Long-Run Effects of Rainfall Shocks

An evaluation of the effects of in utero negative rainfall shocks on children's cognitive ability in rural Kenya

Author: Linn Mattisson
Supervisor: Jan Bietenbeck

NEKP01 Master Essay II
Lund University
Department of Economics
Page total: 29

June 5th 2018

Acknowledgements

First and foremost, my deepest gratitude to Dr Jan Bietenbeck for first suggesting I look into the subject of this essay, and for his patience when answering my many questions.

I would also like to thank my friends in the Lund economics masters – you know who you are – for the past two years. There are no other people with whom I could spend that much time in a basement.

Finally, as always, I owe my family, my Glasgow girls and Thomas.

Abstract

Shocks in utero can have lasting effects on an individual's physical and socio-economic outcomes. By merging direct rainfall data at the district-level with an individual-level dataset on children's education and socio-economic factors I create a proxy for in utero income shock with negative rainfall to evaluate the effect on Kenyan children's cognitive score in rural districts. I find that a drought leads to a decrease of 2.60 % of the standard deviation fall in children's cognitive score. These results are robust for gender and age, as well as the definition of the shock and rural districts. The result suggests that policies aimed at protecting rural households against shocks and alleviating their effects could save large amounts in social costs in the long run.

Keywords: Rainfall, drought, education, fetal origins, Kenya, development

Table of Contents

List of Figures

List of Tables

1 Introduction

The livelihood of many people in developing countries depends crucially on the weather, owing to a large dependence on the primary production sector. In particular, weather shocks pose a significant source of risk to the income of rural, agriculture dependent households. One of the most common weather shocks threatening these households' livelihoods are droughts. Regions such as Sub-Saharan Africa that have low levels of development and where 96% of its cropland is rain-fed (FAO, 2012 quoted in Randell and Gray, 2016) are particularly vulnerable to droughts (Hallegatte et al., 2016). An especially interesting question is whether droughts have long-run effects. There is some evidence that droughts can have long-run negative impacts on individuals' health, cognitive development, and economic well-being (e.g. Dinkelman, 2017; Shah and Steinberg, 2017), although the evidence is still sparse, in particular for Africa. Additionally, as droughts are also expected to increase in frequency and severity with climate change (IPCC, 2007), income shocks caused by droughts will only increase in relevance and scope.

In this paper, I study whether in utero exposure to a drought affects children's long-run cognitive development in Kenya. Several previous studies have shown that negative shocks while a child is in utero can have a lasting effect on children's physical and cognitive development. (e.g. Maccini and Yang, 2009; Burke et al., 2015). Furthermore, cognitive skills have been widely recognised to be a good predictor for labour market success at the individual level (Caruso, 2017) and of economic growth at the country level (Hanushek and Woessmann, 2012).

My empirical analysis exploits variation in rainfall levels within small geographic areas (districts) over time. This introduces credible exogenous variation to my estimation since rainfall is unaffected by human decisions. Specifically, I estimate the impact of being in utero during a year in which rainfall is below the 20^{th} percentile of the long-term mean of a district on cognitive test scores of children aged 8-16. My regressions control for district and age specific effects; thus, the impacts are identified as being solely within-district variation. The test scores measure children's basic numeracy and literacy skills using standardised tests conducted as part of the so-called Uwezo surveys. Since these surveys are representative at the district level it is then possible to combine with district-level rainfall data while retaining the same degree of representation.

I find that a negative rainfall shock leads to a decrease equal to 2.60% of the standard deviation in the cognitive scores. The response to the shock depends neither on age nor gender, which further highlights the persistence of the effect of the shock. The results are as expected more severe when a shock is defined as a year in which rainfall is below the 10^{th} percentile of the long-term mean, and likewise has a stronger effect the more narrowly a rural district is defined. This in turn further supports the belief that the effect on children's cognitive ability is driven by mechanisms related to agrarian production.

This paper is part of a growing body of literature that concerns itself with the impact weather shocks have on child and adult outcomes. Another such paper is that by Shah and Steinberg (2017) who find that children exposed to higher rainfall between the ages 0-2 have consistently higher test scores and are more likely to enrol in school (whereas rainfall shocks in later life leads children to pursue productive work instead of schooling). Maccini and Yang (2009) find that more generous early-life rainfall leads to better health and schooling outcomes for women in Indonesia. Caruso (2017) notes that being exposed to excessive rainfall in a storm in Guatemala has a negative effect on children's years of schooling and health. Furthermore, Kumar et al. (2016) find that in utero droughts are associated with lower childhood nutritional status in India while Dinkelman (2017) likewise shows that in the South African Homelands an exposure to a drought in infancy leads to a higher occurrence of mental and physical capabilities.

What the result of these studies seem to suggest is that severe shocks – such as a cyclone or a drought – while a child is either in utero will have negative effects on children's physical and cognitive development, whereas slightly higher rainfall which leads to more agrarian production, such as the case with Shah and Steinberg (2017) will have positive effects on children's outcomes. Whether rainfall is too much or too little can be said to be determined by the effect it has on income: If the rainfall leads to crop failure and loss of income it has a negative effect on children's outcomes. Thus, droughts can be seen as a proxy for negative rural income shocks. This can be further argued to affect child outcomes since an income-constrained household might have to reduce maternal or infant medical investment, or lose the ability to obtain nutritious food. These constraints can also be explained by the so-called fetal origins hypothesis developed by Barker (1998). Here, Barker argues that negative in utero shocks have negative effects lasting well to adulthood. When maternal health is compromised, babies can suffer from low metabolic function, low birth weight as well as cognitive impairment, all of which can lead to lower human capital potential (Currie, 2009).

The main contribution of this paper is twofold. First, this paper presents one of the first analyses of its kind to Africa, where droughts are particularly important: Kenya as a relatively rural and poor economy in Sub-Saharan Africa is likely to be sensitive to changes in weather. The second contribution of this paper is the use of the Uwezo dataset which allows me to use standardised test scores to evaluate children's cognitive ability. This is in many ways an improvement over more common schooling measurements such as enrolment (used by e.g. Baez et al., 2017) or years of schooling (used by e.g. Randell and Gray, 2016; Alderman et al.; 2006) as these are more indirect measures of children's actual cognitive capability.

This structure of the rest of the paper is as follows: The succeeding section (2) will present the data used. I next move on to explain my the estimation strategy in section 3. Following this, section 4 presents the results. Section 5 provides some robustness checks of the main result. Then, section 6 provides a brief discussion of the results which then closes with the conclusion in section 7.

2 Data

2.1 Data Collection

2.1.1 Kenyan UWEZO Data

The main data are taken from the Kenyan Uwezo survey data. Uwezo started collecting data in 2009 as a part of Twaweza, an independent East African initiative that promotes information and social capabilities in Kenya, Tanzania and Uganda. (Uwezo East Africa at Twaweza, 2014). The data are collected from the IPUMS International database (Minnesota Population Center, 2015). Households participating in the survey are randomly chosen each year in a process in such a way that the sample is representative at both national and district level. This produces a repeated cross-section data set. In this paper I use five survey waves, from 2011 to 2015. In the chosen household, all children aged 6-16 are assessed on core numeracy and literacy skills with complementary socio-economic household characteristics also collected. The test scores are based on a standard grade 2 test in Maths, English and Swahili (where the expected age is 8). This includes being able to, for example, solve an addition problem or read a paragraph.

The Uwezo surveys thus have three characteristics that are key to this paper. First, the scope of the survey means the survey information is representative to the whole country. This means there are a high number of observations in the sample from which information can be drawn, but more importantly it also means all children's test scores are available. Normally, data on schooling only involves those enrolled in school, so this sample is more representative of the whole population. This is quite unusual. Second, these standardised test scores provide a better comparison and more accurate reflection of the actual cognitive ability of children as opposed to the more common years of schooling metric, which is more readily available but far less telling. Third, the household characteristics that are collected also provide a geographic identifier by stating what district the child is surveyed in, and the age of the children are at the time of the survey. This will prove essential when merging the

UWezo dataset with the spatial data.

2.1.2 Spatial Data

The precipitation data are collected from CHIRPS 2.0 Stations (see Funk et al., 2014) on the Kenyan district level. To make the link between the Uwezo and the spatial data, the district borders must be defined the same way. Following the same methodology as Bietenbeck et al. (2017) I define the district borders in the spatial data by the same 2009 census as in the Uwezo survey data.

To find the in utero shock I take the survey year minus the age of the child to find the birth year where the child is aged 0. I then generate a dummy variable which takes the value one if the district experienced a negative rainfall shock in that year. A rainfall shock is defined by a year where the rainfall lies below the 20^{th} percentile of rainfall over the available range in the data, 1980-2016. Two examples of this is provided in Figure 1. For Baringo Central, the blue observations below the lower red line represent the potential shock years. For Bungoma West, the red observations below the upper red line constitute potential shock years. By design, a shock is therefore defined relative to each region's distribution of rainfall. This is intuitive from an agricultural perspective: As crops are adapted to the region's climate, it is deviations from what is normal within the district that will lead to lower crop yield and income. This definition is also used in similar studies (Shah and Steinberg, 2017; Baez et al., 2017; Bietenbeck et al., 2017).

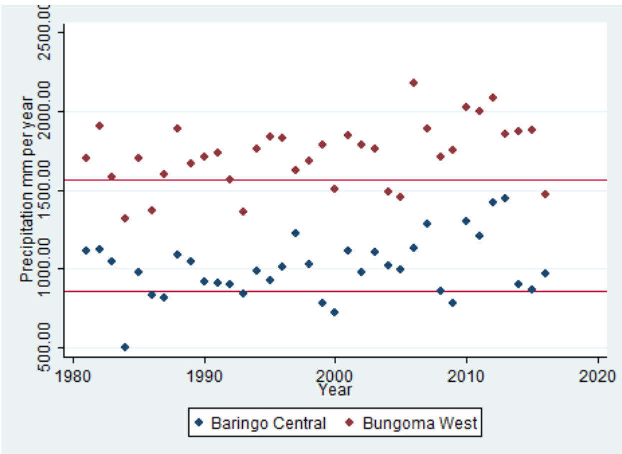

Figure 1: Shock Definitions

Finally, I also collect data on night light density measured from space from the year 2000. I extract the data from the US National Oceanic and Atmospheric Information (see NOAA, 2014). Night light has been shown to be a good proxy for economic growth, productivity, and above all urbanisation (Mellander et al., 2015), which is the purpose of this data in this study.

2.2 Descriptive Statistics

Table 1 summarises the main variables used. Notably, English, Maths and Swahili scores have been standardised by age and survey year to have mean zero and standard deviation one. This is why, in the table, the mean and standard deviation appear to be slightly inaccurate; the table summarises *across* age and wave rather than *within* each age-wave combination. The cognitive score is made by first taking the average of these three outcomes for each child. This is then standardised again according to the survey wave and age.

Next, Basic Maths and Basic English take the value one if the child passed the standard criteria for either subjects, defined as being able to do multiplication or read a paragraph, respectively. Enrollment takes the value 1 if the child has been enrolled at any point. The table also shows that rural districts far outnumber urban ones.

Table 1: Descriptive Statistics

Context	Variable Name	Unit	Mean	SD
Education				
	Cognitive Score	Std Score	-0.0131	1.005265
	Maths Score	Std Score	-0.01129	1.006043
	English Score	Std Score	-0.01344	1.004223
	Swahili Score	Std Score	-0.01211	1.005214
	Basic Maths	Dummy	0.610322	0.487678
	Basic English	Dummy	0.663241	0.472602
	Enrolled	Dummy	0.937703	0.241695
Socio-Economic				
	Age	Count [8,16]	11.66399	2.517605
	Girl	Dummy	0.487072	0.499833
	Urban District	Dummy	0.0280174	0.1650228
	Gender-Ratio (Girls/Boys)	Count ratio	0.9533505	0.0833859
	Mothers Edu	Ordered Ranking	0.88166	0.690096
Spatial				
	Shock (<20p)	Dummy	0.252385	0.434381
	Shock (<10p)	Dummy	0.1228998	0.3283225
	Nightlight density (2000)	% of light pixels	4.002989	10.71007

Rainfall is defined as the mean annual precipitation in each district. Again, an in utero shock is defined as a dummy taking the value one if there was a drought in the year the child is aged zero. Night light density is taken from the year 2000. It is the district-average percentage of light pixels on a night satellite picture from space.

3 Empirical Strategy

3.1 Identification

In this paper I evaluate the effect of in utero exposure to a drought. However, as Glewwe and King (2006) emphasise, it is difficult to extract a causal effect by simply observing children who experience an in utero shock and their later education performance; any effect driven by omitted variables or other endogeneity issues can confound the causal effect. The exception is if a pure experiment or a natural experiment can be identified. Rainfall, by its nature, falls into the category of a natural experiment and can be considered random since it is unaffected by human decisions. This quasi-random allocation of droughts creates treatment and control groups that should not differ systematically other than from the shock itself.

As the rainfall data is on the district level, I cannot strictly compare individuals treated or not treated with a drought. Instead, I compare children living in a district which experienced a drought to children living in districts that did not experience a shock. This will however still succeed to capture a causal effect of a drought since a household is directly affected by a drought in their district. This might be directly in the form of consumption of farm yields, which decrease in the event of a drought. A household could also be affected indirectly if they generate income from farm labour: as crop failure reduces the demand for farm labour, a shock will lead to less income generated which leads to lower consumption (see e.g. Shah and Steinberg, 2017 for more detail). A shock in a rural district is also likely to reduce the supply of food goods available on the formal market, increase food prices, and lead to a more general economic downturn since the dominating production sector (the rural sector) is constrained.

The identification of the causal effect also crucially assumes there is a noticeable effect on children from the rainfall shock. This has frequently been argued and showed empirically in related literature but the intuition in a rural developing country was nicely summarised by Kumar et al. (2016) as a drought leading to "crop failure and a steep decline in household income, which in turn may affect maternal and fetal nutrition through reduced food consumption" (p 54).

3.2 Model Specification

Formally, I employ an Ordinary Least Squares model which generates a difference-in-differences estimate from running the following regression equation as my baseline:

$$CognitiveScore_{icdw} = \alpha + \delta RainfallShock_d + X'_{idc}\beta + \gamma_1 Cohort_c + \gamma_2 District_d + \gamma_3 Wave_w + \varepsilon_{icdw} \quad (1)$$

where i refers to an individual child, c refers to the cohort (birth year), d is district-specific and w refers to the survey wave. The shock is, as discussed in the data section, district-specific in space and time-specific to the individual's birth year. γ_1 captures each cohort fixed effect while γ_2 capture survey wave fixed effects. γ_3 are district dummies capturing district fixed effects.

The specification above can be interpreted as a difference-in-difference model. The fixed effects correct for cohort-, survey wave- and district-invariants. The coefficient δ thus captures the difference in the cognitive score between children who were and were not exposed to an in utero rainfall shock on a within-district level. This allows me to find the causal effect on the cognitive ability of children exposed to a shock.

To make the results more efficient I would also like to include control variables, captured by the term $X'_{idc}\beta$. However, as the shock occurs around the birth of the child, all socio-economic factors measured in the survey are possibly outcomes of the shock itself. Some examples of these are household income, number of siblings and the occupation of the parents. Including these controls could potentially bias the estimation.

I have two candidates for controls that should avoid the bad control problem. The first one is the highest level of education achieved by the mother. This has been shown to correlate with the socio-economic status and human capital of a person (Currie and Moretti, 2003). The other possible control variable is the level of economic growth and urbanisation of a district as proxied by night light density. As night light change is slow and unlikely to correlate with children's cognitive skill, I use district mean density for the year 2000 throughout.

To prepare for estimation I further cut the sample in two ways. Firstly, since the school test scores are designed for grade 2 students (8 year olds) I drop children younger than 8 years. Secondly, I want to improve the link between the rainfall shock and income by only including rural areas as only rural households are directly affected by droughts. The 2013 and 2014 survey years include a categorisation of districts as being either "urban", "semi-arid", "normal" or "arid". I drop the districts categorised as urban, which leaves me with 147 districts. This together with the merging of the spatial data leaves me with a dataset of about 470,000 observations.

4 Results

4.1 Pre-Estimation Results

To get an initial overview of the data characteristics I first perform some preliminary analysis of the summary statistics. I compare the mean of the cognitive score in the sample for children who experienced a shock against those who did not experience a shock. I then compare the means for the individual subjects. As is clear in Table 2 there is a visible difference in means, where the children who experienced a shock have a lower mean than the control sample for all four outcomes. The same results are shown visually for cognitive score in Figure 2. The blue observations represent district means of the control sample, and the red observations the district means of the treatment group. While there is a generous spread across regions, there is clearly a trend visible. This is furthered captured by the two mean lines within the figure. As the descriptive analysis indicates there is a difference in the means of the cognitive score, I move on to my intended estimation.

Table 2: Non Parametric Mean Comparison

Variable	Mean	Mean	Diff-in-Means
Sample:	**No Shock**	**Shock**	-
Cognitive Score	-0.00332	-0.04208	-0.03876
Maths	-0.0012	-0.04114	-0.03994
English	-0.00489	-0.03876	-0.03387
Swahili	-0.00305	-0.03892	-0.03587

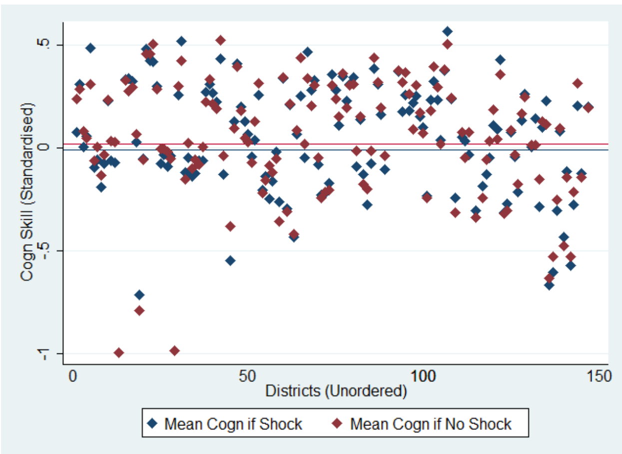

Figure 2: Mean Comparison

Finally, before performing the main regressions I also check for heteroskedasticity using the White Test. The null hypothesis of homoskedasticity is rejected and I conclude that the standard errors would give misleading result. I then cluster standard errors by district and subsequently find that a second White Test does not reject homoskedasticity. Thus, clustered standard errors are used in all regressions.

4.2 Main Results

The regression results from estimating Equation (1) are presented in Table 3. The first column is a stripped specification which only regresses cognitive score on the shock variable of interest and a constant. The subsequent columns add fixed effects as well as controls. When controlling for the fixed effects, the magnitude of the shock goes down but it increases marginally when adding control variables (which also reduces the sample size). The district fixed effects by far has the largest impact of the three fixed effects categories. While the mother's education is highly significant and positive as expected, controlling for night light does not improve the estimation. This could suggest that the economic growth/urbanisation night lights proxy are already captured in the fixed effects since night light density is defined on the district level here.

Table 3: Main Results

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
Rainfall Shock	-0.0388***	-0.0220***	-0.0389***	-0.0609***	-0.0365***	-0.0260***
	(0.0105)	(0.00598)	(0.0104)	(0.0160)	(0.00929)	(0.00626)
Mothers Edu						0.254***
						(0.0101)
Nightlight						6.274
						(4.384)
Constant	-0.00332	0.215***	-0.00134	0.00365	0.182***	-0.00511
	(0.0288)	(0.000973)	(0.0321)	(0.0372)	(0.0244)	(0.00358)
District FE	NO	YES	NO	NO	YES	YES
Wave FE	NO	NO	YES	NO	YES	YES
Cohort FE	NO	NO	NO	YES	YES	YES
Observations	439,915	439,915	439,915	439,915	439,915	392,246
R-squared	0.000	0.106	0.000	0.000	0.107	0.126

Here and henceforth: Robust standard errors clustered by district in parenthesis, *** p<0.01, ** p<0.05, * p<0.1. The outcome is the cognitive score.

The effect of an in utero shock is negative and significant at the 1% level for all six specifications and the size of the coefficient is fairly consistent. The sixth column is the preferred specification since this controls for all fixed effect factors as well as controls. The regression result in this column implies that a shock in utero decreases the cognitive score by 2.60 % of the standard deviation.

Following the main result where the outcome variable is cognitive score, I estimate the effect of an in utero shock on alternative schooling-related outcomes. I regress enrolment (using both a probit and OLS specification) as well as standardised Maths, English and Swahili on the rainfall shock. These results are available in the appendix. The effect on enrolment is weaker than it is for actual performance in terms of test scores, where test scores are of similar magnitude to the cognitive score. An in utero drought is associated with a decrease of 4.38% of the standard deviation in Maths test scores, a 3.43% of the standard deviation decrease in the English test scores with an equivalent 3.64% SD decrease in the Swahili test score.

I also estimate the effect of a rainfall shock on the incidence of basic Maths and English being met. The results imply that being exposed to a shock means your likelihood of meeting basic Maths (English) skills fall by 1.47% (1.57%) of the standard deviation. The smaller magnitude of these two variables than the cognitive score suggests that the link between basic Maths and English and the shock is weaker than for the cognitive score. While the reason for this can be analysed in several ways, I speculate that it is because a big part of the effect on cognitive ability is on those already far below the expected cognitive level. That is, those more affected by the shock are those that were less likely to pass the "basic" criteria from the start. This suggests that the shock might exacerbate the struggles of the already disadvantaged.

4.3 Heterogeneity Analysis

I next analyse heterogeneous effects of the in utero drought on age and gender, starting with the age decomposition analysis. The model (including fixed effects and controls) is first estimated unconditionally using interaction terms for each age in the sample taking the value one for each age. Hence, the first interaction is the shock as previously defined, multiplied by a dummy taking the value one where the child is aged eight. This can be formalised by the equation below:

$$\begin{aligned} CognitiveScore_{icdw} &= \sum_{i=8}^{16} [RainfallShock_d * AgeDummy_{icdw}] \\ &+ MothersEdu_{idcw} + \gamma_1 Cohort_c + \gamma_2 Age_i + \gamma_3 District_d + \gamma_4 Wave_w + \varepsilon_{icdw} \quad (2) \end{aligned}$$

Next, I analyse the age effect when conditioning the sample on the age of the child and then running separate regressions for each age. Hence, in the first estimation the whole sample consists of children aged eight. The regression outputs are given in

Tables 11 and 12 in the appendix. The results of interest are summarised in Figure 3 with the shock coefficient value on the y-axis and the age as according to each specification on the x-axis. The graph shows that there is a fairly constant negative effect of the rainfall shock on cognitive score. The fluctuation across age is small and also statistically insignificant.

Figure 3: Age-Decomposition of Rainfall Effect

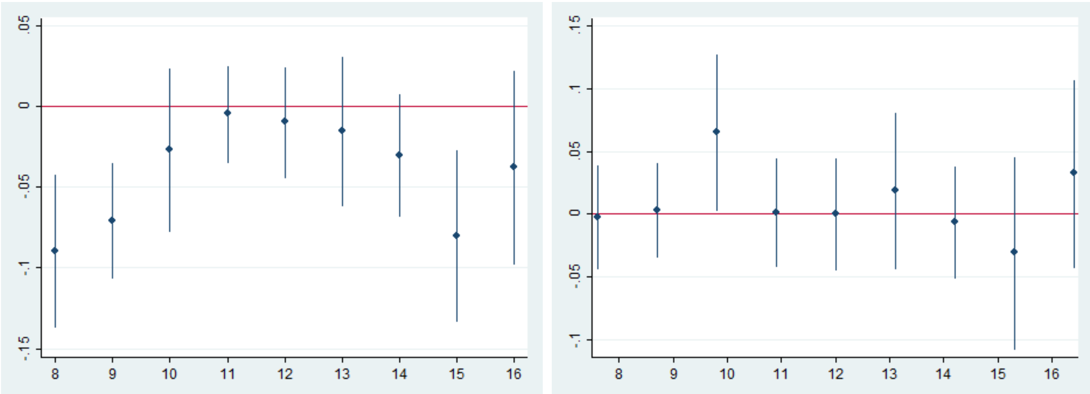

(a) Unconditional (b) Conditional

ALthough some coefficients are significant, is is difficult to attach any meaning to this given that the ages that are significant are somewhat inconsequential in their order and for the two different estimations. These small, mainly insignificant coefficients can be interpreted as the age of the child being unrelated to the effect of the in-utero exposure to the shock. As the age of the child is directly proportional to the time since the shock, it suggests that being exposed to an in utero-shock has a persistent and constant effect on cognitive ability, at least for when the child is school-aged (which is as far as the sample goes).

Moving on to other heterogeneous responses, it is possible that the response to the rainfall shock is different for girls and boys. For example, Maccini and Yang (2009) find that a shock in the first year of the life is only significant for women in explaining adult self-reported health, and Dinkelman (2017) show that a rainfall shock has a significantly larger negative effect for men. The gender decomposition results are given in Table 4. The first column evaluates the gendered effects via an interaction term for the in utero shock and a dummy which takes the value one for girls. Columns (2) and (3) splits the sample into girls and boys and analyse the effects separately.

The first column supports no evidence of a heterogeneous effect on the two genders. The interaction term is insignificant and the coefficient is very close to zero. However, it does suggest that the magnitude of the shock goes up to 3.71% of SD when including the gender as a control. The split-sample analysis in the latter two columns implies that girls' cognitive score is somewhat more sensitive (-4.27% SD)

Table 4: Gender Decomposition

Sample: VARIABLES	(1) Total	(2) Girls	(3) Boys
Rainfall Shock	-0.0371***	-0.0427***	-0.0393***
	(0.0108)	(0.0103)	(0.0106)
Girl	0.0754***		
	(0.0103)		
Shock*Girl	-0.00602		
	(0.00983)		
Observations	392,246	191,447	200,799
R-squared	0.127	0.148	0.111

Note: District, wave and cohort fixed effects used throughout as well as mother's education as a control. The outcome variable is cognitive score.

than boys' cognitive score (-3.93% SD) to in utero droughts. The estimates are negative and significant at the 1% level. In reconciling these results, it is implied that the difference between girls and boys is not large enough to be relevant. One explanation of this is that male-bias is a cultural phenomenon more common in Asia where prior studies found a gender heterogeneous effect of shocks (e.g. India for Shah and Steinberg, 2017 and China for Qian, 2008).

An alternative explanation for the result is that the gender bias takes shape already in the survival rates of the children. It is possible that rainfall shocks affect the survival rate of girls particularly hard. They might, for example, be more likely to be neglected when income is scarce. This difference in survival rates would therefore not be captured by analysing the cognitive scores of (surviving) girls and boys. I hence construct a district-specific gender-ratio and regress this gender-ratio on the shock. I also construct a variable counting the cohort size in each district depending on whether there was an in utero-shock or not. I then regress this on the rainfall shock-variable. These two results are summarised in Table 5 below.

The estimation provides no support for a difference in the gender-ratio originating from a shock. In column 2 on the other hand, the estimation results suggest that the size of the cohort is significantly reduced if a shock occurred. One can thus speculate that that a drought has an impact on infant or early-life mortality, but that this mortality rate is not gendered. This in a way makes sense even if there is a male bias, if the mortality is induced by changes in fetal survival chances, as it is likely the gender of the child is unknown prior to the birth. However, since there are only 140 observations at the district-level the result in both columns can be considered

Table 5: District-Level Outcomes

VARIABLES	(1) Gender Ratio	(2) No. Children
Rainfall Shock	-0.0491	-1,709***
	(0.0650)	(615.4)
Level of Observation	District	District
Observations	140	140
R-squared	0.493	0.270

econometrically weak. Thus, these results indicate more information about the intermediate time period between the birth year and the test scores will provide better insight into the mechanisms for the long-run effect of in utero shocks.

5 Robustness Checks

The results in the preceding section relies on two main identifying assumptions. The first one assumes that shocks are indeed random and create otherwise comparable treatment and control groups. The second assumptions relies on a rainfall shock having a tangible effect on in utero conditions for the child. It is difficult to test either of these conditions but some explorations can be made.

First, I can perform a placebo regression on the time-dimension and see if the in utero timing is accurately capturing the link to children's cognitive development. To do this I take the in utero year and move the shock to $\pm$ 2 or 3 years. This is to account for the fact that some children will be in utero in year "-1" which therefore is not a valid placebo year. I thus create a dummy which takes the value one for two or three years after or before the in utero shock if the district experienced a drought that year. This gives four different specifications. The placebo regressions are summarised in Table 6.

As can be seen, three out of four of the coefficients of the shock are statistically insignificant. The reason for the significant result in column 3 is unclear, since one would expect that the direction of the shock is at least comparable to the shock effect in column 4 which is one year prior. It is also possible that it is significant out of chance, in particularly as the effect is so close to zero. I therefore consider this a spurious result and that the placebo regressions overall show that it is in fact the shock in utero which matters for the cognitive score.

Furthermore, one way to check how accurate rainfall is as a proxy for (rural) income is to regress split sample analyses between urban and non-urban districts. The results are presented in Table 7. As can be seen, the effect on the cognitive

Table 6: Placebo Regressions

VARIABLES	(1)	(2)	(3)	(4)
Shock (birth year +2)	-0.0161			
	(0.0103)			
Shock (birth year +3)		0.00124		
		(0.00916)		
Shock (birth year -2)			-0.0181*	
			(0.00934)	
Shock (birth year -3)				0.00294
				(0.00867)
Constant	0.0359	0.0365	0.0479	0.0357
	(0.0283)	(0.0281)	(0.0294)	(0.0286)
Observations	392,246	392,246	392,246	392,246
R-squared	0.126	0.125	0.126	0.125

Note: The outcome variable in all columns is cognitive score. A constant and control is used in all specifications as well as district, wave and cohort fixed effects.

score is statistically significant at the 1% level for the non-urban sample, while it is insignificant to the urban sample. This lends support to the assumption that a drought only has a direct effect on rural consumption and income.

Table 7: Rainfall-Rural Connection

Sample:	**Urban**	**Non-Urban**
Shock	-0.0496	-0.0388***
	(0.0265)	(0.0105)
Constant	0.467***	-0.00332
	(0.0474)	(0.0288)
Observations	12,646	439,915

Note: The outcome variable in both columns is cognitive score.

Next, I move on to analysis of the robustness of the main results. To see how sensitive the effect of the in utero rainfall shock is to the specification of the shock, I define it more narrowly as a year where the district experienced rainfall below its 10^{th} percentile. This is therefore equivalent to only defining more serious droughts as a shock. I separately regress cognitive score and the subject-specific test scores on the more narrowly defined shock variable. The results are given in the four columns of Table 8. The causal coefficient of the effect of a rainfall shock on the cognitive

Table 8: 10th percentile Rainfall Shock

VARIABLES	(1) Cognitive Score	(2) Maths	(3) English	(4) Swahili
Rainfall Shock (10p)	-0.0444***	-0.0418***	-0.0385***	-0.0436***
	(0.0126)	(0.0116)	(0.0117)	(0.0124)
Observations	392,246	402,852	403,619	399,340
R-squared	0.126	0.119	0.111	0.104

Note: District, wave and cohort fixed effects as well as control and constant used in all columns. The outcome variable is cognitive score.

score is -4.44 % of the standard deviation. This is almost twice the size of the original specification. This higher value is intuitive; since we are now only measuring the severest of droughts, the conditions leading to lower cognitive ability are stronger. It is therefore consistent if these children consequently have an even lower cognitive score.

As the origin and accuracy of the district categorisation variable used to define the urban and rural sample is unknown, I further investigate how sensitive the results are to the way the sample is divided according to the rural-urban criteria. I first define a rural district according to the night light density data since it is a good proxy for urbanisation (Mellander et al., 2015; Storeygard et al., 2014). A rural district is defined as a district where the mean night light density is less than 2.64%, which is the 75^{th} percentile. This means the top 25% of the most night light-dense districts are dropped. To visualise this definition of the sample, see Figure 4. The left figure shows that Kenya has some spots of brighter, more urban districts. In the right figure, these are greyed out in the map and excluded from the estimation. The estimation results are given in column (1) in Table 9.

Additionally, I make a more narrow definition of the rural sample by again utilising the categorisation variable. Recall, that I originally dropped districts classified as urban and kept districts classified as "normal", "semi-arid" and "arid". I now also drop the districts classified as "normal". I then estimate my regression equation as previously and summarise the result in column (2) in Table 9.

As expected, the sample size is substantially reduced, and the magnitude of the rainfall shock on the cognitive score increases. In both columns, the significance level and magnitude are increased compared to the original specifications. When the estimation sample is defined by the survey categorisation, the rainfall shock lowered the cognitive score by 3.46% SD – with the night light defined sample the effect is -4.9% SD. It should be noted that the sample based on categorisation-variable is far smaller than for the night light-based sample. Hence, it is therefore likely that the

Figure 4: Urban districts defined by night light density

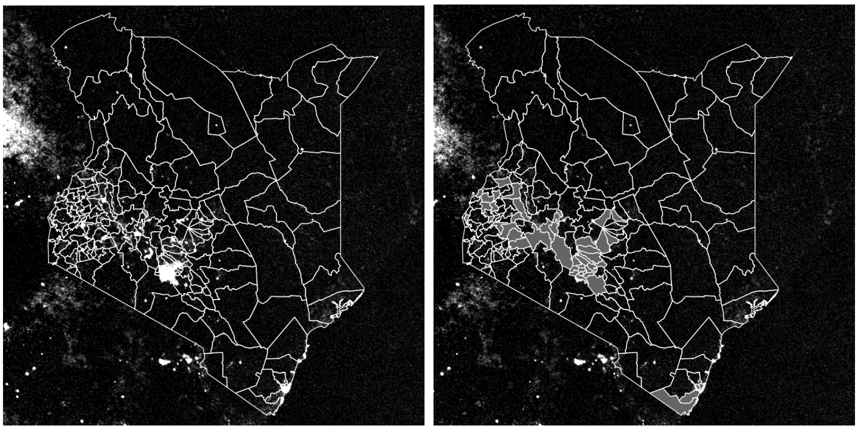

(a) All districts (b) Urban districts (grey) excluded

Table 9: Different rural-defined sample

VARIABLES	(1)	(2)
Rainfall Shock	-0.0346***	-0.0490***
	(0.0112)	(0.0155)
Number of districts	103	60
Observations	300,572	172,036
R-squared	0.124	0.130

Note: The outcome variable is the cognitive score. Column (1) is cut to districts where night light $<$ 2.64 %. Column (2) drops districts classified as 'urban' and 'normal' districts. District, wave and cohort fixed effects are used in both columns.

districts remaining in this more strict definition of a rural district are even more arid. Thus – as expected – a stricter definition of a rural district leads to a stronger effect of a drought. This further legitimises the viewpoint that the connection between the weather shock and cognitive ability is via rural income and consumption.

6 Discussion

6.1 Evaluation

To what extent can the results analysed in this paper be considered causal? Since the shock itself is very plausibly exogenous, I can identify a causal effect. The causal effect identified is the effect of having a shock in one's district which is slightly different from observing a direct income shock to the household. This is, however, likely convergent to a household income shock since households are overwhelmingly rural and highly dependent on rainfall. If anything, the magnitude of the drought effect on the cognitive score is underestimated.

Thus, a complementary estimation strategy could include more localised information on the rainfall shock. Furthermore, an extension to this paper could also more formally look at the mechanisms for how a rainfall shock affects schooling. On one hand, this could be done by incorporating data directly observing rural crop production or income in relation to the shock. On the other hand, to further establish how an income shock affects the cognitive score, data showing if there is any compensatory or intensifying behaviour on for example the parents' side would further shed light on how droughts affect children's cognitive ability.

Another potential source of imprecision is that the in utero shock was defined as the child's birth year i.e. the year the child is aged zero. However this measure is somewhat imprecise as the child's birth date determines how well the in utero period overlaps with "age zero". Because of this imprecision, it is possible that the true causal effect is underestimated. At loss of actual date of birth, the quarter of birth which Kumar et al. (2016) use would also improve the accuracy of identifying the in utero period.

Another factor which might underestimate the effect of the in utero drought is if a drought has a significant effect on infant mortality (as found by e.g. Kudamatsu et al., 2012). If this has a 'culling of the weak' effect, the weakest children will not enter the sample at school age. If these children would be expected to also have weaker cognitive scores, the children that enter the sample are already the better off or stronger children. It is then likely that the observed difference in cognitive score between children who did and did not experience an in utero shock will be underestimated. Incorporating data on infant mortality rates would shed further light on this potential source of bias.

Finally, Caruso (2017) has argued that selective migration can confound the effect of a weather shock on later outcomes. If the district in utero is different than the district of residence noted in the survey it will weaken the connection between the shock and the schooling outcome. Caruso (2017) argues that in his case in Latin

America 7% of individuals would be misclassified (being exposed to a weather shock or not) if their birth location was based on current rather than birth location. However, Baez et al. (2017) and Shah and Steinberg (2017) both find that their results are robust to selective migration. It is therefore difficult to assess the potential bias affecting my results based on prior literature, but intuitively mixing the control and treatment groups should underestimate the results of anything.

I expect that the external validity in terms of Kenya is high since the sample involves a random sample from across the country. Perhaps a more interesting question is to what extent these results are applicable to other countries. It is indeed applicable in the sense that the main results conform with economic theory, intuition as well as conforming with the effects found by studies using similar estimation strategies. This paper does however depart from the result of other country-studies in terms of the gender decomposition, so likely these specific results can be considered region-specific.

6.2 Implications

This paper shows that there is a persistent, negative effect on children's cognitive ability if they have been subjected to a negative rainfall shock. On the household level, this further emphasises the vulnerability of households to the natural elements. It also shows that the effect of a drought should not only be measured in terms of the direct effect on consumption, instability and income, but that the intergenerational effect in terms of to-be-born children should also be included when estimating the welfare loss.

To address these costs to affected families, policy should aim at both preventing and containing the costs of rainfall shocks. The former serves to highlight the importance for institutions to exist that provide functioning insurance against drought and subsequent crop failure. A study by Karlan et al. (2014) in small farms in Ghana found that the largest constraint to farm investment is uninsured risk. This suggests that mitigating risk would have positive effects on the production as well as the productivity of farming.

To mitigate the effects of an income shock resulting from drought it must be countered by some income (or consumption) support which prevent mothers from being exposed to health risks that are passed on to their in utero children. One such way is for the government to develop a functioning social support system. However, such a welfare system involves large costs and institutions that are slow to respond. A less comprehensive policy is to facilitate intra-economic trade if shocks are localised which would weaken the dependence on local rural production. This means a shock in one district can be offset by a good crop yield in a different district.

A connection can also be made to wider education policy. The results in terms of an educated labour force can be sub-optimal because of these underlying, long-run causes of lower cognitive ability. This can be the case even if other measures to improve education are taken, such as high enrollment or cash transfers. A more holistic view when investing into education is therefore necessary.

Finally, as extreme weather events and natural disasters are expected to be more common as the effects of climate change grow, this study provides some insight into possible costs of climate change in the future. As Dell et al. (2014) write, more frequent weather shocks like drought will lead to either adaptation or intensification. Adaptation involves providing the institutions and infrastructure necessary to offset the socio-economic costs to droughts. Intensification means that more frequent shocks would overburden the economic system and have an even greater net effect than the sum of individual shocks.

Adapting to more frequent droughts can involve everything from providing better irrigation, less rain-dependent crops or insurance against crop failure. This study thus provides further evidence for the costs of the market failure insurance markets in developing countries as this can have large bearings on the costs of failed rural production.

7 Conclusion

This paper finds that there are far-reaching effects of a negative in utero rainfall shock on children's cognitive ability. I find that the effect of a negative rainfall shock on cognitive ability is 2.60 percent of the standard deviation. The age and gender decomposition effect found no significant heterogeneous responses to a drought. I take this to imply that a shock has a persistent effect which means children carry the effects of an in utero shock at least through to their schooling years. The results are robust to several specifications.

The benefit of incorporating direct rainfall data means the shock can be reliably quantified, and above all provides the exogenous variation important to establish a causal effect. Owing to this causal interpretation, the results further shed light on the importance of aligning educational policy with the intrinsic qualities and risks of rural residence. Furthermore, the findings also imply that to tackle the costs associated with more frequent adverse weather in future generations, a long-run perspective today can better allow us to meet these challenges.

References

Alderman, H., Hoddinott, J., and Kinsey, B. (2006). Long-term consequences of early childhood malnutrition. *Oxford Economic Papers*, 58(3):450–47.

Baez, J. E., Lucchetti, L., Genoni, M. E., and Salazar, M. (2017). Gone with the storm: Rainfall shocks and household wellbeing in guatemala. *Journal of Development Economics*, 53(8):1253–127.

Barker, D. J. P. (1998). Mothers, Babies and Health in Later Life. *Edinburgh: Livingstone*, pages 13–42.

Bietenbeck, J., Ericsson, S., and Wamalwa, F. M. (2017). Pre-School Attendance, School Progression and Cognitive Skills in East Africa. *Cape Town: SALDRU, UCT. (SALDRU Working Paper Number 222*.

Burke, M., Gong, E., and Jones, K. (2015). Income shocks and HIV in Africa. *Economic Journal*, 125(585):1157–1189.

Caruso, G. D. (2017). The legacy of natural disasters: The intergenerational impact of 100 years of disasters in Latin America. *Journal of Development Economics*, 127(March):209–233.

Currie, J. (2009). Healthy, Wealthy, and Wise: Socioeconomic Status, Poor Health in Childhood, and Human Capital Development. *Journal of Economic Literature*, 47(1):87–122.

Currie, J. and Moretti, E. (2003). Mother's Education and the Intergenerational Transmission of Human Capital: Evidence from College Openings. *The Quarterly Journal of Economics*, 118(4):1495–1532.

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

Dinkelman, T. (2017). Long-run Health Repercussions of Drought Shocks: Evidence from South African Homelands. *Economic Journal*, 127(604).

FAO (2012). Fao (food and agricultural organization of the un) statistical yearbook 2012. *http://www.fao.org/docrep/015/i2490e/i2490e00.htm*.

Funk, C. C., Peterson, P. J., Landsfeld, M. F., Pedreros, D. H., Verdin, J. P., Rowland, J. D., Romero, B. E., Husak, G. J., Michaelsen, J. C., and Verdin, A. P. (2014). A Quasi-Global Precipitation Time Series for Drought Monitoring. *U.S. Geological Survey Data Series*, 832:4.

Glewwe, P. and King, E. (2006). The impact of early childhood nutritional status on cognitive development: does the timing of malnutrition matter? Number 4 in Handbook of Development Economics, chapter 56. Elsevier.

Hallegatte, S., Vogt-Schilb, A., Bangalore, M., and Rozenberg, J. (2016). Unbreakable: Building the Resilience of the Poor in the Face of Natural Disasters.

Hanushek, E. A. and Woessmann, L. (2012). Do better schools lead to more growth? Cognitive skills, economic outcomes, and causation. *Journal of Economic Growth*.

IPCC (2007). *Climate change 2007: The Physical Science Basis : contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, volume 2010.

Karlan, D., Osei, R., Osei-Akoto, I., and Udry, C. (2014). Agricultural decisions after relaxing credit and risk constraints *. *The Quarterly Journal of Economics*, 129(2):597–652.

Kudamatsu, M., Persson, T., and Strömberg, D. (2012). Weather and Infant Mortality in Africa. *CEPR Discussion Paper No. DP9222*.

Kumar, S., Molitor, R., and Vollmer, S. (2016). Drought and Early Child Health in Rural India. *Population and Development Review*, 42(1):53–68.

Maccini, S. and Yang, D. (2009). Under the Weather : Health , Schooling , and Economic of Early-Life Rainfall Consequences. *The American Economic Review*, 99(3):1006–1026.

Mellander, C., Lobo, J., Stolarick, K., and Matheson, Z. (2015). Night-time light data: A good proxy measure for economic activity? *PLoS ONE*, 10(10).

Minnesota Population Center (2015). Integrated Public Use Microdata Series, International: Version 6.4.

NOAA (2014). National Oceanic and Atmospheric Administration. Version 4 DMSP-OLS Nighttime Lights Time Series.

Qian, N. (2008). Missing Women and the Price of Tea in China: The Effect of Sex-Specific Earnings on Sex Imbalance. *Quarterly Journal of Economics*, 123(3):1251–1285.

Randell, H. and Gray, C. (2016). Climate variability and educational attainment: Evidence from rural Ethiopia. *Global Environmental Change*, 41:111–123.

Shah, M. and Steinberg, B. M. (2017). Drought of opportunities: Contemporaneous and long-term impacts of rainfall shocks on human capital. *Journal of Political Economy*, 125(2):527–561.

Storeygard, A., Henderson, J. V., and Deichmann, U. (2014). 50 years of urbanization in Africa : examining the role of climate change. (June):1–44.

Uwezo East Africa at Twaweza (2014). Are our children learning? Literacy and numeracy across East Africa 2013.

Appendix

Table 10: Alternative Outcomes

VARIABLES	(1) Probit Enrolment	(2) OLS Enrolment	(3) OLS Maths	(4) OLS English	(5) OLS Swahili	(6) OLS Basic Math	(7) OLS Basic English
Rainfall Shock	-0.121***	-0.00480**	-0.0438***	-0.0343***	-0.0364***	-0.0147***	-0.0157***
	(0.0190)	(0.00188)	(0.00835)	(0.00896)	(0.00955)	(0.00338)	(0.00418)
Observations	428,857	428,857	402,852	403,619	399,340	428,857	428,857
R-squared	-	0.153	0.120	0.111	0.104	0.220	0.208

Note: District, Wave and Cohort Fixed Effects used throughout as well as a constant and the mother's education level as a control.

Table 11: Sample Restriction by Age

SAMPLE:	8	9	10	11	12	13	14	15	16
Rainfall Shock	-0.00204	0.00355	0.0659**	0.00158	0.000523	0.0191	-0.00647	-0.0308	0.0326
	(0.0209)	(0.0193)	(0.0315)	(0.0220)	(0.0227)	(0.0314)	(0.0227)	(0.0387)	(0.0379)
Observations	51,110	43,099	56,554	35,821	53,060	42,449	44,047	33,799	32,307
R-squared	0.152	0.163	0.143	0.140	0.133	0.122	0.126	0.136	0.160

Note: The dependent variable in both columns is cognitive skill. District, wave and cohort fixed effects used throughout as well as constant and control.

Table 12: Age Decomposition

VARIABLES	(1)
Shock*Dummy(8)	-0.0892***
	(0.0239)
Shock*Dummy(9)	-0.0705***
	(0.0182)
Shock*Dummy(10)	-0.0265
	(0.0255)
Shock*Dummy(11)	-0.00424
	(0.0152)
Shock*Dummy(12)	-0.00954
	(0.0172)
Shock*Dummy(13)	-0.0150
	(0.0235)
Shock*Dummy(14)	-0.0300
	(0.0193)
Shock*Dummy(15)	-0.0800***
	(0.0269)
Shock*Dummy(16)	-0.0373
	(0.0304)
District FE	YES
Wave FE	YES
Cohort FE	YES
Observations	392,246
R-squared	0.126

Note: The dependent variable is cognitive score.