Uruguay
Venezuela
Zimbabwe

Countries added in Sample B

Bulgaria
Hungary
Kuwait
Panama
Paraguay

Poland

Qatar

Saudi Arabia

List of regions

Eastern Europe and Central Asia

Latin America and the Caribbean

Middle East and North Africa

South and East Asia and the Pacific

Sub-Saharan Africa

Western Europe and offshoots

Essay 2

Temperature shocks, growth and poverty thresholds: evidence from rural Tanzania

Abstract

Growing interest in the impacts of climate change in poor countries has sparked attention on the relationship between temperature and micro growth dynamics. Using the LSMS-ISA Tanzania National Panel Survey by the World Bank, we study the relationship between rural household consumption growth and temperature shocks over the period 2008 – 2013.

The main finding is a sharp heterogeneity: temperature shocks have a negative and significant impact on household growth only if their initial consumption lies below a critical threshold, i.e. they slow the convergence process among households. The main transmission channels appear to be agricultural yields and labour productivity. Extrapolating from weather to climate, these findings support the Schelling Conjecture: economic development would be the best solution to cope with climate change for poor farming households in rural areas, and closing the yield gap and modernizing agriculture remains crucial for adaptation to the negative impacts of global warming in Sub-Saharan Africa.

1 Introduction

Poorer countries are generally found to be more vulnerable to climate change and weather variability, but research is concentrated in richer countries. Many would suspect that poorer people are more vulnerable too, but research is scarce. We here estimate the impact of weather shocks in a poor country, Tanzania, and differentiate that impact across the income spectrum.

We use the empirical tools and models of development economics to examine the link between short-term household welfare dynamics and temperature shocks in rural Tanzania, to provide an empirical answer to the following questions: what is the micro relationship between temperature shocks, poverty and economic growth? Is the idea of a climate-induced poverty trap plausible to describe the growth dynamics of farmer households in a rural developing context?

The justification for such an exercise stems from the lack of within-country works on the relationship between temperature and economic growth, as emphasized by Tol (2016): "The pattern of vulnerability that is seen between countries, is likely to hold within countries as well. This would strengthen the worries about climate change, but there has hardly been any research on the quantification of the intra-country distributional implications of the impacts of climate change."

This article thus speaks to two well distinct strands of research: the development literature on poverty traps, that investigates the issues of poverty persistence, growth divergence and multiple equilibria; and the emerging climate-economy literature that studies short-run elasticities of weather shocks impacts on growth to infer about future impacts of climate change.

Tanzania is an appropriate setting for such a study for a number of reasons. It is by now commonly accepted that the future impacts of climate change will disproportionately affect poorer and hotter countries (Tol, 2015), and especially people living in rural, remote and scarcely populated areas, whose main source of income is agriculture. Sub-Saharan Africa, in particular, has been identified as one of the most vulnerable parts of the world to the threats posed by climate change (IPCC, 2014). Tanzania is a poor and hot Sub-Saharan country, where in 2015 68% of the population lived in rural areas[1]. It is constantly classified as a country under high risk from the impacts of future climate change: temperatures in the country are predicted to rise 2–4 °C by 2100, "with warming more concentrated during the dry season and in the interior regions of the country" (Rowhani, Lobell, Linderman & Ramankutty, 2011). Ahmed et al. (2011) underline the importance of agriculture for

[1] http://data.worldbank.org/indicator/SP.RUR.TOTL.ZS?locations=TZ

the Tanzanian economy: "The importance of agriculture to the poor is particularly true for Tanzania, where agriculture accounts for about half of gross production, and employs about 80 percent of the labour force. Agriculture in Tanzania is also primarily rain-fed, with only two percent of arable land having irrigation facilities—far below the potentially irrigable share".

Tanzania is also a country which exhibits quite large climatic diversity, as noted by Rowhani et al. (2011): "on the Indian Ocean, the United Republic of Tanzania possesses a complex landscape, formed by the western and eastern branches of the East African Rift, resulting in substantial spatial variability in climate within the nation. The country's climate varies from tropical at the coast to temperate in the highlands".

Last but not least, data availability: we use the Living Standard Measurement Survey (LSMS) – Integrated Survey on Agriculture (ISA) Tanzania National Panel Survey by the World Bank, a three-wave household longitudinal dataset covering the period 2008 – 2013.

We employ a micro-growth model borrowed from the standard growth literature, and test for convergence among households and for the significance of weather shocks as determinants of growth, while controlling for heterogeneity. Then, we test for the presence of consumption thresholds with regard to the impacts of temperature shocks. Finally, guided by previous theoretical and empirical literature, we test for a set of transmission channels that can justify potential heterogeneity of impacts and explain the lack of consumption smoothing behaviours, namely: health expenditure, labour productivity, average crop yields and asset growth.

However, let us be clear at the outset.

Obviously, given the short-run nature of this dataset, our capacity to assess convergence is limited, and we can only cautiously infer about long-run convergence. Also, we do not directly test for the presence of multiple equilibria and hence for the existence of a poverty trap. Under a classic 'poverty trap' threshold, households are trapped in an equilibrium with permanently low income, while here we only check whether there is a consumption threshold above which temperature impacts turn insignificant, i.e. whether impacts disappear as households grow richer. Deceleration is not bifurcation, as noted by Dercon (2004) and Jalan and Ravallion (2002).

What emerges is a sharp and striking heterogeneity: temperature-induced consumption shocks only affect the poorest households. The observed growth of rural households suffers from a negative and significant contemporaneous impact of temperature shocks only if their initial consumption level lies below a critical threshold. In other words, positive temperature shocks slow convergence among households, and enhance inequalities. The main transmission channels responsible for this heterogeneity appear to be agricultural yields and labour productivity, due to the fact that agricultural yields, technologies and assets, as well as income sources, differ substantially across consumption quartiles. Additionally, no impact on asset growth is found, suggesting asset smoothing is probably taking place and that poorest households choose to voluntarily destabilize their consumption in order not to sell their assets, or that more simply they do not have enough assets to sell to cope with the income reduction caused by temperature shocks.

The contributions of this essay are the following. First, it complements aggregate growth - climate empirics with available micro panel data, by providing evidence on the (short-run) micro causal relationship between weather anomalies, poverty and growth. Second, it links the weather-economic growth literature with the development literature on poverty traps, by applying the tools and models of the latter to the research questions of the former. Third, it contributes to the development literature, by testing for consumption *vs* asset smoothing, which has been rarely been done according to Carter and Lybbert (2012)[2]; and by showing that, when controlling for temperature shocks (often ignored in development literature), precipitation impacts are insignificant and close to zero.

The rest of this essay is arranged as follows. Section 2 reviews the relevant literature. Section 3 illustrates the empirical framework and the identification strategy. Section 4 describes data and provides introductory descriptive statistics. Section 5 shows and comments the results of the empirical analysis. Section 6 conducts a host of robustness checks. Section 7 investigates the channels of the heterogeneity of impacts. Section 8 wraps up, illustrates the policy implications of the analysis with regard to climate change, reminds some *caveats* and concludes.

2 Literature review

The recent and growing body of empirical works focusing on the climate-economy relationship and its channels stems from the interest to try to understand and quantify the future impacts of climate change on human welfare.

Dell, Jones and Olken (2014) review this literature and notice how old cross-sectional works (Dell, Jones, & Olken, 2009; Gallup, Sachs, & Mellinger, 1999; Nordhaus, 2006), whose validity is challenged by the risk of endogeneity and omitted variable bias, have recently been replaced by more appropriate and robust panel methods, both at the macro (Bansal & Ochoa, 2011; Burke, Hsiang, &

[2] "Unfortunately, much of the empirical literature has not tested consumption smoothing against a theoretically well-defined alternative".

Miguel, 2015; Dell, Jones, & Olken, 2012; Hsiang, 2010; Hsiang & Jina, 2014) and micro (Cachon, Gallino, & Olivares, 2012; Cachon et al., 2012; Graff Zivin & Neidell, 2014; Heal & Park, 2015; Niemelä, Hannula, Rautio, Reijula, & Railio, 2002; Sudarshan & Tewari, 2013) level, which isolate the exogenous effect of weather variables on the economic outcome of interest.

The main findings of this emerging literature are that weather affects economic activity and growth through a wide range of channels (agriculture, health and mortality, labour productivity, energy, conflict and political stability, among the others) and that these impacts are substantially bigger and significant in poor countries.

These panel estimates have then been employed and calibrated *ad hoc* in simulation studies on the impacts of future climate change (Lemoine & Kapnick, 2015; Moore & Diaz, 2015) to provide empirically-grounded impact estimates to be used in Integrated Assessment Models (IAMs), and overcome the critiques about the arbitrary choice of crucial parameters like the damage function and climate sensitivity (Pindyck, 2012, 2013; Stern, 2013; Weitzman, 2009, 2010).

However, these works only estimate short-run elasticities, whereas climate change is by definition a long-run phenomenon, which cannot be captured by empirical works because of the intrinsic difference between 'climate' and 'weather': a 1°C shock in a given year and place is not the same of a permanent 1°C global increase.

Furthermore, short-run panels can capture neither the possibility of adaptation, since responses from economic agents to year-to-year variations may well be different from responses to gradual and long-run changes, nor potential intensifying effects from phenomena which lie outside the range of historical experience (massive sea level rise, a thermohaline circulation slowdown, the release of methane from melting permafrost, etc.), both of which could eventually take place in the long-run and drastically change not just the magnitude, but even the nature of the current short-run elasticities.

On the other hand, as noted by Dell, Jones and Olken (2014), since climate change is not a sudden shock, but a stochastic warming process along an upward trend, recent historical experience is an appropriate setting to study warming effects. Furthermore, the use of econometric techniques such as, for instance, the inclusion of lags in the regressions, long differences and the analysis of persistence of impacts of past shocks, can alleviate these concerns and partially fill the conceptual gap between short- and long-run impacts.

Thus, despite the external validity issue, it is reasonable to assume that this empirical literature can be informative about the structure of the damage function for climate change.

In the development literature, however, the econometric identification of the impacts of weather shocks (especially rainfall) on human welfare is standard practice. Weather shocks, in fact, represent the exogenous shock *par excellence*, and many development studies have employed them both as independent regressors, in the context of the search for the determinants of growth, and as instruments for the estimation of impacts of other variables of interest. In her pioneering work, Paxson (1992) used rainfall shocks to construct estimates of transitory income, and found that unexpected income shocks did not have serious welfare consequences for Thai farm households, because they used savings and dissavings to buffer consumption from income shocks. The partial insurance strategies adopted by poor farmers against a temporary shock could indeed imply a reduction in crop yields with potentially negative impacts on consumption growth. Indeed, there is empirical evidence that self-insurance mechanisms only partially succeed (Morduch, 1995; Townsend, 1995), that households might not be able to smooth their consumption in response to income fluctuations due to credit or liquidity constraints (Hirvonen, 2016; Morduch, 1995; Rosenzweig & Wolpin, 1993), and that uninsured risk may well be a cause of poverty, due to two distinct mechanisms, one *ex ante* or behavioural and one *ex post* (Dercon, 2004). The first can be explained as follows: since poorer farmers are generally risk-averse, uninsured risk determines *ex-ante* changing in behaviour that implies precautionary saving and/or other optimal strategies to avoid profitable but risky opportunities at the expenses of mean returns (Dercon, 1996, 2004; Elbers, Gunning, & Kinsey, 2007). To this end, Ligon and Schechter (2003) provide a micro-founded household level vulnerability measure by applying the so-called Jensen inequality. Dercon (1996), analysing, through a theoretical model of risk-taking behaviours, the relationship between risk, crop choice and savings in rural Tanzania, finds that wealthier households engage in more risky but higher return activities than households with a poor asset base, and notes that "this evidence is to some extent disconcerting if one is interested in rural growth without exacerbating rural inequality, since it shows the existence and the mechanisms of a 'poverty trap'". The *ex post* impact, instead, is the one that materializes after a 'bad' state (Dercon, 2004): in this respect weather shocks are shown to have an impact on *ex-post* poverty too. In such a context, several theoretical models underline the issues of persistence to highlight that temporary shocks can affect long-term outcomes such as the process of income convergence among households (Carter, Little, Mogues & Negatu; Little, Mogues & Negatu, 2007; Reis, 2009). This permanent effect of temporary shocks has been typically explained by asset smoothing (Barrett & Carter, 2013; Carter & Barrett, 2006; Carter Little, Mogues & Negatu, 2007) or by the conservative behaviour of risk-averse households that shy away from investing in profitable but risky technologies (Reis, 2009).

Indeed, this is what has emerged from many empirical studies on household welfare dynamics: Fafchamps, Udry and Czukas (1998), using panel data for farming households in Burkina Faso, test the hypothesis that households keep livestock as a buffer stock to insulate their consumption from income fluctuations, but only find evidence for very limited consumption smoothing. Dercon (2004) himself, using panel data from Ethiopia during the period 1989 – 1997, finds that rainfall shocks had a substantial contemporaneous impact on food consumption growth, and also shows persistence of impacts, suggesting that rainfall shocks may have a long-lasting effect which goes beyond the welfare cost of short-term consumption fluctuations. His subsequent works in the same setting confirmed these results (Dercon & Christiaensen, 2011; Dercon, Hoddinott, & Woldehanna, 2005; Dercon & Krishnan, 2000). Carter et al. (2007) explore the asset dynamics of Ethiopian and Honduran households in the wake of environmental shocks, and find that household growth can be hit not just in the immediate aftermaths but also in the long-run, and that coping strategies adopted are costly and can be a source of divergence among households. Hirvonen (2016), using the Kagera Health and Development Survey (KHDS), spanning the period 1991-2009, shows how household consumption co-moves with temperature, and then uses temperature shocks as a proxy for income shocks to study long-term migration decisions in Tanzania.

Other studies have instead focused on the possibility of long-run impacts on household welfare from weather shocks. Hoddinott and Kinsey (2001) first, reviewing literature on household responses to weather-related shocks, note how what emerges is that "[...] some, but not all households can smooth consumption. In particular, households facing liquidity constraints have limited smoothing ability. For these households, therefore, income fluctuations will generate a welfare loss". Then, drawing on a panel dataset in Zimbabwe, they try to determine whether these shocks have only transitory or also permanent effects, by examining growth in the heights of young children. They discover droughts have a long-lasting impact on child growth, and that this impact is heterogeneous, i.e. greatest amongst children living in poor households. They notice how this points to the possibility of the intergenerational transmission of poorer health status resulting from drought shocks. Alderman, Hoddinott and Kinsey (2006) follow this path and explore the long-term consequences of shocks on individuals, starting from the observation that where temporary shocks have long-lasting impacts, utility losses may be higher, and finding analogous results.

The amount of evidence of both short-run and long-run impacts of weather shocks on household welfare, and the limited evidence for precautionary saving and consumption smoothing, has been the spark for the development of another strand of literature, based on the concept of "poverty traps".
The concept of poverty traps has been proposed both in macro- as well as in microeconomics and is closely related to the idea of convergence in neoclassical economics. The assumption of diminishing returns is a crucial one in neoclassical economic growth: essentially, it implies that the incomes of poorer countries (households) will eventually 'catch up' over time with those of richer countries (households). But, following empirical evidence on macro growth which contradicted the assumed convergence hypothesis between countries, as Carter and Barrett (2006) describe, "within the macro

growth literature, two alternatives to the neoclassical growth model have emerged to account for the observed pattern of divergence", namely the idea of club convergence (Baumol, 1986; De Long, 1988; Quah, 1996, 1997) and the concepts of thresholds and multiple equilibria (Azariadis & Drazen, 1990; Hansen, 2000; Murphy, Shleifer, & Vishny, 1989).

At the micro level, as Carter and Barrett (2006) argue, it may be that "As with nations, individuals may also have intrinsic characteristics (skills, savings propensities, discount rates, and geographic locations) that condition their desired level of accumulation and ultimate equilibrium level of well-being. However, there may also be analogues to the locally increasing returns to scale that generate multiple equilibria and thwart the ability of initially poor households to catch up and converge with their wealthier neighbours".

Starting from this hypothesis, an empirical literature has developed to try and detect the presence of thresholds and multiple equilibria at the micro level. The task is hard, as noticed by Barrett and Carter (2013), Carter and Barrett (2006) and Jalan and Ravallion (2002), due to the lack of sufficiently long panels at the household level in developing countries, which contrasts with the fact that convergence among households, as well as post-shock recovery, are long-run processes.

While it is thus difficult to empirically detect the presence of multiple equilibria, several studies have attempted to do so, and have provided evidence of at least significant persistency of poverty. These works can be divided in two categories. The first has focused on income and consumption growth as indicators of household welfare (Dercon, 2004; Jalan & Ravallion, 2002, 2004). Dercon (2004) only tests for, and discovers, persistence of shocks, but he cannot assert the existence of a poverty trap, as he explicitly states: "This is not the same as testing for the existence of a 'poverty trap' in the sense of the investigation of the threshold, below which there is a tendency to be trapped in permanently low income, from which no escape is possible except for by large positive shocks. Persistence within the time period of the data does not exclude permanent effects, but does not imply them either". Jalan and Ravallion (2002; 2004) draw from the standard growth literature to derive micro-based growth models and explicitly test for divergence due to spatial factors and geographic externalities, finding evidence which supports the notion of "geographic poverty traps", i.e. the idea that, *ceteris paribus*, the welfare of a household living in a well-endowed area grows while the one of an otherwise identical household living in an unfavourable geographic area stagnates.

The other, the so-called 'asset-based' approach, taking cue from the theoretical underpinnings provided by Barrett and Carter (2006; 2013), focuses on asset growth as the dependent variable of interest, arguing that looking at assets makes it possible to distinguish persistent structural poverty from poverty that passes naturally with time thanks to the growth process. This second empirical current is mainly represented by the works of Carter et al. (2007), who show that the idea of asset-based poverty traps is consistent with the post-shock growth experience in Honduras after Hurricane

Mitch, and in Ethiopia after the drought of the late 1990s, while also providing empirical support for the concept of “asset smoothing” (opposed to the hypothesis of consumption smoothing), according to which poorer households with very low assets (typically, livestock), choose to voluntarily destabilize consumption not to sell assets and be caught in a poverty trap from which it would be almost impossible to recover; Carter and Lybbert (2012), who test the two alternative hypothesis of consumption and asset smoothing, and using a panel dataset from West Africa they apply threshold estimation techniques which provide support for asset, and not consumption, smoothing in response to external shocks; Barrett et al. (2006), who examine welfare dynamics in rural Kenya and Madagascar and again, mixing quantitative and qualitative evidence, find that poor households defend their critical asset levels through asset smoothing, even if this comes at the cost of an immediate reduction in consumption.
Finally, Barrett and Swallow (2006) try to unify macro and micro literature on poverty traps by providing the theoretical framework of “fractal poverty traps”, in which multiple dynamic equilibria, caused by endogenous and / or exogenous conditions, exist simultaneously at multiple scales (micro, meso and macro) and are self-reinforcing through feedback effects.

The idea of poverty traps has also been proposed and tested for in the context of the debate on the long-run determinants of growth and development. The two main currents in this literature are the geography hypothesis, which draws from the hypothesis of environmental determinism put forward in Diamond (1999) and Huntington (1922), namely that climate and geography are the fundamental drivers of development, and has found qualified empirical support in the works of Alsan (2014), Andersen, Dalgaard, and Selaya (2016), Gallup, Sachs and Mellinger (1999), and Olsson and Hibbs (2005); and the institution hypothesis (Acemoglu, Johnson, & Robinson, 2000, 2001; Easterly & Levine, 2003; Rodrik, Subramanian, & Trebbi, 2004), which conversely endorses institutional determinism and stresses the primacy of institutions over geography as a determinant of long-run growth. As Dell, Jones and Olken (2014) observe, the fact that geographic characteristics and institutional quality are highly correlated makes it challenging to definitely settle the debate. In this context, Bloom, Canning and Sevilla (2003), Bonds, Keenan, Rohani, and Sachs (2010), and Strulik (2008) provide both theoretical underpinnings and empirical evidence for the idea of ‘climate-induced’ poverty traps, while Tol (2011) explores the long-run mechanisms (diseases, infant mortality, fertility, education) through which climate and climate change could widen or deepen poverty traps or even cause intergenerational poverty traps. This large body of literature notwithstanding, Tol (2015) notes: “The literature on the impact of climate (change) on development has yet to reach firm conclusions. Climate change could moderate the rate of economic growth, but estimates range from high to low. More people may be trapped in poverty because of climate, but this effect could be large or small."

3 Empirical framework and identification strategy

Our empirical framework belongs to the strand of the literature that looks at growth in developing countries by using micro-level data, drawing in particular on the works of Carter et al. (2007); Dercon (2004); Jalan and Ravallion (2002). We assess convergence by using a standard empirical growth model, in a framework borrowed from the macro literature Barro & Sala-i-Martin, 1990; Solow, 1956), where growth rates are assumed to be negatively related to the initial income levels:

$$(1) \quad ln\, Y_{it} - ln\, Y_{it-1} = \alpha ln\, Y_{it-1} + \beta \Delta Temp_{gt} + \gamma \Delta Pre_{gt} + \Omega Z_{it} + \omega X_{it} + \mu_i + q_{it} + w_t + \theta_{rt} + \varepsilon_{it}$$

In this equation, the left-hand side variable is the annualised growth rate in annual household per adult-equivalent[3] consumption between t and t-1, and $ln\, Y_{it-1}$ is household per adult-equivalent lagged consumption[4]. The coefficient α, if negative and statistically significant, would indicate, on average, convergence among households.

In all our specifications, Y_{it} either denotes food consumption or total consumption.

The reason why we use two different dependent variables is that looking only at food consumption growth one may confound the impact of weather shocks with the effects of relative price changes. In fact, due to changes in the ratio between food *vs* non-food prices, food consumption may follow a different growth path from total consumption. While Dercon (2004), due to lack of data availability for non-food expenditure, had to largely limit his analysis to food consumption growth, we employ both to address this concern.

The inclusion of lagged consumption level as an independent regressor may raise concerns about endogeneity. However, endogeneity tests, based on the difference of two Sargan-Hansen statistics – one for the equation with the smaller set of instruments, where lagged consumption is treated as endogenous and instrumented with asset and education levels at t-1, and one for the equation with the larger set of instruments, where lagged consumption is treated as exogenous – do not reject the assumption of exogeneity of this variable (see Table A.1). Furthermore, the core findings do not

---

[3] We use an adult-equivalent scale that was already included in the dataset instead of a per capita measure, since per capita measures would underestimate the welfare of households with children with respect to families with no children, and the welfare of large households with respect to small households, as stressed in the Basic Information Document of the original LSMS-ISA surveys. Basic Information Documents for the surveys are available at the following link: http://econ.worldbank.org/WBSITE/EXTERNAL/EXTDEC/EXTRESEARCH/EXTLSMS/0,,contentMDK:23635561~pagePK:64168445~piPK:64168309~theSitePK:3358997,00.html

[4] Given the household fixed-effects model, we could not include initial consumption levels because they are time-invariant. Hence the choice of including lagged levels, which in a panel with only three waves is in practice very similar.

change when we use other estimation methods (see Section 6) which treat lagged consumption level as endogenous.

This basic empirical growth model is augmented to investigate the potential impacts of weather shocks.

$\Delta Temp_{gt}$ and ΔPre_{gt} are temperature and precipitation shocks, where 'shocks' mean 'anomalies' in the sense defined by Dell, Jones and Olken (2014), i.e. our weather variables are calculated as the level difference between their average values in the period between interviews and the long-run means, divided by the long-run standard deviation[5]. This means we assume that level changes matter not only in an absolute sense but also, more importantly, in terms of deviation from their long-run averages. Given we have a short-run panel and only limited climatic variation, this choice of the weather functional form suits better the nature of our data.

A common practice in the development literature on the relationship between growth and shocks is the fact almost all these works only include rainfall shocks in the empirical analysis, neglecting the potential role of temperature as a determinant of household growth.

Indeed, climate literature (Auffhammer, Hsiang, Schlenker & Sobel, 2013); Dell, Jones, & Olken, 2014) has warned against the risk of omitted variable bias when dealing with the effects of weather regressors, and recommends to always include at least both temperature and precipitation as independent variables. Since the two are closely correlated, excluding temperature, as commonly done in many empirical development works, may mean attributing to precipitation shocks an impact which could be actually due to temperature. We avoid this risk by including both.

To capture potential heterogeneity of impacts, we also interact weather shocks with dummies for being "poor" and for living in "hot" areas, as well as with dummies for consumption quartiles.

Other than weather shocks, we include two sets of control variables.

ΩZ_{it} is a vegetation time series which includes variables providing data on the start of the wettest quarter, average changes in greenness, and onsets of greenness increase and decrease.

These vegetation variables were already included in the original World Bank data as part of the Integrated Survey on Agriculture (ISA); we chose to add them in the regression following the advice in Auffhammer et al. (2013) and Dell, Jones and Olken (2014): it is important to include a rich set of climatic variables in the regression (when available), given the risk of omitted variable bias due to the fact climatic variables are always highly correlated.

Household controls include household size, the square of household size, the age of the household

[5] The subscript *g* indicates temperature and precipitation variabes are observed at the grid level.

head and its squared term, a dummy for the gender of the household head, average years of education among adults, the number of infants (i.e. less than 5-year old) and dummies capturing a variety of self-reported shocks, both idiosyncratic (illness and deaths of household members) and covariate (e.g. market) shocks. The inclusion of control variables reduces the risk of omitted variable bias and provides smaller standard errors in the estimates.

As for the other elements in the equation, μ_i are household fixed effects; q_{it} are quarter of year dummies to capture when the interview took place; w_t are wave dummies; θ_{rt} are region-year fixed effects, to allow for differentiated time trends in different regions and capture idiosyncratic local shocks, as suggested by Dell, Jones and Olken (2012); ε_{it} are error terms clustered simultaneously at the Enumeration Areas (EAs) and wave levels, following the two-way clustering recommended by Cameron, Gelbach and Miller (2011). EAs are the main stratification level in the NPS surveys and also the closest unit to the grid level where temperature and precipitation are observed; furthermore, in most rural areas, EAs are defined by village boundaries[6].

After finding heterogeneity, we try to detect a critical consumption threshold for the significance of temperature impacts. In order to do so, we employed the Hansen (2000) threshold estimator following the approach by Carter et al. (2007). This model distinguishes two impact regimes conditional to a critical value of lagged (pre-shock) consumption level:

$$(2) \quad ln\, Y_{it} - ln\, Y_{it-1} = \begin{cases} \alpha ln\, Y_{it-1} + \beta^{l} \Delta Temp_{gt} + \gamma \Delta Pre_{gt} + \Omega Z_{it} + \omega X_{it} + \mu_i + q_{it} + w_t + \theta_{rt} + \varepsilon_{it} & if\ ln\, Y_{it-1} \leq \sigma \\ \alpha ln\, Y_{it-1} + \beta^{u} \Delta Temp_{gt} + \gamma \Delta Pre_{gt} + \Omega Z_{it} + \omega X_{it} + \mu_i + q_{it} + w_t + \theta_{rt} + \varepsilon_{it} & if\ ln\, Y_{it-1} > \sigma \end{cases}$$

Where the superscripts *l* and *u* on the coefficient β indicate, respectively, the lower and upper regime of temperature impacts, conditional on lagged consumption level.

4 Data and descriptive statistics

A. Data

The data used in this work are taken from two different sources.

Household data

---

[6] In their works on Tanzania, Hirvonen (2016) clusters standard errors at the village level, Bengtsson (2010) at the "cluster"-level, i.e. the main stratification unit and the level at which rainfall is observed. Given the absence of village location data due to confidentiality reasons, EA coordinates were the most appropriate choice for the clustering level.

Household data come from the Tanzania National Panel Surveys, part of the World Bank collection of household surveys known as Living Standards Measurement Study – Integrated Survey on Agriculture (LSMS – ISA). In particular, this panel consists of three surveys: 2008 – 2009; 2010-2011; 2012-2013[7]. These three surveys have been cleaned and aggregated using household identification numbers to build a three-round panel. All the monetary values in the surveys have been deflated, in order to convert nominal values in real/constant values, using the Consumer Price Index (CPI) for Tanzania by the World Bank[8], and they are expressed in Tanzanian shillings at 2013 monetary values. Importantly, we only selected rural households in building the panel, dropping urban households for which confounding factors would have been more likely. After cleaning the data, we are left with a balanced panel of 1,585 georeferenced households. This panel includes data on household and, as part of the ISA questionnaire, vegetation time series and geographic variables, as well as data on crops and agriculture.

Finally, data on the monetary value of total crop production and other agricultural characteristics used in Section 5 have been developed by the FAO Rural Income Generating Activities (RIGA) Team starting from the household data contained in the survey questionnaires.

Weather data

Weather data are taken from NASA's Modern-Era Retrospective analysis for Research and Applications, Version 2 (MERRA-2), which is a global, gridded data set based on retrospective analysis of historical weather data obtained from satellite images and weather stations (Rienecker et al., 2011). The dataset provides daily temperature measures aggregated into grids that are 1/2° in latitude x 2/ 3° in longitude (which corresponds roughly to 55 km x 75 km at the equator). As with all re-analysis products, the data set is a combination of observed and imputed data points, using observation where and when available, and physics-based interpolation where and when needed.

We aggregated in two ways. First, we computed long-run averages over the period 1980 – 2015. Second, we built average measures of weather variability during the period between interviews for each household. However, to better catch the weather conditions during the growing season, as suggested by Hirvonen (2016), we chose to exclude the summer months from the computations of both averages (namely, June, July, August and September)[9].

---

[7] These data are available at: http://econ.worldbank.org/WBSITE/EXTERNAL/EXTDEC/EXTRESEARCH/EXTLSMS/0,,contentMDK:23635561~pagePK:64168445~piPK:64168309~theSitePK:3358997,00.html

[8] http://data.worldbank.org/indicator/FP.CPI.TOTL?page=1

[9] See http://www.geog.ox.ac.uk/research/climate/projects/undp-cp/UNDP_reports/TanzaniaTanzania.lowres.report.pdf, where it is stated that "the 'short' rains [take place] in October to December and the long rains in March to May, whilst the southern, western and central parts of the country experience one wet season that continues October through April or

Hence, temperature at time *t* is average monthly growing season temperature in the period between t and t-1, expressed in degree Celsius. Precipitation at time *t*, instead, is calculated as average monthly growing season precipitation (in millimetres) in the period between t and t-1. Long-run average temperature and precipitation represent respectively long-run average monthly growing season temperature and long-run average monthly growing season precipitation. Finally, as already specified above, temperature and precipitation *shocks* (or anomalies) at time *t* are defined as the level difference between their values at *t* and their long-run averages, divided by the long-run standard deviation.

We used latitude and longitude coordinates to link these gridded weather data to household data. Unfortunately, for confidentiality reasons we did not have access to the exact location of households, but only to the average of household GPS coordinates in each enumeration area (EA), for which a random offset within a 5-km range was applied for rural households. Such an offset range, anyway, is not an issue of concern for us given the medium resolution of our weather data.

Furthermore, given the risk of incorrect inference when dealing with historical weather data, emphasized by Auffhammer et al. (2013), as a robustness check we also run a sensitivity analysis for our results by using a different source of weather data, where temperature data come from the *CRUCY Version 3.23* by the Climatic Research Unit (CRU) of the University of East Anglia (CRU, 2016), and have a resolution of 1/2° in latitude x 1/2° in longitude, and rainfall data come from the same NPS Dataset as part of the ISA module, and they contain data on total rainfall in the wettest quarter within 12-month periods starting in July previous to each round.

B. Descriptive statistics

Table 1 provides descriptive statistics for the main variables employed in the empirical analysis. Annualised average total and food consumption growth rates are both negative: they decreased on average by about 1.4 and 1.7 percentage points each year. However, the standard deviation is large for both variables, indicating heterogeneity in the growth paths experienced by rural households. Both temperature and precipitation anomalies were, on average, positive in the timespan considered, but for them as well it is worth noting the huge standard deviation, suggesting substantial heterogeneity in the weather conditions experienced by households living in different geographical areas.

5 Baseline results

Tables 2 and 3 report the results from estimating Equation (1). First, the hypothesis of convergence among households is confirmed: growth rates are negatively related to ‘initial’ consumption levels,

May”. In this way, given the intrinsic difficulty in exactly identifying rainy seasons months for households scattered across the whole country, we excluded the summer months which are not part of any rainy season in Tanzania.

i.e. poorer household grow faster. As for the weather variables, Column 1 shows that, on average and *ceteris paribus*, temperature (precipitation) shocks have a slightly negative (positive) but not significant impact on growth.

Column 2 controls for heterogeneity of impacts, by interacting both temperature and precipitation with a dummy for being "poor", i.e. a dummy with value 1 for households with below median initial food (in Table 2) or total (in Table 3) consumption. Defining a household as "poor" is of course a relative concept in a context like rural Tanzania. We essentially check for heterogeneity of impacts with respect to the poorest amongst the poor. Including these interactions qualitatively changes the results: temperature shocks now have a positive and weakly significant impact for the "non-poor" households, but a large, negative and significant (at the 5 percent level) impact on household growth for "poor" households. Interpreting these results with respect to the within-standard deviation of temperature shocks (0.237), one standard deviation increase in temperature anomalies decreases household per-adult equivalent food consumption growth by about 2.76 %, and household per-adult equivalent total consumption growth by approximately 2.21 %, *ceteris paribus*, for households defined as "poor". Rainfall impacts are insignificant.

Given the presence of heterogeneity with respect to initial consumption, in Column 3 we also check for heterogeneity by interacting shocks with a dummy for living in "hot" areas, which takes value 1 for households living in an area with above mean long-run average monthly growing season temperature. Although the interaction between temperature shocks and the dummy for "poor" households stays unchanged in sign, magnitude and significance, the total effect of temperature shocks on poorest households is now slightly smaller and less significant. The interaction between temperature shocks and a dummy for households living in hotter areas is small and negligible, and so the total effect. Living in a hot area has a positive (and significant, but only in Table 3) impact on growth, but this is very likely due to the fact the hottest areas in Tanzania (coastal regions and Zanzibar) are also the richest ones.

Temperature impacts on growth are always larger on food consumption growth than on total consumption growth, consistently with the fact that most households are subsistence farming households. This will be additionally addressed in Section 7, where the channels of the heterogeneity will be investigated.

Finally, Column 4 in both Tables 2 and 3 explores more in detail the relationship between consumption levels, temperature shocks and their impact on growth, by interacting the lagged consumption term (food consumption in Table 2, total consumption in Table 3) with temperature shocks. The results are consistent with the previous findings: the process of convergence is unaltered, the coefficient for temperature shocks is negative and statistically significant, the interaction between lagged consumption and temperature shocks is positive and statistically significant at the 1 percent level, suggesting that the impacts from temperature shocks tend to decrease as households grow richer.

Figures 1 and 2 show the implications of the results in Column 4 for, respectively, Table 3 and 4. They show the marginal effect of temperature shocks at different lagged consumption levels. When households have a sufficiently high level of pre-shock consumption, impacts from temperature shocks turn first zero and then eventually positive.

Tables 4 and 5 take a closer look, by interacting weather shocks not with a dummy for being "poor", but with dummies for initial consumption quartiles[10]. The results, consistent between tables, reveal even further heterogeneity: as can be seen in Column 1 of both tables, households belonging to the poorest initial quartile suffer from a large, negative and statistically significant impact of temperature shocks, while the second and third quartiles do not, and growth for households in the upper initial quartile is positively and significantly affected, revealing heterogeneity in sign rather than size.

This core finding is not altered when including the interaction for living in a "hot" area, as shown in Column 2 of both tables. Finally, impacts due to precipitation shocks are always insignificant.

In sum, depending on initial conditions, the impacts of temperature shocks on household growth is sharply heterogeneous across quartiles, and poorest households are the only ones to be significantly and negatively affected.

This contrasts with the implications of the negative and significant coefficient of the lagged consumption term: while there seems to be an ongoing process of convergence among households, temperature shocks go in the opposite direction, slowing growth of the poorest households while bolstering growth for the richest ones.

However, we have not precisely identified thresholds of consumption that entail regime changes for temperature shocks. We just interacted shocks with dummies that capture heterogeneity, but these choices are arbitrary. They are not driven by the data.

To overcome this drawback, on the wake of Carter et al. (2007), we present the results for a panel threshold model using the so-called Hansen (2000) estimator, as implemented in a fixed-effect setting by Wang (2015).

Threshold models identify the jumping character or structural break in the relationship between variables. In our context, we are looking for thresholds of pre-shock consumption above or below which there is a structural break in the impact of temperature shocks, as illustrated in Equation (2). Temperature shocks are the regime-dependent variable.

[10] Although we considered the possibility of a quantile regression model as an alternative specification, we ruled out this option because when quantile regression is combined with panel data and a fixed-effect setting, identification and estimation become complicated, since the quantiles of the difference are not equal to the difference in quantiles, and the issue gets even worse when the number of time periods is small (as in our case) (Ponomareva, 2010).

Looking at the previous regressions, it appears there is not just one threshold, but two separate and distinct thresholds. The first is the threshold above which impacts turn negative but statistically insignificant; the second the one above which impacts turn positive and significant. We are therefore looking for two, and not just one, consumption level thresholds.

In Table 6 we present the results for this double threshold model using the Hansen estimator.

In Column 1 the dependent variable is food consumption growth, in Column 2 total consumption growth. As hypothesized, we find two thresholds and three regimes: a first threshold below which impacts of temperature shocks are negative and strongly significant, and above which they turn insignificant; and a second threshold from which impacts turn to being positive and strongly significant. Although the positive impact above the upper threshold is much bigger than the negative impact below the lower threshold, the percentage of observations falling below the lower threshold is much higher (47 % and 24 % , respectively, for food and total consumption) than the percentage of observations above the upper threshold (around 13 % in both cases), revealing it is a smaller group of better-off households that drives the significance of the positive impact for the upper quartile.

Both thresholds, for both dependent variables, are statistically significant at the 1 percent level, as reported in the threshold tests.

After re-converting logs into monetary values, for food consumption we find a lower threshold of approximately 483594 Tanzanian shillings or, expressed at 2013 Purchasing Power Parity (PPP) values[11], 803 dollars; and an upper threshold of approximately 917126 Tanzanian shillings, i.e. about 1523 dollars; for total consumption, instead, the two thresholds are approximately 2434956 Tanzanian shillings, approximately 723 dollars, and 1219559 Tanzanian shillings, or about 2026 dollars.

Temperature shocks, in sum, slow the convergence process, and are a source of divergence. This has strong distributional implications and raises the issue of which channels and transmission mechanisms could be responsible for such a sharp heterogeneity of impacts. These questions are addressed in Section 7 but, first, Section 6 conducts a number of tests to assess the robustness of our results to different sensitivity analyses, and make sure our findings are not driven by the chosen identification strategy or by properties of the data used.

6 Robustness checks

We explore the robustness of our results with respect to spatial autocorrelation, different weather data and different estimation strategies.

[11] For the PPP conversion factor in 2013: https://data.worldbank.org/indicator/PA.NUS.PPP?locations=TZ .

A. Conley (1999) standard errors

It is well known that both economic growth and temperature are spatially autocorrelated. One could thus argue that confidence in our results are inflated because we fail to take this into account. We therefore re-run the quartile regressions from Tables 4 and 5 correcting for Conley (1999) standard errors, which are robust to both spatial autocorrelation and heteroskedasticity. The computation of the Conley standard errors is based on a weighing matrix that places greater weight on observations that are closer to each other, and the weights decay to zero after a pre-specified distance cut-off is met. We use the following cut-off points: 50, 75 and 100 km. These regressions are reported in Table 7: in Column 1 the dependent variable is food consumption growth, in Column 2 is total consumption growth. The core results are basically unchanged: our findings are not weakened when correcting for spatial autocorrelation and spatially-robust standard errors.

B. Different weather data

Results could be driven by properties of the weather data, the selection of weather stations, the homogenization of the data, and the imputation of missing observations. Auffhammer et al. (2013) highlight the risk of using reanalysis data, since reanalysis is conducted with models that, like economic models, are imperfect and contain systematic biases. Moreover, they recommend to always check that results also hold when using a different data source.

For temperature data, we use the *CRUCY Version 3.23* by the Climatic Research Unit (CRU) of the University of East Anglia (CRU, 2016), a gridded dataset which has a resolution of 1/2° in latitude x 1/2° in longitude. While the MERRA-2 Reanalysis data combine information from ground stations, satellites, and other sources with a physical climate model to create gridded weather data products, CRU data are gridded data, statistically interpolated from ground stations (Dell, Jones and Olken, 2014). Table A.2 in the Appendix provides descriptive statistics for the CRU temperature data. ΔTemp is on average almost 5 times bigger compared to average temperature shocks in Table 1. Despite this, the correlation between the two temperature series is more than 90 %.

As for rainfall, we use precipitation data that come from the NPS Dataset as part of the ISA module, and our variable is now average total rainfall in the wettest quarter before the interview. These data were taken from the NOAA datasets on African Rainfall Climatology (ARC) data. ARC data blend rain gauge measurements and InfraRed (IR) satellite information to render a daily, high resolution (0.1°x0.1°) gridded estimate covering the Africa continent.[12] Since data on the long-run standard deviation are not included, we simply define rainfall shocks as level differences from the long-run

[12] Data can be found at: ftp://ftp.cpc.ncep.noaa.gov/fews/newalgo_est_dekad/ .

average.

The results are in Table 8. The pattern of heterogeneity holds, and the effect size is similar, both for the negative impacts on households belonging to the poorest quartile and for the positive impacts for households belonging to the richest quartile. Precipitation shocks are now often significant, and seem to point to heterogeneity as well, but they are also quite sensitive to specification, and since we detect no significant precipitation impacts on crop yields using the same data source (see Section 7), we conclude their significance here is likely incidental.

In sum, our main findings hold when using a different source of weather data.

C. Hausman – Taylor regressions

Following Dercon (2004), we repeat our empirical analysis using the Hausman - Taylor (1981) model, which involves partitioning the time-invariant and time-varying vector of variables in two groups each, of which one group of variables is assumed to be uncorrelated with the fixed effect.

The Hausman-Taylor model, being a random-effect model for panel data allows us to include time-invariant variables in our regressions. In particular, in addition to region dummies[13], we add distance to the nearest major road and long-run averages for our weather variables. Given the strong partitioning assumptions implied by this estimation strategy, we adopt a cautious approach, following Dercon (2004): lagged consumption terms and all household controls (with the exception of self-reported covariate shocks) are treated as time-varying endogenous variables; dummies for consumption quartiles are treated as time-invariant endogenous; all weather and geographic variables, both time-varying and time-invariant, are treated as exogenous.

Results can be found in Table 9 for food consumption growth (Column 1) and total consumption growth (Column 2)[14]. Despite stark differences between estimation strategies, the overall picture is consistent with the results from the fixed-effect specification: the convergence process is confirmed, and temperature shocks only harm poorest households, although here also the second poorest quartile is negatively and significantly affected. Interestingly, while the coefficient for the upper quartile is still positive, its magnitude has decreased and its significance has disappeared in Column (1) and diminished in Column (2). This will be further addressed in the next robustness check. As above, there is no statistically discernible effect of rainfall shocks, while both long-run temperature and precipitation have a positive impact on both food and total consumption growth.

---

[13] Region dummies were included separately from year dummies because the estimation of Hausman-Taylor regressions requires the presence of time-invariant exogenous variables.

[14] Incidentally, although not reported in Table 7, distance from the nearest major road always has a large and significant effect on growth, consistently with what found by Dercon (2004) in rural Ethiopia, hinting at public infrastructure as another source of divergence among households.

D. Two-Step Difference GMM

As a third, and last, estimation strategy we employ the two-step difference GMM, first proposed by Arellano and Bond (1991). This estimation method is especially recommended for dynamic panels which exhibit the following characteristics (Roodman, 2006): "1) "small *T*, large *N*" panels, meaning few time periods and many individuals; 2) a linear functional relationship; 3) one left-hand-side variable that is dynamic, depending on its own past realizations; 4) independent variables that are not strictly exogenous, meaning they are correlated with past and possibly current realizations of the error; 5) fixed individual effects; and 6) heteroskedasticity and autocorrelation within individuals but not across them". Arellano–Bond estimation transforms all regressors by differencing, and uses the generalized method of moments (GMM) as the estimation method. Importantly, it adjusts for the potential bias caused by the inclusion of a lagged dependent variable as a regressor. The Hansen-J tests reported ensure the specification is valid, and the standard errors are corrected using Windmeijer (2005) adjustment procedure.

In distinguishing between endogenous and exogenous variables, we followed Dercon (2004) and Jalan and Ravallion (2002): lagged consumption terms and all household controls are treated as endogenous, and weather shocks and vegetation time series as exogenous.

The results for the two-step Arellano-Bond GMM estimation are reported in Table 10.

They are consistent with the fixed-effect and Hausman-Taylor regressions discussed above: heterogeneity of impacts from temperature shocks is confirmed, with a strong and significant impact only for households belonging to the poorest initial quartile.

Similarly to the Hausman-Taylor model, temperature impacts for households in the richest quartiles are still positive, but much smaller and not significant anymore. This means that the significance of the positive impact detected using the fixed-effect model is not robust to different estimation strategies, and should be interpreted with extreme caution. Finally, precipitation is insignificant.

7 Transmission channels and mechanisms

Having demonstrated robustness, we now explore why there is such a sharp heterogeneity of impacts and indeed a *change in sign* of impacts on household growth depending on initial consumption. We shed light on this question by investigating three main channels: health expenditure, labour productivity and agricultural yields. These three transmission channels find their theoretical and empirical grounding in previous literature. Furthermore, to explain if and how households cope with impacts of temperature shocks, we investigate impacts on asset growth, to test for the potential presence of asset-smoothing in contrast to consumption smoothing behaviours.

A. *Health expenditure channel*

As summarized by Dell, Jones and Olken (2014) and Heal and Park (2015), several studies have examined the impact temperature can have on morbidity and mortality, which in turn affect labour productivity and income (and *vice versa*).

Empirical works such as, among the others, Barreca (2012), Burgess et al. (2011), Deschênes and Greenstone, (2011) and Goldberg et al. (2011) have documented the effects of temperature and heat waves on health, particularly mortality, using panel methods. From a theoretical point of view, instead, the long-run relationship between health, environment/climate and poverty traps has been explored by Raffin (2012), Strulik (2008) and Tol (2011).

In our framework, it could be temperature shocks on consumption growth appear, at least partially, through the health channel: temperature could affect health and hence productivity, and this in turns may affect income and subsequently consumption.

We test this mechanism by using the baseline specification set out in Equation (1) with a different dependent variable: instead of consumption growth, we now use as Y_{it} the ratio between health expenditure and total expenditure[15]. The expected sign of the relationship is the opposite: in response to temperature shocks, the growth rate of the ratio should increase. Table 11, Column 1 partially supports our hypothesis: temperature shocks have a positive (but not significant) impact on the growth rate of the health expenditure / total expenditure ratio. Furthermore, to justify the pattern of heterogeneity, one would expect this ratio to increase significantly more for households belonging to the poorest quartile. As reported in Column 2, this is not the case: the impact is small and insignificant for all quartiles, and the sign is not the expected one. Hence, either the health channel is not responsible for the heterogeneity we find, or there is a transmission mechanism which is ongoing but cannot be caught given the limitations and short-run nature of our data.

Column 3 shows that living in a hot area has a large, positive and significant effect on the growth rate of the ratio of health to total expenditure. In other words, if the weather is anomalously hot, people spend more on health care.

B. *Labour productivity channel*

A recent but already large micro literature (Cachon, Gallino, & Olivares, 2012; Cachon et al., 2012; Graff Zivin & Neidell, 2014; Heal & Park, 2015; Niemelä et al., 2002; Park, 2017; Sudarshan & Tewari, 2013) has found vast and significant effects of temperature increases on the productivity of workers, especially on those who work outdoor.

---

[15] To calculate the growth rate of this ratio, we increased both per a.e. health and total expenditures by the same small increment (the equivalent of a US dollar) for all households.

In a context like rural Tanzania, a large share of workers is involved in outdoor work, primarily in farming. Outdoor work is more exposed to heat waves, and agriculture in Tanzania is still largely traditional and thus still involves a lot of manual labour. These characteristics make workers in rural areas vulnerable to stress from temperature shocks, but there could also be significant differences in farmers' characteristics that entail heterogeneity. Labour productivity may thus help explaining the heterogeneous impacts on consumption growth.

We created a rough measure of agricultural labour productivity by dividing the monetary value of household total crop production (taken from the FAO Rural Income Generating Activities (RIGA) Team Database[16]) in the 12 months before the interview by the number of family members engaged in agricultural activities in the 12 months before the interview.

We are aware this measure represents a rough and only approximate proxy of (agricultural) labour productivity, stemming from one of the more general definitions of labour productivity as the ratio between total output and number of employed persons, but it is also the only one that we could get[17]. Consequently, our left-hand side variable is the growth rate of (agricultural) labour productivity between t and t-1[18]. Analogously to Equation (1), we regress this dependent variable on lagged agricultural labour productivity, temperature and precipitation shocks as well as controls and fixed effects.

Since preliminary endogeneity tests (see Table A.3) rejected the assumption of exogeneity of the lagged dependent variable, the model was estimated using two-step difference GMM.

Results are reported in Table 12. Column 1 shows average impacts. Temperature anomalies have a large and significant at the 5 level impact on the growth rate agricultural labour productivity. One within-standard deviation increase in temperature shocks decreases agricultural labour productivity growth by approximately 5.61 %, on average, *ceteris paribus*. Column 2 disentangles this aggregate impact across initial consumption quartiles: there is a large and significant negative effect on the poorest quartile, while impacts are not significant for the other quartiles. Precipitation shocks are insignificant and close to zero. This overall picture is consistent with the consumption growth regressions, and confirms labour productivity as one of the transmission channels responsible for the heterogeneity of impacts, but not for the sign change.

Why is there such a discrepancy of impacts on agricultural labour productivity growth across quartiles? Tables A.4 reports some descriptive statistics that can help clarifying this issue. It shows the average Agricultural Wealth Index for the four initial consumption quartiles. The Agricultural

[16] The FAO-RIGA Database can be found at: http://www.fao.org/economic/riga/riga-database/en/.

[17] Another shortcoming is that we only investigate the aggregate impact, without disentangling the impacts between labour supply and labour demand. Unfortunately, such refinements go beyond the limitations of our data.

[18] We added a small amount (the equivalent of a US dollar) to labour productivity values of all households not to lose observations with zeros when calculating growth rates.

Wealth Index was again taken from the FAO-RIGA Database, and is a specific aggregated index based on a factor analysis of the agricultural assets and technologies used by rural households in the sample. In this context this is useful because it also proxies for the use of technologies that decrease the need for manual labour. The average index is more than three times higher for the upper quartile compared to the poorest quartile, although oddly very low for the third quartile.

Additionally, Table A.5 reports the percentage of households, across quartiles, for which farming was not the main source of income in at least two waves. According to our hypothesis above, the less households depend on farming activities, the less they work outdoors, and the lower the impact on labour productivity. Farming was the main source of income for about 81% of households in the poorest quartile. This share falls and, for the richest quartile, two-thirds of households depend on farming as the main source of income.

This further enhances the influence of weather variability on the labour productivity of poorest households compared to that of the wealthier households.

Aware of the limitations of our labour productivity measure, we find heterogenous effects on the growth rate agricultural labour productivity, which partially explains heterogeneity of impacts on consumption growth.

This impact on labour productivity may have directly affected income or also entailed an indirect effect through crop yields, as Sudarshan & Tewari (2013) hypothesize: "Observed productivity losses in agriculture that have been attributed by default to plant growth responses to high temperatures may in fact be partly driven by lower labor productivity". Of course, the opposite may also be true: impacts on agricultural labour productivity may be actually driven by losses in crop yields.

C. *Agricultural yield channel*

Several studies have investigate the relationship between crop yields and weather variability, starting from the very plausible assumption that extreme temperatures and rainfall above or below a certain threshold may have damaging consequences on crop yields, especially in developing countries whose agriculture is less modernized (Challinor, Wheeler, Craufurd & Slingo., 2005; Feng, Krueger, & Oppenheimer, 2010; Guiteras, 2009; Levine & Yang, 2006; Li et al., 2010; Porter & Semenov, 2005; Rowhani et al., 2011; Schlenker & Lobell, 2010; Welch et al., 2010).

Other works have provided theoretical underpinnings to explain how low crop yields and yield gaps could be one of the reasons why smallholder farmers are trapped in chronic poverty (Barrett & Swallow, 2006; Sachs, 2008; Tittonell & Giller, 2013).

On the wake of this literature, we investigate the agricultural yield channel to explain heterogeneity of impacts on consumption growth. Crop yields are defined as quantity produced (in kilograms) divided per hectare of cultivated land. Thanks to the ISA module in the original dataset, we had access to crop data for the two rainy seasons (long and short) preceding the interview month. In investigating the impacts of weather shocks on crops, we must also take into account the possibility of non-linear effects, given the apparent inverted-U and non-linear relationship between temperature and plant growth (Dell, Jones & Olken, 2014; Hirvonen, 2016; Schlenker & Roberts, 2009). In order to do so, we draw from Ahmed et al. (2011), Hirvonen (2016), and Rowhani et al. (2011) works on Tanzania and adopt a specific temperature measure, the number of growing degree days (GDDs) (Schlenker & Roberts, 2009) in the twelve months preceding the interview month. Following the procedure implemented by Hirvonen (2016), we took daily minimum and maximum temperatures from the MERRA-2 data and approximated the diurnal temperature distribution by interpolating between the minimum and maximum temperature values using a sinusoidal curve. Growing degree days are then measured by the time of exposure to a certain temperature range. As in Hirvonen (2016), we set the lower bound to 8°C and the upper bound to 34°C. Exposure to temperatures above 34°C is considered harmful for agricultural yields[19]. In our regressions we use a spline function of the GDD variable. The first part of this variable captures temperature exposure between 8-34°C and the second exposure to temperatures above 34°C. We included average total precipitation during the two wettest quarters before the interview and its squares, using the alternative ARC rainfall data (cf. Tables 8 and A2), because they use the actual household plot location.

Table 13 reports the results for this specification. The dependent variable is average crop yield during the previous two rainy seasons. In Column 1 we only look at the aggregate impact. The estimates suggest that it is exposure to extreme temperatures (above 34°C) which is harmful for crop yields. , In Column 2 we check whether this negative effect mainly comes through maize and paddy, two of the most important crops in the country, as suggested by previous literature on the impacts of temperature on crop yields in Tanzania (Ahmed et al. 2011; Rowhani et al., 2011).

Therefore, we include interactions with a dummy for 'Maize & paddy non-specializers', a dummy with value 1 for households in which maize and paddy account for less than 50% of total crop production in a given wave[20]. As expected, negative effects on crop yields from extreme temperatures are driven by impacts on maize and paddy, and disappear if households are not specialized in the cultivation of these two crops.

In Column 3 we decompose the aggregate impact of GDDs by looking at impacts across initial consumption quartiles. Rainfall impacts are close to zero and insignificant. Impacts of GDDs between 8-34°C is essentially zero for all four quartiles. Exposure to extreme temperatures (above 34°C) has negative and strongly significant impact on crop yields of households in the poorest quartile, a

[19] Descriptive statistics on GDDs can be found in the Appendix, Table A.6.
[20] See Table A.7 for descriptive statistics of this dummy.

negative and insignificant impact on crop yields of households in the second and third quartiles, and a positive but insignificant impact on crop yields of households in the upper quartile. These results are consistent with the pattern of heterogeneity of temperature shocks on consumption growth.

Why are there such big differences in the impacts from extreme temperatures on crop yields across quartiles?

Table A.8 reveals that richer households produce more crops (Column 1) and have more productive plots (Column 2). The heterogeneity of impacts can thus be explained by the fact that richer households are advantaged by better agricultural assets, technologies and soil quality, which make them less vulnerable to the negative impacts entailed by temperature shocks, which conversely have serious welfare consequences for poorest households.

We have yet to account for the sign change for the upper quartile. The use of irrigation is still very limited (Table A.9) and so the use of inorganic fertilizers (Table A.10), but richer households show better conditions. Tables A.11-A.14 in the Appendix show data taken from the ISA module on the use of 'improved' seeds for maize and paddy. Improved seeds are more drought-resistant and can mitigate the negative impacts of extreme temperatures. Tables A.11 and A.12 show that the use of improved maize seeds sharply differ across consumption quartiles. Tables A.13 and A.14 reveal the same pattern with regard to the use of improved paddy seeds.

D. *Testing for asset smoothing*

We have established that the main channels that account for the heterogeneity of impacts on consumption growth are agricultural yields and labour productivity.

But we did not explain yet why households are not smoothing consumption by drawing on their assets. Drawing from previous theoretical and empirical literature (Barrett et al., 2006; Barrett & Carter, 2013; Carter & Barrett, 2006; Carter et al., 2007; Carter & Lybbert, 2012), we test the two alternative hypothesis of consumption *vs* asset smoothing by repeating the baseline specification but using, as an alternative dependent variable, asset growth instead of consumption growth. Our measure of assets is Tropical Livestock Units (TLUs), again taken from the FAO-RIGA Dataset. Descriptive statistics for TLUs is reported in Table A.15: the gap in TLUs per adult-equivalent across quartiles is evident.

The dependent variable, therefore, is now annualised percentage change in (ln) per a.e. household TLUs between t and t-1[21]. Table 14 reports the results. In Column 1 we can see that, while convergence among households is confirmed, temperature shocks have, on average, a negative but not significant impact on asset growth. In Column 2, where we decompose the impacts by

[21] To calculate asset growth and use logarithms, since many households have no assets at all and this implied the presence of many zeroes in the data, we followed the method implemented in Carter et al. (2007) and increase all livestock assets per adult-equivalent by the same small increment (namely the minimum value in the sample above zero).

consumption quartiles, impacts are always negative but we do not find any significance.
These findings imply several considerations. First, it was a good choice to look at consumption growth instead of asset growth, following the reasoning in Carter et al. (2007), who argued that in the context of weather shocks such as droughts, characterized by a gradual onset and a prolonged effect (differently from the immediate disruption entailed by environmental shocks such as hurricanes or typhoons), impacts on welfare growth could appear through consumption and not through assets. Indeed, had we chosen asset growth as the dependent variable, we would have found no impacts at all. Second, poorest households in our sample could be performing asset-smoothing, i.e. they might be choosing of voluntarily destabilize consumption and stubbornly hold on to their livestock, in order not to sell them and then fall in a poverty trap from which there could be no recovery. This is consistent with what Carter et al. (2007) find for Ethiopia, where they note that "poor households seek to defend their assets in the face of successive droughts rather than liquidate them and perhaps limit their subsequent chances of recovery.".
Therefore, we are prone to assert that, for the poorest households in our sample, asset smoothing is probably taking place, while the choice of using assets as buffer stocks, one of the main risk-coping strategy hypothesized in literature, was either not adopted or not effective during the survey period (Kazianga & Udry, 2006; Morduch, 1995).

8 Discussion and conclusion

Using the LSMS-ISA Tanzania Panel Surveys by the World Bank, we find a causal relationship between temperature shocks, household growth and poverty in rural Tanzania. There is heterogeneity of impacts from temperature shocks, which affect household consumption growth only if initial consumption levels lie below critical thresholds, explained by the impacts of temperature anomalies on two interrelated transmission channels: labour productivity and, more importantly, on crop yields. Richer and poorer households differ not only in that the former have more diversified income sources and are less engaged in outdoor farming activities, but especially to the existence of a 'yield gap' and differences in crucial agricultural characteristics. Such differences among households may also be related to *ex-ante* risk-managing behaviours (Dercon, 2004), e.g. the conservative behaviour of the poorer risk-averse households that shy away from investing in profitable but risky technologies (such as modern agricultural inputs) and stick to low-risk, low-return activities, as indeed Dercon (1996) himself argues it is the case in the context of rural Tanzania, or, more simply, they cannot have access to these technologies because of credit and liquidity constraints.
Importantly, while the negative effect for households below the lower threshold has proved to be robust, the significance of the positive impact above the second upper threshold disappeared when using different estimation methods such as the Hausman-Taylor random-effect model and two-step

difference GMM, and the magnitude itself of such effect positive substantially decreased. Furthermore, the analysis of the transmission channels found no evidence whatsoever of a significantly positive impact. While it may be that richest households are indirectly taking advantage from the negative impacts on poorest households and earning more from their crop activities, this explanation only goes so far, and is not supported by sufficiently robust empirical evidence.
In any case, temperature shocks have a heterogeneous *ex-post* impact which slows the process of convergence and enhances inequalities. These micro results are consistent with what found on the relationship between growth, temperature shocks and poverty by macro studies (Dell, Jones & Olken, 2012; Letta and Tol, 2016)

However, these findings must be interpreted with caution for a number of reasons, the first being the nature and limitations of the data. We use a six-year panel with only three rounds, so we are only estimating a short-run elasticity between temperature shocks and growth in a given time span. It may well be that our period of study saw relatively mild weather variability, which could explain the absence of a significant average impact. Severe droughts may well entail much more pervasive consequences. However, even such extreme scenarios are unlikely to overturn the core finding that it is the poorest households who suffer more from the negative impacts of temperature shocks.
Second, convergence is a long-run process. Even though we observe convergence in this short-run panel, we can only infer about long-run convergence, but not directly test for it.
In the future, the availability of longer household-level panels for developing countries could alleviate these issues, enabling further research to test whether these findings, emerged from short-run elasticities, also hold in the medium or long run.
Third, the consumption thresholds we detected, other than being intrinsically relative and data-driven, are not thresholds in the sense of the existence ‘poverty traps’, below which households are permanently trapped in low income. Temperature shocks have a diverging effect which enhances inequalities and slows the convergence process, but does not reverse it. Making all households reach the critical threshold level above which impacts turn insignificant, would make this source of divergence disappear. There are no multiple equilibria, but rather different regimes of impacts separated by pre-shock consumption thresholds. Rather than a climate-induced poverty trap, whose potential existence was the research question at the heart of this work, if anything we could define this relationship a poverty-induced climate trap.

Extrapolating from weather to climate, such a qualitative finding is particularly relevant to climate change policy. Sub-Saharan Africa is one of the most vulnerable parts of the world to the threats posed by climate change (IPCC, 2014). Climate change is a serious threat to crop productivity (Knox, Hess, Daccache & Wheeler, 2012) and food security (Challinor, Wheeler, Garforth, Craufurd & Kassam, 2007) in Africa, and Tanzania is no exception (Ahmed et al., 2011; Hirvonen, 2016; Rowhani et al., 2011) .

The so-called Schelling Conjecture (Schelling, 1992 & 1995), i.e. that economic development would reduce vulnerability to climate change, and Schelling's point that the need for greenhouse gas abatement cannot be separated from the developing world's need for immediate development (Schelling, 1997), find empirical support in the results of this work. Tol (2015) illustrates the Schelling Conjecture with regard to Africa: "In the worst projections, climate change could cut crop yields in Africa by half (Porter et al., 2014). At present, subsistence farmers often get no more from their land than one-tenth what is achieved at model farms working the same soil in the same climate (Mueller et al., 2012). The immediate reason for the so-called yield gap is a lack of access to irrigation, high-quality seeds, pesticides, fertilizers, tools, and things like that. The underlying causes include a lack of access to capital and product markets due to poor roads and insecure land tenure (Dorward et al., 2004; Foley et al., 2011). [...] Indeed, modernizing agriculture in Africa would also make it less *vulnerable* to climate change (Howden et al., 2007; Mendelsohn & Dinar, 1999)." More broadly, these results increase the concerns over the issue of the distributional implications of future impacts, because they show that inequalities of impacts hold at the micro level as they do at the macro level, as already guessed by Tol (2016). If the impacts of temperature shocks decrease as households grow richer, growth is the key for rural Tanzanian households: diversifying income sources, reducing outdoor work, modernizing agriculture, closing the yield gap and using drought-resistant seeds would all make households less vulnerable to the negative impacts of weather shocks, and less dependent on climate.

However, a note of caution is in order with regard to these considerations. External validity is an issue: weather variations are *not* climate variations: climate change is a long-run phenomenon in which other factors, as intensification of impacts, global non-linear effects and adaptation, could completely alter the nature and magnitude of the current elasticities (Dell, Jones and Olken 2014). This is true for the Schelling Conjecture as well, which may overturn or break down in the tail for high warming above 3°C (Tol, 2016).

These *caveats* notwithstanding, we reckon that development and poverty reduction should be key and paramount elements of any climate policy, especially in vulnerable contexts like rural Tanzania, and that inequality of impacts will be, within countries other than between countries, the first and foremost challenge posed by climate change.

Bibliography

Acemoglu, D., Johnson, S., & Robinson, J. A. (2000). The colonial origins of comparative development: An empirical investigation. *American Economic Review*, 91(5), 1369-1401.

Acemoglu, D., Johnson, S., & Robinson, J. A. (2001). Reversal of fortune: Geography and institutions in the making of the modern world income distribution. *The Quarterly Journal of Economics*, 117(4), 1231-1294.

Ahmed, S. A., Diffenbaugh, N. S., Hertel, T. W., Lobell, D. B., Ramankutty, N., Rios, A. R., & Rowhani, P. (2011). Climate volatility and poverty vulnerability in Tanzania. *Global Environmental Change*, *21*(1), 46–55.

Alderman, H., Hoddinott, J., & Kinsey, B. (2006). Long term consequences of early childhood malnutrition. *Oxford economic papers*, *58*(3), 450–474.

Alsan, M. (2014). The effect of the tsetse fly on African development. *The American Economic Review*, *105*(1), 382–410.

Andersen, T. B., Dalgaard, C.-J., & Selaya, P. (2016). Climate and the Emergence of Global Income Differences. *The Review of Economic Studies*, rdw006.

Arellano, M., & Bond, S. (1991). Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *The review of economic studies*, *58*(2), 277–297.

Auffhammer, M., Hsiang, S. M., Schlenker, W., & Sobel, A. (2013). Using weather data and climate model output in economic analyses of climate change. *Review of Environmental Economics and Policy*, ret016.

Azariadis, C., & Drazen, A. (1990). Threshold externalities in economic development. *The Quarterly Journal of Economics*, *105*(2), 501–526.

Bansal, R., & Ochoa, M. (2011). *Temperature, aggregate risk, and expected returns*. National Bureau of Economic Research.

Barreca, A. I. (2012). Climate change, humidity, and mortality in the United States. *Journal of Environmental Economics and Management*, *63*(1), 19–34.

Barrett, C. B., & Carter, M. R. (2013). The economics of poverty traps and persistent poverty: empirical and policy implications. *The Journal of Development Studies, 49*(7), 976–990.

Barrett, C. B., Marenya, P. P., McPeak, J., Minten, B., Murithi, F., Oluoch-Kosura, W., … Wangila, J. (2006). Welfare dynamics in rural Kenya and Madagascar. *The Journal of Development Studies*, *42*(2), 248–277.

Barrett, C. B., & Swallow, B. M. (2006). Fractal poverty traps. *World development*, *34*(1), 1–15.

Barro, R. J., & Sala-i-Martin, X. (1990). *Economic growth and convergence across the United States*. National Bureau of Economic Research.

Baumol, W. J. (1986). Productivity growth, convergence, and welfare: what the long-run data show. *The American Economic Review*, 1072–1085.

Bengtsson, N. (2010). How responsive is body weight to transitory income changes? Evidence from rural Tanzania. *Journal of Development Economics*, *92*(1), 53–61.

Bloom, D. E., Canning, D., & Sevilla, J. (2003). Geography and poverty traps. *Journal of Economic Growth*, *8*(4), 355–378.

Bonds, M. H., Keenan, D. C., Rohani, P., & Sachs, J. D. (2010). Poverty trap formed by the ecology of infectious diseases. *Proceedings of the Royal Society of London B: Biological Sciences*, *277*(1685), 1185–1192.

Burgess, R., Deschenes, O., Donaldson, D., & Greenstone, M. (2011). Weather and death in India. *Cambridge, United States: Massachusetts Institute of Technology, Department of Economics. Manuscript*.

Burke, M., Hsiang, S. M., & Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature* (527): 235-239.

Cachon, G., Gallino, S., & Olivares, M. (2012). Severe weather and automobile assembly productivity. *Columbia Business School Research Paper*, (12/37).

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, *29*(2), 238–249.

Carter, M. R., & Barrett, C. B. (2006). The economics of poverty traps and persistent poverty: An asset-based approach. *The Journal of Development Studies*, *42*(2), 178–199.

Carter, M. R., Little, P. D., Mogues, T., & Negatu, W. (2007). Poverty traps and natural disasters in Ethiopia and Honduras. *World development*, *35*(5), 835–856.

Carter, M. R., & Lybbert, T. J. (2012). Consumption versus asset smoothing: testing the implications of poverty trap theory in Burkina Faso. *Journal of Development Economics*, *99*(2), 255–264.

Challinor, A. J., Wheeler, T. R., Craufurd, P. Q., & Slingo, J. M. (2005). Simulation of the impact of high temperature stress on annual crop yields. *Agricultural and Forest Meteorology*, *135*(1), 180–189.

Challinor, A., Wheeler, T., Garforth, C., Craufurd, P., & Kassam, A. (2007). Assessing the vulnerability of food crop systems in Africa to climate change. *Climatic change*, *83*(3), 381–399.

Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of econometrics*, *92*(1), 1–45.

CRU (2016), Climatic Research Unit, University of East Anglia. Retrieved 9 July 2016, from: http://www.cru.uea.ac.uk/data

De Long, J. B. (1988). Productivity growth, convergence, and welfare: comment. *The American Economic Review*, *78*(5), 1138–1154.

Dell, M., Jones, B. F., & Olken, B. A. (2009). *Temperature and income: reconciling new cross-sectional and panel estimates*. National Bureau of Economic Research.

Dell, M., Jones, B. F., & Olken, B. A. (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics*, *4*(3), 66–95.

Dell, M., Jones, B. F., & Olken, B. A. (2014). What do we learn from the weather? The new climate–economy literature. *Journal of Economic Literature*, *52*(3), 740–798.

Dercon, S. (1996). Risk, crop choice, and savings: Evidence from Tanzania. *Economic development and cultural change*, *44*(3), 485–513.

Dercon, S. (2004). Growth and shocks: evidence from rural Ethiopia. *Journal of Development Economics*, *74*(2), 309–329.

Dercon, S., & Christiaensen, L. (2011). Consumption risk, technology adoption and poverty traps: Evidence from Ethiopia. *Journal of development economics*, *96*(2), 159–173.

Dercon, S., Hoddinott, J., & Woldehanna, T. (2005). Shocks and consumption in 15 Ethiopian villages, 1999-2004. *Journal of African economies*.

Dercon, S., & Krishnan, P. (2000). Vulnerability, seasonality and poverty in Ethiopia. *The Journal of Development Studies*, *36*(6), 25–53.

Deschênes, O., & Greenstone, M. (2011). Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US. *American Economic Journal: Applied Economics*, *3*(4), 152–185.

Diamond, J. (1999). *Guns, germs, and steel: The fates of human societies*. WW Norton & Company.

Dorward, A., Kydd, J., Morrison, J., & Urey, I. (2004). A policy agenda for pro-poor agricultural growth. *World development*, *32*(1), 73–89.

Easterly, W., & Levine, R. (2003). Tropics, germs, and crops: how endowments influence economic development. *Journal of monetary economics*, *50*(1), 3–39.

Elbers, C., Gunning, J. W., & Kinsey, B. (2007). Growth and risk: Methodology and micro evidence. *The World Bank Economic Review*, *21*(1), 1–20.

Fafchamps, M., Udry, C., & Czukas, K. (1998). Drought and saving in West Africa: are livestock a buffer stock? *Journal of Development economics*, *55*(2), 273–305.

Field, C. B., Barros, V. R., Mastrandrea, M. D., Mach, K. J., Abdrabo, M.-K., Adger, N., … others. (2014). Summary for policymakers. In *Climate change 2014: impacts, adaptation, and vulnerability. Part A: global and sectoral aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* (pagg. 1–32). Cambridge University Press.

Foley, J. A., Ramankutty, N., Brauman, K. A., Cassidy, E. S., Gerber, J. S., Johnston, M., … others. (2011). Solutions for a cultivated planet. *Nature*, *478*(7369), 337–342.

Gallup, J. L., Sachs, J. D., & Mellinger, A. D. (1999). Geography and economic development. *International regional science review*, *22*(2), 179–232.

Goldberg, M. S., Gasparrini, A., Armstrong, B., & Valois, M.-F. (2011). The short-term influence of temperature on daily mortality in the temperate climate of Montreal, Canada. *Environmental research*, *111*(6), 853–860.

Graff Zivin, J., & Neidell, M. (2014). Temperature and the allocation of time: Implications for climate change. *Journal of Labor Economics*, *32*(1), 1–26.

Hansen, B. E. (2000). Sample splitting and threshold estimation. *Econometrica*, *68*(3), 575–603.

Hausman, J. A., & Taylor, W. E. (1981). Panel data and unobservable individual effects. *Econometrica: Journal of the Econometric Society*, 1377–1398.

Heal, G., & Park, J. (2015). *Goldilocks economies? temperature stress and the direct impacts of climate change*. National Bureau of Economic Research.

Hirvonen, K. (2016). Temperature Changes, Household Consumption, and Internal Migration: Evidence from Tanzania. *American Journal of Agricultural Economics*, aaw042.

Hoddinott, J., & Kinsey, B. (2001). Child growth in the time of drought. *Oxford Bulletin of Economics and statistics*, *63*(4), 409–436.

Howden, S. M., Soussana, J.-F., Tubiello, F. N., Chhetri, N., Dunlop, M., & Meinke, H. (2007). Adapting agriculture to climate change. *Proceedings of the national academy of sciences*, *104*(50), 19691–19696.

Hsiang, S. M. (2010). Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America. *Proceedings of the National Academy of sciences*, *107*(35), 15367–15372.

Hsiang, S. M., & Jina, A. S. (2014). *The causal effect of environmental catastrophe on long-run economic growth: Evidence from 6,700 cyclones*. National Bureau of Economic Research.

Huntington, E. (1922). *Civilization and climate*. Yale University Press.

Jalan, J., & Ravallion, M. (2002). Geographic poverty traps? A micro model of consumption growth in rural China. *Journal of applied econometrics*, *17*(4), 329–346.

Kazianga, H., & Udry, C. (2006). Consumption smoothing? Livestock, insurance and drought in rural Burkina Faso. *Journal of Development Economics*, *79*(2), 413–446.

Knox, J., Hess, T., Daccache, A., & Wheeler, T. (2012). Climate change impacts on crop productivity in Africa and South Asia. *Environmental Research Letters*, *7*(3), 034032.

Lemoine, D., & Kapnick, S. (2015). A top-down approach to projecting market impacts of climate change. *Nature Climate Change* 6(5):514-519.

Letta, M., Tol, R.S.J, (2016). *Weather, climate and total factor productivity* (Working Paper Series No.102-2016). Department of Economics, University of Sussex.

Li, S., Wheeler, T., Challinor, A., Lin, E., Ju, H., & Xu, Y. (2010). The observed relationships between wheat and climate in China. *Agricultural and Forest Meteorology*, *150*(11), 1412–1419.

Ligon, E., & Schechter, L. (2003). Measuring vulnerability. *The Economic Journal*, *113*(486), C95–C102.

Mendelsohn, R., & Dinar, A. (1999). Climate change, agriculture, and developing countries: does adaptation matter? *The World Bank Research Observer*, *14*(2), 277–293.

Moore, F. C., & Diaz, D. B. (2015). Temperature impacts on economic growth warrant stringent mitigation policy. *Nature Climate Change*, *5*(2), 127–131.

Morduch, J. (1995). Income smoothing and consumption smoothing. *The journal of economic perspectives*, *9*(3), 103–114.

Mueller, N. D., Gerber, J. S., Johnston, M., Ray, D. K., Ramankutty, N., & Foley, J. A. (2012). Closing yield gaps through nutrient and water management. *Nature*, *490*(7419), 254–257.

Murphy, K. M., Shleifer, A., & Vishny, R. W. (1989). Industrialization and the big push. *Journal of political economy*, *97*(5), 1003–1026.

Niemelä, R., Hannula, M., Rautio, S., Reijula, K., & Railio, J. (2002). The effect of air temperature on labour productivity in call centres—a case study. *Energy and Buildings*, *34*(8), 759–764.

Nordhaus, W. D. (2006). Geography and macroeconomics: New data and new findings. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(10), 3510–3517.

Olsson, O., & Hibbs, D. A. (2005). Biogeography and long-run economic development. *European Economic Review*, *49*(4), 909–938.

Park, J. (2017). *Will We Adapt? Labor Productivity and Adaptation to Climate Change*. Cambridge, Massachusetts, USA: Harvard Environmental Economics Program.

Paxson, C. H. (1992). Using weather variability to estimate the response of savings to transitory income in Thailand. *The American Economic Review*, 15–33.

Pindyck, R. S. (2012). Uncertain outcomes and climate change policy. *Journal of Environmental Economics and management*, *63*(3), 289–303.

Pindyck, R. S. (2013). Climate change policy: What do the models tell us? *Journal of Economic Literature*, *51*(3), 860–872.

Ponomareva, M. (2010). Quantile regression for panel data models with fixed effects and small T: Identification and estimation. *University of Western Ontario*.

Porter, J. R., & Semenov, M. A. (2005). Crop responses to climatic variation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *360*(1463), 2021–2035.

Porter, J. R., Xie, L., Challinor, A. J., Cochrane, K., Howden, S. M., Iqbal, M. M., … Travasso, M. I. (2014). *Chapter 7: Food security and food production systems*. Cambridge University Press.

Quah, D. T. (1996). Twin peaks: growth and convergence in models of distribution dynamics. *The economic journal*, 1045–1055.

Quah, D. T. (1997). Empirics for growth and distribution: stratification, polarization, and convergence clubs. *Journal of economic growth*, *2*(1), 27–59.

Raffin, N. (2012). Children's environmental health, education, and economic development. *Canadian Journal of Economics/Revue canadienne d'économique*, *45*(3), 996–1022.

Rienecker, M. M., Suarez, M. J., Gelaro, R., Todling, R., Bacmeister, J., Liu, E., … others. (2011). MERRA: NASA’s modern-era retrospective analysis for research and applications. *Journal of climate*, *24*(14), 3624–3648.

Reis, R. (2009). The Time-Series Properties Of Aggregate Consumption: Implications For The Costs Of Fluctuations. *Journal of the European Economic Association*, *7*(4), 722‒753.

Rodrik, D., Subramanian, A., & Trebbi, F. (2004). Institutions rule: the primacy of institutions over geography and integration in economic development. *Journal of economic growth*, *9*(2), 131–165.

Roodman, D. (2006). How to do xtabond2: An introduction to difference and system GMM in Stata. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=982943

Rosenzweig, M. R., & Wolpin, K. I. (1993). Credit market constraints, consumption smoothing, and the accumulation of durable production assets in low-income countries: Investments in bullocks in India. *Journal of political economy*, *101*(2), 223–244.

Rowhani, P., Lobell, D. B., Linderman, M., & Ramankutty, N. (2011). Climate variability and crop production in Tanzania. *Agricultural and Forest Meteorology*, *151*(4), 449–460.

Sachs, J. (2008). The end of poverty: economic possibilities for our time. *European Journal of Dental Education*, *12*(s1), 17–21.

Schelling, T. C. (1992). Some economics of global warming. *The American Economic Review*, *82*(1), 1–14.

Schelling, T. C. (1995). Intergenerational discounting. *Energy policy*, *23*(4–5), 395–401.

Schelling, T. C. (1997). The cost of combating global warming: Facing the tradeoffs. *Foreign Affairs*, 8–14.

Schlenker, W., & Lobell, D. B. (2010). Robust negative impacts of climate change on African agriculture. *Environmental Research Letters*, *5*(1), 014010.

Schlenker, W., & Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to US crop yields under climate change. *Proceedings of the National Academy of sciences*, *106*(37), 15594–15598.

Schmidhuber, J., & Tubiello, F. N. (2007). Global food security under climate change. *Proceedings of the National Academy of Sciences*, *104*(50), 19703–19708.

Stern, N. (2013). The structure of economic modeling of the potential impacts of climate change: grafting gross underestimation of risk onto already narrow science models. *Journal of Economic Literature*, *51*(3), 838–859.

Strulik, H. (2008). Geography, health, and the pace of demo-economic development. *Journal of Development Economics*, *86*(1), 61–75.

Sudarshan, A., & Tewari, M. (2013). *The economic impacts of temperature on industrial productivity: Evidence from indian manufacturing*. Working Paper. http://icrier.org/pdf/working_paper_278.pdf

Tittonell, P., & Giller, K. E. (2013). When yield gaps are poverty traps: The paradigm of ecological intensification in African smallholder agriculture. *Field Crops Research*, *143*, 76–90.

Tol, R. S. J. (2011). *Poverty Traps and Climate Change* (Paper No. WP413). Economic and Social Research Institute (ESRI). https://ideas.repec.org/p/esr/wpaper/wp413.html

Tol, R. S. J. (2015). *Economic impacts of climate change* (Working Paper Series No. 7515). Department of Economics, University of Sussex. https://ideas.repec.org/p/sus/susewp/7515.html

Tol, R. S. J. (2016). *Dangerous Interference With The Climate System: An Economic Assessment* (Working Paper Series No. 100-2016). Department of Economics, University of Sussex. https://ideas.repec.org/p/sus/susewp/10016.html

Townsend, R. M. (1995). Consumption insurance: An evaluation of risk-bearing systems in low-income economies. *The Journal of Economic Perspectives*, *9*(3), 83–102.

Wang, Q., & others. (2015). Fixed-effect panel threshold model using Stata. *Stata Journal*, *15*(1), 121–134.

Weitzman, M. L. (2009). On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics*, *91*(1), 1–19.

Weitzman, M. L. (2010). What Is The" Damages Function" For Global Warming—And What Difference Might It Make? *Climate Change Economics*, *1*(01), 57–69.

Welch, J. R., Vincent, J. R., Auffhammer, M., Moya, P. F., Dobermann, A., & Dawe, D. (2010). Rice yields in tropical/subtropical Asia exhibit large but opposing sensitivities to minimum and maximum temperatures. *Proceedings of the National Academy of Sciences*, *107*(33), 14562–14567.

Windmeijer, F. (2005). A finite sample correction for the variance of linear efficient two-step GMM estimators. *Journal of econometrics*, *126*(1), 25–51.

Zivin, J. G., & Shrader, J. (2016). Temperature extremes, health, and human capital. *The Future of Children*, *26*(1), 31–50.

Table 1

Descriptive statistics

	Mean	Var	sd	Obs
Food consumption growth rate	-1.696	992.409	31.503	3168
Total consumption growth rate	-1.441	901.549	30.026	3170
Food consumption	13.117	0.318	0.564	4754
Total consumption	13.384	0.334	0.578	4755
∆Temp	0.083	0.105	0.324	3170
∆Pre	0.051	0.023	0.153	3170
Temp	23.755	7.260	2.694	3170
Pre	117.998	589.714	24.284	3170
Long-run average temperature	23.658	6.924	2.631	4755
Long-run average precipitation	114.747	576.907	24.019	4755
Household size	5.659	10.029	3.167	4755
Number of infants (< 5 years)	0.918	1.147	1.071	4755
Adult education level	4.593	8.338	2.888	4750
Age of the household head	49.615	241.137	15.529	4755
Gender of the household head	0.239	0.182	0.426	4755

Notes:

Food consumption growth rate is the annualised percentage change in household per adult equivalent food consumption between t and t-1. Total consumption growth rate is the annualised percentage change in household per adult equivalent consumption between t and t-1. Food consumption is the natural logarithm of household per adult-equivalent food consumption, expressed in Tanzanian shillings. Total consumption is the natural logarithm of household per adult-equivalent total consumption, expressed in Tanzanian shillings. ∆Temp is the difference between average monthly growing season temperature in the period between interviews and long-run (1980-2015) average monthly growing season temperature, divided by long-run (1980-2015) standard deviation, and expressed in degree Celsius. ∆Pre is the difference between average monthly growing season precipitation in the period between interviews and long run (1980-2015) average monthly growing season precipitation, divided by long-run (1980-2015) standard deviation, expressed in mm. Temp is average monthly growing season temperature in the period between interviews. Pre is average monthly growing season precipitation in the period between interviews. Long-run average temperature is the average monthly growing season temperature over the period 1980-2015, expressed in degree Celsius. Long-run average precipitation represents average monthly growing season precipitation over the period 1980-2015, expressed in mm. Adult education level represents the average years of education among adults, where adult means > 15 year old.

Table 2

FE regressions – Food consumption

Dependent variable: food consumption growth rate	(1)	(2)	(3)	(4)
L1.Food	-72.965***	-75.796***	-75.808***	-74.281***
	(1.219)	(1.299)	(1.304)	(1.326)
ΔTemp	-1.895	9.925*	11.093**	-338.600***
	(4.750)	(5.332)	(5.449)	(44.868)
Poor x ΔTemp		-21.588***	-21.460***	
		(4.537)	(4.541)	
Hot x ΔTemp			-2.653	
			(3.718)	
ΔPre	0.839	3.259	2.113	-4.941
	(6.673)	(8.386)	(9.339)	(6.622)
Poor x ΔPre		-8.758	-8.482	
		(9.620)	(9.673)	
Hot x ΔPre			2.127	
			(10.264)	
Hot			4.032	
			(3.689	
L1.Food x ΔTemp				25.713***
				(3.438)
Obs	3,164	3,164	3,164	3,164
Adj. R^2	0.831	0.835	0.835	0.841
Vegetation time series	Yes	Yes	Yes	Yes
Household controls	Yes	Yes	Yes	Yes
Total temperature effect for poorest households		-11.663**	-10.366*	
		(5.091)	(5.308)	
Total temperature effect for households in hot areas			8.441	
			(5.748)	
Total temperature effect for poorest households in hot areas			-13.019**	
			(5.482)	
Total precipitation effect for poorest households		-5.499	-6.329	
		(7.742)	(8.387)	
Total precipitation effect for households in hot areas			4.240	
			(10.401)	

Notes: All specifications include households FE, wave dummies, region x year FE and quarter of year dummies. Vegetation time series includes data on changes in crop greenness within growing season and onset of greenness increase and decrease. Household controls include household size, squared household size, age of the household head, squared age of the household head, gender of the household head, number of infants, adult education level and dummies for self-reported idiosyncratic and covariate shocks. Food consumption growth rate is the annualised percentage change in (ln) household per a.e. food consumption between t and t-1. L1.Food is lagged household per a.e. (ln) food consumption. ΔTemp is the difference between average monthly growing season temperature in the period between interviews and long-run (1980-2015) average monthly growing season temperature, divided by long-run (1980-2015) standard deviation, and expressed in degree Celsius. ΔPre is the difference between average monthly growing season precipitation in the period between interviews and long run (1980-2015) average monthly growing season precipitation, divided by long-run (1980-2015) standard deviation, expressed in mm. Poor is a dummy with value 1 for households with below median initial food consumption. Hot is a dummy with value 1 for households living in an area with an above mean long-run average monthly growing season temperature. Standard errors are in parentheses and are clustered at the EA and wave levels.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3

FE regressions – Total consumption

Dependent variable: total consumption growth rate	(1)	(2)	(3)	(4)
L1.Cons	-71.193***	-73.532***	-73.618***	-72.671***
	(1.299)	(1.380)	(1.387)	(1.338)
ΔTemp	-0.328	8.494*	9.199*	-319.134***
	(4.198)	(4.478)	(4.736)	(39.811)
Poor x ΔTemp		-17.813***	-17.565***	
		(3.748)	(3.739)	
Hot x ΔTemp			-1.645	
			(3.268)	
ΔPre	0.695	1.777	0.217	-6.080
	(5.848)	(7.452)	(8.279)	(5.597)
Poor x ΔPre		-5.771	-4.890	
		(8.412)	(8.495)	
Hot x ΔPre			2.370	
			(8.380)	
Hot			13.687***	
			(2.855)	
L1.Cons x ΔTemp				23.868***
				(2.988)
Obs	3,166	3,166	3,166	3,166
Adj. R^2	0.830	0.833	0.833	0.840
Vegetation time series	Yes	Yes	Yes	Yes
Household controls	Yes	Yes	Yes	Yes
Total temperature effect for poorest households		-9.319**	-8.366*	
		(4.694)	(4.897)	
Total temperature effect for households in hot areas			7.553	
			(4.846)	
Total temperature effect for poorest households in hot areas			-10.012**	
			(5.117)	
Total precipitation effect for poorest households		-3.994	-4.673	
		(6.747)	(7.235)	
Total precipitation effect for households in hot areas			2.587	
			(8.879)	

Notes: All specifications include households FE, wave dummies, region x year FE and quarter of year dummies. Vegetation time series includes data on changes in crop greenness within growing season and onset of greenness increase and decrease. Household controls include household size, squared household size, age of the household head, squared age of the household head, gender of the household head, number of infants, adult education level and dummies for self-reported idiosyncratic and covariate shocks. Total consumption growth rate is the annualised percentage change in (ln) household per a.e. total consumption between t and t-1. L1.Cons is lagged household per a.e. (ln) food consumption. ΔTemp is the difference between average monthly growing season temperature in the period between interviews and long-run (1980-2015) average monthly growing season temperature, divided by long-run (1980-2015) standard deviation, and expressed in degree Celsius. ΔPre is the difference between average monthly growing season precipitation in the period between interviews and long run (1980-2015) average monthly growing season precipitation, divided by long-run (1980-2015) standard deviation, expressed in mm. Poor is a dummy with value 1 for households with below median initial consumption. Hot is a dummy with value 1 for households living in an area with an above mean long-run average monthly growing season temperature. Standard errors are in parentheses and are clustered at the EA and wave levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4

FE initial quartile regressions – Food consumption

Dependent variable: Food consumption growth rate	(1)	(2)
L1.Food	-77.172***	-77.224***
	(1.344)	(1.351)
q1 x ΔTemp	-19.847***	-19.157***
	(5.164)	(5.338)
q2 x ΔTemp	-5.693	-4.985
	(5.332)	(5.403)
q3 x ΔTemp	4.604	5.234
	(5.659)	(5.944)
q4 x ΔTemp	16.115***	16.784***
	(5.844)	(5.909)
Hot x ΔTemp		-1.386
		(3.677)
q1 x ΔPre	-6.451	-8.752
	(10.031)	(10.497)
q2 x ΔPre	-4.833	-7.239
	(8.634)	(9.354)
q3 x ΔPre	5.244	2.913
	(9.904)	(10.943)
q4 x ΔPre	-2.776	-5.841
	(10.418)	(11.452)
Hot x ΔPre		7.024
		(10.337)
Hot		3.525
		(3.702)
Obs	3,164	3,164
Adj. R^2	0.837	0.837
Vegetation time series	Yes	Yes
Household controls	Yes	Yes

Notes: All specifications include households FE, wave dummies, region x year FE and quarter of year dummies. Vegetation time series includes data on changes in crop greenness within growing season and onset of greenness increase and decrease. Household controls include household size, squared household size, age of the household head, squared age of the household head, gender of the household head, number of infants, adult education level and dummies for self-reported idiosyncratic and covariate shocks. Food consumption growth rate is the annualised percentage change in (ln) household per a.e. food consumption between t and t-1. L1.Food is lagged household per a.e. (ln) food consumption. q1, q2, q3, q4 are initial food consumption quartiles. ΔTemp is the difference between average monthly growing season temperature in the period between interviews and long-run (1980-2015) average monthly growing season temperature, divided by long-run (1980-2015) standard deviation, and expressed in degree Celsius. ΔPre is the difference between average monthly growing season precipitation in the period between interviews and long run (1980-2015) average monthly growing season precipitation, divided by long-run (1980-2015) standard deviation, expressed in mm. Hot is a dummy with value 1 for households living in an area with above mean long-run average monthly growing season temperature. Standard errors are in parentheses and are clustered at the EA and wave levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5

FE initial quartile regressions – Total consumption

Dependent variable: Total consumption growth rate	(1)	(2)
L1.Cons	-75.155***	-75.297***
	(1.378)	(1.387)
q1 x ΔTemp	-14.965***	-15.279***
	(5.068)	(5.098)
q2 x ΔTemp	-3.732	-3.738
	(5.504)	(5.666)
q3 x ΔTemp	1.483	1.034
	(4.734)	(5.323)
q4 x ΔTemp	18.664***	18.436***
	(5.565)	(5.624)
Hot x ΔTemp		0.780
		(3.451)
q1 x ΔPre	-3.016	-5.158
	(9.118)	(9.555)
q2 x ΔPre	-6.526	-7.999
	(8.921)	(9.254)
q3 x ΔPre	3.671	0.846
	(8.803)	(9.925)
q4 x ΔPre	-5.478	-8.184
	(10.307)	(10.928)
Hot x ΔPre		6.415
		(8.563)
Hot		14.725***
		(2.894)
Obs	3,166	3,166
Adj. R^2	0.837	0.837
Vegetation time series	Yes	Yes
Household controls	Yes	Yes

Notes: All specifications include households FE, wave dummies, region x year FE and quarter of year dummies. Vegetation time series includes data on changes in crop greenness within growing season and onset of greenness increase and decrease. Household controls include household size, squared household size, age of the household head, squared age of the household head, gender of the household head, number of infants, adult education level and dummies capturing self-reported idiosyncratic and covariate shocks. Total consumption growth rate is the annualised percentage change in (ln) household per a.e. consumption between t and t-1. L1.Cons is lagged household per a.e. (ln) consumption. q1, q2, q3, q4 are initial consumption quartiles. ΔTemp is the difference between average monthly growing season temperature in the period between interviews and long-run (1980-2015) average monthly growing season temperature, divided by long-run (1980-2015) standard deviation, and expressed in degree Celsius. ΔPre is the difference between average monthly growing season precipitation in the period between interviews and long run (1980-2015) average monthly growing season precipitation, divided by long-run (1980-2015) standard deviation, expressed in mm. Hot is a dummy with value 1 for households living in an area with above mean long-run average monthly growing season temperature. Standard errors are in parentheses and are clustered at the EA and wave levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6

Double threshold model – Hansen Estimator

Dependent variable:	(1) ΔFood	(2) ΔCons
L1.Food	-74.698***	
	(1.256)	
L1.Cons		-72.326***
		(1.295)
ΔPre	-2.645	-5.367
	(6.709)	(5.809)
ΔTemp_Lower regime	-14.682***	-18.347***
	(4.878)	(4.863)
ΔTemp_Medium regime	5.340	1.383
	(4.846)	(4.386)
ΔTemp_Upper regime	29.135***	28.953***
	(6.811)	(6.638)
Obs	3,168	3,170
Adj. R^2	0.775	0.770
Vegetation time series	Yes	Yes
Household controls	Yes	Yes

Notes: All specifications include households FE, wave dummies, region x year FE and quarter of year dummies. Vegetation time series includes data on changes in crop greenness within growing season and onset of greenness increase and decrease. Household controls include household size, squared household size, age of the household head, squared age of the household head, gender of the household head, number of infants, adult education level and dummies capturing self-reported idiosyncratic and covariate shocks. ΔFood is the annualised percentage change in (ln) household per a.e. food consumption between t and t-1. ΔCons is the annualised percentage change in (ln) household per a.e. consumption between t and t-1. L1.Food is lagged household per a.e. (ln) food consumption. L1.Cons is lagged household per a.e. (ln) consumption. ΔTemp is the difference between average monthly growing season temperature in the period between interviews and long-run (1980-2015) average monthly growing season temperature, divided by long-run (1980-2015) standard deviation, and expressed in degree Celsius. ΔPre is the difference between average monthly growing season precipitation in the period between interviews and long run (1980-2015) average monthly growing season precipitation, divided by long-run (1980-2015) standard deviation, expressed in mm. Standard errors are in parentheses and are clustered at the EA and wave levels.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Threshold Confidence intervals and effect tests

<u>Column (1) – Food consumption</u>

1) Threshold estimator (level = 95):

Model	Threshold	Lower	Upper
Th-1	13.089	13.086	13.093
Th-21	13.089	13.084	13.093
Th-22	13.729	13.709	13.733

2) Threshold effect test (bootstrap = 300 300):

Threshold	RSS	MSE	Fstat	Prob	Crit10	Crit5	Crit1
Single	5.12e+05	161.770	141.92	0.000	17.715	22.171	27.298
Double	5.04e+05	159.234	50.43	0.000	20.140	22.664	26.723

3) Percentage of observations in each regime:

Lower regime:	47.16 %
Medium regime:	39.90 %
Upper regime:	12.94 %

Column (2) – Total consumption

1) Threshold estimator (level = 95):

Model	Threshold	Lower	Upper
Th-1	13.297	13.285	13.300
Th-21	12.983	12.979	12.991
Th-22	14.014	14.005	14.024

2) Threshold effect test (bootstrap = 300 300):

Threshold	RSS	MSE	Fstat	Prob	Crit10	Crit5	Crit1
Single	4.72e+05	148.891	113.50	0.000	16.957	19.678	26.294
Double	4.61e+05	145.622	73.09	0.000	18.415	22.431	29.031

3) Percentage of observations in each regime:

Lower regime:	23.56 %
Medium regime:	63.47 %
Upper regime:	12.97 %

Table 7

FE regressions with spatially-robust SEs

Dependent variable:	(1) ΔFood	(2) ΔCons
L1.Food	-77.224	
Conley(1999), 50 km cut-off	(0.911)***	
Conley(1999), 75 km cut-off	(0.914)***	
Conley(1999), 100 km cut-off	(0.943)***	
L1.Cons		-75.297
Conley(1999), 50 km cut-off		(1.007)***
Conley(1999), 75 km cut-off		(1.037)***
Conley(1999), 100 km cut-off		(1.087)***
q1 x ΔTemp	-19.157	-15.279
Conley(1999), 50 km cut-off	(3.679)***	(3.246)***
Conley(1999), 75 km cut-off	(3.744)***	(3.255)***
Conley(1999), 100 km cut-off	(3.823)***	(3.252)***
q2 x ΔTemp	-4.985	-3.738
Conley(1999), 50 km cut-off	(3.473)	(3.572)
Conley(1999), 75 km cut-off	(3.392)	(3.485)
Conley(1999), 100 km cut-off	(3.370)	(3.322)
q3 x ΔTemp	5.324	1.034
Conley(1999), 50 km cut-off	(3.704)	(3.466)
Conley(1999), 75 km cut-off	(3.658)	(3.427)
Conley(1999), 100 km cut-off	(3.632)	(3.390)
q4 x ΔTemp	16.784	18.436
Conley(1999), 50 km cut-off	(3.572)***	(3.539)***
Conley(1999), 75 km cut-off	(3.492)***	(3.501)***
Conley(1999), 100 km cut-off	(3.409)***	(3.454)***
Hot x ΔTemp	-1.386	0.780
Conley(1999), 50 km cut-off	(2.280)	(2.118)
Conley(1999), 75 km cut-off	(2.293)	(2.124)
Conley(1999), 100 km cut-off	(2.306)	(2.080)
q1 x ΔPre	-8.752	-5.158
Conley(1999), 50 km cut-off	(7.473)	(6.665)
Conley(1999), 75 km cut-off	(7.185)	(6.487)
Conley(1999), 100 km cut-off	(7.167)	(6.459)
q2 x ΔPre	-7.239	-7.999
Conley(1999), 50 km cut-off	(6.245)	(5.942)
Conley(1999), 75 km cut-off	(6.116)	(5.752)
Conley(1999), 100 km cut-off	(6.288)	(5.596)
q3 x ΔPre	2.913	0.846
Conley(1999), 50 km cut-off	(6.898)	(6.512)

Conley(1999), 75 km cut-off	(6.956)	(6.436)
Conley(1999), 100 km cut-off	(7.128)	(6.434)
q4 x ΔPre	-5.841	-8.184
Conley(1999), 50 km cut-off	(7.080)	(7.142)
Conley(1999), 75 km cut-off	(7.085)	(6.999)
Conley(1999), 100 km cut-off	(7.023)	(6.908)
Hot x ΔPre	7.024	6.415
Conley(1999), 50 km cut-off	(6.520)	(5.557)
Conley(1999), 75 km cut-off	(6.527)	(5.616)
Conley(1999), 100 km cut-off	(6.686)	(5.758)
Hot	3.525	14.725
Conley(1999), 50 km cut-off	(6.588)	(5.201)***
Conley(1999), 75 km cut-off	(6.586)	(5.207)***
Conley(1999), 100 km cut-off	(6.513)	(5.244)***
Obs	3,164	3.166
Adj. R^2	0.768	0.765
Vegetation time series	Yes	Yes
Household controls	Yes	Yes

Notes: All specifications include households FE, wave dummies, region x year FE and quarter of year dummies. Vegetation time series includes data on changes in crop greenness within growing season and onset of greenness increase and decrease. Household controls include household size, squared household size, age of the household head, squared age of the household head, gender of the household head, number of infants, adult education level and dummies for self-reported idiosyncratic and covariate shocks: ΔFood is the annualised percentage change in (ln) household per a.e. food consumption between t and t-1. ΔCons is the annualised percentage change in (ln) household per a.e. consumption between t and t-1.. L1.Food is lagged household per a.e. (ln) food consumption. L1.Cons is lagged household per a.e. (ln) consumption. q1, q2, q3, q4 are initial food consumption quartiles in Column (1) and initial consumption quartiles in Column(2). ΔTemp is the difference between average monthly growing season temperature in the period between interviews and long-run (1980-2015) average monthly growing season temperature, divided by long-run (1980-2015) standard deviation, and expressed in degree Celsius. ΔPre is the difference between average monthly growing season precipitation in the period between interviews and long run (1980-2015) average monthly growing season precipitation, divided by long-run (1980-2015) standard deviation, expressed in mm. Hot is a dummy with value 1 for households living in an area with above mean long-run average monthly growing season temperature. Conley (1999) standard errors are in parentheses and are robust to both spatial and temporal autocorrelation.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8

FE initial quartile regressions - Alternative weather data

Dependent Variables:	(1) ΔFood	(2) ΔFood	(3) ΔCons	(4) ΔCons
L1.Food	-76.191***	-76.234***		
	(1.343)	(1.347)		
L1.Cons			-74.291***	-74.270***
			(1.392)	(1.396)
q1 x ΔTemp	-14.205***	-14.636***	-10.985**	-11.147**
	(4.602)	(4.736)	(4.622)	(4.786)
q2 x ΔTemp	-5.339	-5.963	-3.778	-3.964
	(5.507)	(5.559)	(5.031)	(5.073)
q3 x ΔTemp	-0.051	-0.649	-1.797	-2.111
	(5.768)	(5.838)	(4.809)	(4.931)
q4 x ΔTemp	15.130**	14.725**	19.063***	18.940***
	(5.897)	(5.906)	(5.058)	(5.155)
Hot x ΔTemp		2.090		2.623
		(2.453)		(2.342)
q1 x ΔPre	-0.009	-0.002	-3.819***	-0.004
	(0.011)	(0.012)	(1.295)	(0.011)
q2 x ΔPre	-0.001	0.007	1.124	0.003
	(0.010)	(0.010)	(0.763)	(0.011)
q3 x ΔPre	0.019**	0.027**	0.407	0.017
	(0.009)	(0.011)	(1.169)	(0.011)
q4 x ΔPre	0.025**	0.036***	5.007***	0.030**
	(0.010)	(0.013)	(1.250)	(0.014)
Hot x ΔPre		-0.021*		-0.018
		(0.012)		(0.011)
Hot		2.193		10.960***
		(3.445)		(3.198)
Obs	3,164	3,164	3,166	3,166
Adj. R^2	0.835	0.836	0.835	0.836
Vegetation time series	Yes	Yes	Yes	Yes
Household controls	Yes	Yes	Yes	Yes

Notes: All specifications include households FE, wave dummies, region x year FE and quarter of year dummies. Vegetation time series includes data on changes in crop greenness within growing season and onset of greenness increase and decrease. Household controls include household size, squared household size, age of the household head, squared age of the household head, gender of the household head, number of infants, adult education level and dummies capturing self-reported idiosyncratic and covariate shocks. ΔFood is the annualised percentage change in (ln) household per a.e. food consumption between t and t-1. ΔCons is the annualised percentage change in (ln) household per a.e. consumption between t and t-1. L1.Food is lagged household per a.e. (ln) food consumption. L1.Cons is lagged household per a.e. (ln) consumption. q1, q2, q3, q4 are initial food consumption quartiles in Column (1) and initial consumption quartiles in Column (2). ΔTemp is the difference between average monthly growing season temperature in the period between interviews and long-run (1983-2015) average monthly growing season temperature, divided by long-run (1983-2015) standard deviation, and expressed in degree Celsius. ΔPre is the difference between total precipitation during the previous wettest quarter and long-run average (2001 – 2013) total precipitation during the wettest quarter, expressed in mm. Hot is a dummy with value 1 for households living in an area with above mean long-run average monthly growing season temperature. Standard errors are in parentheses and are clustered at the EA and wave levels.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9

Hausman – Taylor regressions

Dependent variables:	(1) ΔFood	(2) ΔCons
L1.Food	-75.877***	
	(1.302)	
L1.Cons		-74.520***
		(1.277)
q1 x ΔTemp	-21.797***	-18.784***
	(3.888)	(3.625)
q2 x ΔTemp	-9.955***	-9.270**
	(3.818)	(4.064)
q3 x ΔTemp	-1.894	-4.931
	(4.615)	(3.942)
q4 x ΔTemp	6.179	10.206**
	(4.372)	(4.248)
q1 x ΔPre	-10.441	-7.986
	(7.424)	(7.642)
q2 x ΔPre	-3.261	-7.862
	(8.020)	(7.216)
q3 x ΔPre	2.289	-0.584
	(8.345)	(7.064)
q4 x ΔPre	-0.020	-1.665
	(8.364)	(8.918)
Long-run average temperature	1.049*	1.246**
	(0.557)	(0.588)
Long-run average precipitation	0.132**	0.129*
	(0.067)	(0.069)
Obs	3,164	3,166
Vegetation time series	Yes	Yes
Household controls	Yes	Yes

Notes: All specifications include wave, region, year and quarter of year dummies. All household controls are treated as time-varying endogenous variables with the exception of self-reported covariate shocks. Distance (in KMs) to nearest major road is included and treated as time-invariant exogenous. ΔFood is the between-wave percentage change in (ln) household per a.e. food consumption. ΔFood is the annualised percentage change in (ln) household per a.e. food consumption between t and t-1. L1.Food is lagged household per a.e. (ln) food consumption and is treated as endogenous. ΔCons is the annualised percentage change in (ln) household per a.e. consumption between t and t-1. L1.Cons is lagged household per a.e. (ln) consumption and is treated as endogenous. q1, q2, q3, q4 are food consumption quartiles in Column (1) and total consumption quartiles in Column (2); they are all treated as time-invariant, endogenous variables. standard deviation, expressed in mm. ΔTemp is the difference between average monthly growing season temperature in the period between interviews and long-run (1980-2015) average monthly growing season temperature, divided by long-run (1980-2015) standard deviation, and expressed in degree Celsius. ΔPre is the difference between average monthly growing season precipitation in the period between interviews and long run (1980-2015) average monthly growing season precipitation, divided by long-run (1980-2015) standard deviation, expressed in mm. All the weather variables are treated as exogenous. Standard errors are in parentheses and are clustered at the household level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10

Two-step Difference GMM

Dependent variables:	(1) ΔFood	(2) ΔCons
L1.Food	-70.120***	
	(7.108)	
L1.Cons		-74.439***
		(5.701)
q1 x ΔTemp	-19.993***	-20.437***
	(6.929)	(6.065)
q2 x ΔTemp	-9.166	-7.303
	(5.769)	(5.928)
q3 x ΔTemp	-7.351	-1.323
	(6.051)	(6.254)
q4 x ΔTemp	4.081	10.417
	(8.389)	(7.461)
q1 x ΔPre	0.806	-2.193
	(9.327)	(9.242)
q2 x ΔPre	-3.949	-0.300
	(10.617)	(10.181)
q3 x ΔPre	8.584	12.033
	(12.051)	(10.431)
q4 x ΔPre	2.755	-3.414
	(12.770)	(12.652)
Obs	1,581	1.533
Vegetation time series	Yes	Yes
Household controls	Yes	Yes
Hansen – J test (p)	0.584	0.510

Notes: All specifications include households FE, wave dummies, year FE and quarter of year dummies. Region x time FE are used as additional instruments. All household controls are treated as endogenous. ΔFood is the annualised percentage change in (ln) household per a.e. food consumption between t and t-1. L1.Food is lagged household per a.e. (ln) food consumption and is treated as endogenous. ΔCons is the annualised percentage change in (ln) household per a.e. consumption between t and t-1. L1.Cons is lagged household per a.e. (ln) consumption and is treated as endogenous. q1, q2, q3, q4 are initial food consumption quartiles in Column (1) and initial total consumption quartiles in Column (2). ΔTemp is the difference between average monthly growing season temperature in the period between interviews and long-run (1980-2015) average monthly growing season temperature, divided by long-run (1980-2015) standard deviation, and expressed in degree Celsius. ΔPre is the difference between average monthly growing season precipitation in the period between interviews and long run (1980-2015) average monthly growing season precipitation, divided by long-run (1980-2015) standard deviation, expressed in mm. Weather variables and the vegetation time series variables are treated as exogenous. Robust standard errors are in parentheses and are corrected using Windmeijer's procedure. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11

Health expenditure channel

Dependent variable: Share of health expenditure growth rate	(1)	(2)	(3)
L1.Share of health expenditure	-73.730*** (1.235)	-73.599*** (1.232)	-73.531*** (1.274)
ΔTemp	0.869 (20.100)		
ΔPre	-0.248 (27.421)		
q1 x ΔTemp		-4.911 (20.097)	-8.074 (21.462)
q2 x ΔTemp		0.789 (22.606)	-1.778 (23.226)
q3 x ΔTemp		15.652 (23.579)	11.926 (25.551)
q4 x ΔTemp		-6.589 (22.664)	-9.452 (23.241)
Hot x ΔTemp			7.174 (14.605)
q1 x ΔPre		8.047 (36.340)	6.092 (38.827)
q2 x ΔPre		-15.644 (33.183)	-16.615 (34.962)
q3 x ΔPre		7.444 (36.515)	3.828 (40.909)
q4 x ΔPre		9.122 (39.249)	7.140 (41.776)
Hot x ΔPre			7.824 (43.651)
Hot			23.312** (2.153)
Obs.	2,952	2,952	2,952
Adj. R^2	0.820	0.820	0.821
Vegetation time series	Yes	Yes	Yes
Household controls	Yes	Yes	Yes

Notes: All specifications include households FE, wave dummies, region x year FE and quarter of year dummies. Vegetation time series includes data on changes in crop greenness within growing season and onset of greenness increase and decrease. Household controls include household size, squared household size, age of the household head, squared age of the household head, gender of the household head, number of infants, adult education level and dummies capturing self-reported idiosyncratic and covariate shocks. q1, q2, q3, q4 are initial consumption quartiles. Dependent variable is (ln) per a.e. between-wave percentage of the health expenditure / total expenditure ratio. L1.Share of health expenditure is lagged ln per a.e. health expenditure / total expenditure ratio. ΔTemp is the difference between average monthly growing season temperature in the period between interviews and long-run (1980-2015) average monthly growing season temperature, divided by long-run (1980-2015) standard deviation, and expressed in degree Celsius. ΔPre is the difference between average monthly growing season precipitation in the period between interviews and long run (1980-2015) average monthly growing season precipitation, divided by long-run (1980-2015) standard deviation, expressed in mm. Hot is a dummy with value 1 for households living in an area with above mean long-run average monthly growing season temperature. Standard errors are in parentheses and are clustered at the EA and wave levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 12

Labour productivity channel – Two-step Difference GMM

Dependent variable:	(1) ΔALP	(2) ΔALP
L1.ALP	-70.206*** (4.677)	-70.657*** (4.785)
ΔTemp	-23.658** (11.942)	
ΔPre	-14.961 (14.656)	
q1 x ΔTemp		-32.790** (12.933)
q2 x ΔTemp		-20.770 (13.975)
q3 x ΔTemp		-18.107 (18.216)
q4 x ΔTemp		-17.618 (14.574)
q1 x ΔPre		-30.274 (27.740)
q2 x ΔPre		1.864 (20.903)
q3 x ΔPre		-18.424 (27.794)
q4 x ΔPre		-14.348 (26.282)
Obs	1,130	1,130
Vegetation time series	Yes	Yes
Household controls	Yes	Yes
Hansen – J test (p)	0.235	0.247

Notes: All specifications include households FE, wave dummies, year FE and quarter of year dummies. Region x time FE and month of interview dummies are used as additional instruments. All household controls are treated as endogenous with the exception of self-reported covariate shocks. ΔALP is agricultural labour productivity growth between t and t-1. L1.ALP is lagged (ln) agricultural labour productivity, instrumented using lagged assets and education levels at t-1. q1, q2, q3, q4 are initial total consumption quartiles. ΔTemp is the difference between average monthly growing season temperature in the period between interviews and long-run (1980-2015) average monthly growing season temperature, divided by long-run (1980-2015) standard deviation, and expressed in degree Celsius. ΔPre is the difference between average monthly growing season precipitation in the period between interviews and long run (1980-2015) average monthly growing season precipitation, divided by long-run (1980-2015) standard deviation, expressed in mm. Weather variables and the vegetation time series variables are treated as exogenous. Robust standard errors are in parentheses and are corrected using Windmeijer's procedure. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 13

Agricultural yield channel - Crop yields

Dependent variable: Crop yield	(1)	(2)	(3)
Number of GDDs (8-34 °C)	0.000	0.000	
	(0.001)	(0.001)	
Number of GDDs (34 + °C)	-0.020**	-0.022**	
	(0.010)	(0.010)	
Precipitation	-0.000	-0.000	
	(0.002)	(0.002)	
$(\text{Precipitation})^2$	0.000	0.000	
	(0.000)	(0.000)	
Maize & paddy non-specializers x Number of GDDs (8-34 °C)		-0.000	
		(0.000)	
Maize & paddy non-specializers x Number of GDDs (34 + °C)		0.026	
		(0.021)	
Maize & paddy non-specializers		0.460	
		(1.099)	
q1 x Number of GDDs (34 + °C)			-0.052***
			(0.016)
q2 x Number of GDDs (34 + °C)			-0.020
			(0.015)
q3 x Number of GDDs (34 + °C)			-0.017
			(0.011)
q4 x Number of GDDs (34 + °C)			0.011
			(0.021)
q1 x Precipitation			0.001
			(0.003)
q2 x Precipitation			-0.000
			(0.004)
q3 x Precipitation			-0.000
			(0.002)
q4 x Precipitation			-0.002
			(0.004)
q1 x $(\text{Precipitation})^2$			-0.000
			(0.000)
q2 x $(\text{Precipitation})^2$			-0.000
			(0.000)
q3 x $(\text{Precipitation})^2$			0.000
			(0.000)
q4 x $(\text{Precipitation})^2$			0.000
			(0.000)
Obs	3,537	3,537	3,537
Adj. R^2	0.595	0.599	0.599
Vegetation time series	Yes	Yes	Yes

Total effect of Number of GDDs (8-34 °C) for households not specialized in maize and paddy production	-0.000 (0.001)
Total effect of Number of GDDs (34 + °C) for households not specialized in maize and paddy production	0.005 (0.021)

Notes: All specifications include households FE, wave dummies, region x year FE and quarter of year dummies. Vegetation time series includes data on changes in crop greenness within growing season and onset of greenness increase and decrease. Crop yield is average crop yield (kg / ha) during the previous two rainy seasons. 'Maize & paddy non-specializers' is a dummy with value 1 for households in which maize and paddy account for less than 50% of total crop production in a given wave. q1, q2, q3, q4 are initial total consumption quartiles. Standard errors are in parentheses and are clustered at the EA and wave levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 14
FE regressions – Asset growth as the dependent variable

Dependent variable: Asset growth	(1)	(2)
L1.Assets	-74.762***	-75.053***
	(1.834)	(1.832)
ΔTemp	-5.823	
	(22.094)	
ΔPre	-27.314	
	(32.146)	
q1 x ΔTemp		-2.823
		(24.355)
q2 x ΔTemp		-3.731
		(25.099)
q3 x ΔTemp		-16.042
		(29.640)
q4 x ΔTemp		-4.402
		(28.642)
q1 x ΔPre		66.468
		(44.547)
q2 x ΔPre		-75.504*
		(42.217)
q3 x ΔPre		-80.426*
		(48.702)
q4 x ΔPre		-29.418
		(59.022)
Obs	2,223	2,223
Adj. R^2	0.800	0.804
Vegetation time series	Yes	Yes
Household controls	Yes	Yes

Notes: All specifications include households FE, wave dummies, region x year FE and quarter of year dummies. Vegetation time series includes data on changes in crop greenness within growing season and onset of greenness increase and decrease. Household controls include household size, squared household size, age of the household head, squared age of the household head, gender of the household head, number of infants, adult education level and dummies capturing self-reported idiosyncratic and covariate shocks. Asset growth is the annualised percentage change in (ln) household per a.e. household Tropical Livestock Units (TLUs) between t and t-1. L1.Assets is lagged household per a.e. (ln) asset level (TLUs). q1, q2, q3, q4 are initial consumption quartiles. ΔTemp is the difference between average monthly growing season temperature in the period between interviews and long-run (1980-2015) average monthly growing season temperature, divided by long-run (1980-2015) standard deviation, and expressed in degree Celsius. ΔPre is the difference between average monthly growing season precipitation in the period between interviews and long run (1980-2015) average monthly growing season precipitation, divided by long-run (1980-2015) standard deviation, expressed in mm. Standard errors are in parentheses and are clustered at the EA and wave levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 1

Marginal effect of ΔTemp on food consumption growth

at different lagged food consumption levels

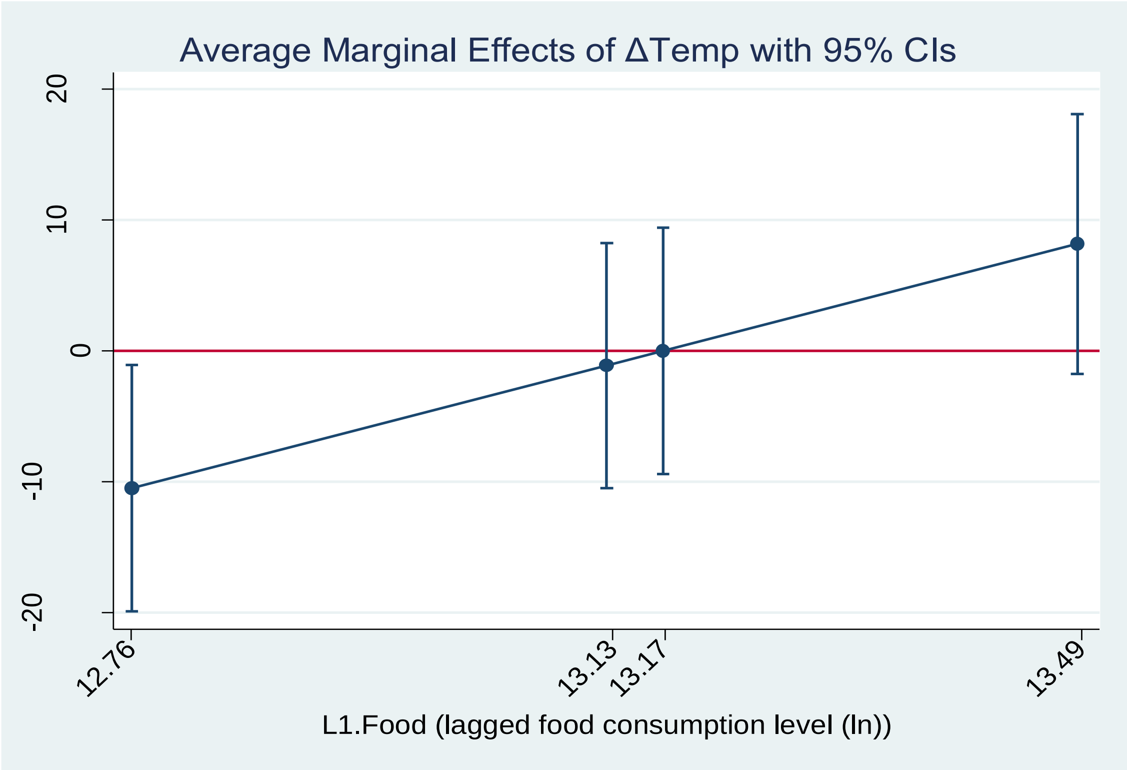

Figure 2

Marginal effect of ΔTemp on total consumption growth at different lagged total consumption levels

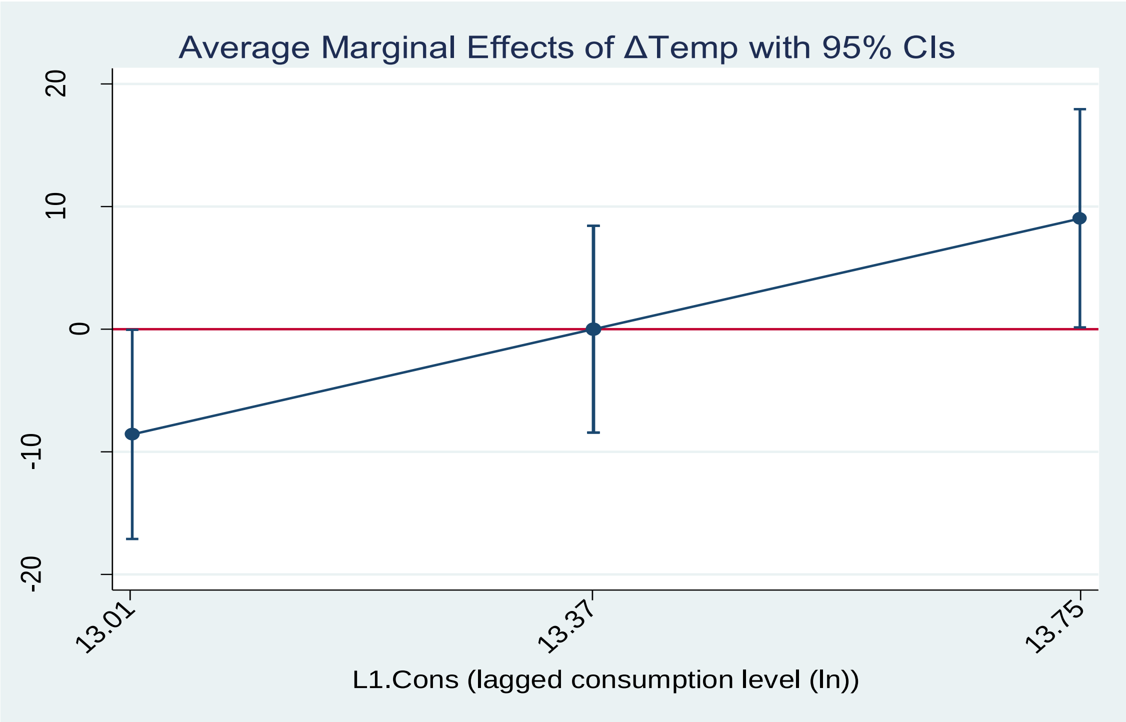

Appendix

Table A.1

Instrumented FE regressions – Endogeneity tests

Dependent variable:	(1) ΔFood	(2) ΔCons
L1.Food	-98.481*	
	(51.761)	
L1.Cons		-91.758***
		(29.374)
ΔTemp	2.846	2.607
	(7.394)	(4.352)
ΔPre	2.440	-0.306
	(6.178)	(7.209)
Observations	3092	3094
Adjusted R-squared	0.304	0.342
Vegetation time series	Yes	Yes
Household controls	Yes	Yes

Notes: L1.Food is lagged household per a.e. (ln) food consumption, instrumented using lagged assets and education levels at t-1. L1.Cons is lagged household per a.e. (ln) total consumption, instrumented using lagged assets and education levels at t-1. ΔTemp is the difference between average monthly growing season temperature in the period between interviews and long-run (1980-2015) average monthly growing season temperature, divided by long-run (1980-2015) standard deviation, and expressed in degree Celsius. ΔPre is the difference between average monthly growing season precipitation in the period between interviews and long run (1980-2015) average monthly growing season precipitation, divided by long-run (1980-2015) standard deviation, expressed in mm Standard errors are in parentheses and are clustered at the household and wave levels .
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Endogeneity tests:

Regressor	Test	p-value
L1.Food	0.074	0.786
L1.Cons	0.423	0.515

Table A.2

Descriptive statistics – Alternative weather data

	Mean	Var	sd	Obs
ΔTemp	0.405	0.131	0.363	3170
ΔPre	-21.565	8585.501	92.658	4755
Long-run average temperature	23.948	4.362	2.089	4755
Long-run average precipitation	502.203	19198.690	138.559	4755

Notes:

ΔTemp is the difference between average monthly growing season temperature in the period between interviews and long-run (1983-2015) average monthly growing season temperature divided by long-run (1983-2013) standard deviation, and expressed in degree Celsius. ΔPre is the difference between total precipitation during the previous wettest quarter and long-run average (2001 - 2013) total precipitation during the wettest quarter divided by average decadal (2001 - 2013) standard deviation, expressed in mm. Long-run average temperature is the average monthly growing season temperature over the period 1983-2015, expressed in degree Celsius. Long-run average precipitation represents long-run average (2001 - 2013) total precipitation during the wettest quarter. Data source is the *CRUCY Version 3.23* by the University of East Anglia for temperature data, and the Tanzania LSMS-ISA NPS surveys for rainfall data.

Table A.3

Labour productivity channel – Endogeneity test

Dependent variable:	(1) ΔALP
L1.ALP	-227.889
	(220.885)
ΔTemp	-58.596
	(92.139)
ΔPre	-15.059
	(71.783)
Observations	2260
Vegetation time series	Yes
Household controls	Yes

Notes: ΔALP is agricultural labour productivity growth between t and t-1. L1.ALP is lagged (ln) agricultural labour productivity, instrumented using lagged assets and education levels at t-1. ΔTemp is the difference between average monthly growing season temperature in the period between interviews and long-run (1980-2015) average monthly growing season temperature, divided by long-run (1980-2015) standard deviation, and expressed in degree Celsius. ΔPre is the difference between average monthly growing season precipitation in the period between interviews and long run (1980-2015) average monthly growing season precipitation, divided by long-run (1980-2015) standard deviation, expressed in mm Standard errors are in parentheses and are clustered at the household and wave levels .
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Endogeneity test:

Regressor	Test	p-value
L1.ALP	7.611	0.0058

Table A.4

Descriptive statistics –Agricultural Wealth Index

Variable: Agricultural Wealth Index				
	Mean	Var	sd	Obs
q1	0.066	1.151	1.073	905
q2	0.097	1.054	1.027	981
q3	0.018	0.841	0.917	931
q4	0.228	1.878	1.370	836

Notes: q1, q2, q3, q4 are initial consumption quartiles.

Agricultural Wealth Index is from the FAO Rural Income Generating Activities (RIGA) Team.

Table A.5

Descriptive statistics – Main source of income

Variable: Main source of income is not farming (in at least two periods) - % of households		
	Yes	No
Whole sample	24.61	75.39
q1	19.40	80.60
q2	20	80
q3	25.25	74.75
q4	33.75	66.25

Notes: q1, q2, q3, q4 are initial consumption quartiles.

Table A.6

Descriptive statistics – Growing degree days

	Mean	Var	sd	Obs
Number of GDDs (8-34 °C)	3905.047	389495.400	624.096	4755
Number of GDDs (34 + °C)	3.280	46.273	6.802	4755

Table A.7

Descriptive statistics – Maize and paddy as a share of total crop production

	Maize and paddy account for 50% or more of total crop production - % of households	
	Yes	No
q1	50.59	49.41
q2	58.44	41.46
q3	51.60	48.40
q4	47.81	52.19

Notes: q1, q2, q3, q4 are initial consumption quartiles.

Table A.8

Descriptive statistics – Average crop yield and quantity produced

	(1) Mean quantity (kg)	(2) Mean crop yield (kg / ha)	Obs
q1	1268.625	715.602	876
q2	1452.362	1033.638	965
q3	1479.123	1225.526	903
q4	1762.087	1201.825	793

Notes: q1, q2, q3, q4 are initial consumption quartiles.

Table A.9

Descriptive statistics – Irrigation

	Use of irrigation in the previous long rainy season - % of households	
	Yes	No
q1	1.95	98.05
q2	3.30	96.70
q3	3.84	96.16
q4	6.05	93.95

Notes: q1, q2, q3, q4 are initial consumption quartiles.

Table A.10

Descriptive statistics – Inorganic fertilizers

	Use of inorganic fertilizers in the previous long rainy season - % of households	
	Yes	No
q1	17.65	82.35
q2	19.10	80.81
q3	25.25	74.75
q4	23.46	76.54

Notes: q1, q2, q3, q4 are initial consumption quartiles.

Table A.11

Descriptive statistics – Use of improved maize seeds on at least one plot

Variable: Use of improved maize seeds on at least one plot across waves - % of households		
	Yes	No
q1	34.16	65.84
q2	41.24	58.76
q3	46.48	53.52
q4	53.46	46.54

Notes: q1, q2, q3, q4 are initial consumption quartiles.

Table A.12

Descriptive statistics – Use of improved maize seeds on at least half plots

Variable: Use of improved maize seeds on at least half of the household plots across all waves - % of households		
	Yes	No
q1	8.77	91.23
q2	10.65	89.35
q3	18.79	81.21
q4	22.08	77.92

Notes: q1, q2, q3, q4 are initial consumption quartiles.

Table A.13

Descriptive statistics – Use of improved paddy seeds on at least one plot

Variable: Use of improved maize seeds on at least one plot across waves - % of households		
	Yes	No
q1	19.35	80.65
q2	24.76	75.24
q3	27.03	72.97
q4	27.15	72.85

Notes: q1, q2, q3, q4 are initial consumption quartiles.

Table A.14

Descriptive statistics – Use of improved paddy seeds on at least half plots

Variable: Use of improved paddy seeds on at least half of the household plots across all waves - % of households		
	Yes	No
q1	4.27	95.73
q2	6.27	93.73
q3	6.61	93.39
q4	16.49	83.51

Notes: q1, q2, q3, q4 are initial consumption quartiles.

Table A.15

Descriptive statistics –Tropical Livestock Units per adult-equivalent

Variable: TLU level p.a.				
	Mean	Var	sd	Obs
Whole sample	0.436	1.328	1.152	3653
q1	0.257	0.337	0.580	926
q2	0.424	1.031	1.016	963
q3	0.410	1.152	1.073	937
q4	0.680	2.890	1.700	827

Essay 3

Household growth, weather shocks and resilience thresholds in rural Tanzania: a long-run assessment

Abstract

The study of the dynamics of household resilience, i.e. the ability to withstand shocks, is becoming more and more important in development economics. We contribute to this literature by exploring the relationship between household growth, weather shocks and resilience dynamics in rural Tanzania. By building a synthetic panel we are able to study a longitudinal dataset covering the time span 2000 – 2013 and make up for the lack of long micro panels which hamper the analysis of growth in developing countries. Our measure of resilience is an ad hoc FAO index which we show to be a factor of significant heterogeneity in the impacts of temperature shocks on growth. Our main contribution is the estimation of "resilience thresholds" above which households are immune to the negative effects of temperature shocks. The existence of resilience thresholds is a significant finding for policy-making and especially relevant for the crucial issue of adaptation to climate change in developing countries.

Introduction

Development resilience is closely related to the idea of stochastic poverty traps, since it concerns the likelihood that adverse outcomes do not persist for an extended period in a dynamic setting (Barrett & Constas, 2014). Furthermore, it looks at poverty dynamics by explicitly considering both the impact of possible stressors and shocks and household capacity to cope with them.

Consistently with the original nature of resilience as the ability to respond to ecological stress, our aim is to empirically test household resilience to food insecurity in the presence of observable weather shocks. To this end, inspired mainly by Dercon (2004) and the standard growth literature, we look at the relationship between household food consumption growth, the correlated likelihood of falling into poverty traps [1] due to a source of growth divergence represented by weather shocks and the protective role of resilience capacity.

While household welfare will always be higher in a shock-free environment, the effect on household consumption is theoretically ambiguous. On the one hand, according to the precautionary saving literature (Caballero, 1990; Carroll & Kimball, 2001; Carroll & Samwick, 1998; Deaton, 1992) lower consumption levels are, on average and *ceteris paribus*, compatible with a higher consumption growth (Paxson, 1992). On the other hand, the literature on poverty traps (Carter & Barrett, 2006; Carter, Little, Mogues, & Negatu, 2007; Carter & Lybbert, 2012; Zimmerman & Carter, 2003) highlights the opportunity of alternative coping strategies, such as voluntarily destabilizing consumption to asset-smooth and consequently avoid the risk of falling into poverty traps.

Looking at (food) consumption growth can provide useful insights on household welfare dynamics and 'resilience' traps, since we can both identify potential deviations from positive consumption growth paths of convergence in presence of shocks (weather shocks in our case) as well as recovery dynamics, growth divergence and multiple regimes of shocks due to resilience heterogeneity.

Unfortunately, robust empirical evidence on such issues has been hampered so far by the lack of long micro panels, measurement error and attrition (Antman & McKenzie, 2007). Short panels, in fact, are unable to disentangle between short-term transition movements and the long-term analysis of shocks and recovery paths, including identification of thresholds and multiple equilibria (Carter & Barrett, 2006). Because of the chronic scarcity of long panels of households in developing countries, but also due to the inherent inconsistency of the concept of "households" in a long-term perspective (household splits are characterized by short-term frequencies), the use of pseudo-panels can be

[1] This means we explicitly focus on one out of the four dimensions of food security agreed by the international community: access. Hence, this empirical analysis is not able to provide insights on the availability, stability and utilization dimensions.

considered as a workable solution. Synthetic panels can create longer time span compared to genuine panels, minimize attrition and smooth individuals' response errors (Deaton, 1985).
By matching the Living Standard Measurement Survey (LSMS) – Integrated Survey on Agriculture (ISA) Tanzania National Panel Survey by the World Bank, which provides a three-wave household longitudinal dataset covering the period 2008 – 2013[2], and the National Household Surveys 2000 and 2007 by the Tanzanian Department of Statistics, we build up integrated pseudo-panels for Tanzania covering the period 2000 - 2013.

We then propose the following identification strategy: i) in a standard empirical stochastic micro-growth model augmented with weather shocks and controlling for households and geographical heterogeneity, we test for the relevance of a set of household characteristics relevant to resilience to food insecurity, gathered into a measure of resilience developed by the Food and Agriculture Organization (FAO), the Resilience Capacity Index (from here on RCI), in determining a stable and positive food consumption growth; ii) we look for potential heterogeneity of impacts from temperature and precipitation shocks on household food consumption growth with respect to pre-shock RCI iii) we test for the presence of critical "resilience thresholds", in order to check for bifurcation of impacts from weather shocks due to different resilience capacity regimes.

The choice of Tanzania is not only driven by data availability.
Tanzania is a poor Sub-Saharan African country, which exhibits high average temperatures and a large climatic diversity. It is constantly classified as a country under high risk from the impacts of climate change: temperatures in the country are predicted to rise 2–4 °C by 2100 (Rowhani, Lobell, Linderman & Ramankutty, 2011). Moreover, it is still a predominantly rural country where agriculture accounts for about half of gross production, employs about 80 percent of the labour force and is primarily rainfed (Ahmed et al., 2011).
Finally, several empirical works have documented a causal relationship between welfare dynamics and weather shocks in the country and particularly in rural areas (Ahmed et al., 2011; Arndt et al., 2012; Bengtsson, 2010; Dercon, 2004; Hirvonen, 2016; Rowhani et al., 2011).

The essay is organized as follows: Section 1 reviews the literature on resilience to food security; Section 2 provides information on data, methodology for the estimation of the RCI, pseudo-panel creation and descriptive statistics; Section 3 illustrates the identification strategy and the empirical model; Section 4 reports the outcomes of the empirical analysis; Section 5 presents robustness checks;

[2] The new NPS 2014-2015 wave, recently released (in August 2017) by the Word Bank, was not yet available when this work was undertaken.

Section 6 discusses and concludes.

Section 1
Resilience: review of the literature

1.A. Resilience definition and measurement

Measuring resilience - i.e., people's ability to withstand shocks - has become an urgent task as climate change trends, ecosystem fragility and geo-political instability have led to increasingly unpredictable risks. The well-being of the world's poor is now subject to a more challenging series of shocks and stressors (Constas, Frankenberger, & Hoddinott, 2014). The response to assist households in dealing with these shocks depends, in turn, on accurate identification and measurement of their resilience determinants. Definitions and measurement efforts aiming at assessing resilience in development, and specifically with reference to food insecurity, appeared much earlier than a solid theoretical framework has been advanced.

Following the seminal paper from Pingali, Alinovi and Sutton (2005), Alinovi et al. (2010) and Alinovi, Mane and Romano (2010) were among the first authors who tried to define and measure household resilience to food insecurity. In their framework, the household is the entry level of analysis because it is the decision-making unit where the most important decisions are made on how to manage risks, both *ex ante* and *ex post,* including those affecting food security. To measure resilience, these studies construct a resilience index as a latent variable (unobservable) through a two-stage factor analysis based on observables.

Vaitla et al. (2012) define resilience as "the ability of an individual, a household, a community or an institution to withstand a shock or setback of some type and recover, or as a "bounce back, after a setback". The authors present a livelihood change approach to measuring resilience focusing on how assets held by a household or other social unit are used in various livelihood strategies to achieve certain outcomes.

Smith & Frankenberger (2017) define a household as resilient if "able to maintain its well-being even in the face of shocks and stressors" and identify the determinants of household resilience through three types of capacities - absorptive, adaptive and transformative. In such a framework, they investigate the relationship between food security, household resilience and the impact of the 2014 flooding in Northern Bangladesh. By including in their fixed-effect regressions a continuous interaction between resilience capacity and shock exposure, they detect evidence that resilience mediates the relationship between food security and shocks and mitigated the negative impact of the flooding.

FAO (2016) proposes RIMA, the Resilience Index Measurement and Analysis, where resilience is measured as a latent variable through structural equation models that include a host of relevant variables which act as proxies for the natural environment, social inclusion and well-being perception[3].

Alfani et al. (2015), after defining resilience as “the ability of countries, communities and households to manage change, by maintaining or transforming living standards in the face of shocks or stresses - such as earthquakes, droughts or violent conflicts - without compromising their long-term prospects”, propose an empirical alternative to all previous approaches that depend on longitudinal data to estimate household resilience, which makes use of readily available cross-sectional data, but is weakened by the well-known issues of omitted variable bias and confounding risk.

Cissé and Barrett (2015) instead, adopt a definition of resilience that was originally proposed by Barrett and Constas (2014): “Development resilience is the capacity over time of a person, household or other aggregate unit to avoid poverty in the face of various stressors and in the wake of myriad shocks.” Consequently, they develop an innovative moment-based approach to measure resilience by taking into account stochastic and possibly nonlinear well-being dynamics. An additional important contribution of their work is the derivation of a decomposable resilience measure based on the Foster-Greer-Thorbecke class of poverty measures, which makes possible the comparison of resilience capacities of various sub-populations of interest.

In 2013, the Resilience Measurement Technical Working Group (RMTWG), comprised of experts in resilience measurement, was established under the auspices of the Food Security Information Network[4], in order “to secure consensus on a common analytical framework and guidelines for food and nutrition security resilience measurement, and to promote adoption of agreed principles and best practices”.

Thanks to the work of the RMTWG, a sound consensus has been recently established around some key elements of resilience measurement. First, resilience is defined as the ability to respond to shocks as well as stressors, which can be analysed at different levels of aggregation (individuals, households, communities, organizations, systems or even entire states), and applied to both short-term and long-term consequences of shocks and stressors. Second, resilience has to be indexed against an outcome, e.g. food security. Third, since resilience is a genuinely dynamic concept, resilience measurement necessarily requires a dynamic analytical framework, in order to capture all possible pathways to achieve resilience, by accounting for agents’ heterogeneity in gaining their own livelihoods in a risky environment.

[3] See Section 2.2.

[4] The Food Security Information Network (FSIN) is a global initiative co-sponsored by FAO, WFP and IFPRI to strengthen food and nutrition security information systems for producing reliable and accurate data to guide analysis and decision-making. See more at http://www.fsincop.net/home/en/ .

In short, resilience measurement should be based on a theoretical framework focusing on the stochastic dynamics of individual and collective human well-being, especially on the avoidance of and escape from irreversible negative states over time (e.g. chronic poverty or food insecurity), and allowing for nonlinearities and multiple equilibria that can generate bifurcated welfare dynamics shocks, as recently proposed by Barrett and Constas (2014).

Consistently, it is now largely accepted that an indicator which aims to capture resilience capacity should be: a) aggregative of the multiple dimensions of resilience; b) able to vary within a clear and determined range; c) field tested; d) able to detect changes over time; e) subject to tests of validity and reliability[5].

1.B. Resilience Thresholds

The concept of resilience thresholds was first introduced in the early 1970s in seminal papers on ecology (Holling, 1973; May, 1977). Following Groffman et al. (2006), a threshold can be defined as the point at which there is an abrupt change in a quality, property or phenomenon or where small changes in a driver may produce large responses in a system.

Thresholds are fundamental in order to evaluate alternate states, i.e. the phenomenon whereby a system can change from one status to another (Walker & Meyers, 2004). An alternate state may also be considered as "stable", although this does not mean that it can be classified as good or bad. It is simply a steady state which, at a certain point, can be shifted as the result of a shock.

A regime shift involving alternate stable states occurs when a threshold level is passed, resulting in a change of direction (the trajectory) of the system itself. In some cases, crossing the threshold brings about a sudden, large, and dramatic change in some (or many) of the indicators (or variables) that identify the system.

While the concept of threshold is linked to the inner nature of resilience (Holling, 1973), there are no internationally accepted normative standards on how to define and identify thresholds.

Nevertheless, resilience is particularly important when a system approaches a threshold that is critical for regime shifts (Barrett & Constas, 2014; Holling, 1973).

The incorporation of the concept of thresholds and non-linear dynamics in development resilience analysis stems from the work of Barrett and Constas (2014) who, on the wake of the literature on poverty traps and multiple equilibria (Barrett & Carter, 2013; Carter & Barrett, 2006; Carter et al., 2007), consider the potential existence of thresholds when the expected path is not respected by the actual resilience capacity behaviour. In their example, two thresholds exist that separate three distinct

[5] As emphasized by Hoddinott (2014)"proposed measure should be subjected to tests of validity and reliability; in the case of measures of resilience capacity, we are also interested in understanding their predictive power".

regimes / stable states: humanitarian disaster; a chronic poverty zone where people will still be able to recover; a non-poor zone within which people are reasonably expected to recover from shocks.

Well-being is intrinsically - and, some would argue, increasingly - stochastic, affected by a range of exogenous events. In nearly every resilience measurement approach these events tend to be summarized in the disturbance term of an estimating regression. If one is interested in identifying thresholds above which a household will not suffer long-term development consequences of its food security, there are basically two alternative strategies: find a unique resilience threshold, or break down resilience in its main components and seek thresholds for each of them. Resilience literature is inherently multi-dimensional and, indeed, the biggest added value of a resilience perspective is the reciprocal influence that resilience components have each other. In this respect, the first strategy, i.e. the identification of a single threshold with respect to a unique resilience indicator, is more faithful to the origins and aims of resilience analysis.

Finally, a resilience threshold should be contextualized with respect to a specific shock, posing the concept of resilience thresholds as intrinsically relative.

Unfortunately, empirical efforts to analyse the potential existence and nature of such thresholds have been hampered so far by a lack of readily available empirical data (Walker & Meyers, 2004), by the inherent complexity of nonlinear dynamics and by multiple factor controls that operate at diverse spatial and temporal scales (Groffman et al., 2006).

As a consequence of these operative issues, most attempts have been limited to the investigation of exogenous thresholds unilaterally set up without any reference to real life or validate norms (Barrett & Constas, 2014; Upton, Cissé, & Barrett, 2016) whose internal (other than external) validity remains questionable.

Section 2
Data, RCI and Pseudo-panels

This section illustrates the three operative steps which preceded the empirical analysis: data matching, estimation of the RCI and construction of the synthetic panels.

2.1 Data

Our household dataset consists of five surveys: two repeated cross-sections, namely the 2000-2001 and 2007 Tanzania Household Budget Surveys (HBS)[6] by the Tanzanian National Bureau of Statistics (NBS); and the 2008-2009, 2010-2011 and 2012-2013 waves of the National Panel Survey

[6] These surveys can be found at: http://www.nbs.go.tz/ .

(NPS) by the World Bank, part of the World Bank collection named Living Standards Measurement Study (LSMS) – Integrated Survey on Agriculture (ISA)[7]. Both collections include extensive data on household expenditure, which made possible the calculation of annual measures of household per capita food consumption, and on household assets and characteristics, but only the LSMS-ISA waves include specific questionnaires on agriculture.

These household data have been cleaned, aggregated and integrated with weather data.

Weather data include time series on precipitation and land surface temperature.

Monthly time series (covering the period 1983 - 2016) of precipitation data (in millilitres) come from the Climate Hazards Group InfraRed Precipitation with Station data (CHIRPS) (Funk et al., 2015). CHIRPS data have a very high spatial resolution of 0.05° latitude by 0.05° longitude.

Monthly time series (for the period 1983 - 2015) temperature data are taken from the Centre for Climatic Research (CRU) at the University of East Anglia with a 0.5° latitude by 0.5° longitude spatial resolution.

While the LSMS-ISA households are geo-referenced (although using a random offset of 5 KMs), HBS households are not, thus allowing only matching with weather data at the district level.

Consistently, we aggregated both weather time series at the district-monthly level using geo-spatial software before merging them with household data[8].

We explicitly restrict our analysis to rural households, since we are interested in the impacts from weather shocks which primarily affect farming households.

Our aggregated dataset consists of 16,190 original observations.

Starting from this pooled dataset, the RCI was estimated for each household before the creation of pseudo-panels.

2.2 The Resilience Capacity Index (RCI)

Resilience capacity is a (i) multidimensional and (ii) unobservable construct.

Therefore, a common approach in the literature is to look at resilience by (i) unpacking and estimating each pillar (or principle, as in Pingali, Alinovi and Sutton (2005), or capacity, as in Béné, Frankenberger, and Nelson (2015) that contributes to resilience and (ii) by employing latent variable models to estimate a proxy indicator of the unobservable resilience capacity.

---

[7]The complete LSMS-ISA dataset collection is available at the following link: http://econ.worldbank.org/WBSITE/EXTERNAL/EXTDEC/EXTRESEARCH/EXTLSMS/0,,contentMDK:23617057~pagePK:64168445~piPK:64168309~theSitePK:3358997,00.html .

[8] Other than temperature and precipitation data, other biophysical variables, namely slope, elevation and length of the growing period, available from the FAO Database, were included in the dataset. The FAO Database is available at: http://www.fao.org/geonetwork/srv/en/metadata.show?id=48025&currTab=simple.

Among the available statistical techniques for latent variables, namely Principal Component Analysis (PCA), Factor Analysis (FA), Structural Equation Models (SEMs) and Multiple Indicator and Multiple Causes (MIMIC), FA allows for the reduction of a set of observed variables (used as proxy indicators for the latent variable), in a single variable, the component of interest. The data reduction mechanism relies on finding cross-correlations between the observed variables, identifying the number of (unobservable) factors reflected in the correlations, and predicting the latent outcome as a linear combination of underlying factors.
On the contrary, PCA reduces the dimensionality of a data set by deriving a new set of variables – called principal components – that are uncorrelated, retain most of the sample information, and are ordered by the fraction of the total information that each component explains. However, PCA cannot be employed to create a latent variable that is linearly correlated with the underlying dimensions, as it does not consider linear relations during the estimation process. In addition, the components are calculated using the variance of the observable variables, and the total variance appears in the solution (Costello & Osborne, 2005). On top of this, the literature has widely recognized the inadequacy of such a method when using categorical variables (Kolenikov & Angeles, 2004).
Conversely, SEMs assume that the set of measured variables is an imperfect measure of the latent variable of interest. SEMs use a factor analysis-type model to measure the latent variable *via* observed variables, while simultaneously using a regression-type model to identify relationships among the underlying variables (Bollen, 1989). Differently from SEMs, FA assumes that the residual errors (i.e. unique factors) are uncorrelated with each other and with the common (i.e. latent) variable.
Among the SEMs, the MIMIC model allows including a measurement part of the estimation, as described below.

Taking into account its multidimensional and unobservable nature, this work estimates an index of the resilience capacity at household level, the so-called Resilience Capacity Index (RCI), through a two-step procedure. In the first step, the pillars of resilience are estimated through Factor Analysis (FA) from observed variables (except in case of data constraints, see below). The procedure allows for the reduction of the set of variables, used as proxy indicators for the latent variable, in a single variable, the pillar of interest. The pillars of resilience considered in this study are: Access to Basic Services (ABS), Assets (AST), Social Safety Nets (SSN) and Adaptive Capacity (AC). This choice follows the FAO-Resilience Index and Measurement Analysis (RIMA-II) conceptual framework, that extensively explains the definition of the four pillars of resilience (FAO, 2016).
All the observed variables used to estimate the pillars are listed in Table A.1 in the Appendix along with summary statistics. The choice of the variables adopted for estimating each pillar is based on

literature review (FAO, 2016), data availability, context analysis and statistical properties of the variables. The factors considered for the estimation of each resilience pillar are only those able to explain at least 95 percent of the variables' variance.[9] The first-step results of the three FA employed for estimating the ABS, AST and AC pillars are reported in Tables A.2-A.4 in the Appendix. Conversely, due to data constraints, the SSN pillar was created by simply adding public and private transfers received by households for all waves, except for the HBS 2007 survey, for which data on transfers were not available and were replaced with a dummy capturing household participation in a saving / credit group.

In the second step, a MIMIC model is estimated (Bollen & Davis, 2009). Specifically, a system of equations is constructed, specifying the relationships between an unobservable latent variable (RCI), a set of outcome indicators (food security indicators), and a set of attributes (four resilience pillars). The food security indicators employed in the MIMIC model are food expenditure and the Simpson index of dietary diversity. Table A.1 in the Appendix reports the summary statistics of the food security indicators.

The MIMIC model is made up of two components, namely the measurement equation (1) – reflecting that the observed indicators of food security are imperfect indicators of resilience capacity – and the structural equation (2), which correlates the estimated attributes to resilience capacity:

$$\begin{bmatrix} Food\ expenditure\ pc \\ Simpson\ index \end{bmatrix} = [\Lambda_1, \Lambda_2] \times [RCI] + [\varepsilon_1, \varepsilon_2] \quad (1)$$

$$[RCI] = [\beta_1, \beta_2, \beta_3, \beta_4] \times \begin{bmatrix} ABS \\ AST \\ SSN \\ AC \end{bmatrix} + [\varepsilon_3]\ . \quad (2)$$

Therefore, the RCI is jointly estimated by its causes, the pillars, and its indicators, the food security variables. This procedure ensures a proper link between the estimated resilience capacity and household food security.

The MIMIC results present a good fit of the data, as shown in Table A.5 in the Appendix. All the pillar coefficients, except for Social Safety Nets, are positive and statistically significant.

The RCI is not anchored to any scale of measurement. Therefore, a scale has been defined setting the coefficient of the food expenditure loading (Λ1) as equal to 1, meaning that one standard deviation increase in RCI implies an increase of one standard deviation in food expenditure. This also defines

[9] When more than one factor explains the 95 percent of the variable variance, the resilience pillar is constructed as a weighted sum of the factors estimated by FA. The weights employed in the sum are represented by the percentages of variance explained by each factor.

the unit of measurement for the other outcome indicator (Λ2) and for the variance of the two food security indicators.

The use of FA in the first step and MIMIC in the second one is supported by the technical literature and a validation test. In fact, to obtain robust results, the RCI has been estimated by employing alternative techniques for both the pillar estimation - FA, both with the iterated principal-factor method (ipf) and the principal-component factor (pcf) - and the RCI estimation, namely employing the general form of the SEM instead of the MIMIC model in the second step.
Using the general SEM means estimating the RCI by employing only Equation (2), without considering food security indicators. Therefore, a total of six models have been estimated by combining MIMIC or SEM, in the second step, with FA (ipf), FA (pcf) and PCA in the first one.
A validation index allows to directly compare the goodness of fit achieved by the different models. The index is based on the model Chi-squared but takes also into account the degree of freedom as well as the number of observations. A smaller value of the index ensures a better fit of the model.
By comparing the validation indexes[10], it emerges that among the three MIMC models, the use of FA in the first-step ensures smaller value of the index, without major differences between FA (ipf) and FA (pcf). The smallest validation index is obtained by combining PCA in the first-step and SEM in the second-step. Finally, only minor differences are detected in the validation indexes of the MIMIC versus SEM models by FA type.
In conclusion, as stated above, the use of MIMIC model *versus* a general SEM guarantees the consistency with the estimation of the RCI linked to household food security. Additionally, according to the literature and the validation test, the FA is the preferred technique for pillar estimation when combined in a second step with the MIMIC model.

2.3 Pseudo-panels

Since the seminal work of Deaton (1985), the pseudo-panel approach has been used in several empirical applications to estimate changes over time. It is essentially a hybrid between repeated cross-sections and genuine panel data, but presents some additional advantages with respect to the latter: i) attrition and non-response problems are minimized, and individuals' response errors smoothed (Verbeek, 2008); ii) larger sample dimension (N) and time span (T) can be created by combining different sources; iii) longer-term dynamics can be studied than is usually possible with the existing panels since repeated cross-sectional surveys are more common than genuine panels (especially in developing countries) (Antman & McKenzie, 2007).

---

[10] Results are available upon request.

The main drawback is of course that the same individuals cannot be followed over time, so that individual welfare dynamics over time cannot be observed and are not available for constructing instruments or transforming a model to first-differences or in deviations from individual means (Verbeek, 2008).
The fundamental assumption of the pseudo-panel framework is that the mean cohort behaviour reproduces the form of an individual behaviour in that specific cohort.
A 'cohort' is defined as a group with fixed membership, individuals of which can be identified as they show up in the surveys (Deaton, 1985).
Operationally, it means to group individuals sharing some common characteristics into cohorts, after which the averages within these cohorts are treated as observations in a pseudo panel (Verbeek, 2008).

For the construction of our cohorts, we paid careful attention to the following recommendations underlined by the relevant literature: i) cohorts can be more or less broadly defined, but on the basis of *time-invariant* variables that are observed for all individuals and should allow sufficient temporal and cross-sectional variation in the true cohort mean for parameter identification (Verbeek & Nijman, 1992). Possible choices include variables like date of birth, race, gender, region etc. (Verbeek, 2008) which may also be interacted to increase the number of cohorts C. McKenzie (2004) stresses the importance of allowing for inter-cohort parameter heterogeneity; ii) each individual must be a member of exactly one cohort, which stays the same for all T; iii) in constructing cohort samples, there is a trade-off between n_c and C, that is between the accuracy of each cohort mean (n_c) and the number of observations (C) of the pseudo-panel. A large n_c minimizes measurement error but increases the efficiency loss. By skimming the literature on pseudo-panel applications, it may be argued that the size of n_c is more important than the size of C. Hence, the challenge is to find a balance between these two dimensions, where the optimal choice would be the one that minimizes the heterogeneity within each cohort (internally homogeneous) but maximizes the heterogeneity among them. Verbeek and Nijman (1992) argue that with a sufficiently large cohort size ($n_c \geq 100$) measurement error issues and the cohort nature of the data can safely be ignored, but this will come at the cost of reducing the number of cohorts C; iiii) according to Verbeek & Vella (2005), the fixed effects estimator based on the pseudo-panel of cohort averages may provide an attractive choice, even when a lagged dependent variable is included in the model.

As already explained, our dataset consists of five surveys: two repeated cross-sections and a three-wave panel. Aggregating these five waves in a pseudo-panel allows us to expand the timespan of our empirical analysis up to fourteen years, from 2000 to 2013, more than a decade of empirical observation.

Our baseline analysis is conducted on two different pseudo-panels (Versions 1 and 2) whose core characteristics are shown in Table A.6.
In Version 1 we use, as group variables - i.e. variables interacted with each other to create cohorts - quintiles of long-run average temperature and year-of-birth of the household head quintiles.
We employ long-run temperature quintiles as a proxy for climatic areas, while the use of year-of-birth bands follows standard practice in pseudo-panel literature. Furthermore, they are meant as an exogenous factor to our resilience index and its four pillars.
Similarly, the only difference in Version 2 is the use of long-run precipitation quintiles instead of long-run temperature quintiles as cohort variables. This because using a different proxy for climatic areas leads to different groupings of households into cohorts and consequently could drastically change the results of the empirical analysis. This alternative version is used to perform robustness checks (see Section 5).
As reported in Table A.6, in both pseudo-panels, Versions 1 and 2, the number of cohorts C is 25 and the average number of observation per cohort, $n_{c,}$ is 648[11].

Tables A.7 and A.8 shows descriptive statistics for, respectively, Version 1 and Version 2: food consumption growth is very high on average (more than 6 % in both versions), but this is likely due to the fact these synthetic panels are based on data coming from two separate sources (HBS and NPS). Overall, the standard deviation is quite high for all the reported variables, pointing to a good degree of between-cohort variation and heterogeneity in terms of growth paths, resilience capacity and climatic conditions.
Although most variables exhibit similar distributions across the two pseudo-panels, please note that the mean of temperature shocks for Version 2 is almost three times higher than in Version 1, i.e. temperature shocks are, on average, three times stronger in the former than in the latter dataset. Precipitation shocks, instead, are almost identical.
As will be shown below, the difference in the magnitude of temperature shocks between the two pseudo-panels will be important in explaining part of the empirical results.

Section 3
Identification strategy and empirical model

Our benchmark identification strategy is derived from the empirical work of Dercon (2004), Carter et al. (2007), Jalan and Ravallion (2002, 2004).

[11] Note also that our waves are different in terms of observations. HBS 2000 includes data on over 7000 households, HBS 2007 on about 3000 households and the three NPS waves only approximately 2000 households. This implies that the average number of observations per cohort-wave differ across waves.

In particular, we build our model on the wake of the macro-literature (Solow, 1956) , and augment a standard empirical growth model, where cohort food consumption growth rates are conditional on lagged consumption levels, with a set of additional explanatory variables:

$$\Delta Y_{it} = \alpha + \beta_1 ln\, Y_{it-1} + \beta_2 RCI_{it-1} + \beta_3 \Delta RCI_{it} + \beta_4 \Delta Temp_{dt} + \beta_5 \Delta Pre_{dt} + \beta_6 X_{it} + \mu_i + \theta_t + \varepsilon_{it} \quad [3]$$

In Equation (3), the dependent variable is the annualised growth rate in real cohort monthly per capita food consumption[12], and $ln\, Y_{it-1}$ is lagged cohort monthly per capita consumption[13]. β_1 represents the convergence term: a negative sign of the coefficient β_1 would indicate the presence of an ongoing convergence process, i.e. that poorest household tend to catch up over time with richer households.

The use of consumption instead of income as a proxy for welfare is motivated by the setting: in rural Tanzania most households depend on self-employed agriculture and consequently measuring income is subject to large errors (Deaton, 1997; Hirvonen, 2016); in contexts such as this consumption has been viewed as providing a more reliable measure of welfare (Deaton & Grosh, 2000).

We then amend the standard empirical growth model by including additional variables which allow us to explicitly look at the role of resilience capacity, weather shocks and their dynamic interactions. First, the role played by resilience capacity is identified by two distinct variables.

RCI_{it-1} captures the impact of household resilience at time t-1, i.e. pre-'treatment', on food consumption growth. Hence, a positive sign of the coefficient β_2 would indicate that, on average and *ceteris paribus*, being more resilient to food insecurity can boost household growth.

We then include ΔRCI_{it}, which stands for the annualised between-wave change in the RCI[14]. This variable is added for two reasons: i) being the RCI a catch-all index for a full set of relevant household characteristics and assets which are supposed to mimic the latent Tanzanian farmers' resilience to food insecurity, its annualised change captures a wide range of time-varying variables, thus offsetting many potential sources of omitted variable bias; ii) the sign of its coefficient β_3 can be informative about the importance of resilience *dynamics*, i.e. the evolution over time of household ability to withstand shocks, in determining the consumption growth trend.

Second, the impacts of weather shocks.

We test for possible deviations from the assumed convergence path of consumption growth by

[12] Food consumption is measured as food expenditure.

[13] Given our fixed-effect setting, we could not directly include initial consumption levels since they are time-invariant. Hence the choice of including lagged consumption levels to assess potential convergence, which follows standard practice in this micro literature.

[14] Explicitly, it means: ΔRCI_{it}, = $RCI_{it} - RCI_{it-1}$.

including temperature $[\beta_4 \Delta Temp_{dt}]$ and precipitation $[\beta_5 \Delta Pre_{dt}]$ shocks, observed at the district level (hence the subscript d). As for the functional form of weather shocks, we followed Dercon (2004), i.e. our weather variables are calculated as the difference in logarithms between their values at t and t-1, both scaled by long-run means.[15] This because we assume that level changes matter not only in an absolute sense but also, more importantly, in proportion to an area's usual, long-run, variation (Dell, Jones, & Olken, 2014). The parameters associated to weather shocks test whether they can slow or derail the assumed convergence process in the average household growth path. We include both temperature and precipitation in our regressions, because climate literature (Auffhammer, Hsiang, Schlenker and Sobel, 2013; Dell, Jones, & Olken, 2014) has warned against the risk of omitted variable bias and incorrect inference when dealing with the effects of weather regressors that are always highly correlated with each other, like temperature and precipitation usually are.

Since this literature also recommends to always include all the available climatic and biophysical variables, we also add to our model other biophysical variables, X_{it}, which act as additional controls, namely elevation, slope and length of the growing period (LGP) for each household.

Third, the dynamic relationship between resilience capacity and climatic shocks: we are interested in assessing whether there is heterogeneity in the impacts of weather shocks with respect to pre-shock resilience capacity. To test for this prediction, we augment the baseline specification in Equation (3) by including several interactions between lagged or initial RCI level and weather variables[16].

Finally, μ_i are cohort fixed effects, θ_t are wave fixed effects and ε_{it} are error terms clustered at the cohort level.

This is the baseline specification to look for the aggregate impacts of the variables of interest on food consumption growth.

As a second step of our analysis, we employ a fixed-effect panel threshold model in order to potentially detect a critical RCI threshold or tipping point that entails a change of regime of the impacts from temperature shocks, on the wake of the literature on poverty traps.

In this way, we can test whether there is heterogeneity of impacts from temperature shocks with respect to the RCI, and if resilience to food insecurity plays a role in counterbalancing or eliminating altogether the assumed diverging effects entailed by weather shocks.

In doing so, we follow the approach by Carter et al. (2007) and opt for the use of the Hansen (2000) threshold estimator, as implemented in a fixed-effect setting by Wang (2015).

---

[15] Since we are working with synthetic panels, weather data were aggregated for each household in the original sample starting from their interview month and going backwards for the average total number of months between each survey.

[16] Additionally, we also include interactions between weather shocks and a dummy to control for heterogeneity of impacts for households living in hotter-than-average areas. See Section 4.

This model replicates Equation (3) but distinguishes two impact regimes conditional to a critical value of the RCI:

$$\Delta Y_{it} = \begin{cases} \alpha + \beta_1 ln\, Y_{it-1} + \beta_2 RCI_{it-1} + \beta_3 \Delta RCI_{it} + \beta_4^l \Delta Temp_{dt} + \beta_5 \Delta Pre_{dt} + \beta_6 X_{it} + \mu_i + \theta_t + \varepsilon_{it} \;\; if\; RCI_{it-1} \leq \omega \\ \alpha + \beta_1 ln\, Y_{it-1} + \beta_2 RCI_{it-1} + \beta_3 \Delta RCI_{it} + \beta_4^u \Delta Temp_{dt} + \beta_5 \Delta Pre_{dt} + \beta_6 X_{it} + \mu_i + \theta_t + \varepsilon_{it} \;\; if\; RCI_{it-1} > \omega \end{cases} \quad (4)$$

In Equation (4) the superscripts l and u on the coefficient β_4 indicate, respectively, the lower and upper regime of temperature impacts, conditional on the RCI value at time t-1, i.e. the pre-shock period, in order to capture heterogeneity with respect to *ex ante*, and not *ex post*, household resilience capacity to food insecurity.

Section 4
Empirical outcomes and the estimation of "resilience thresholds"

Table 1 shows the results of the baseline identification strategy set in Equation (3).
First, some general remarks.
Lagged consumption level has a negative and significant at the 1 percent level effect on growth. This points to the plausibility of the assumption of convergence among households, i.e. poorer households tend to have higher growth rates than richer households, leading to a catch-up growth process.
Lagged RCI has a positive and strongly significant effect on growth: net of the convergence process, being more resilient to food insecurity enhances growth.
Analogously, as one would expect given the high correlation with food consumption growth, the coefficient of ΔRCI is positive, large and strongly statistically significant, providing evidence that an increase in resilience capacity over time helps households to withstand shocks and stressors and boosts growth.
Finally, the adjusted R-squared is extremely high, but this too was expected given the large set of covariates and fixed effects we chose to include. While this may raise concerns about overidentification, we reckon this is actually an advantage and not a drawback: there is almost no residual variation, so we are confident we can identify the impacts from weather shocks with a good degree of precision.
Please also note that these general findings will hold in all the following regressions.

Let us now turn to the impacts of weather shocks on food consumption growth.
In Column (1) we only look at aggregate impacts. Both temperature and rainfall shocks have, on average and *ceteris paribus*, a negative but insignificant effect on the growth rate of food consumption.

However, these aggregate results could be hiding heterogeneity of impacts conditional on household resilience capacity, i.e. on their ability to withstand weather shocks. Following the path of Smith and Frankenberger (2017), our hypothesis is that given the occurrence of the same weather shock, its impact will be stronger for the least resilient household.

In Column (2) we start investigating whether this is the case, by including two separate interactions of both temperature and precipitation shocks with a 'Low average pre-shock' RCI dummy, which takes value 1 for cohorts having an average pre-shock RCI below the 25th percentile. This dummy identifies cohorts who were on average the least resilient ones *before* shocks occurred. Additionally, as a further control and also to check for potential heterogeneity with respect to different climatic areas (it may be impacts are stronger in hotter-than-average areas, due to intensification effects, or smaller, due to adaptation over time, as emphasized by Dell, Jones and Olken (2012)) we also include an interaction between weather shocks and a 'hot' dummy, which takes value 1 for cohorts living in areas with an above median long-run average temperature. The results show a sharp heterogeneity of impacts with respect to resilience capacity: temperature shocks now have a negative, larger and statistically significant effect for the least resilient cohorts, but negative, smaller and insignificant effect for the rest of cohorts. In particular, a one within-standard deviation increase in temperature shocks (0.018) entails a 0.63 % decrease in food consumption growth for cohorts with a low average pre-shock resilience capacity index. The pattern is analogous in hot areas: while there seems to be an intensifying effect due to living in hotter-than-average areas (even though they are always insignificant, impacts on average more than double in hot areas), only the least resilient cohorts are significantly affected by temperature shocks. Therefore, rather than heterogeneity with respect to climatic areas, we detect heterogeneity with respect to pre-shock resilience capacity. As for precipitation, there seems to be a negative and statistically significant impact for the most resilient households, and a positive and weakly impact of precipitation shocks in hot areas; however, precipitation shocks will prove to be highly sensitive to specification in the following regressions[17] so this result should be interpreted with extreme caution.

In Columns (3) and (4) we explore different definitions of 'least resilient cohorts and check whether heterogeneity of impacts from temperature shocks persists.

In Column (3) the previous RCI interactions are now replaced by interactions between weather shocks and 'Low pre-shock RCI', a dummy taking value 1 for cohorts with a lagged RCI below the 25th percentile in a given wave. This dummy captures heterogeneity not with respect to *average* pre-shock RCI levels, but with respect to the RCI level before each shock.

[17] Due to space reasons, in Table 1 as well as in Table 4 in Section 6 the total impacts of precipitation shocks for the least resilient households and for households living in hot areas are not reported. This because they are always either statistically insignificant or only weakly statistically significant. Consequently, our focus is on temperature shocks.

The findings detected in Column (2) still hold and are actually reinforced: the impacts of temperature shocks are larger and statistically significant at the 1 % level for the least resilient cohorts. A one within-standard deviation in temperature shocks entails a decrease of approximately 1 percentage point in food consumption growth for cohorts with a below the 25th percentile pre-shock RCI level, while the effect is positive and insignificant for the rest of cohorts. The pattern is again the same for cohorts in hot areas, but now there is no intensifying effect from living in hot areas but rather weak evidence of adaptation, represented by the positive and statistically significant at the 10 % level impact of the interaction between the 'hot' dummy and temperature shocks.

It may be that, more than pre-shock conditions in resilience capacity, what matters is resilience capacity at the beginning of the study period. To explore this possibility, in Column (4) we change again our definition of 'least resilient cohorts' and interact weather shocks with a 'Low initial RCI' dummy, which takes value 1 for cohorts with an *initial* RCI below the 25th percentile.

The previous results are confirmed: cohorts with a low initial RCI are the only ones to be negatively and significantly affected by temperature shocks: a one within-standard deviation now reduces the growth rate of food consumption by 0.74 %. The intensifying effect of living in hotter-than-average areas is now again evident and also statistically significant.

Finally, in Column (5) we replace the interactions with dummies with interactions between temperature and precipitation shocks and the RCI at time t-1 (pre-shock). These *continuous* interactions can shed light on whether the negative impacts of shocks disappear as households become more resilient over time, i.e. they can provide empirical evidence on the *dynamic* relationship between resilience and weather shocks.

While this regression confirms rainfall shocks are insignificant and so their relationship with resilience, for temperature shocks the causal relationship is supported: temperature shocks have a negative and statistically significant impact on cohort growth which gradually gets smaller and eventually disappears as cohorts get more resilient, thanks to the dampening effect resilience has on temperature shocks, indicated by the positive and statistically significant continuous interaction.

Such sharp and persisting heterogeneity of temperature impacts with respect to resilience capacity is worth being investigated in further detail.

It raises the hypothesis of "resilience thresholds" which entail a bifurcation of impacts from weather shocks on growth.

The RCI dummies we included in the interactions only represent subjective classifications between who is more and who is less resilient to food insecurity. Indeed, any dummy would set an arbitrary threshold, imposed by us and not by the true distribution and nature of the data.

To overcome this drawback and let the data speak for us, we use Equation (4) which implements the

fixed-effect panel threshold model as adapted by Wang (2015) from Hansen (2000). Thanks to the use of this estimator not only potential thresholds are data-driven and not arbitrarily chosen, but we can also test for the statistical significance of such thresholds by using the bootstrap procedure.
Table 2 shows the empirical results.
There is a bifurcation of impacts from temperature shocks: in the lower regime, i.e. below the threshold, the effect is large and strongly statistically significant. Above the threshold, in the upper regime, the impact is small and negligible. For cohorts below the threshold a one within-standard deviation increase in temperature shocks reduces growth by about 1.1 percentage points, while for cohorts above the threshold only by approximately 0.1 %.
Table 3 shows the threshold value, confidence intervals and tests. Bifurcation of impacts occurs at a pre-shock RCI value of 54.609, far below the mean.
This threshold is statistically significant at the 1 percent level, as indicated by the bootstrap test.
One may ask whether there is more than one resilience threshold. We investigated if multiple thresholds could be detected, but their significance was always rejected and the single threshold is by far the one which suits better the nature of the data.

Section 5
Robustness checks

A pseudo-panel is made of cohort data and not composed by 'true' households whose growth dynamics are followed over time. Choices involving the creation of pseudo-panels are arbitrary, although we tried to stick to the usual prescriptions in the relevant literature. Therefore, one could be concerned that our results may be driven by the specific nature of this pseudo-panel and the cohort variables used. This is a valid point which we acknowledge.
Consequently, we think it is important to show that our core findings are not altered when making different choices about the creation of the synthetic panel. This is why our robustness check is the repetition of the same analysis conducted above using a different pseudo-panel. As introduced in Section 2, we call this alternative longitudinal dataset Pseudo-Panel Version 2.
The unique but relevant difference compared to our baseline version is that we substitute, as a cohort-variable, long-run average precipitation quintiles to long-run average temperature quintiles.
The core characteristics of Version 2 in terms of C and n_c are exactly the same of Version 1, as shown in Table A.6.
Tables 4-6 the outcomes of the empirical analysis conducted on Version 2.
Table 4 replicates the same regressions of Table 1.
The pattern of sharp and significant heterogeneity of temperature impacts on growth conditional on pre-shock resilience capacity is confirmed in all cases, although physiologically different in

magnitude. Temperature shocks slow the convergence process. There seems to be a stronger evidence that living in hot areas amplifies the negative impacts of temperature shocks on growth, and there is no evidence of adaptation, i.e. smaller impacts in hot areas, whatsoever.
Conversely, precipitation shocks, again, are almost always insignificant and very sensitive to specification.
Table 5 illustrates the results from the threshold model.
Again, cohorts below the resilience thresholds are strongly and significantly affected by temperature shocks, whereas the effect for cohorts above the threshold is slightly positive and insignificant. A one within-standard deviation (0.023) in temperature shocks entails a decrease in the growth rate of food consumption of 0.37 % for households in the lower regime and an increase of 0.05 % for households in the upper regime.
Table 6 shows threshold value, confidence intervals and effect tests. The value of the threshold is higher compared to Table 3: bifurcation in this panel occurs at a lagged RCI level of 60.796, and the threshold is significant at the 5 % level.
Again, the hypothesis of multiple thresholds was rejected.

While these results are qualitatively equivalent to those emerged using the baseline Version 1 pseudo-panel dataset, one may ask why they do differ quantitatively, in particular with regard to the threshold value.
Our answer is that the value of the RCI threshold is conditional on the magnitude of temperature shocks: on average and *ceteris paribus*, if temperature shocks are stronger households must be more resilient to be immune to the impacts of such shocks.
In this Version 2, as stressed above (cf. Table A.8), temperature shocks are on average three times higher than in the baseline version, and we reckon this explains why the threshold value is higher in Pseudo-Panel Version 2 compared to Pseudo-Panel Version 1.

Section 6
Discussion and conclusion

We are aware that there are several limitations to this empirical exercise which go beyond the precise quantification of thresholds.
We do not claim our resilience thresholds are absolute thresholds for households living in the rural Tanzanian setting, let alone for other countries.
While absolute thresholds would be the perfect tool for policymakers, since they would indicate what is the resilience level that households should be helped to reach in order to be immune from shocks,

this is a utopia: resilience thresholds are intrinsically *relative*, so some *caveats* are necessary.
First, our resilience thresholds only apply to weather shocks and specifically to temperature shocks. Second, these thresholds are limited to the rural Tanzanian context, while in other countries or settings the picture could be completely different. Third, such thresholds are data-driven, and they may change significantly if using other datasets.
Fourth, even though we take inspiration from the literature on poverty traps, we can only talk of resilience thresholds, not of resilience 'traps': households below the RCI threshold are not permanently trapped in a low equilibrium from which there is no escape; rather, their growth paths are negatively affected by temperature shocks and this slows down, but does not reverse, the growth convergence process. However, climatic shocks of exceptional magnitude (at least compared to the temperature shocks experienced in our sample) could hit households so hard to entail not just a slowdown, but even a reversal of the convergence process.
Fifth, our RCI is only one of the possible ways of measuring resilience which is, as stressed in Section 2, an inherently latent and unobservable variable. Therefore, the use of other empirical approaches to resilience may well lead to different conclusions.

Still, despite the above remarks, we reckon the qualitative contribution is meaningful.
Smith and Frankenberger (2017) found suggestive evidence that resilience capacity reduced the negative impact of floods on household food security in Northern Bangladesh.
Here, instead, we show that resilience capacity can neutralize the diverging effect of temperature shocks on food consumption growth in rural Tanzania, and detect the existence of a critical resilience threshold.
In this particular case, such threshold is embodied by a value of the RCI, our measure of resilience, comprised between 54.6 and 60.8.
Going back to the original aggregate dataset, this means we find that between around 25 % and 47 % of households in our sample are below the resilience threshold and consequently vulnerable to temperature shocks.

Extrapolating with respect to global warming, with the usual caution about external validity (Dell, Jones, & Olken, 2014), this conclusion is especially relevant in view of the impacts of climate change in developing countries in the 21st century, which we know will be disproportionately bigger in poorer and hotter countries such as Tanzania (Tol, 2015).
Being above the critical resilience threshold neutralizes the diverging effect by temperature shocks: the logical implication is that adaptation strategies and policy interventions in developing countries will need to focus on boosting household resilience, so to empower all households to eventually reach

the threshold and cope with the consequences of climate change. But resilience remains a multi-dimensional concept: in this sense, while in this study we only estimated thresholds with respect to an aggregated measure of resilience, the identification of the main resilience determinants mediating the temperature-economic growth relationship, as well as the disaggregated exploration of different temperature regimes with respect to single drivers or dimensions of resilience, are necessary and complementary research avenues for the next future.

More generally, despite the remarks made above about the very limited, if any, external validity of our resilience thresholds, many fundamental questions stem from this work: there may potentially be resilience thresholds for households not just with respect to weather shocks, but for any shock; not just in rural Tanzania, but in many developing countries.

The pioneeristic work by Barrett and Constas (2014) laid the theoretical foundations for the hypothesis of non-linear paths, regime shifts and critical thresholds for development resilience, but empirical evidence on this crucial topic is still scant.

Given the self-evident importance of such issues from a policy-making perspective, the dynamic relationship between development resilience and shocks should be considered a top priority by future research.

References

Alinovi L., D'Errico, M., Mane, E., & Romano, D. (2010). Livelihoods strategies and household resilience to food insecurity: an empirical analysis to Kenya. *European Report on Development*. Paper prepared for the conference on "Promoting resilience through social protection in Sub-Saharan Africa", organised by the European Report of Development in Dakar, Senegal 28-30 June 2010. Available at http://erd.eui.eu/media/BackgroundPapers/Alinovi-Romano-D%27Errico-Mane.pdf

Alinovi, L., Mane, E., & Romano, D. (2010). Measuring household resilience to food insecurity: application to Palestinian households. *Agricultural Survey Methods (eds. R. Benedetti, M. Bee, G. Espa and F. Piersimoni)*, John Wiley & Sons, Ltd, Chichester, UK. doi: 10.1002/97.

Antman, F., & McKenzie, D. (2007). Poverty traps and nonlinear income dynamics with measurement error and individual heterogeneity. *The journal of development studies*, *43*(6), 1057–1083.

Auffhammer, M., Hsiang, S. M., Schlenker, W., & Sobel, A. (2013). Using weather data and climate model output in economic analyses of climate change. *Review of Environmental Economics and Policy*, ret016.

Barrett, C. B., & Constas, M. A. (2014). Toward a theory of resilience for international development applications. *Proceedings of the National Academy of Sciences*, *111*(40), 14625–14630.

Bollen, K.A. (1989) Structural equations with latent variables. New Jersey: John Wiley & Sons, Inc.

Bollen, K.A., Bauer, D.J., Christ, S.L., and Edwards, M.C. 2010. Overview of structural equation models and recent extensions. In Kolenikov, S., Steinley, D., and Thombs, L. (eds.), Statistics in the social sciences: Current methodological developments. Hoboken, NJ, Wiley. 37-79.

Caballero, R. J. (1990). Consumption puzzles and precautionary savings. *Journal of monetary economics*, *25*(1), 113–136.

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, *29*(2), 238–249.

Carroll, C. D., & Kimball, M. S. (2001). *Liquidity constraints and precautionary saving*. National Bureau of Economic Research. Available at: http://www.nber.org/papers/w8496 .

Carroll, C. D., & Samwick, A. A. (1998). How important is precautionary saving? *Review of Economics and*

Statistics, *80*(3), 410–419.

Carter, M. R., & Barrett, C. B. (2006). The economics of poverty traps and persistent poverty: An asset-based approach. *The Journal of Development Studies*, *42*(2), 178–199.

Carter, M. R., Little, P. D., Mogues, T., & Negatu, W. (2007). Poverty traps and natural disasters in Ethiopia and Honduras. *World development*, *35*(5), 835–856.

Carter, M. R., & Lybbert, T. J. (2012). Consumption versus asset smoothing: testing the implications of poverty trap theory in Burkina Faso. *Journal of Development Economics*, *99*(2), 255–264.

Cissé, J.D., & Barrett, C.B. (2016). Estimating development resilience: a conditional moments-based approach. *African Development Bank Group*, Working Paper No. 236.

Costello, A. & Osborne, J. (2005). Best practices in exploratory factor analysis: Four recommendations for getting the most from your analysis. *Practical Assessment, Research & Evaluation* 10(7):1-9.

Deaton, A. (1985). Panel data from time series of cross-sections. *Journal of econometrics*, *30*(1–2), 109–126.

Deaton, A. (1992). *Understanding consumption*. Oxford University Press.

Dell, M., Jones, B. F., & Olken, B. A. (2014). What do we learn from the weather? The new climate–economy literature. *Journal of Economic Literature*, *52*(3), 740–798.

Dercon, S. (2004). Growth and shocks: evidence from rural Ethiopia. *Journal of Development Economics*, *74*(2), 309–329.

Food and Agriculture Organization (FAO). (2016). Resilience Index Measurement and Analysis (RIMA) - II. Rome: FAO. Available at: http://www.fao.org/3/a-i5665e.pdf .

Hansen, B. E. (2000). Sample splitting and threshold estimation. *Econometrica*, *68*(3), 575–603.

Holling, C. S. 1973. Resilience and stability of ecological systems. *Annual Review of Ecological Systems* 4:1nualMay, R. M. 1977. Thresholds and breakpoints in ecosystems with a multiplicity of stable states. *Nature* 269:471–477

Kolenikov S., Angeles G. 2004. *The use of discrete data in pca: theory, simulations, and applications to socioeconomic indices*. Chapel Hill: Carolina Population Center, University of North Carolina, 1-59. URL: http://www.cpc.unc.edu/measure/resources/publications/wp-04-85.

Jalan, J., & Ravallion, M. (2002). Geographic poverty traps? A micro model of consumption growth in rural

China. *Journal of applied econometrics*, *17*(4), 329–346.

McKenzie, D. J. (2004). Asymptotic theory for heterogeneous dynamic pseudo-panels. *Journal of Econometrics*, *120*(2), 235–262.

Morduch, J. (1995). Income smoothing and consumption smoothing. *The journal of economic perspectives*, *9*(3), 103–114.

Paxson, C. H. (1992). Using weather variability to estimate the response of savings to transitory income in Thailand. *The American Economic Review*, 15–33.

Pingali, P., Alinovi, L., & Sutton, J. (2005). Food security in complex emergencies: enhancing food system resilience. *Disasters*, 29(1), s5-s24.

Propper, C., Rees, H., & Green, K. (2001). The demand for private medical insurance in the UK: a cohort analysis. *The Economic Journal*, *111*(471), 180–200.

Smith, L. C., & Frankenberger, T. R. Does Resilience Capacity Reduce the Negative Impact of Shocks on Household Food Security? Evidence from the 2014 Floods in Northern Bangladesh. *World Development*. Available at: http://dx.doi.org/10.1016/j.worlddev.2017.07.003 .

Townsend, R. M. (1995). Consumption insurance: An evaluation of risk-bearing systems in low-income economies. *The Journal of Economic Perspectives*, *9*(3), 83–102.

Upton, B.J., Cissé, D. J., & Barrett, B.C. (2016). Food security as resilience: reconciling definition and measurement. *Agricultural Economics*, 47(1), 135-147.

Verbeek, M. (2008). Pseudo-panels and repeated cross-sections. In *The econometrics of panel data*., 369–383. Springer.

Verbeek, M., & Nijman, T. (1992). Can cohort data be treated as genuine panel data? *Empirical economics*, *17*(1), 9–23.

Verbeek, M., & Vella, F. (2005). Estimating dynamic models from repeated cross-sections. *Journal of econometrics*, *127*(1), 83–102.

Wang, Q., & others. (2015). Fixed-effect panel threshold model using Stata. *Stata Journal*, *15*(1), 121–134.

Zimmerman, F. J., & Carter, M. R. (2003). Asset smoothing, consumption smoothing and the reproduction of inequality under risk and subsistence constraints. *Journal of Development Economics*, *71*(2), 233–

260.

Table 1

Impacts on food consumption growth – Pseudo-panel Version 1

Dependent variable: ΔFood	(1)	(2)	(3)	(4)	(5)
L1.Food	-51.948***	-44.487***	-46.986***	-46.577***	-51.423***
	(5.238)	(4.702)	(5.438)	(6.006)	(4.569)
L1.RCI	3.776***	3.195***	3.098***	3.306***	3.658***
	(0.427)	(0.370)	(0.427)	(0.468)	(0.358)
ΔRCI	7.389***	7.418***	7.426***	7.366***	7.325***
	(0.101)	(0.122)	(0.112)	(0.129)	(0.106)
ΔTemp	-12.998	-10.697	4.966	-13.141	-255.308**
	(17.367)	(14.247)	(16.323)	(15.812)	(110.287)
Low average pre-shock RCI x ΔTemp		-24.367**			
		(10.476)			
Hot x ΔTemp		-12.371	-1.879	-15.511**	-9.262
		(8.172)	(7.035)	(7.195)	(7.260)
ΔPre	-1.959	-3.182**	-4.881**	-5.141***	20.469
	(1.494)	(1.422)	(1.933)	(1.646)	(25.403)
Low average pre-shock RCI x ΔPre		-0.999			
		(3.400)			
Hot x ΔPre		3.976*	4.180*	4.511**	6.674**
		(2.054)	(2.218)	(1.944)	(2.596)
Low pre-shock RCI x ΔTemp			-60.694***		
			(15.635)		
Low pre-shock RCI x ΔPre			6.632		
			(3.110)		
Low pre-shock RCI			0.112		
			(0.285)		
Low initial RCI x ΔTemp				-27.988***	
				(9.247)	
Low initial RCI x ΔPre				-0.506	
				(3.081)	
L1.RCI x ΔTemp					3.802**
					(1.663)
L1.RCI x ΔPre					-0.404
					(0.408)
Constant	-63.728***	-51.077***	-56.872***	-52.344***	-62.922***
	(14.159)	(10.358)	(12.971)	(12.784)	(11.543)
Observations	100	100	100	100	100
Adjusted R-squared	0.993	0.994	0.994	0.994	0.994
Biophysical controls	Yes	Yes	Yes	Yes	Yes

Total temperature effect for cohorts with low average pre-shock RCI	-35.064**		
	(16.550)		
Total temperature effect for cohorts living in hot areas	-23.068	3.087	-28.653
	(17.946)	(19.995)	(19.477)
Total temperature effect for cohorts with low average pre-shock RCI living in hot areas	-47.435**		
	(20.963)		
Total temperature effect for cohorts with low pre-shock RCI		-55.728***	
		(16.864)	
Total temperature effect for cohorts with low average pre-shock RCI living in hot areas		-57.608***	
		(19.609)	
Total temperature effect for cohorts with low pre-shock RCI			-41.129**
			(18.034)
Total temperature effect for cohorts with low average pre-shock RCI living in hot areas			-56.641**
			(22.040)

Notes: All specifications include cohort and time fixed effects. Biophysical controls include slope, elevation and length of the growing period. ΔFood is food consumption growth rate, i.e. the average annual percentage change in (ln) cohort monthly per capita food consumption between t and t-1. L1.Food is lagged (ln) cohort monthly per capita food consumption. L1.RCI is the lagged Resilience Capacity Index, scaled from 1 to 100. ΔRCI is the annualised change in the RCI Index between t and t-1. ΔTemp is the difference in logarithms of average temperature levels at and t-1, both scaled by long-run means. ΔPre is the difference in logarithms of average total precipitation levels at t and t-1, both scaled by long-run means. Low average pre-shock RCI is dummy with value 1 for cohorts having an average pre-shock RCI below the 25th percentile. Low pre-shock RCI is a dummy taking value 1 for cohorts with a lagged RCI below the 25th percentile in a given wave. Low initial RCI is a dummy with value 1 for households with an initial RCI below the 25th percentile. Hot is a dummy taking value 1 for cohorts living an area with an above mean long-run average temperature.

Standard errors are in parentheses and are clustered at the cohort level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2

Threshold model – Pseudo-panel Version 1

	(1)
Dependent variable:	ΔFood
L1.Food	-52.711***
	(5.651)
L1.RCI	3.777***
	(0.445)
ΔRCI	7.363***
	(0.099)
ΔPre	-3.075*
	(1.496)
ΔTemp_Lower regime	-61.361***
	(20.920)
ΔTemp_Upper regime	-5.661
	(16.505)
Constant	-64.476***
	(14.319)
Observations	100
Adjusted R-squared	0.994
Biophysical controls	Yes

Notes: All specifications include cohort and time fixed effects. Biophysical controls include slope, elevation and length of the growing period. ΔFood is food consumption growth rate, i.e. the average annual percentage change in (ln) cohort monthly per capita food consumption between t and t-1. L1.Food is lagged (ln) cohort monthly per capita food consumption. L1.RCI is the lagged Resilience Capacity Index, scaled from 1 to 100. ΔRCI is the annualised change in the RCI Index between t and t-1. ΔTemp is the difference in logarithms of average temperature levels at and t-1, both scaled by long-run means. ΔPre is the difference in logarithms of average total precipitation levels at t and t-1, both scaled by long-run means. Hot is a dummy taking value 1 for cohorts living an area with an above mean long-run average temperature.
Standard errors are in parentheses and are clustered at the cohort level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3

Threshold tests and confidence intervals – Pseudo-panel Version 1

RCI Threshold*

Model	Threshold	Lower	Upper
RCI	54.609	53.130	55.139

* The threshold value of RCI is at time t-1.

Threshold effect test (bootstrap = 300):

Threshold	RSS	MSE	Fstat	Prob	Crit10	Crit5	Crit1
Single	13.5417	0.1411	21.08	0.0033	13.2478	15.5232	20.0510

Table 4

Impacts on food consumption growth – Pseudo-panel Version 2

Dependent variable: ΔFood	(1)	(2)	(3)	(4)	(5)
L1.Food	-42.068***	-45.379***	-47.292***	-42.862***	-46.258***
	(4.448)	(3.917)	(4.360)	(3.481)	(3.957)
L1.RCI	2.996***	3.293***	3.394***	3.033***	3.258***
	(0.352)	(0.308)	(0.337)	(0.269)	(0.290)
ΔRCI	7.146***	7.266***	7.230***	7.208***	7.212***
	(0.129)	(0.145)	(0.131)	(0.118)	(0.127)
ΔTemp	-5.067	-7.918	-6.398	-5.666	-185.110**
	(4.619)	(5.566)	(7.760)	(4.323)	(70.249)
Low average pre-shock RCI x ΔTemp		-7.593			
		(9.612)			
Hot x ΔTemp		-12.772	-12.212	-18.023**	-16.717**
		(7.541)	(7.708)	(6.488)	(7.552)
ΔPre	-1.944	0.691	1.387	0.402	-42.163
	(2.519)	(2.942)	(2.686)	(2.514)	(34.675)
Low average pre-shock RCI x ΔPre		-4.326			
		(4.911)			
Hot x ΔPre		-6.972**	-6.954**	-6.859**	-6.502**
		(2.918)	(3.061)	(2.845)	(2.735)
Low pre-shock RCI x ΔTemp			-11.755		
			(8.460)		
Low pre-shock RCI x ΔPre			-4.829		
			(5.190)		
Low pre-shock RCI			0.382		
			(0.371)		
Low initial RCI x ΔTemp				-19.087***	
				(5.819)	
Low initial RCI x ΔPre				-1.262	
				(3.499)	
L1.RCI x ΔTemp					2.938**
					(1.180)
L1.RCI x ΔPre					0.692
					(0.556)
Constant	-53.564***	-65.786***	-65.701***	-58.154***	-61.601***
	(12.066)	(10.825)	(11.398)	(10.096)	(9.386)
Observations	100	100	100	100	100
Adjusted R-squared	0.996	0.997	0.997	0.997	0.997
Biophysical controls	Yes	Yes	Yes	Yes	Yes

Total temperature effect for cohorts with low average pre-shock RCI	-15.511* (9.033)		
Total temperature effect for cohorts living in hot areas	-20.690** (7.879)	-18.610** (8.831)	-23.689*** (7.328)
Total temperature effect for cohorts with low average pre-shock RCI living in hot areas	-28.283** (9.964)		
Total temperature effect for cohorts with low pre-shock RCI		-18.153*** (6.413)	
Total temperature effect for cohorts with low average pre-shock RCI living in hot areas		-30.364*** (9.138)	
Total temperature effect for cohorts with low pre-shock RCI			-24.753*** (6.689)
Total temperature effect for cohorts with low average pre-shock RCI living in hot areas			-42.776*** (9.341)

Notes: All specifications include cohort and time fixed effects. Biophysical controls include slope, elevation and length of the growing period. ΔFood is food consumption growth rate, i.e. the average annual percentage change in (ln) cohort monthly per capita food consumption between t and t-1. L1.Food is lagged (ln) cohort monthly per capita food consumption. L1.RCI is the lagged Resilience Capacity Index, scaled from 1 to 100. ΔRCI is the annualised change in the RCI Index between t and t-1. ΔTemp is the difference in logarithms of average temperature levels at and t-1, both scaled by long-run means. ΔPre is the difference in logarithms of average total precipitation levels at t and t-1, both scaled by long-run means. Low average pre-shock RCI is dummy with value 1 for cohorts having an average pre-shock RCI below the 25th percentile. Low pre-shock RCI is a dummy taking value 1 for cohorts with a lagged RCI below the 25th percentile in a given wave. Low initial RCI is a dummy with value 1 for households with an initial RCI below the 25th percentile. Hot is a dummy taking value 1 for cohorts living an area with an above mean long-run average temperature.
Standard errors are in parentheses and are clustered at the cohort level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5

Threshold model – Pseudo-panel Version 2

	(1)
Dependent variable:	ΔFood
L1.Food	-45.248***
	(4.568)
L1.RCI	3.211***
	(0.360)
ΔRCI	7.134***
	(0.113)
ΔPre	-1.822
	(2.261)
ΔTemp_Lower regime	-16.094***
	(4.811)
ΔTemp_Upper regime	2.288
	(5.901)
Constant	-60.330***
	(12.500)
Observations	100
Adjusted R-squared	0.997
Biophysical controls	Yes

Notes: All specifications include cohort and time fixed effects. Biophysical controls include slope, elevation and length of the growing period. ΔFood is food consumption growth rate, i.e. the average annual percentage change in (ln) cohort monthly per capita food consumption between t and t-1. L1.Food is lagged (ln) cohort monthly per capita food consumption. L1.RCI is the lagged Resilience Capacity Index, scaled from 1 to 100. ΔRCI is the annualised change in the RCI Index between t and t-1. ΔTemp is the difference in logarithms of average temperature levels at and t-1, both scaled by long-run means. ΔPre is the difference in logarithms of average total precipitation levels at t and t-1, both scaled by long-run means. Hot is a dummy taking value 1 for cohorts living an area with an above mean long-run average temperature.
Standard errors are in parentheses and are clustered at the cohort level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6

Threshold tests and confidence intervals – Pseudo-panel Version 2

RCI Threshold*

Model	Threshold	Lower	Upper
RCI	60.796	60.379	60.835

* The threshold value of RCI is at time t-1.

Threshold effect test (bootstrap = 300):

Threshold	RSS	MSE	Fstat	Prob	Crit10	Crit5	Crit1
Single	16.528	0.1722	15.78	0.0267	11.6289	13.7316	19.6328

Appendix

Table A.1

RCI Variables - Descriptive statistics

	Mean	Var	sd	Obs
Dwelling Index	-0.150	0.459	0.677	16190
Distance from hospital (inverse)	0.764	5.656	2.378	16190
Distance from primary school (inverse)	17.627	264.651	16.268	16190
Wealth Index	0.008	0.427	0.653	16190
Agricultural Wealth Index	0.151	1.508	1.228	16190
Tropical Livestock Units (per capita)	0.296	1.436	1.198	16190
Land owned (per capita)	0.424	1.533	1.238	16190
Public Transfers (per capita)	1.163	294.673	17.166	16190
Private transfers (per capita)	12.186	2017.016	44.911	16190
Participation in a saving group	0.044	0.042	0.204	10254
Average years of education	4.724	8.730	2.955	16190
Dependency ratio (inverse)	2.056	0.878	0.937	16190
Farming is not the main source of income	0.268	0.196	0.443	16190
Monthly per capita food expenditure (usd)	28.671	588.273	24.254	16190
Simpson Index	0.605	0.026	0.161	16190

Table A.2

Pillar factor loadings for ABS

Variable	Factor 1	Factor 2	Uniqueness
Dwelling Index	0.154	0.101	0.966
Distance from hospital (inverse)	0.276	0.009	0.924
Distance from primary school (inverse)	0.216	-0.084	0.946

Table A.3

Pillar factor loadings for AST

Variable	Factor 1	Factor 2	Factor 3	Uniqueness
Wealth Index	0.708	-0.148	0.053	0.475
Agricultural Wealth Index	0.764	-0.025	-0.044	0.413
Tropical Livestock Units (per capita)	0.322	0.276	-0.042	0.818
Land owned (per capita)	0.133	0.261	0.070	0.909

Table A.4

Pillar factor loadings for AC

Variable	Factor 1	Factor 2	Uniqueness
Average years of education	0.403	0.144	0.817
Dependency ratio (inverse)	-0.122	0.228	0.933
Farming is not the main source of income	0.463	-0.066	0.782

Table A.5

MIMIC results

ABS	0.048 (22.13)**
AST	0.021 (5.24)**
SSN	0.002 (16.94)
AC	0.071 (22.77)**
Food_expenditure_usd (log)	1 --
Simpson Index	0.105 (21.59)**
N	16,190
Chi2	58.106
Prob > Chi2	0.000
CFI	0.988
TLI	0.965
RMSEA	0.034
RCI	61.292
St.dev.	10.057
Min	0
Max	100

* p<0.05; ** p<0.01

Table A.6

Characteristics of the pseudo-panels

Version	Variables used for cohort construction	Number of cohorts (C)	Average N of observations per cohort (n_c)
1	Long-run average temperature quintiles*Year-of-birth of the household head quintiles	25	647.6
2	Long-run average precipitation quintiles*Year-of-birth of the household head quintiles	25	647.6

Table A.7

Descriptive statistics – Pseudo-panel Version 1

	Mean	Var	sd	Obs
Food consumption level	32.855	88.180	9.390	125
ΔFood	6.019	28.349	5.324	100
RCI	63.378	17.644	4.201	125
ΔRCI	0.789	0.576	0.759	100
Temperature	23.197	3.430	1.852	125
Precipitation	60.088	26.546	5.152	125
Long-run temperature	23.054	3.446	1.856	125
Long-run precipitation	60.348	13.046	3.612	125
ΔTemp	0.002	0.000	0.018	100
ΔPre	0.033	0.004	0.064	100

Notes:
Food consumption is cohort monthly per capita food consumption expressed in dollars at 2010 Purchasing Power Parity (PPP). ΔFood is the annualised food consumption growth rate between t and t-1, i.e. the average annual percentage change in (ln) cohort per capita food consumption. RCI is the Resilience Capacity Index, scaled from 1 to 100. ΔRCI is the annualised change in the RCI Index between t and t-1. Temperature is average monthly temperature in the years between t and t-1, expressed in degree Celsius. Precipitation is average monthly precipitation in the years between t and t-1, expressed in mm. ΔTemp is the difference in logarithms of average temperature levels at and t-1, both scaled by long-run means. ΔPre is the difference in logarithms of average precipitation levels at t and t-1, both scaled by long-run means. Long-run temperature is average monthly temperature during the period 1981-2014. Long-run precipitation is average monthly precipitation during the period 1983-2016.

Table A.8

Descriptive statistics – Pseudo-panel Version 2

	Mean	Var	sd	Obs
Food consumption level	33.162	82.441	9.080	125
ΔFood	6.691	60.786	7.797	100
RCI	63.496	16.392	4.049	125
ΔRCI	0.887	1.222	1.106	100
Temperature	23.256	0.980	0.990	125
Precipitation	60.141	114.408	10.696	125
Long-run temperature	23.108	0.776	0.881	125
Long-run precipitation	60.470	88.037	9.383	125
ΔTemp	0.006	0.001	0.024	100
ΔPre	0.032	0.005	0.071	100

Notes:
Food consumption is cohort monthly per capita food consumption expressed in dollars at 2010 Purchasing Power Parity (PPP). ΔFood is the annualised food consumption growth rate between t and t-1, i.e. the average annual percentage change in (ln) cohort per capita food consumption. RCI is the Resilience Capacity Index, scaled from 1 to 100. ΔRCI is the annualised change in the RCI Index between t and t-1. Temperature is average monthly temperature in the years between t and t-1, expressed in degree Celsius. Precipitation is average monthly precipitation in the years between t and t-1, expressed in mm. ΔTemp is the difference in logarithms of average temperature levels at and t-1, both scaled by long-run means. ΔPre is the difference in logarithms of average precipitation levels at t and t-1, both scaled by long-run means. Long-run temperature is average monthly temperature during the period 1981-2014. Long-run precipitation is average monthly precipitation during the period 1983-2016.