NHH

Norwegian School of Economics
Bergen, Spring 2019

After the rain - Exploring the link between rainfall shocks and early childhood development

Akshay Moorthy
Supervisor: Vincent Somville

Master thesis, Economics and Business Administration
Major: Economics

NORWEGIAN SCHOOL OF ECONOMICS

This thesis was written as a part of the Master of Science in Economics and Business Administration at NHH. Please note that neither the institution nor the examiners are responsible – through the approval of this thesis – for the theories and methods used, or results and conclusions drawn in this work.

Contents

Abstract

This thesis studies the link between exogenous and unforeseen variations in household income and the development level of very young children in Uganda using rainfall shocks as an instrument for income variation. The analysis links household data and child development measures from 2336 households from 9 Ugandan districts with 28 years of rainfall data to look at the effects of rainfall shocks in-utero and in early childhood upon measures of the motor, early literacy, early numeracy and social-emotional development of children between the ages of 3 and 5. A simple model of childhood development is developed to illustrate possible causal channels and challenges associated with studying this relationship. Reduced form OLS estimates indicate the existence of links between early life rainfall shocks and a child's non-cognitive development level. There is no evidence of a significant link between rainfall shocks in-utero and our measures of child development. Heterogeneity analysis reveals differential links along gender, education and asset ownership dimensions. Decomposition of yearly rainfall deviation into binary and seasonal shocks indicates different effects of shocks in the context of agriculture in Uganda. Rainfall shocks in the Ugandan harvest season may have the opposite effect of rainfall shocks in the planting season. The results are highly sensitive to the choice of inference calculation. Beyond education and health programs, insuring households against income risks from climate events as well as mitigating the source of these risks could play an important role in meeting early child development goals.

Acknowledgments

This thesis would not have been possible without the support and guidance of Vincent Somville. I am grateful to Vincent Somville, Kjetil Bjorvatn and Selim Gulesci for providing the idea that led to this paper. The opportunity to work with and learn from Denise Ferris, Lore Vandewalle, Arne Nasgowitz, Nansubuga Robinah, Balakaita Joseph, Mulondo Charles and many others associated with the larger project has been invaluable.

1 Income shocks, rainfall and early childhood development

Around 219 million children worldwide are at the risk of not being able to reach their maximum development level which might manifest in a 20% loss in their adult income (Black et al. 2017). Meeting the sustainable development goal (SDG) of "ensuring that all girls and boys have access to quality early childhood development, care and pre-primary education so that they are ready for primary education" by 2030 (included in the SDGs in 2015) will require focused investment and effort by many stakeholders. Policymakers and investors ought to be ready and willing to prioritise early childhood development as research shows that, even in the USA, the rate of return on investments in early childhood development (henceforth ECD) could be as high as 7-10% (Heckman et al. 2010). Worse, not investing in ECD could lead to the inter-generational persistence of economic and educational backwardness (Grantham-McGregor et al. 2007). Therefore, understanding both the need for and the potential impact of programs targeted towards early childhood development is important.

The effect of a well designed program might be nullified by unpredictable events such as wars or disease outbreaks or even by predictable events (albeit with unpredictable effects) such as the looming and very real spectre of climate change. In high income countries, infrastructure (physical, financial and social) helps dampen the effect of these uncertainties (Kochar 1995). But the people of low- and middle-income countries (LMICs) are often at the mercy of the "five horsemen of the apocalypse" - state failure, climate change, famine, disease and migration (Morris 2010). The links between unforeseen risks and early childhood development in LMICs is therefore an important research problem because "a child's brain is not born, it is built" (Britto 2017). Nutrition, shelter from conflict and disease, and the love of a caring adult are some of the key ingredients that are crucial in the early stages of a child's life (Walker et al. 2011). Being deprived of these resources leads to an enormous waste of potential as children in such situations are restrained from reaching their maximum potential. Children in LMICs are especially susceptible to this as there are many kinds of risks that children and families in such countries regularly face.

There are two hurdles that have to be overcome in order to study the relationship between income and ECD properly. First, collecting the exhaustive household and child development data set that is necessary to describe an accurate structural relationship comes with significant practical challenges and high costs, especially in a low-income setting. Second, there is the problem of the endogeneity of income on child development - beyond directly affecting childhood development, income can affect other things that may modify the level of childhood development. The importance of side stepping this endogeneity problem accentuates the need to conduct field experiments or to look for naturally occurring experiments (Glewwe and Miguel 2007). Naturally occurring experiments have significantly lower costs than field studies or randomised controlled trials and provide a way around the endogeneity problem. Many researchers have used the occurrence of rainfall shocks as a proxy for income in LMICs to study various economic and social phenomena. E. Miguel (2005) and Paxson (1992) are good examples of this approach, using rainfall shocks to study income-related effects on violence and savings respectively.

The identifying assumption used in this thesis is that unforeseen rainfall shocks will cause variations in household income that can be allocated towards human capital - in our case, investments that will aid the development of a young child in the household. Rain that falls during an inopportune time, such as right after sowing or during the middle of a harvest period may lead to significant losses in agricultural output and farmer welfare, especially where technological adoption/access is low and risk mitigation avenues are minimal or non-existent. An unexpected loss in income may lead to lower investments in human capital which could affect the health and educational outcomes of children (Jensen 2000). Conversely, unforeseen additional rainfall at the right time could boost household income unexpectedly. This is a situation that will vary across (and possibly, even within) countries as local environmental conditions, farming practices and individual preferences play important roles in this complex relationship. Consequently, studying this relationship is demanding in terms of data requirements and sample selection; many of these factors are difficult to observe within a sample and it may be difficult to collect representative members of each sub group in the observed sample.

Despite these challenges, several researchers have studied the effect of rainfall shocks on an array of childhood and adult development indicators in recent years. Some of the outcomes examined in this context are health outcomes (e.g. height and weight), educational attainment (years of schooling, entry delay into schooling), cognitive skills (test scores) and non-cognitive skills. Maccini and Yang (2009) look at the long-term effect of early-life rainfall shocks on adult outcomes in rural Indonesia. They find significant positive links between instrumented rainfall shocks in the year of birth of women and several (self-reported and measured) health and educational attainment indicators. The Indonesian context was further examined by Cornwell and Inder (2015) who also find some effects of early life shocks (especially during the time when the child is in-utero) on health outcomes through a nutrition based analytical approach. Thai and Falaris (2014) document the disadvantageous effects of negative income shocks on health and schooling outcomes when families are dependent on agriculture in rural Vietnam. Shah and Steinberg (2017) investigate the effect of early life rainfall (income) shocks on children's test scores in rural India and find that scores and years of schooling are higher when children and their families are affected by increases in income in early childhood. Leight, Glewwe, and Park (2015) identify a negative effect of adverse rainfall shocks on cognitive outcomes but no impact on non-cognitive outcomes (such as measures of social skills) in China's Gansu province. In Uganda, Björkman-Nyqvist (2013) finds gender differences in the response of human capital investment (measured by schooling) in the presence of negative income shocks caused by below average rainfall. Girls appear to play second fiddle to boys when it comes to human capital investment, as is the case in so many places in the world.

This thesis contributes to the literature studying the effects of income shocks on early childhood development in 5 ways: a) Using data collected with a relatively new ECD assessment instrument to evaluate short run effects, b) investigating the non-cognitive skill development of young children in a low income environment, c) heterogeneous responses to rainfall shocks along rainfall timing, household and child-specific dimensions, d) exploring the different effects of shocks occurring in different seasons (in the context of Ugandan agriculture), and e) exploring the impact of using different strategies to calculate inference statistics. Using a quasi-experimental setup to look at the effect of rainfall shocks faced by families on the cognitive and non-cognitive capabilities of very young Ugandan children, this thesis studies the reduced form relationship between early life/in-utero rainfall shocks and measures of early learning development using a freshly collected data set of measures of the development level of children aged between 3 and 5 years from 9 Ugandan districts. The results suggest that there may be a significant link between early life income shocks and the social-emotional (non-cognitive) development of young Ugandan children.

1.1 The Ugandan context

The Ugandan economy is highly dependent on agriculture which employs 73.6% of its labour force (Uganda Bureau of Statistics 2016). The Uganda National Panel Survey (2015/16), which was conducted on a representative sample of the Ugandan population, found that just 0.6% of Ugandan farmers were using irrigation as a part of their agricultural activity. This indicates that almost all Ugandan farmers and, consequently, a large number of Ugandan households are extremely susceptible to variations in rainfall.

Uganda is blessed with ample and fairly consistent rainfall at levels high enough to sustain two harvests every year and without the irregularities that characterise the rainfall patterns of many African countries (Gommes and Petrassi 1996). March-May and September/October to November/December are typically considered the “rainy seasons” (Mubiru et al. 2012; Jury 2018) which means that most sowing/planting happens early in these seasons. The remaining months of the year can be considered as “harvest” months.

Recent research on the links between rainfall variability and measures of income or consumption in Uganda has generally found a positive relationship between above average rainfall and income. Asiimwe and Mpuga (2007) find that above average rainfall in the first planting period (March-May) results in lower household income, but above average rainfall in harvest periods lead to higher incomes. Björkman-Nyqvist (2013) finds a positive relationship between above average (yearly) rainfall and both agricultural production as well as national income. These Ugandan studies and many other studies (such as Levine and Yang (2014), who supplement the findings in Maccini and Yang (2009), by identifying a positive correlation between rainfall shocks and agricultural output) generally find (and exploit) a positive relationship between rainfall variability and household income to construct reduced form models to study various economic outcomes. The theoretical model that we will construct in section 2 of this thesis does not require a directional assumption of income on our outcomes of interest. In section 5, we attempt to decompose the seasonal effect of a rainfall shock and additionally look at the different effects of surpluses and deficits of rainfall when compared to the long-term averages.

The data set and analytical techniques used in this thesis do not however account for the equilibrium effects of rainfall shocks. Given the primarily agrarian nature of the population being studied here, households are quite likely to be both producers and consumers (indeed, about half of the households in our sample derive some income from farming their own plots of land). They may either consume what they produce or trade from the market to smooth out surpluses/deficits. In the presence of differential income shocks brought about by an uneven distribution of rainfall shocks, households that do not receive the direct beneficial effects of a good shock to their income (induced by a production gain because of a favourable rainfall shock) may receive an indirect benefit due to a reduction in prices of some consumable goods because of the productivity enhancing effect of favourable rainfall shocks on other households.

For example, if many households that produce rice receive a favourable rainfall shock, the supply of rice increases and prices may drop. Other households who were not able to produce as much rice as they need for their consumption due to adverse rainfall shocks will now benefit slightly from being able to purchase some rice from the market at prices that are lower than otherwise. On the other hand, producers of rice now earn less per unit of rice than they did before. The equilibrium state thus depends on several factors such as the sensitivity of prices to supply fluctuations, the availability of alternate income-generating opportunities, the elasticity of wages in response to adverse conditions such as droughts (Jayachandran 2006) and the extent to which markets are integrated and individual households are able to trade their goods on these markets. In the absence of primary or secondary data to account for these factors, equilibrium effects remain unexplored in this thesis. It must be noted that these effects do seem to matter in the LMIC context - for example, Aker (2012) finds that food markets in Niger become more tightly integrated under drought conditions.

1.2 Early Childhood Development

Traditionally, early childhood development has been measured either at the individual level or at the population level. Individual measures are more useful when evaluating the impact of a program or intervention on treatment groups as they provide more detailed information on the development status of a particular child (D. C. McCoy et al. 2016). Measures such as the rate of stunting, malnourishment and years of schooling have also been used to look at the level of child development (Black et al. 2017). The inclusion of early childhood development in the UN SDGs in 2015 has spurred the development of tools to measure childhood development, several of which are described in Fernald et al. (2017). Selecting the correct tool given the needs of the research project (what is to be measured and in what context) and the resources available (technical expertise, costs, inter-cultural applicability) is very important as each tool/instrument has a set of advantages and disadvantages.

Measures of childhood development are constructed using the International Development and Early Learning Assessment (IDELA) survey instrument (described in greater detail in section 3). The IDELA allows for the examination of different causal channels as it provides measures of four distinct developmental domains - motor skills, early literacy, early numeracy and social-emotional skills. Notably, our outcomes of interest capture non-cognitive development levels (specifically, the social-emotion and motor skill components), marking a departure from the literature - typically, population level measurements or cognitive test score achievement are used as ECD indicators. Acquiring skills of this nature during early childhood has been found to have a positive effect on long-run outcomes such as labour income (Heckman, Stixrud, and Urzua 2006). Investigating the sensitivity of this dimension of early childhood development to unforeseen shocks in household income in a LMIC setting contributes towards filling a gap in the existing literature.

The data set used in this thesis allows the investigation of the short run impact of plausibly exogenous income variation on childhood development. Most of the studies mentioned above look at the impact of early childhood shocks on adult/teenage outcomes - i.e. medium and long run effects of early childhood income shocks. Such approaches have the clear advantage of being able to identify outcomes that get affected down the line. However, by their very nature the findings are subject to scepticism owing to the long time periods between causes and effects which may attenuate (or even misrepresent) the "treatment" effect. Our data set consists of first-hand information from households and an assessment of the development level of one child (aged 3 to 5) from each selected household. The trade-off is that while this data set may produce stronger evidence of short term effects, the question of whether effects in early childhood can be attenuated naturally (i.e. without any special programs or alternate treatments) over time cannot be explored.

The IDELA has been shown to be internally valid (Wolf et al. 2017; P. F. Halpin et al. 2019) and has been used in many different settings (Pisani, Borisova, and Dowd 2018; Children 2018). This analysis is potentially comparable across other countries and settings provided that the same (or similar - Pisani, Borisova, and Dowd (2018) shows that the IDELA is similar to other assessment tools commonly used in this domain) assessments are used. However, P. F. Halpin et al. (2019) find that the between-country comparative validity of the IDELA is limited. Some of the reasons cited in that study are differences in the groups studied (they are generally not nationally representative, but are selected as part of a program or initiative) and cross-cultural differences (which may impact the administration of the assessment) and different developmental trajectories in different countries. The IDELA is simple to administer, yet is complex in what it aims to accomplish, which may be a disadvantage when trying to compare ECD achievement across countries.

A disadvantage of using these measures is the lack of quantifiable interpretations of any effects found. The economic significance of IDELA scores has not been established yet - the tool is very new, and (to our knowledge) no correlations have been established between childhood development as measured by the IDELA and economic outcomes. The effect of indicators such as years of schooling or stunting measurements has been quantified in several different ways; for example, finding a significant effect of rainfall shocks on years of schooling can be quantified into a dollar loss in adult income based on earlier estimates from several studies as in Maccini and Yang (2009) or Thai and Falaris (2014). Since we do not have the ability to do this, we interpret the IDELA scores (and effects on these scores) primarily as an indicator of the level of child development that is comparable across subjects in this study.

In section 4, we define the empirical approach and discuss various considerations and potential problems with the specification. In concordance with recent literature, the analysis utilises different techniques of calculating inference statistics in an attempt to arrive at more reliable estimates in the presence of choices made during the research and sample design process. Our analysis finds that inferences vary quite dramatically depending upon the choice of clustering methods and control variable inclusion.

In section 5, we present the results from estimating various specifications across different measures of rainfall and outcome variables. We also attempt to decompose the effects of rainfall by using various measures of rainfall to try and understand the relative importance of the type (above or below average rainfall) and timing (planting and harvest seasons) along the lines of Asiimwe and Mpuga (2007). The analysis also looks at the question of whether certain types of households and children are affected in different ways by rainfall shocks by exploring the interaction of rainfall shocks with various control variables. We conclude the analysis in section 6 by briefly summarising the key findings and discussing the implications and limitations of this thesis.

2 Conceptual model

In order to understand the channels through which income shocks can affect child development, we develop a simple model along the lines of Glewwe and Miguel (2007). Consider a simple household consisting of two people - one adult and one child. The adult earns and allocates income Y_t in time period t. Income comes from multiple streams such as agriculture, labour, business ownership, return on investments, etc. Income can be spent on consuming goods and services (C_{adult}, C_{child}), leisure time (L_{adult}) or investment (I_t). The child has no say in the allocation process and is dependent on the preferences of the adult as well as the absolute level of income for its utility in any period. The household's budget constraint in any given time period is:

$$Y_t + (1 + r)I_{t-1} = C_{adult} + C_{child} + L_{adult} + I_t$$

The adult wants to maximise his/her utility U_{adult} from various sources such personal consumption (C_{adult}), various unobserved sources (X) and his/her perception of the child's utility ($g(U_{child}, J)$) and leisure (L_{adult}).g() is a function that captures the adult's perception of their child's utility, reflecting their beliefs on what is good for the child and how "happy" the child is. This can depend upon many things, plausibly including factors such as their education level, cultural perspectives and individual preferences which are captured by J, a vector of adult characteristics and preferences).

$$U_{adult} = f(C_{adult}, g(U_{child}, J), L_{adult}, X)$$

In utero, the effects on the child are largely driven by the mother's (represented in this case by the adult) health and nutrition intake (Walker et al. 2011). The development level of the child is increasing in health, nutrition, other factors (M) and leisure time (Blau and Grossberg 1990). The production function in-utero can be defined as:

$$Dev_{Utero} = m(Health_{adult}, Nutrition_{adult}, L_{adult}, M, \alpha)$$

where α is the "endowment", perhaps genetic or biological in nature. Health and Nutrition are assumed to be functions of adult consumption, so we can simplify and write the above as:

$$Dev_{Utero} = m(l(C_{adult}), L_{adult}, M, \alpha)$$

The child's development level in the current period T, which is the subject of our investigation, also depends on events that happen in early childhood (i.e. the past) as well as in-utero events. Post birth, the effects on the child are driven by many factors including consumption focused on the child, the family and home environment (functions of adult consumption, $l(C_{adult})$) and and other factors (N e.g. exposure to government/NGO programs (transfers, micro-nutrients, vaccinations, etc)). ψ is an endowment, similar to α above. The development level of the child (post-birth) is increasing in consumption, family environment (which can be defined as a function of non-child consumption and leisure allocation), health and external transfers.

$$Dev_t = h(C_{child}, l(C_{adult}), N, \psi, Dev_{utero}, \sum_{n=1}^{t-1} Dev_n)$$

The child's utility in a time period t depends upon the child's consumption and the child's development level:

$$U_{child,t} = i(C_{child}, Dev_t)$$

Using the above relations, the adult's utility function can be rewritten in terms of the child's development level as:

$$U_{adult} = f(C_{adult}, g(i(C_{child}, Dev_t), J), L_{adult}, X)$$

Inverting the above, the development level of the child can be written in terms of the utility of the adult as:

$$Dev_t = i_p(g_p(f_p(U_{adult}, C_{adult}, L_{adult}, X), J), C_{child})$$

The relationship above was derived using many simplifying assumptions about childhood development, household income and household preferences. It's complexity illustrates the difficulty of studying the structural problem. All of the terms in the budget constraint appear here, nested within multiple functions. Additionally, the vectors X and J may also be affected by income. Thus, variations in income will affect this equation in various ways as $C_{adult}, L_{adult}, C_{child}$ are highly correlated.

In the presence of income shocks - such as one induced by unforeseen variations in rainfall - allocation decisions will vary depending upon the preferences of each individual adult. Changes in the amount allocated to each of the components of the individual production functions will lead to changes in the level of early childhood development. However, the direction of the change is not easy to determine even in this simple model. For example, increasing the allocation in the post-birth period to C_{Child} will have a positive effect on development post birth. But reduced $C_{adult} + L + I$, which are also a part of the development production function could crowd out the effect of an increase in C_{child} as reductions in adult leisure time that are detrimental to early childhood development could nullify the positive impact of increasing the quantity of resources devoted to the child (e.g. better food, better bed, clothes, etc). Similarly, increasing the allocation to L during the in-utero period will increase the in-utero component of childhood development, but reduced resources available for allocation to nutrition and health could lead to a reduction in the level of childhood development.

An interesting interpretation of the above is that the utility of the child is maximised when the adult is able to make efficient allocations based upon perfect knowledge of the influences of home environments, leisure, health, nutrition etc. both in-utero and in every time period. In the absence of perfect knowledge, each function begins to diverge from a pareto optimal allocation and ends up leading to sub-optimal child development over time. The longer the time period, the strong is the divergence effect from optimum, so perhaps a discount multiplier is required on the accumulation term in the previous equation to make the function have diminishing returns to this form of information asymmetry.

The pseudo-accumulation of early development further compounds the difficulty of defining and measuring the structure and chain of causality that affects childhood development. Income shocks in the past may have effects far into the future. Fully defining the state of development of a child in order to estimate the exact impact of such an income shock requires complete knowledge of the quantity of every single resource and investment channel that the child may have ever crossed paths with. Thus, most (if not all) studies of this nature use a reduced form approach to explore the effect of income shocks on health/education/development outcomes.

In order to derive a reduced form that could help recover some information about the structural parameters (Blinder 1973), consider a simplified version of the model described above. The first simplifying assumption is that childhood development is a linear function of income and a vector of household characteristics.

$$Dev_t = \alpha_1 + \beta_1 Y_t + \gamma_1 \widehat{H}_t + u_1$$

Y_t is endogenous as it can affect multiple household characteristics and household characteristics can also affect the income in (and across) time periods. Adapting the model and reasoning used by Paxson (1992), we assume that variations in permanent income are smoothed over all time periods. Income is thus a linear function of this permanent income and additional income arising from exogenous, unpredictable shocks such as, rainfall shocks.

$$Y_t = \alpha_2 + \beta_2 Y_P + \delta \widehat{Shock}_t + u_2$$

Then the reduced form representation can be obtained by eliminating Y_t:

$$Dev_t = \alpha_3 + \beta_3 \widehat{Shock}_t + \beta_4 Y_P + \gamma_1 \widehat{H}_t + u_3$$

The coefficient β_3 does not indicate the direct effect of a rainfall shock. It is the combination of the shock's effect on income as well as the effect of income on development both with and without the shock. The advantage of using this approach is that we recover some information on the larger structure by sacrificing fidelity in establishing a chain of causality. Perturbations in the $Shock_t$ vector will lead to changes in temporary income which is unanticipated and thus could not be smoothed over all time periods. The mechanisms discussed earlier in this section mean that the direction of the effect of a shock on development in this equation is ambiguous and will depend strongly upon household preferences (which are not observed in this equation) and perhaps household characteristics (H_t). Thus, while this model can give us some information about the effect of a rainfall shock on development, it is quite likely that the u_3 term in the equation also picks up a lot of the variation. However, this model ought to yield a certain amount of explanatory power on the direction of the effect (given that temporary income can strongly influence household allocation outcomes as discussed above). It is this equation that forms the basis for the empirical model used in this thesis.

3 Data

3.1 Early Childhood Development

Measures of early childhood development come from a survey conducted as a part of a larger research project involving a randomised control trial to study the effect of providing improved access to childcare services and/or unconditional cash transfers to poor Ugandan families (targeted specifically at mothers). The final data set that is available for analysis consists of test scores from 2336 children between 3 and 5 years old from low income households across 9 Ugandan districts. The districts were not chosen randomly - choices were made to spread districts evenly across 3 Ugandan regions, prioritising accessibility and availability of field staff. Within each district, villages and households were first selected for a large census on the economic condition of these households. Based on this, a smaller sample of ~3000 households was constructed by selecting households which had both a female caregiver as well as a child between 3 and 5 years of age from which 2467 unique children were finally assessed. A few observations had to be removed from the raw sample set due to factors such as multiple assessment entries, enumerator malfeasance, problems in the field and administrative errors leading to a final, usable sample consisting of 2336 observations.

The survey uses the International Development and Early Learning Assessment (IDELA) instrument developed by Save The Children, an NGO focused on childhood development. The tool consists of a battery of questions and tests that aim to measure the level of competency or mastery that children possess across four domains - motor skills, early literacy, early numeracy and socio-emotional skills (Pisani, Borisova, and Dowd 2018). Table 1 provides an overview of some of the skills that are evaluated as a part of each domain.[1] This illustrates one of the major advantages of the IDELA tool - the multiple dimensions are measured separately which provides a granular assessment of early development. From an economic perspective, this will allow researchers to disentangle child development into multiple outcomes which could facilitate the identification of specific causal chains given a known source of variation.

[1] More information is available in Appendix C.

Table 1: The IDELA domains

Domain	Feature	
Gross and fine motor skills	Hopping on one foot Copying a shape	Drawing a human figure Folding a piece of paper
Emergent literacy and language	Print awareness Letter identification Phonemic awareness	Expressive vocabulary Emergent writing Listening
Emergent numeracy	Measurement and comparison Number identification One to one correspondence Simple problem solving	Classification and sorting Shape identification Simple arithmetic
Social - Emotional development	Peer relationships Empathy Conflict resolutions	Emotional awareness and regulation Self-awareness

The IDELA was designed specifically for use in LMICs with an emphasis on reducing the costs of administration and on removing the necessity for specialised enumerators and/or trainers - problems that afflict other tools in this area (Pisani, Borisova, and Dowd 2018). The IDELA has been tried and tested in many developing countries including Uganda and has been found to be internally valid and consistent (Wolf et al. 2017) and externally (in comparison to other commonly used tools, Pisani, Borisova, and Dowd (2018)). The tool is freely available for use and does not require expert knowledge to administer, factors which made both training and evaluation significantly cheaper and easier. The IDELA is also available in the major Ugandan languages which provides ease of communication and improves the chances that the assessments are accurate. The risk of biased assessments in favor of children of higher literacy (P. F. Halpin et al. 2019) who are often better placed to understand and answer questions that are not accurately translated to the local context is mitigated to an extent.

The survey was facilitated by BRAC Uganda who trained a team of enumerators to conduct the IDELA survey. After the training, the enumerators conducted a small pilot survey in order to assess the efficacy of the training and ability to use the survey equipment provided to them. After ironing out the issues and questions identified during the pilot phase, the team conducted the survey on the entire sample identified for the purpose of the project. The results we use in this thesis were captured by this team using mobile devices to record the results of each individual test.[2] Table 2 shows the gender and age distribution of cases across Uganda.

[2] Some assessments were initially conducted using paper surveys. However, these were later converted into assessments captured on the mobile phone to eliminate instances of enumerator malfeasance. Most of the troublesome cases have been removed from the final sample used in the analysis.

Table 2: Gender composition in each district

District	Boys	Girls	Total
Central			
Masaka	127	126	253
Mukono	111	130	241
Mityana	121	132	253
Eastern			
Iganga	155	165	320
Mbale	149	150	299
Jinja	136	148	284
Western			
Kabarole	141	125	266
Kasese	94	84	178
Kyenjojo	120	122	242
Total			
All Districts	1154	1182	2336

The survey consists of a battery of questions, most of which are coded as SUCCESS, FAILURE, or SKIP. The results of the survey were processed using standard instructions provided by the developers of the IDELA tool. Importantly, they advise the conversion of all SKIP results to FAILURE. Table 3 shows the summary of the test scores on a scale of 0-1. The scores can be interpreted as percentage values indicating, for example, that the average child in our sample scores about 36% on the overall test. Scores increase as children grow older, which is a desirable feature of an effective childhood development measurement tool (Fernald et al. 2017). Interestingly, girls score higher than boys in every category across all age groups.

Table 3: IDELA Summary - by gender and age (raw scores)

		Age = 3 (n: 836)		Age = 4 (n: 868)		Age = 5 (n: 632)		All ages (n: 2336)	
Category	Gender	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Motor skills	Male	0.27	0.20	0.50	0.26	0.71	0.21	0.48	0.28
	Female	0.33	0.23	0.59	0.25	0.77	0.18	0.54	0.28
	Total	0.30	0.22	0.55	0.26	0.73	0.20	0.51	0.29
Early Literacy	Male	0.17	0.12	0.27	0.18	0.44	0.22	0.28	0.21
	Female	0.19	0.13	0.33	0.19	0.51	0.24	0.33	0.22
	Total	0.18	0.13	0.31	0.19	0.47	0.23	0.30	0.22
Early Numeracy	Male	0.22	0.12	0.33	0.16	0.46	0.19	0.33	0.18
	Female	0.23	0.13	0.37	0.17	0.51	0.19	0.35	0.19
	Total	0.22	0.13	0.34	0.16	0.48	0.19	0.34	0.19
Socio-emotional	Male	0.20	0.12	0.29	0.16	0.37	0.17	0.28	0.17
	Female	0.21	0.13	0.31	0.16	0.39	0.17	0.30	0.17
	Total	0.20	0.13	0.30	0.16	0.38	0.17	0.29	0.17
Total	Male	0.21	0.12	0.35	0.16	0.50	0.16	0.34	0.18
	Female	0.24	0.13	0.40	0.16	0.54	0.17	0.38	0.19
	Total	0.23	0.12	0.37	0.16	0.52	0.17	0.36	0.19

Note:
n(male) = 1154, n(female) = 1182

The IDELA score measures reported here are compared to the scores reported from other similar surveys in table 4 below (IDELA network 2019). However, P. F. Halpin et al. (2019) find that the IDELA in its current form is not well suited to between-country comparisons owing to idiosyncratic differences between countries (as discussed in the first section of this thesis). Based on the selection of literature available for access online (IDELA network 2019), most studies that have used the IDELA tool use it to look at program impact within-country, rather than using it as a baseline for comparison across countries. In that sense, even though the IDELA tool appears to be reliable and robust, it's true application could be to look at differences over time relative to past levels rather than the absolute levels across countries.

The IDELA tool is also relatively new and has only been used in the field since 2014/2015. Therefore, there are no studies on the medium- or long-term outcomes using variations in measured IDELA characteristics in early childhood. This makes evaluating the impact of a score or the effect of a change in the score because of an intervention difficult as there are not established links between an IDELA measure and an economic outcome.

Table 4: IDELA scores across countries

Category	Data set	Age		
		3	4	5
Socio-emotional	Uganda (2018)*	0.2	0.30	0.38
	Uganda	-	0.31	0.46
	Burundi	0.16	0.17	0.27
	India	-	0.38	0.49
	Ethiopia	-	0.27	0.30
Motor	Uganda (2018)*	0.3	0.55	0.73
	Uganda (2016)	-	0.50	0.75
	Burundi (2018)	0.21	0.27	0.39
	India (2018)	-	0.64	0.71
	Ethiopia (2018)	-	0.36	0.42

Note:
Taken from individual data sets on the IDELA website. Missing values as samples did not report values for those ages. * = The data set used in this thesis.

Figure 1 shows the kernel density plots of the IDELA socio-emotional score achievement for each age group (3, 4 and 5 years). As visible in table 3, there is a large and significant difference in scores across age groups. As mentioned before, this is a desirable (and expected) feature of an effective childhood development measurement tool/data set (Fernald et al. 2017). The rate of neurological development in children is very high between the ages of 3 and 5 (Black et al. 2017). Variations in the scores in an individual assessment capture this feature of early childhood development. The IDELA scores are also calculated as percentages and not absolute numbers. In order to account for the large differences across age groups and to enable easier interpretation of the scores (and coefficients of interest), the scores were standardised with respect to the mean of each age group (3/4/5) , the IDELA scores used in regression analysis were normalised to have a mean of 0 and a standard deviation of 1.

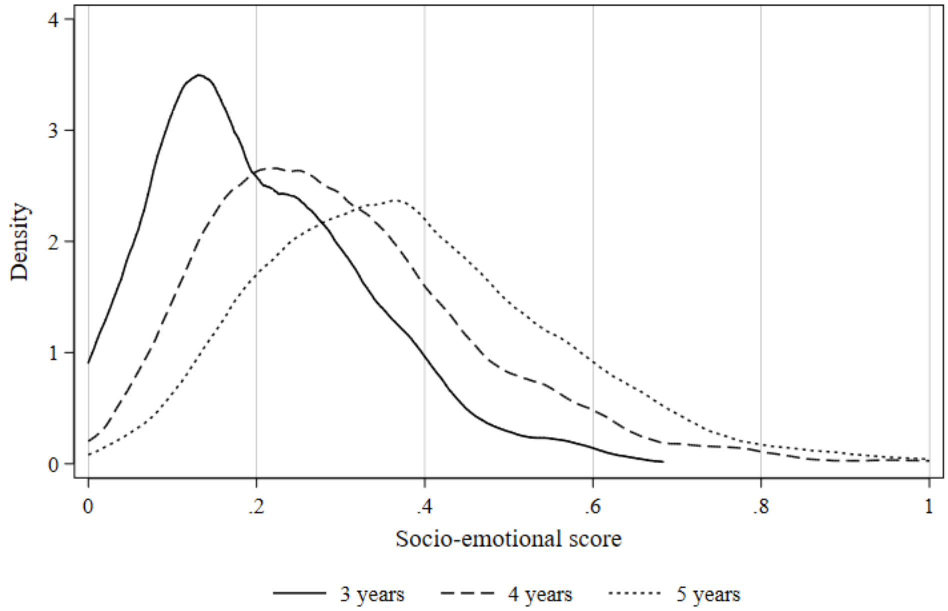

Figure 1: IDELA scores (percentages) across age groups

3.2 Household data

Household information was collected mostly from the baseline survey conducted as a part of the RCT described in the previous section. The survey was designed to collect information in order to study the relationship between child-care services and maternal labour force participation. The survey was conducted in the months of November and December 2018 across approximately 3000 households distributed across 9 districts in Uganda. The objective of the survey was to collect household information such as demographics, educational attainment, asset ownership and measures of preferences and risk aversion amongst others prior to the implementation of a randomised controlled trial. The summary statistics below are a subset of the household survey data for households which consented to be tested with the IDELA tool and for which the data was successfully recorded (i.e, households whose children are a part of the IDELA data set).

Variable	Mean	Median	SD	N
Child age	3.93	4	0.78	2336
Respondent age	35.4	33	10.9	2289
Siblings in HH	2.42	2	1.79	2336
HH size	6.26	6	2.93	2289

Variable	Mean	Median	SD	N
HH Income (USD)	$ 514	$ 212	$ 1200	1125

The household income data point is problematic in several ways. It is calculated as the sum of income earned by a household from several streams in the past 12 months before the survey - businesses, land ownership (rent or agricultural income), livestock, labour income and other sources. A large number of households could not (or did not) provide a specific number for the income earned from a stream. For example, of 1167 (out of a household sample set of 2668) observations, 530 reported that they were not aware of the income generated by their business in the previous year. The total income variable was constructed by excluding any observation where one (or more) streams reported income as "unknown", which leads to income information being available for just 1125 of the 2336 households in which children were assessed using the IDELA tool. Thus, while we report summary statistics for household income in the table above, we do not use it in our empirical analysis.

The survey also captured a large amount of information on asset ownership, business activities, time use and intra-household allocation preferences. In terms of broad measures of asset ownership, 67% of the surveyed households own some land, 43% own one or more businesses and 46% own some livestock. As a measure of the education level (or preferences for education) of a household we classify the respondents and household heads into 4 categories - no education or only nursery, some primary, some secondary and some tertiary/higher/certification. 51% of the respondents have some primary education, and only 5.4% have education beyond the secondary level whereas 49.5% of household heads have some primary education, and 8.9% of household heads have education beyond the secondary level. The major crops grown by land owning households are beans, cassava, sweet potato, ground nuts, maize and matooke[3].

The IDELA scores vary quite a lot along many of the household/demographic dimensions, some of which can be seen quite clearly in Table 6. There are strong differences in scores between genders, with girls outperforming boys by over a quarter of a standard deviation. Educated household heads and respondents (which in almost all cases is the primary caregiver of the evaluated child) imply (perhaps intuitively) higher test score achievement. Land and livestock ownership appear to be negatively correlated with test score achievements.

[3]A delicious and extremely satisfying variant of the plantain, often described as Uganda's national food. The author however strongly believes that the humble Rolex is the true claimant to that title.

Table 6: IDELA score variation

Variable	Category	n	Motor Skills	Early Literacy	Early Numeracy	Socio-emotional	Total IDELA
Gender	Male	1154	-0.151	-0.130	-0.083	-0.0620	-0.135
	Female	1182	0.142	0.127	0.081	0.0600	0.132
District	Masaka	253	-0.142	0.289	0.193	0.2850	0.145
	Mukono	241	0.295	0.349	0.190	0.1020	0.293
	Mityana	253	0.157	0.543	0.359	0.2130	0.367
	Iganga	320	-0.236	-0.492	-0.159	-0.2900	-0.350
	Mbale	299	0.184	0.195	0.273	0.5710	0.351
	Jinja	284	0.038	-0.385	-0.269	-0.2080	-0.229
	Kabarole	266	-0.163	-0.201	-0.065	-0.1560	-0.169
	Kasese	178	-0.228	0.106	-0.265	-0.3480	-0.219
	Kyenjojo	242	0.079	-0.210	-0.312	-0.2720	-0.182
Land ownership	No	760	0.130	0.141	0.101	0.0780	0.141
	Yes	1576	-0.063	-0.068	-0.049	-0.0380	-0.068
Business ownership	No	1324	-0.036	-0.037	-0.029	-0.0190	-0.039
	Yes	1012	0.047	0.049	0.038	0.0250	0.051
Livestock ownership	No	1250	0.041	0.025	0.010	0.0090	0.027
	Yes	1086	-0.047	-0.029	-0.011	-0.0100	-0.031
Respondent education	No or Nursery	258	-0.083	-0.212	-0.209	-0.2070	-0.205
	Some Primary	1183	-0.107	-0.097	-0.070	-0.0348	-0.097
	Some Secondary	725	0.136	0.140	0.116	0.0760	0.143
	Some Tertiary	123	0.376	0.556	0.469	0.2930	0.510
HH head education	No or Nursery	216	-0.096	-0.217	-0.179	-0.1640	-0.187
	Some Primary	1067	-0.080	-0.060	-0.027	-0.0050	-0.055
	Some Secondary	680	0.109	0.072	0.027	0.0140	0.072
	Some Tertiary	192	0.261	0.447	0.393	0.2720	0.409

There is strong variation in the IDELA scores across districts; Masaka, Mukono and Mityana (from Uganda's relatively prosperous Central region) and Mbale (a relatively industrialised region bordering Kenya) have significantly higher scores than the other districts.[4] The difference in means was established using a series of t-tests with adjusted p values ($p = 0.05/8 = 0.00625$ as each district is compared against 8 other districts). Table 7 contains the reported p-values from a two tailed t-test with significances calculated as above ($95\% = .00625$). A significant value indicates that the two districts have different mean IDELA total scores. In a large number of cases, the t-test was significant using the adjusted p-value. Some of the district-score components are not normally distributed, which calls into question the validity of the t-test. However, the existence of significant differences between districts which do not fail the Shapiro-Wilk normality test indicate that scores may in fact vary between districts. Similar results are found across all IDELA score components.

[4] Masaka, Mukono and Mityana are fairly close to the capital, Kampala. This provides these districts with better infrastructure and access to various goods and services than the relatively less well-connected Western districts (Kasese, Kabarole and Kyenjojo).

Table 7: T-tests by district: P-values for IDELA total scores (age standardised)

District	Masaka	Mukono	Mityana	Iganga	Mbale	Jinja	Kabarole	Kasese	Kyenjojo
Masaka	-								
Mukono	0.0725	-							
Mityana	0.0057*	0.3578	-						
Iganga	0*	0*	0*	-					
Mbale	0.0249	0.5291	0.8594	0*	-				
Jinja	0*	0*	0*	0.1201	0*	-			
Kabarole	0.0002*	0*	0*	0.0189	0*	0.4691	-		
Kasese	0*	0*	0*	0.0774	0*	0.8962	0.5352	-	
Kyenjojo	0.0002*	0*	0*	0.0356	0*	0.5789	0.8834	0.6542	-

Note:
* = Significance level is 95% if the p-value is <= 0.00625. A significant value indicates that the two districts have different mean IDELA total scores.

GPS coordinates were recorded for each household/child surveyed in the IDELA and household surveys. These coordinates were used to map rainfall data to each household, which is described below.

3.3 Rainfall Data

The rainfall information used in this thesis comes from the National Oceanic and Atmospheric Administration's (NOAA) African Rainfall Climatology - 2 (ARC-2) product. The ARC-2 data set provides estimates of precipitation for each cell of a 0.1 degree longitude by 0.1 degree latitude grid. Relative to other similar data sets, the ARC-2 data set provides higher spatial and temporal resolution combined with a relatively simpler estimation approach that potentially reduces biases in the estimation of precipitation (Novella and Thiaw 2013). Daily precipitation data is available for the region of interest (9 districts in Uganda) from 1983 till the present day. The data was downloaded using the RNOAA package for R (H. Edmund et al. 2014). The R code used to extract this information is included in appendix B of this thesis. Rainfall measures have been constructed using precipitation data between 1991 and 2018 (both inclusive). Data between 1983 and 1991 was not used because of a non-insignificant number of missing observations in this time period, which is less of a problem post 1991. The number of missing observations is less than half a percent over the selected time period. To our knowledge, this is the first study in this literature to use this data set; other Ugandan studies have used rainfall data collected from Ugandan weather monitoring stations at a district level.

The reference point for the calculation of rainfall aggregates is a child's month of birth. In a few cases, this data point was not recorded (in 164 cases, the survey respondent did not know the month in which the child was born). This results in a final data set consisting of 2172 children for whom we have both the IDELA evaluation and rainfall information.

Table 8: Rainfall Summary

Variable	Mean	Median	SD
Precipitation			
Average Rainfall	1266 mm	1280 mm	81.80
Utero year Rainfall	1293 mm	1302 mm	145.80
Year 1 Rainfall	1286 mm	1276 mm	158.90
Year 2 Rainfall	1249 mm	1223 mm	175.40
Deviation			
Utero Deviation	1.70%	2.70%	0.10
Year 1 Deviation	0.90%	1.15%	0.11
Year 2 Deviation	-2.14%	-2.42%	0.12
Good shocks			
Utero year	0.22	0	0.42
Year 1	0.21	0	0.41
Year 2	0.18	0	0.39
Bad shocks			
Utero year	0.14	0	0.35
Year 1	0.16	0	0.37
Year 2	0.27	0	0.44

The first (and most intuitive) measure of rainfall variation used in this analysis is the deviation in precipitation seen at a household relative to a long-term mean for that household. The closest ARC-2 grid point was identified using the recorded GPS coordinates of the household. Long-term average rainfall is the mean of the sum of the total rainfall grouped into 12-month periods. The month and year of birth of a child is used as the reference point for calculating the rainfall measure. The rainfall totals are calculated in the 12 months before the birth of the child, 12 months after the birth of the child and for the 2nd year of the child's life. The primary "rainfall shock" measure is then constructed by taking the difference of the natural logs of each of these totals and the natural log of the long-term average[5] computed using the same reference point.

Figure 2 shows the spatial variation in the magnitude of rainfall shocks in each of the three periods (in-utero, year 1 and year2 relative to the birth of a child) and the average rainfall observed by a household using the nearest available rainfall data point. Since children of

[5]E.g. $Shock_{Year1} = ln(rainfall_{Year1}) - ln(MeanRainfall)$

different ages have been sampled randomly across the chosen villages, some variation in the shocks can be seen between neighbouring households. There does not seem to be any serial correlation on shocks across years (i.e. households do not, generally speaking, face a sequence of positive or negative shocks). The correlation coefficients of the various shock variables are very low, with the largest being a -0.28 correlation between year 1 and year 2 deviations from the long-term mean. In our context, this can be interpreted as the chances of households receiving consecutive years of good or bad rainfall being acceptably low and may not significantly bias any estimates using these as explanatory variables. Interestingly, the coefficients are negative in sign indicating that it is more likely that a good rainfall year follows a bad one and vice-versa.

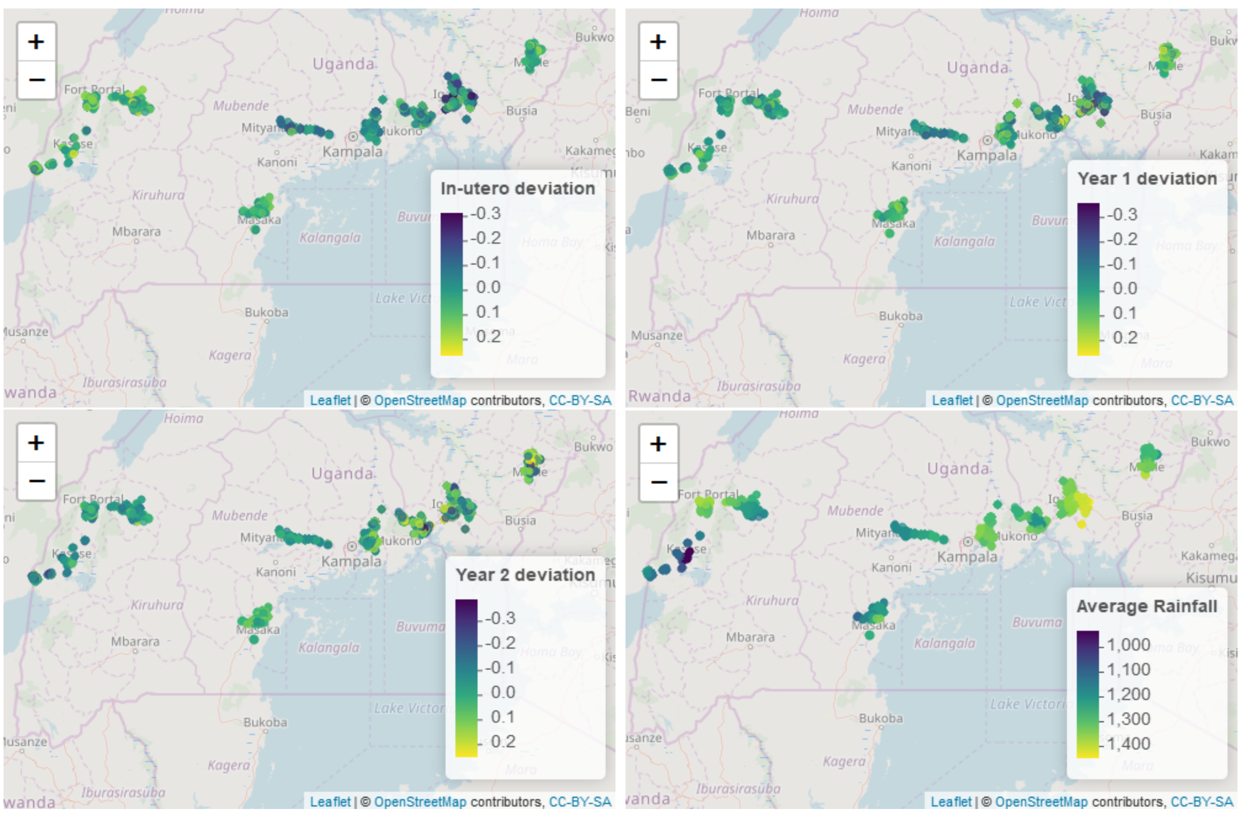

Figure 2: Distribution of rainfall shocks

As a secondary rainfall variation measure, we define binary variables that indicate whether a household received a positive or a negative shock. The threshold for a definitive positive or negative shock is defined as a 10% deviation from the long-term mean for a particular observation. This measure, although less precise than a continuous rainfall shock variable, allows the investigation of heterogeneous responses in the outcome variables to positive and negative rainfall variation. The 10% figure is, at best, an arbitrary choice; meteorological definitions of droughts vary from country to country, with a 15% shortfall from the long-term mean being used in countries like India to define droughts. The 10% figure was chosen as a compromise between a significant deviation and the necessity of having a sufficient number of observations above the shock cut-off threshold.

Lastly, given the possibility of differential impacts of a rainfall shock depending upon the timing of the shock, we construct variables similar to the primary measure but aggregated over two six-month periods (instead of the entire year). Relative to the month of a child's birth, aggregate rainfall in the two harvest seasons and two planting seasons along with the respective long-term means are calculated. In this analysis, we follow the information presented in Asiimwe and Mpuga (2007) and Mubiru et al. (2012) (as well as the recommendations of the BRAC researchers based in Uganda) in defining the *planting* season as the months of March-May and September-November. The *harvest* season is defined as the rest of the year (i.e. June-August and December-February). We also construct binary measures of these rainfall shocks (similar to the ones constructed on yearly rainfall above) to explore the differential impact of different types of rainfall shocks in each season. These are by their very nature fairly loose definitions, and will vary from crop to crop and from household to household. Several factors that might influence the dynamics of responses to these shocks remain unobserved in our analysis.

Unlike the yearly shocks described in Table 8, the harvest shock variables are not centred around 0, which is clearly visible in Figure 3. In absolute terms, harvest season averages are significantly lower than planting season averages - 447mm vs 819mm. Thus, even though the harvest rainfall in our observation set is a lot lower than the long-term average, the overall rainfall is not (as illustrated in table 8).

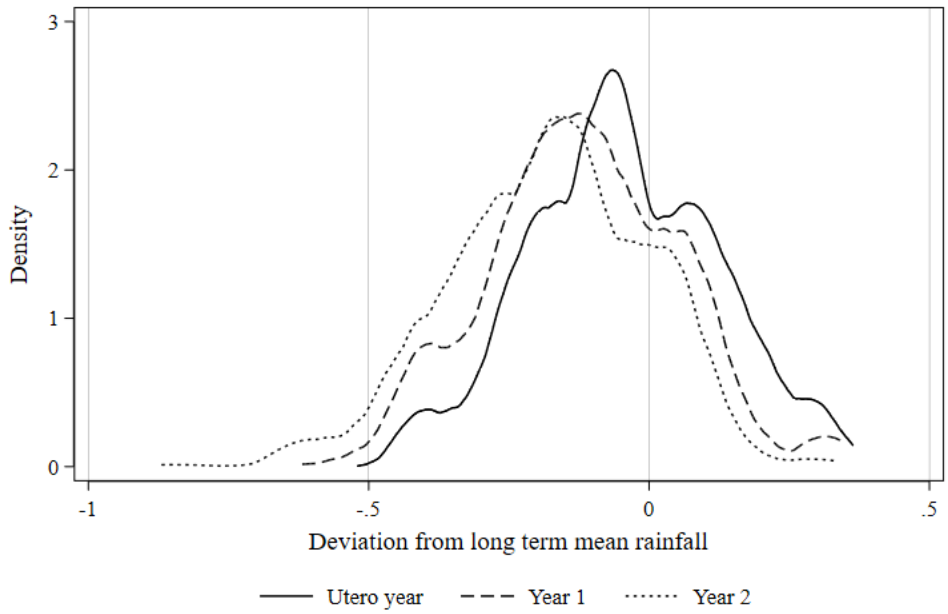

Figure 3: Harvest season rainfall shocks

Despite these uncertainties, investigating the differential effect of a rainfall shock in each season on early childhood development outcomes may reveal some information on the actual effect of rainfall. As our subsequent analysis reveals, rainfall in an economic and social perspective cannot be viewed as merely a continuous variable. Different types of rainfall events may work through different channels and need to be studied as separate phenomena.

4 Empirical Framework

Based on the conceptual model and the data points available for this analysis, the major threats to identification are omitted variables and measurement bias. Not knowing the sources of household income and how they evolve over time (at least during the early childhood period) leaves any estimate susceptible to omitted variable biases as these are important pieces of information. Without these, it is difficult to draw a causal link using channels such as maternal labour supply. Having a sense of time allocation between labour and leisure and being able to control for the vulnerability of household income to the vagaries of the weather gods would facilitate far more reliable estimates.

Secondly, despite rigorous training and randomised monitoring of enumerators[6] in the field,

[6] 19 enumerators did about 120 assessments each. They were assigned to specific districts based on

there are systematic differences between enumerators in how certain questions in the IDELA are scored which could point to a certain amount of incorrect administration of the survey in the field. Owing to practical considerations, enumerators were not randomly assigned across the entire sample set. This means that there is a good chance that the measurement error is not randomly distributed over the available observations. Some of the specifications used in the next section were also estimated with the inclusion of an enumerator level control variable in an attempt to partially alleviate the concern of systematic measurement error in the IDELA scores because of differences in enumerator training uptake and perception. However, the inclusion of this control variable did not significantly change results, and is excluded from most of the analysis.[7]

Under the assumption that permanent (smoothed) income can be subsumed into functions of household characteristics and regional idiosyncrasies, the conceptual model specified earlier can be extended to the following specification:

$$Score_{DHC} = \beta \cdot \widehat{Rainfall} + \gamma \cdot \widehat{Char_C} + \omega \cdot \widehat{Char_H} + \tau \cdot district + \epsilon$$

which is the reduced form specification used to study the link between child development outcomes and exogenous rainfall shocks. $Score_{DHC}$ is the IDELA score (either category or total) of a child C who belongs to a household H in district D. The primary outcome variables - the $Score_{DHC}$ variables - are normalised with respect to the means and standard deviations of each age cohort as there are large differences across age groups as described in the previous section.[8] $\widehat{Rainfall}$ is a vector of rainfall shocks (using one of the measures described in the previous section). $\widehat{Char_C}$ is a vector of child characteristics such as age or gender. $\widehat{Char_H}$ is a vector of household and survey respondent characteristics. This model can be estimated using ordinary least squares regression, with different measures of rainfall shocks and a selection of techniques for calculating the inference statistics.

4.1 Inferences

Assuming that the error terms obtained upon estimating the above model using OLS are uncorrelated could lead to incorrect inferences. The standard errors will be biased downwards unless potential correlations within similar groups of individuals (namely, clusters) are

language and familiarity considerations.

[7]Including enumerators as control variables led to missing F values in the OLS output from STATA. Although the F values were still significant, we chose to exclude this control in our analysis.

[8]$Score_{norm,age} = (Score_{raw} - mean_{age})/SD_{age}$ for each age (3, 4 and 5).

controlled for (Cameron and Miller 2015). In our context, it is quite likely that there are several unobserved dimensions of heterogeneity, some of which could potentially depend upon local factors such as ethnic differences or variation in development levels and/or economic production. The standard approach in this domain (based on the papers surveyed and discussed earlier in this thesis) is to cluster observations using a geographical group as the grouping dimension. Abadie et al. (2017) argue quite convincingly for a more rigorous approach towards the identification of a clustering variable (and indeed, whether clustering is necessary at all).

Following their approach involves looking at clustering as being driven by either sampling or experimental design instead of being driven by the empirical model being used. Our design is quasi-experimental, using random variation in rainfall. Thus, the experiment itself is randomly assigned to individuals - in theory, every observed individual has an equal chance of being affected by a rainfall shock. However, district selection from the population of available districts was not perfectly random. It was influenced by factors such as the availability of field infrastructure, ease of access and the desire to have a good spread across the country. This resembles a non-random draw from a larger sample which may possibly be heterogeneous as districts each have their own individual characteristics that are idiosyncratic but not necessarily normally distributed across all districts. In the Ugandan context, there are several districts that are quite different from any of the nine that were selected for the purpose of this project. Therefore, district level clustering may be necessary to account for the impact of selecting non-randomly from heterogeneous districts even if we include district fixed effects as a control variable in the empirical model (attributed to Arellano (1987) in Abadie et al. (2017)).

Additionally, the process of identifying observations also relies on selecting individuals randomly within each district. Each Ugandan district is divided into sub-counties and parishes. Parishes are very small units; however sub-counties are a potential candidate for clustering. There are 90 unique sub-counties in our sample set which is favourable because the larger the number of clusters, the more efficient the standard error estimates become (Cameron and Miller 2015). However, observations are not distributed evenly across sub-counties and the variance is high with several having fewer than 10 observations whereas some have over a hundred. Clearly, the sampling design did not use sub-counties as a criterion for sample selection which, according to Abadie et al. (2017), is a strong argument against using this as a clustering dimension. Lastly, the proposition introduced by Abadie et al. (2017) (which they support using a simulation approach) is that the presence of heterogeneous treatment effects constitutes a necessary condition to justify clustering. Therefore, we will also explore heterogeneity of our treatment (i.e. rainfall shocks) across districts.

This analysis explores the differences in the inferences using three separate methods. The standard errors for each coefficient of interest are computed using heteroscedasticity-robust clustering (the *robust* option in STATA), clustering by sub-county (using $vce(cluster)$), and clustering by district (using wild bootstrapping[9]). The wild bootstrap method for estimating standard errors is used when clustering by district as it provides reliable inference statistics even when the number of clusters is small (Webb 2013, Roodman et al. (2019)).

4.2 Heterogenous effects

To investigate heterogeneous responses of childhood development outcomes to rainfall shocks we can estimate:

$$Score_{DHC} = \beta \cdot \widehat{Rainfall} + \sum \theta \cdot X \cdot \widehat{Rainfall} + \gamma \cdot \widehat{Char_C} + \omega \cdot \widehat{Char_H} + \tau \cdot district + \epsilon$$

where the summation is across the vector X of dummies (e.g. age = 3/4/5). First, as discussed in the previous section, we estimated the above model with district indicator variables to inform the clustering decision. The deviation in rainfall in 3 time periods (12 months before the child's birth, 12 months immediately after birth and the 2nd year after birth) from the long-term mean were used as the rainfall shocks. For each group, the sum of the respective β and θ will give the effect of a particular rainfall shock on a particular group. For example, when looking at the effects of rainfall shocks on gender, the equation will take the form:

$$\begin{aligned} Score_{total} = {} & \beta_1 Shock_{utero} + \beta_2 Shock_{year1} + \beta_3 Shock_{year2} \\ & + \theta_{1,1} Shock_{utero} Male + \theta_{2,1} Shock_{year1} Male + \theta_{3,1} Shock_{year2} Male \\ & + controls + \epsilon \end{aligned}$$

In this example, if $\beta_1 + \theta_{1,1}$ is significant, then a rainfall shock when a male child is in-utero is linked to a significant effect on the total IDELA score $Score_{total}$.

Several terms of the form $district_D \cdot Rainfall_H + Rainfall_H$ (i.e. the effect of a particular rainfall shock on a household H in district D) were significant with SEs calculated using

[9] The wild bootstrap is implemented in STATA in the *boottest* package. Additionally, as advised by the authors of the package we use the Webb weight option as the number of clusters is small.

robust standard errors. This implies that some rainfall shocks are linked with different early childhood development outcomes in different districts. Districts were explicitly selected during the survey design. Thus, following the approach recommended by Abadie et al. (2017), most of the analysis in this thesis clusters standard errors by districts. Since the number of districts is quite low (9) relative to the recommended number of clusters (30-40), the wild bootstrap technique is used to calculate inference statistics.

4.3 Control variables

District fixed effects are included to absorb any idiosyncratic characteristics of each district. Child controls consist of the age and gender of the child. The age is an integer value between 3 and 5. Survey respondent controls consist of the education (none or nursery, some primary, some secondary, secondary completion and above), occupational status (employed/unemployed) and the age of the survey respondent. Household characteristic controls include the education and occupational status of the household head, the size of the household and the number of children in the household. Asset controls include indicators of whether a household owns land, businesses or livestock. We also include indicators for whether the household has any outstanding loans and whether the household has received any monetary transfers over the past 12 months. Finally, we use the log of the total income as an optional control in some specifications. However, the income variable is quite noisy and the chance of measurement error in this variable is quite high given the local context[10]. Hence, while the income variable is used in a couple of scenarios, the results may not be reliable and are not interpreted or discussed in detail.

5 Results

Estimates of the coefficients of interest (the rainfall shock variables) from the equation described above using age standardised scores and wild bootstrapping of standard errors (clustering at the district level) are shown in Table 9.[11] The only significant coefficient is

[10] "...It really plays into the cultural context. Unless they are people with a standard job they won't get paid regularly and often casual labor is paid based on work done ... Wage will be on how many bundles you complete in a day. Animals on the other hand are often fixed price that doesn't really fluctuate - and its easier to remember how many animals I've sold in a year or month then how many times I worked on my neighbours farm as a casual labourer." - Denise Ferris, researcher at BRAC Uganda.

[11] For brevity, only the major coefficients are shown here. The full regression tables are presented in table 16 in the appendix.

that of the deviation of rainfall from the long-term mean in the 2nd year of a child's life with socio-economic score. A 1% deviation from the long-term mean is linked with a change of .00378 standard deviations in the socio economic score achieved by a child. This coefficient is significant at the 5% level with standard errors calculated using wild bootstrapping, clustered at the district level.

As discussed in the earlier sections of this thesis, the coefficient is not merely the effect of rainfall on scores - it includes other coefficients that are endogenous in the structural equations, and therefore cannot be interpreted as the direct effect of rainfall. No evidence of a significant link between early life rainfall shocks on any other score component is found with this specification. Further, the interpretation depends quite a lot on the choice of inference calculation; when robust standard errors, simple clustering on districts, or clustering on sub counties is used, more significant results emerge. Following the discussion in the previous section, we do not use alternate methods - inference statistics calculated using the wild bootstrap technique are used throughout the rest of this analysis. Table 17 in the appendix compares the standard errors calculated using different methods for this particular regression.[12]

Table 9: Regression results: age-standardised scores, all controls

Variable	Motor Skills	Early Literacy	Early Numeracy	Socio-emotional	Total IDELA
In-utero	0.424	0.436	0.181	-0.02	0.332
	0.56	0.56	0.47	0.42	0.51
Year 1	0.173	0.635	0.467	0.403	0.517
	0.61	0.77	0.61	0.28	0.61
Year 2	0.194	0.621	0.413	0.378**	0.466
	0.63	0.66	0.53	0.16	0.56
n	2007	2007	2007	2007	2007
R-squared	0.09	0.19	0.11	0.12	0.14
Controls	All	All	All	All	All
Inference	Wild Bootstrap	Wild Bootstrap	Wild Bootstrap	Wild Bootstrap	Wild Bootstrap

Note:
* = 90%, ** = 95%, *** = 99% significance

The signs of the coefficients are worthy of note - a rainfall shock (which in the Ugandan context would typically imply a positive rather than an adverse effect on household income) has a positive link on some areas of early childhood development. One-sided hypothesis tests reject the null that the coefficients that are found to be significant are negative. In our setting, this implies that the children assessed in our study may either benefit from a positive

[12]Interestingly, SEs calculated using robust errors and clustered by sub counties are similar.

Table 10: Regression results: socio-emotional std. scores, incremental controls

Variable	1	2	3	4	5	6	7
In-utero	0.12	-0.002	-0.007	-0.024	-0.02	0.074	-0.42
	0.35	0.48	0.45	0.42	0.42	0.35	0.76
Year 1	0.527**	0.538*	0.444*	0.432*	0.403	0.503**	0.047
	0.24	0.25	0.22	0.26	0.28	0.24	0.38
Year 2	0.312**	0.371***	0.363***	0.399***	0.378**	0.409**	0.171
	0.13	0.15	0.15	0.16	0.16	0.17	0.32
Settings							
n	2172	2172	2132	2007	2007	2007	966
District controls	Y	Y	Y	Y	Y	Y	Y
HH controls			Respondent	All	All	All	All
Child controls		Y	Y	Y	Y	Y	Y
Asset index					Y	Y	Y
HH Income							Y
Enumerator						Y	
R-squared	0.1	0.1	0.11	0.12	0.12	0.29	0.17

Note:
* = 90%, ** = 95%, *** = 99% significance. Specification description: (1): district controls only, (2): (1) + child age and gender, (3): (2) + respondent age, education and occupational status, (4): (3) + household head education and occupational status , (5): (4) + land, business, livestock and loan ownership , (6) - (5) + enumerator controls, (7): (5) + income controls

rainfall shock or be harmed by a negative rainfall shock; the shock variables and specification do not throw any light on the relative importance or contribution of the direction of a rainfall shock.

Furthermore, as discussed before, the IDELA scores are not yet linked to long-term outcomes. The analysis therefore tells us that rainfall shocks might have an effect on non-cognitive components of ECD but does not yield the economic significance of this result. It is therefore interpretable solely as a link to a development outcome that has some validity because of other studies that have explored the validity of the IDELA tool in comparison to accepted tools used in the field of ECD.

Table 10 looks at the differences in the point estimates and wild bootstrap inference statistics for the standardised socio-emotional score (which we found to be significant above) by incrementally adding sets of control variables. The addition of respondent and household head controls leads to the loss of 165 observations, i.e. a total of 2007 observations in the full specification (reported in table 9). The coefficient on rainfall shocks in the first year of a child's life loses significance as controls are added. The coefficient on rainfall shocks in the second year of a child's life remains significant even after the full set of household controls are added. The magnitude of the coefficient is also fairly similar across specifications 2 to 5. This indicates a certain degree of robustness in our results; the lack of sensitivity to the addition or exclusion of control variables covering a large chunk of our observation set points to the existence of a link between early life rainfall shocks and the acquisition of non-cognitive skills (along the social-emotional dimension).

Specification 6 includes enumerator controls to account for the possibility of differences in how individual enumerators might have administered the assessment. Year 1 shocks do have a significant effect in this specification. The coefficient on the year 2 shock is quite similar to the coefficient in specification 5. There are a few minor variations in other IDELA variables as well, but as no additional significant effects are found when this control is included we do not use this as a control variable in the rest of our analysis.

5.1 Heterogeneity

We further extend our analysis of the impact of rainfall shocks on early childhood development outcomes by investigating the differential effects of shocks on a few heterogeneous groups. After estimating the interaction effect model defined in the previous section.[13] the joint significance of the relevant coefficients was tested using the wild bootstrap method (clustering by districts).

The results are presented in Table 11. There are strong signs of differential responses to some rainfall shocks along dimensions such as the gender of the child, land ownership and the education level of the respondent. These findings correspond to the strong differences in means that we saw in table 6.

Rainfall shocks have different effects on the socio-emotional score achievement of girls compared to boys - a 1% rainfall shock in the first year is linked with a score improvement of 0.0045 standard deviations (significant at the 1% level), and a 1% rainfall shock in the second year is linked with a score improvement of 0.0042 standard deviations (significant at the 5% level). There is no evidence of a significant effect of rainfall shocks on the development of boys aged 3-5. This agrees to some extent with the results found by Björkman-Nyqvist (2013) who found that rainfall shocks affect education investments in girls as their human capital could be more elastic to household income; families may be prioritising the development

[13] Heterogeneity across groups is studied using a model of the form:

$$Score_{DHC} = \beta \cdot \widehat{Rainfall} + \sum \theta \cdot X \cdot \widehat{Rainfall} + \gamma \cdot \widehat{Char_C} + \omega \cdot \widehat{Char_H} + \tau \cdot district + \epsilon$$

For each group, the sum of the respective β and θ will give the effect of a particular rainfall shock on a particular group.

Table 11: Heterogeneity analysis with age-standardised scores and all controls

		Motor Skills			Early Literacy			Early Numeracy			Socio-Emotional			Total IDELA		
Variable	Dimension	Utero	Year 1	Year 2	Utero	Year 1	Year 2	Utero	Year 1	Year 2	Utero	Year 1	Year 2	Utero	Year 1	Year 2
Gender	Male	0.292	0.173	0.27	0.124	0.536	0.714	0.305	0.649	0.393	-0.073	0.354	0.339	0.208	0.53	0.528
	Female	0.541	0.167	0.129	0.358	0.731	0.555	0.079	0.286	0.417	0.015	.449***	.418**	0.447	0.499	0.433
Age	3	1.176	0.648	0.142	0.982*	0.952	0.946	0.659	0.154	0.446	-0.201	0.454	0.496	0.881	0.753	0.573
	4	0.29	-0.615	0.326	0.481	0.463	0.82	0.249	0.959	0.464	-0.219	0.273	0.287	0.264	0.195	0.573
	5	0.199	0.588	-0.026	0.115	1.048	-0.306	-0.12	0.871	0.119	0.441	0.42	0.113	0.176	0.931	-0.04
Respondent education	None or nursery	0.142	-0.599	0.089	-0.106	0.312	-0.252	-0.209	-0.676	-0.252	-0.862	0.309	-0.609	-0.313	-0.25	-0.273
	Some primary	0.555	-0.005	0.014	0.36	0.808	0.387	0.004	0.673	0.298	-0.064	0.59**	.418**	0.316	0.598	0.322
	Some secondary	0.194	0.647	0.401	0.863**	0.997	1.386**	0.74**	0.794	0.81*	0.325	0.261	0.6**	0.598*	0.878	0.927**
	Some tertiary	1.342	0.859	1.141	1.576	-0.749	0.922	0.495	0.064	1.325	1.468	0.574	1.295**	1.484	0.282	1.355
Land ownership	No	0.191	-0.434	-0.128	0.484	0.1	0.682	0.424	-0.232	0.505	0.709	-0.13	0.535**	0.544	-0.211	0.403
	Yes	0.536	0.46	0.335	0.473	0.933	0.579	0.158	0.853	0.359	-0.282	0.683*	0.32*	0.296	0.892	0.491
Outstanding loans	No	0.162	-0.334	0.106	0.008	0.262	0.487	-0.065	0.167	0.374	-0.483	0.213	0.257	-0.076	0.087	0.352
	Yes	0.719	0.931	0.254	0.947*	1.142	0.716	0.497	0.912	0.408	0.533	0.603*	0.449*	0.959	0.27	0.193

Note:
* = 90%, ** = 95%, *** = 99% significance, SE's calculating using Wild Bootstrap

needs of boys ahead of girls in the Ugandan setting in case of income shortfalls or surpluses wrought by the occurrence of unpredictable rainfall shocks.

The analysis also shows differences between households who own land and those who do not. A 1% positive rainfall shock in the second year of a child's life is linked with an improvement in the socio-emotional IDELA score by .005 standard deviations when the child's family does not own any land. When the child's family owns land, shocks in year 1 and year 2 have significant links - a 1% shock linked to .007 and .003 standard deviation score changes respectively. This may indicate that families owning land and who may be more dependent on agricultural income - only a small number of landowning households rent out or sharecrop their land - are affected to a greater extent by rainfall shocks, which manifests in higher early child development on the non-cognitive dimension.

There is some evidence of an interaction effect between the education level of the respondent and rainfall shocks. Socio emotional, early literacy, early numeracy and the total test scores of children whose primary caregiver (i.e. the respondent) has studied at the secondary level or above are affected differently by rainfall shocks in the second year of a child's life. The socio-emotional scores of children whose parents have only undergone some education at the primary level have different links with rainfall shocks in the first and second years of the child's life. As we would expect, the impact on children where the respondent had secondary education (or above) was higher than the impact on children where the respondent only had some primary education. In the context of the model defined earlier, this could be indicative of education having an influence on parental preferences for child utility/well-being. It is plausible that well educated parents place a higher value on ECD and have a better sense of what can improve ECD, which leads to both increased quantity and efficiency of income allocations towards goods and services that could improve ECD.

Our data set also contains information on the loan liabilities of the surveyed households - 40% of the households in which a child was assessed using the IDELA have outstanding loans. The existence of an outstanding loan could plausibly be used as an indicator of an ability and wherewithal to use financial instruments. This could translate into an ability to navigate adverse income shocks more smoothly than families who do not have the access or ability to risk mitigation instruments such as loans. We find that the socio-emotional IDELA score of children whose families who have outstanding loans at the time of assessment are linked to rainfall shocks in the first year of life (.006 standard deviation change on a 1% rainfall shock, significant at the 10% level). The socio-emotional score is also linked with a rainfall shock in the 2nd year of life - .004 standard deviation change on a 1% rainfall increase from the long-term mean, significant at the 10% level.

In relation to the earlier results on the link between rainfall shocks in early childhood and socio-emotional score achievement, some important information can be gleaned from this section of the analysis. First, rainfall shocks seem to have an effect on girls rather than boys. Second, the link is stronger the more educated the household member (usually the mother of the surveyed child) is. Third, children belonging to families that do not own land see stronger links. Lastly, the effects seem to be in the same direction across groups and IDELA components and generally speaking appear to be in accordance with what we would have expected. However, this analysis does not reveal whether positive or negative rainfall shocks have equal effects but in opposite directions, or whether the effects are stronger in either of the possible directions of a shock.

5.2 Rainfall shock decomposition

A positive rainfall shock and a negative rainfall shock may have different outcomes because they do not necessarily imply direct complementarity of income allocation. It is plausible that the effects of a negative shock may be crowded out if parents increase their efforts in an attempt to try and give their children a better life. Similarly, a positive shock may not have any impact on early childhood development if the parents are lazy and/or do not value the child's utility (and development level) enough to take advantage of the income opportunity afforded by a positive rainfall shock. To investigate this angle, we use a binary measure of rainfall shocks (defined as a +/- 10% deviation from the long-term mean) to decompose the effects of positive and negative shocks. Table 12 shows the results of estimating the model using indicator variables for rainfall shocks instead of a continuous measure of rainfall deviation from the long-term mean using the wild bootstrap technique to calculate the inference statistics.

Positive rainfall shocks in the 2nd year of a child's life have an effect on the socio-emotional score component. A positive shock (10% or greater rainfall than the long-term mean) in that year is linked to a .097 standard deviation increase in test score achievement, significant at the 90% level. This is in line with the findings from the previous set of results. The coefficient on the negative rainfall shock in year 2 is not significant. Additionally, there is a significant (90% level) coefficient on the negative rainfall shock in the first year of the child's life - a bad rainfall shock in the first year of a child's life is linked to a subsequent decrease of 0.056 standard deviations in the socio-emotional test score.

This analysis also finds some links between certain types of rainfall shocks and other developmental domains. Positive rainfall shocks in the utero year are linked with a 0.161 standard deviation increase in motor skills test score achievement (significant only at the 10% level). A positive shock in the second year of a child's life is linked with an increase of 0.128 standard deviations in motor skills test score achievement, an effect that is significant at the 5% level. A negative shock in the first year of a child's life is found to be linked with a decrease of 0.169 standard deviations in early literacy test score achievement (significant at the 5% level).

Table 12: Regression results: age-standardised scores, Binary shocks, all controls

Variable	Motor Skills	Early Literacy	Early Numeracy	Socio-emotional	Total IDELA
Positive shocks					
In-utero	0.161*	0.042	0.085	-0.02	0.09
Year 1	0.067	0.059	0.094	0.048	0.085
Year 2	0.128**	0.066	0.064	0.097*	0.111
Negative shocks					
In-utero	-0.014	-0.02	0.061	-0.025	0.004
Year 1	-0.004	-0.09	-0.074	-0.056*	-0.067
Year 2	-0.022	-0.169*	-0.105	-0.116	-0.114
n	2007	2007	2007	2007	2007
Settings					
HH controls	Y	Y	Y	Y	Y
District controls	Y	Y	Y	Y	Y
Child controls	Y	Y	Y	Y	Y
HH Income	N	N	N	N	N
SE method	Wild bootstrap	Wild bootstrap	Wild bootstrap	Wild bootstrap	Wild bootstrap
Standardised	Y	Y	Y	Y	Y
R-squared	0.09	0.19	0.11	0.12	0.15

The signs of the coefficients are consistent - positive shocks have positive coefficients, negative shocks have negative coefficients. The link between positive rainfall shocks and motor skills could perhaps be indicative of a direct effect of an unforeseen boost to income (caused by a positive rainfall shock) on nutrition or leisure allocation which could lead to healthier children and a happier home, represented in our results by higher non-cognitive score achievement (motor skills and socio-emotional development).

5.3 The effect of multiple shocks of the same type

One of the direct implications of the model is that children whose households have been exposed to income shocks more frequently could have amplified effects on their development outcomes. A child whose family has been the beneficiary of two rainfall shocks out of the 3 periods under analysis should be affected significantly more than children whose households have faced just one shock. To study this, we construct indicator variables for all households who have faced two (or three) positive or negative shocks in the three periods (i.e. the utero year, first year and second year of a child's life). These new indicator variables are then added to the previous specification (that uses binary shock values) and estimated using OLS. The hope is that if there are stronger effects due to the accumulative nature of income shocks, the new indicator variables will have values that are significantly different from 0. Inference statistics were calculated using the wild bootstrap technique (with district level clustering).

134 households have faced 2 or 3 negative shocks and 128 have faced 2 or 3 positive shocks. The multiple shock indicator variables are not significant for any of the five IDELA measures. The socio-emotional score effect discussed in the binary shock section retains its significance - a single negative shock has a negative link with score achievement (reducing the scores by 0.1 of a standard deviation, significant at the 10% level). Additionally, two positive shocks are linked positively with motor skills, increasing the score achievement by 0.289 SDs (significant at the 10% level). Overall, the analysis does not find very strong evidence for a link between multiple shocks of the same type (in utero and in early childhood) and ECD of the children in our sample as measured by IDELA score achievement.

Table 13: Regression results: age-standardised scores, aggregated binary shocks, all controls

Variable	Motor Skills	Early Literacy	Early Numeracy	Socio-emotional	Total IDELA
No. of positive shocks					
1	0.106	0.041	0.037	0.017	0.069
	0.09	0.11	0.1	0.09	0.11
2	0.289*	0.151	0.24	0.058	0.238
	0.15	0.23	0.16	0.25	0.16
No. of negative shocks					
1	0.017	-0.103	-0.053	-0.1*	-0.062
	0.11	0.12	0.08	0.06	0.11
2	-0.076	-0.217	-0.112	-0.12	-0.154
	0.17	0.21	0.18	0.09	0.18
n	2007	2007	2007	2007	2007
R-squared	0.09	0.19	0.11	0.12	0.15
Controls	Y	Y	Y	Y	Y
Inference	Wild bootstrap	Wild bootstrap	Wild bootstrap	Wild bootstrap	Wild bootstrap

Note:
* = 90% significance, ** = 95% significance

5.4 Seasonal Decomposition

Table 14 shows the results when the model is estimated using seasonal rainfall shocks. The purpose of estimating this specification is to investigate whether the timing of a rainfall shock matters or not. Asiimwe and Mpuga (2007) in Uganda and Leight, Glewwe, and Park (2015) in rural China are examples of similar studies that use this approach in an attempt to identify seasonal differences in the impact of rainfall shocks. Unpredictable rainfall may have different effects if the shock occurs at a time that is beneficial or harmful to agricultural productivity. Inference statistics are calculated using the wild bootstrap technique with district level clustering.

The results are strikingly different - Harvest season rainfall shocks during the year when the child is in utero and in the first year of a child's life are significant across all score categories except the socio-emotional component. A 1% rainfall shock in the first harvest season of a child's life is linked with a .0036 standard deviation increase in the child's socio-emotional score achievement. In contrast, the only planting season component that is significant is the effect of a rainfall shock during the planting season in the first year of a child's life. Here, the results indicate a negative link between a rainfall shock in year 1 and motor skill score achievement. A 1% increase in rainfall (relative to the long-term mean for the planting season) is linked to a .0036 standard deviation decrease in the motor skill domain score (significant

Table 14: Regression results: age-standardised scores, Seasonal shocks

Variable	Motor Skills	Early Literacy	Early Numeracy	Socio-emotional	Total IDELA
Planting Season					
In-utero	-0.171	-0.281	-0.328	-0.277	-0.305
	0.53	0.45	0.47	0.41	0.52
Year 1	-0.356*	-0.096	0.138	0.138	-0.06
	0.23	0.47	0.33	0.25	0.29
Year 2	0.191	0.356	0.219	0.244	0.306
	0.31	0.31	0.31	0.36	0.32
Harvest Season					
In-utero	0.822***	0.872***	0.536***	0.298	0.789***
	0.22	0.24	0.16	0.3	0.21
Year 1	0.845***	1.03***	0.489*	0.36**	0.854***
	0.3	0.19	0.25	0.17	0.21
Year 2	0.161	0.403	0.277	0.185	0.298
	0.4	0.3	0.19	0.13	0.29
n	2007	2007	2007	2007	2007
R-squared	0.1	0.21	0.11	0.12	0.16

Note:
* = 90%, ** = 95%, *** = 99% significance

at the 10% level). The results indicate that rainfall shocks may have an impact on ECD if the shock happens in the harvest season, but there is no evidence of any effect if the shock happens in the planting season.

However, the data set used in this analysis and the literature reviewed earlier do not provide any clear channels that could explain this. Some crops (such as cotton or wheat) are more susceptible to destruction if unexpected rainfall occurs just as they are ripening (i.e early in the harvest season in our context). Excessively rainy weather can promote the growth of harmful pests and fungi that can damage maturing crops. Similarly, harvested crops can be ruined by unseasonal rainfall if they are not stored securely in a dry space or transported without sufficient waterproofing, both of which are quite possible in rural Uganda where such infrastructure simply may not exist. However, some crops may be at risk if extremely heavy rainfall occurs soon after the seeds are planted as well, as flooded soil does not contain enough oxygen to enable the seeds to germinate.

On the other hand, additional rainfall that falls at the proper time in the proper amount during planting/growing can crop boost yields and crop quality. Thus, beyond possessing different effects, certain types of shocks in certain seasons may have differing effects depending upon the direction and magnitude of this shock. Therefore, the next step in our analysis is to investigate this using an approach similar to the model used in Table 12. Table 15 presents the estimated results using binary indicators of rainfall shocks in each season (greater or lesser than 10% of the long-term seasonal mean) in an attempt to decompose the differential effects of positive and negative seasonal rainfall shocks. Again, there are a few significant results in the harvest season but very few in the planting season. More interestingly, positive shocks in the harvest season seem to have negative effects on IDELA score achievement. However, the direction of the relationship is flipped around in the harvest season. The signs of these results are consistent with the signs of the coefficients estimated using the continuous rainfall measure (Table 14).

Table 15: Regression results: age-standardised scores, Binary seasonal shocks

Shock type	Shock timing	Motor Skills	Early Literacy	Early Numeracy	Socio-emotional	Total IDELA
Planting Season						
	In-utero	0.009	-0.039	-0.007	-0.088*	-0.032
Positive	Year 1	-0.109*	-0.13	-0.029	-0.069	-0.105
	Year 2	0.078	0.091	0.061	-0.1*	0.045
	In-utero	0.265*	0.086	0.22	0.121	0.279
Negative	Year 1	0.134	0.086	0.094	-0.056	0.085
	Year 2	0.088	0.076	0.081	-0.11	.-79
Harvest Season						
	In-utero	0.154*	0.154**	0.058	0.047	0.129*
Positive	Year 1	0.148	0.192*	0.09	0.12	0.16
	Year 2	0.298	0.08	0.168	0.193	0.234
	In-utero	-0.102	-0.094	-0.083	-0.073	-0.109
Negative	Year 1	-0.104*	-0.137**	-0.067	-0.096**	-0.129**
	Year 2	0.117	0.022	0.011	-0.087**	0.032
	n	2007	2007	2007	2007	2007
	R-squared	0.11	0.21	0.11	0.13	0.16

Note:
SE's calculated using Wild Bootstrap. * = 90% significance, ** = 95% significance

This part of the analysis is potentially biased by the skewed distribution of rainfall shocks in the harvest and planting season. For example, around 55% of households have a negative shock in the first year of the assessed child's life in the harvest season whereas only 10% of households have a negative shock in the planting season. The results do however point to the potential differences that may emerge depending upon the choice of rainfall variable. The results are are in line with the results of Asiimwe and Mpuga (2007) who find similar relationships between income/consumption and seasonal rainfall shocks in Uganda.

6 Discussion of Results & Conclusion

This analysis attempted to study the link between exogenous and unforeseen variations in household income and the development level of very young children in Uganda using rainfall shocks as an instrument for household income variation. The question is an important one - understanding this link will inform effective policy design and charitable non-governmental efforts targeted towards achieving the SDG targets in the realm of ECD in LMICs.

6.1 Limitations

It goes without saying that there are several limitations to the approach used in this thesis. The results obtained using this analytical approach may have limited external validity as the reduced form technique yields limited information about the underlying structure. Specifically, while we find some links between rainfall shocks and measures of early childhood development (ECD), these may be suitable only for the context being studied - a selection of economically constrained families in selected Ugandan districts. The reduced form would be far more interesting if repeated measurements of the outcomes of interest were available. If the same children are evaluated at ages 3, 4 and 5, we could use a difference-in-differences approach in addition to the exogenous variation induced by rainfall shocks to extract more robust findings which would be less susceptible to omitted variable biases.

The lack of a quantifiable economic significance of the early childhood measures of development provided by the IDELA restricts us from providing an estimate of the potential impact of rainfall shocks on ECD. This is true of most instruments that measure ECD at a granular, individual level - these are new tools, developed to address problems that have come into the spotlight relatively recently.

Other limitations of this thesis stem from the context and construction of the data set used here. The inability of our analysis to account for equilibrium effects beyond including district fixed effects further restricts generalisability of the results. Further, a large number of children and households were surveyed from several Ugandan districts, but they are somewhat homogeneous in the sense that they were selected based on specific economic and demographic criteria. Since we do not observe the behaviour of all sections of society and from different walks of life, the analysis and findings remain restricted within the bounds of the selected sample.

We discussed omitted variables and measurement errors as large threats to our identification strategy. The complexity of the relationship to be studied and the huge diversity of economic, cultural, technological and social characteristics that go into defining the relationship (some of which are discussed in section 2) are not addressable by the data at our disposal. Also, despite the efforts of the people involved in the design and implementation of the surveys, the challenges of implementing a reliable and cost-effective individual level child assessment within resource and manpower constraints in Uganda are manifold, and a certain amount of measurement error in the IDELA measurements is inevitable.

Lastly, specifications using quadratic forms of the continuous rainfall shock variable could yield more information about the directional effect of both seasonal and yearly rainfall shocks. While the binary indicators used in this analysis do find some interesting results, using a polynomial model for rainfall shocks may uncover more information about the links between rainfall shocks and ECD measurements. Indicatively, a simple quadratic version of the seasonal continuous rainfall shock variable on socio-emotional scores saw an increase in the magnitude of the linear coefficients for harvest seasons, and an even larger positive effect on the quadratic term for the second year harvest season. Thus, using this approach could improve the quality of this sort of analysis.

6.2 Discussion of results

In this analysis, we have looked at two dimensions of rainfall - seasonal against yearly aggregate, and rainfall as a monolithic economic entity against large positive and negative rainfall deviations as different economic entities. The common finding across all the specifications and definitions used is that, at least in our data set, a variation in rainfall in the early years of a child's life is linked with the non-cognitive skill development of that child. When rainfall is defined at the yearly level, a positive deviation in the rainfall received at a household level from the long-term mean seems to be linked with an improvement in non-cognitive skill development as measured by the IDELA socio-emotional and motor skill scores. There is evidence of links between non-cognitive ECD and large (10% or greater) shocks in either direction from the long-term mean of rainfall in the early years of a child's life. The evidence is not very strong for a link between rainfall shocks in-utero and ECD (cognitive and non-cognitive), which is not in line with the expectations of the foetal origins hypothesis and related research in this field. For example, Leight, Glewwe, and Park (2015) find that rainfall shocks in utero have stronger effects than other early life rainfall shocks.

When rainfall is decomposed into two separate periods (i.e. planting and harvest seasons in the local context), the relationship between rainfall shocks in the harvest season and non-cognitive skill development appears to differ from the relationship in the planting season. In the harvest season, a positive deviation from the long-term mean is linked with a positive change in ECD and vice-versa. Indeed, we find stronger evidence for effects of negative shocks than we do for positive ones, especially on socio-economic score achievement which yields the most significant results across all specifications and all our rainfall measures. However, the relationship seems to flip around in the planting season - a positive rainfall shock in the planting season is linked to negative effects in non-cognitive skill development measures. This suggests that different economic pathways may exists in each season, possibly driven by agricultural considerations that dictate the influence of rainfall on agricultural productivity. Additionally, this decomposition yields some evidence of links between rainfall shocks both in-utero and in the early years of a child's life on various measures of non-cognitive skills (and limited evidence on early literacy skills).

Additionally, our heterogeneity analysis revealed that a) the non-cognitive skill development of girls may be more strongly linked to early life rainfall shocks, b) children of educated parents/primary caregivers may receive stronger boosts to non-cognitive development from unforeseen income shocks wrought by rainfall shocks, and c) owning land or having access to loans and financial instruments could amplify the effects of shocks. Although the direction of rainfall shocks and these effects have not been analysed, which prevents drawing any significant conclusions from the points mentioned above, some of these findings resonate with results from similar studies in the literature. However, the identification of these differential effects along household and individual characteristics is useful in that it points to potential pivots for developmental programs focused on ECD - there may be a need to focus on girls, on households where the adults are not educated, households who own land and are dependent on its productivity for their survival, and perhaps on households who are at risk of defaulting on loans because of a bad harvest.

Another implication of our results is that in the event of sudden and rapid climate change - where income shocks caused by rainfall shocks may become far more common and severe - the socio-emotional development of children can be impacted negatively. Children in underdeveloped countries such as Uganda without the physical and financial infrastructure to insulate at-risk farmers and families from the effects of these shocks could be impacted negatively.

The channel linking income shocks to changes in socio-emotional (or other non-cognitive) scores is unclear. Effects on motor skill score achievement can perhaps be linked to poorer health and nutrition forced by a negative shock to income. The determinants of the level socio-emotional skills in children need to be explored further in order to tease out causality. A potential links may be reduced leisure time for mothers (who may have to look for additional income sources in times of scarcity), and leisure time has been linked to improved cognitive development in children (Blau and Grossberg 1990). However, the lack of an effect on cognitive outcomes does not strengthen this hypothesis. Another link could be that adversity leads to unhappier households as stress in parent-child interactions could have negative effects during early childhood. There is some evidence in the psychology and medical literature (for example Shonkoff et al. (2012)), and this could be a fruitful area for future research.

From an econometric perspective, the various specifications used in this analysis and the different inferences obtained by using different methods of calculating inference statistics point to the need for caution in interpreting the results of ordinary least squares regressions in experimental settings similar to this one. The results vary drastically depending upon the clustering method, and there does not seem to be any clear consensus on which method is the most appropriate in settings of this nature where the sample set was specifically constructed for a different purpose. The absence of repeated observations over time and a clearly defined control group makes reliable inference even trickier. However, an advantage of this research design against a few of the studies discussed earlier is that looking at short term outcomes (3-5 years since birth) removes a large amount of uncertainty and attenuation by unobserved factors in the long term. It is slightly easier to construct a causal link between early childhood income shocks and early childhood development than it is to explore the same link for adult outcomes.

In conclusion, the effect of non-cognitive skill development/acquisition in childhood on adult outcomes in a LMIC context is not well studied in the literature. The literature on this in developing countries indicate that the effect of non-cognitive skills are large and manifest in many different ways - Heckman, Stixrud, and Urzua (2006) find an impact on labour outcomes, whereas Carneiro, Crawford, and Goodman (2007) find many different links such as years of schooling, adolescent misbehaviour, etc. Our results, if generalised, indicate that large numbers of children in LMICs living in conditions similar to our sample may not realise their maximum potential in non-cognitive skill development if their families are subjected to unexpected negative income shocks during early childhood. The IDELA measures socio-emotional skills by eliciting the responses of children to hypothetical situations covering scenarios such as identification of conflicts, emotional awareness and control, and most importantly, self-awareness. These are crucial skills and falling behind in terms of both domestic and global competitiveness could reduce the advancement potential of countries such as Uganda.

Program and policy designers must consider the impact of unpredictable shocks of this nature as such events might nullify the impact of a program or investment. Providing adequate risk mitigation facilities such as loans or crop insurance and creating infrastructure to reduce dependence on rainfall for agricultural production and providing the tools and technology required to move away from a dependence on natural rainfall could improve childhood development outcomes by a significant amount. Building awareness amongst families on the importance of child development and guidelines on the efficient use of time and income to boost ECD may prove to be fruitful as well. Beyond education and health programs, insuring households against income risks because of climate events could play an important role in meeting early child development goals.

References

Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey Wooldridge. 2017. “When Should You Adjust Standard Errors for Clustering?” National Bureau of Economic Research.

Aker, Jenny C. 2012. “Rainfall Shocks, Markets and Food Crises: The Effect of Drought on Grain Markets in Niger.” *Center for Global Development, Working Paper*.

Asiimwe, John Bosco, and Paul Mpuga. 2007. “Implications of Rainfall Shocks for Household Income and Consumption in Uganda.” African Economic Research Consortium (AERC).

Avitabile, Ciro, Irma Clots-Figueras, and Paolo Masella. 2014. “Citizenship, Fertility, and Parental Investments.” *American Economic Journal: Applied Economics* 6 (4): 35–65.

Björkman-Nyqvist, Martina. 2013. “Income Shocks and Gender Gaps in Education: Evidence from Uganda.” *Journal of Development Economics* 105. Elsevier: 237–53.

Black, Maureen M, Susan P Walker, Lia CH Fernald, Christopher T Andersen, Ann M DiGirolamo, Chunling Lu, Dana C McCoy, et al. 2017. “Early Childhood Development Coming of Age: Science Through the Life Course.” *The Lancet* 389 (10064). Elsevier: 77–90.

Blau, Francine D, and Adam J Grossberg. 1990. “Maternal Labor Supply and Children’s Cognitive Development.” National Bureau of Economic Research.

Blinder, Alan S. 1973. “Wage Discrimination: Reduced Form and Structural Estimates.” *Journal of Human Resources*. JSTOR, 436–55.

Britto, Pia Rebello. 2017. *Early Moments Matter for Every Child*. ERIC.

Cameron, A Colin, and Douglas L Miller. 2015. “A Practitioner’s Guide to Cluster-Robust Inference.” *Journal of Human Resources* 50 (2). University of Wisconsin Press: 317–72.

Carneiro, Pedro, Claire Crawford, and Alissa Goodman. 2007. “The Impact of Early Cognitive and Non-Cognitive Skills on Later Outcomes.” Centre for Economics of Education.

Children, Save the. 2018. “Beyond Access: Exploring Equity in Early Childhood Learning and Development.” *Journal of Applied Meteorology and Climatology* 52 (3): 588–606.

Cornwell, Katy, and Brett Inder. 2015. “Child Health and Rainfall in Early Life.” *The Journal of Development Studies* 51 (7). Taylor & Francis: 865–80.

Edmund, Hart, Scott Chamberlain, Karthik Ram, and Maintainer Hart Edmund. 2014.

"Package 'Rnoaa'." Citeseer.

Fernald, Lia CH, Elizabeth Prado, Patricia Kariger, and Abbie Raikes. 2017. "A Toolkit for Measuring Early Childhood Development in Low and Middle-Income Countries." World Bank.

Glewwe, Paul, and Edward A Miguel. 2007. "The Impact of Child Health and Nutrition on Education in Less Developed Countries." *Handbook of Development Economics* 4. Elsevier: 3561–3606.

Gommes, R, and F Petrassi. 1996. "Rainfall Variability and Drought in Sub-Saharan Africa." *SD Dimensions, FAO*.

Grantham-McGregor, Sally, Yin Bun Cheung, Santiago Cueto, Paul Glewwe, Linda Richter, Barbara Strupp, International Child Development Steering Group, and others. 2007. "Developmental Potential in the First 5 Years for Children in Developing Countries." *The Lancet* 369 (9555). Elsevier: 60–70.

Halpin, Peter F, Sharon Wolf, Hirokazu Yoshikawa, Natalia Rojas, Sarah Kabay, Lauren Pisani, and Amy Jo Dowd. 2019. "Measuring Early Learning and Development Across Cultures: Invariance of the Idela Across Five Countries." *Developmental Psychology* 55 (1). American Psychological Association: 23.

Heckman, James J, Seong Hyeok Moon, Rodrigo Pinto, Peter A Savelyev, and Adam Yavitz. 2010. "The Rate of Return to the Highscope Perry Preschool Program." *Journal of Public Economics* 94 (1-2). Elsevier: 114–28.

Heckman, James J, Jora Stixrud, and Sergio Urzua. 2006. "The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior." *Journal of Labor Economics* 24 (3). The University of Chicago Press: 411–82.

IDELA network. 2019. "IDELA Data Sets." *IDELA*. Save the Children. Accessed May 20. https://idela-network.org/data/.

Jayachandran, Seema. 2006. "Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries." *Journal of Political Economy* 114 (3). The University of Chicago Press: 538–75.

Jensen, Robert. 2000. "Agricultural Volatility and Investments in Children." *American Economic Review* 90 (2): 399–404.

Jury, Mark R. 2018. "Uganda Rainfall Variability and Prediction." *Theoretical and Applied*

Climatology 132 (3-4). Springer: 905–19.

Kochar, Anjini. 1995. “Explaining Household Vulnerability to Idiosyncratic Income Shocks.” *The American Economic Review* 85 (2). JSTOR: 159–64.

Leight, Jessica, Paul Glewwe, and Albert Park. 2015. “The Impact of Early Childhood Shocks on the Evolution of Cognitive and Non-Cognitive Skills.” *Gansu Survey of Children and Families Papers*.

Levine, David I, and Dean Yang. 2014. “The Impact of Rainfall on Rice Output in Indonesia.” National Bureau of Economic Research.

Maccini, Sharon, and Dean Yang. 2009. “Under the Weather: Health, Schooling, and Economic Consequences of Early-Life Rainfall.” *American Economic Review* 99 (3): 1006–26.

McCoy, Dana Charles, Maureen M Black, B Daelmans, and T Dua. 2016. “Measuring Development in Children from Birth to Age 3 at Population Level.” *The Hague, the Netherlands: Bernard van Leer Foundation*, 34–39.

Miguel, Edward. 2005. “Poverty and Witch Killing.” *The Review of Economic Studies* 72 (4). Wiley-Blackwell: 1153–72.

Morris, Ian. 2010. *Why the West Rules-for Now: The Patterns of History and What They Reveal About the Future*. Profile books.

Mubiru, Drake N, Everline Komutunga, Ambrose Agona, Anne Apok, and Todd Ngara. 2012. “Characterising Agrometeorological Climate Risks and Uncertainties: Crop Production in Uganda.” *South African Journal of Science* 108 (3-4). Academy of Science of South Africa: 108–18.

Novella, Nicholas S, and Wassila M Thiaw. 2013. “African Rainfall Climatology Version 2 for Famine Early Warning Systems.” *Journal of Applied Meteorology and Climatology* 52 (3): 588–606.

Paxson, Christina H. 1992. “Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand.” *The American Economic Review*. JSTOR, 15–33.

Pisani, Lauren, Ivelina Borisova, and Amy Jo Dowd. 2018. “Developing and Validating the International Development and Early Learning Assessment (Idela).” *International Journal of Educational Research* 91. Elsevier: 1–15.

Roodman, David, Morten Ørregaard Nielsen, James G MacKinnon, and Matthew D Webb. 2019. “Fast and Wild: Bootstrap Inference in Stata Using Boottest.” *The Stata Journal* 19

(1). SAGE Publications Sage CA: Los Angeles, CA: 4–60.

Shah, Manisha, and Bryce Millett Steinberg. 2017. "Drought of Opportunities: Contemporaneous and Long-Term Impacts of Rainfall Shocks on Human Capital." *Journal of Political Economy* 125 (2). University of Chicago Press Chicago, IL: 527–61.

Shonkoff, Jack P, Andrew S Garner, Benjamin S Siegel, Mary I Dobbins, Marian F Earls, Laura McGuinn, John Pascoe, et al. 2012. "The Lifelong Effects of Early Childhood Adversity and Toxic Stress." *Pediatrics* 129 (1). American Academy of Pediatrics: e232–e246.

Thai, Thuan Q, and Evangelos M Falaris. 2014. "Child Schooling, Child Health, and Rainfall Shocks: Evidence from Rural Vietnam." *Journal of Development Studies* 50 (7). Taylor & Francis: 1025–37.

Uganda Bureau of Statistics. 2016. "Uganda National Panel Survey - Wave V Report." Uganda Bureau of Statistics.

Walker, Susan P, Theodore D Wachs, Sally Grantham-McGregor, Maureen M Black, Charles A Nelson, Sandra L Huffman, Helen Baker-Henningham, et al. 2011. "Inequality in Early Childhood: Risk and Protective Factors for Early Child Development." *The Lancet* 378 (9799). Elsevier: 1325–38.

Webb, Matthew D. 2013. "Reworking Wild Bootstrap Based Inference for Clustered Errors." Queen's Economics Department Working Paper.

Wolf, Sharon, Peter Halpin, Hirokazu Yoshikawa, Amy Jo Dowd, Lauren Pisani, and Ivelina Borisova. 2017. "Measuring School Readiness Globally: Assessing the Construct Validity and Measurement Invariance of the International Development and Early Learning Assessment (Idela) in Ethiopia." *Early Childhood Research Quarterly* 41. Elsevier: 21–36.

Wooldridge, Jeffrey M. 2003. "Cluster-Sample Methods in Applied Econometrics." *American Economic Review* 93 (2): 133–38.

A Appendix A - Additional Tables

Table 16: Full regression tables, heteroskedastic robust SE's

Variable	Motor Skills	Early Literacy	Early Numeracy	Socio-emotional	Total IDELA
lnuterodiff	0.424	0.436	0.181	-0.0200	0.332
	(0.34)	(0.31)	(0.32)	(0.31)	(0.32)
lnyear1	0.173	0.635*	0.467	0.403	0.517
	(0.29)	(0.26)	(0.29)	(0.28)	(0.28)
lnyear2	0.194	0.621*	0.413	0.378	0.466
	(0.26)	(0.24)	(0.25)	(0.25)	(0.26)
1.district	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)
2.district	0.392***	-0.0284	-0.0759	-0.250*	0.0696
	(0.10)	(0.10)	(0.11)	(0.11)	(0.10)
3.district	0.316**	0.280*	0.171	-0.101	0.234*
	(0.12)	(0.12)	(0.12)	(0.11)	(0.12)
4.district	-0.0270	-0.729***	-0.338**	-0.553***	-0.439***
	(0.11)	(0.10)	(0.12)	(0.11)	(0.11)
5.district	0.238*	-0.236*	-0.0324	0.176	0.0716
	(0.10)	(0.09)	(0.10)	(0.10)	(0.10)
6.district	0.139	-0.696***	-0.497***	-0.537***	-0.414***
	(0.11)	(0.10)	(0.11)	(0.10)	(0.11)
7.district	-0.0131	-0.500***	-0.291**	-0.471***	-0.329**
	(0.11)	(0.10)	(0.10)	(0.10)	(0.10)
8.district	-0.0708	-0.189	-0.473***	-0.662***	-0.370***
	(0.11)	(0.10)	(0.10)	(0.11)	(0.10)
9.district	0.208	-0.492***	-0.510***	-0.569***	-0.329**
	(0.11)	(0.10)	(0.10)	(0.10)	(0.10)
3.age1	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)

Table 16: Full regression tables, heteroskedastic robust SE's *(continued)*

Variable	Motor Skills	Early Literacy	Early Numeracy	Socio-emotional	Total IDELA
4.age1	0.0149	0.00712	-0.0303	-0.0344	-0.00971
	(0.06)	(0.06)	(0.06)	(0.06)	(0.06)
5.age1	-0.0339	0.0168	-0.00396	0.0330	0.00316
	(0.07)	(0.06)	(0.07)	(0.06)	(0.06)
sex	0.272***	0.250***	0.166***	0.109*	0.254***
	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)
0.resp_educ	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)
1.resp_educ	-0.00843	0.0969	0.120	0.145	0.0986
	(0.09)	(0.08)	(0.08)	(0.08)	(0.08)
2.resp_educ	0.181*	0.275**	0.303***	0.237**	0.296***
	(0.09)	(0.09)	(0.08)	(0.09)	(0.09)
3.resp_educ	0.310*	0.557***	0.461**	0.332*	0.492**
	(0.15)	(0.15)	(0.16)	(0.14)	(0.15)
0.resp_occ	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)
1.resp_occ	0.0261	0.0588	0.0418	0.0465	0.0518
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
resp_age	0.00533*	0.00672**	0.00364	0.00235	0.00562*
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
0.HH_head_occ	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)
1.HH_head_occ	-0.137*	-0.0932	-0.0589	-0.0809	-0.116
	(0.07)	(0.06)	(0.07)	(0.07)	(0.07)
0.HH_head_education	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)
1.HH_head_education	0.0375	0.0941	0.0949	0.0260	0.0710
	(0.09)	(0.08)	(0.08)	(0.08)	(0.08)
2.HH_head_education	0.135	0.148	0.0538	0.00565	0.105
	(0.10)	(0.09)	(0.09)	(0.09)	(0.09)

Table 16: Full regression tables, heteroskedastic robust SE's *(continued)*

Variable	Motor Skills	Early Literacy	Early Numeracy	Socio-emotional	Total IDELA
3.HH_head_education	0.212	0.361**	0.313*	0.215	0.325*
	(0.13)	(0.12)	(0.12)	(0.13)	(0.13)
0.sib	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)
1.sib	-0.0736	-0.0177	-0.0260	-0.0324	-0.0465
	(0.08)	(0.07)	(0.08)	(0.08)	(0.08)
2.sib	0.0141	0.0353	0.0471	-0.00159	0.0292
	(0.09)	(0.08)	(0.08)	(0.08)	(0.08)
3.sib	-0.167	-0.0385	-0.0779	-0.171	-0.138
	(0.13)	(0.12)	(0.13)	(0.13)	(0.12)
hh_size	0.0248	0.0147	0.0189	0.0255	0.0253
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
0.business_ownership	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)
1.business_ownership	0.0670	0.0796	0.0554	0.0647	0.0846
	(0.05)	(0.04)	(0.05)	(0.04)	(0.05)
0.land_ownership	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)
1.land_ownership	-0.169**	-0.141**	-0.106*	-0.0752	-0.157**
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)
0.livestock_ownership	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)
1.livestock_ownership	-0.0388	0.00637	0.0321	0.00791	0.000463
	(0.05)	(0.04)	(0.05)	(0.05)	(0.05)
0.loans_outstanding	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)
1.loans_outstanding	0.0152	0.0402	0.0514	0.00345	0.0338
	(0.05)	(0.04)	(0.05)	(0.05)	(0.05)
2.loans_outstanding	0.0631	0.0277	-0.0150	0.0922	0.0570

Table 16: Full regression tables, heteroskedastic robust SE's *(continued)*

Variable	Motor Skills	Early Literacy	Early Numeracy	Socio-emotional	Total IDELA
	(0.10)	(0.08)	(0.10)	(0.10)	(0.10)
0.transfers	0	0	0	0	0
	(.)	(.)	(.)	(.)	(.)
1.transfers	0.00400	-0.0223	-0.00439	0.0414	-0.0000639
	(0.05)	(0.04)	(0.05)	(0.05)	(0.04)
_cons	-0.516**	-0.345*	-0.312	-0.0758	-0.419**
	(0.17)	(0.16)	(0.17)	(0.17)	(0.16)
N	2007	2007	2007	2007	2007
adj. R-sq	0.072	0.179	0.094	0.105	0.129
Standard errors in parentheses					
* p<0.05	** p<0.01	*** p<0.001			

Standard errors calculated using different methods - standardised scores, continuous shocks

Table 17 shows the standard errors calculated using heteroscedasticity robust errors, wild bootstrapping (which was used in the table above) and clustering at the sub-county level. The coefficients on year 1 and year 2 rainfall shocks on early literacy scores are significant when using the robust option and when clustering standard errors at the sub-county level. The coefficient on year 2 shocks on the total literacy score is also significant across these methods. Note however that these coefficients are not significant when standard errors are calculated using wild bootstrapping. In general, the standard errors are much larger with the wild bootstrap method in comparison to the standard errors with the exception of the year 2 variable's standard errors which reduce in magnitude in comparison to other clustering methods.

Table 17: Comparison of inference statistics using different methods

Clustering method	Variable	Motor Skills	Early Literacy	Early Numeracy	Socio-emotional	Total IDELA
Robust	In-utero	0.34	0.31	0.32	0.31	0.32
	Year 1	0.29	.26**	0.29	0.28	.28*
	Year 2	0.26	.24***	0.25	0.25	.26*
Wild Bootstrap	In-utero	0.56	0.56	0.47	0.42	0.51
	Year 1	0.61	0.77	0.61	0.28	0.61
	Year 2	0.61	0.66	0.53	.16**	0.56
Sub counties	In-utero	0.39	0.32	0.32	0.31	0.34
	Year 1	0.37	.36*	0.39	0.29	0.38
	Year 2	0.29	.27**	0.24*	0.23	.47*

Note:
* = 90%, ** = 95%, *** = 99% significance

Standard errors calculated using different methods - standardised scores, binary shocks

Table 18: Standard errors - binary shock indicators on age standardised scores

Clustering method	Shock Type	Variable	Motor Skills	Early Literacy	Early Numeracy	Socio-emotional	Total IDELA
Robust	Good shocks	In-utero	0.06***	0.06	0.06	0.06	0.06
		Year 1	0.06	0.06	0.06	0.07	0.06
		Year 2	0.07*	0.07	0.07	0.07	0.07
	Bad shocks	In-utero	0.08	0.07	0.08	0.07	0.07
		Year 1	0.06	0.06	0.07	0.06	0.06
		Year 2	0.06	0.06***	.06*	.05**	.06**
Wild Bootstrap	Good shocks	In-utero	.09*	0.08	0.09	0.11	0.08
		Year 1	0.14	0.15	0.12	0.11	0.14
		Year 2	.07**	0.09	0.1	0.05*	0.08
	Bad shocks	In-utero	0.14	0.14	0.09	0.06	0.11
		Year 1	0.07	0.11	0.1	0.04*	0.09
		Year 2	0.1	0.1*	0.07	0.09	0.09
Sub counties	Good shocks	In-utero	.06***	0.06	0.06	0.07	0.06
		Year 1	0.08	0.08	0.07	0.08	0.08
		Year 2	.07*	0.07	0.06	0.07	0.07
	Bad shocks	In-utero	0.1	0.08	0.06	0.06	0.08
		Year 1	0.07	0.06	0.09	0.06	0.07
		Year 2	0.06	.06***	.06*	0.07	.06*

Unstandardised scores, comparison with skips counted as NA's

In the table below, the first row reports point estimates and significances using the heteroscedasticity-robust option in STATA. The second row reports the significance & SE (95% CI) using the Wild Bootstrap technique to find SE's that can be more accurate given the heteroscedasticity of unknown form and a small number of clusters. Note that the point estimates remain the same, but the SE changes quite a bit with the exception of one case - socio-emotional scores, which are the only results that retain their significance after the more rigorous wild-bootstrap test.

The effects of rainfall shocks on socio-emotional IDELA scores are small but significant (at the 5% level for year 1 shocks and at the 10% level for year 2 shocks). A 1% positive rainfall shock in year 1 of a child's life seems to cause the socio emotional score to increase by 0.1 percentage points. Given that the median score on this section of the IDELA is .27 (max: 1), this points to a 3.7% increase in the score if there was a rainfall shock of 1% in the first year of a child's life.

Table 19: Regression results

						Consider skip as NA				
Variable	Motor	Early Literacy	Early Numeracy	Socio-emotional	Total IDELA	Motor	Early Literacy	Early Numeracy	Socio-emotional	Total
In-utero	0.119	0.081	0.021	0	0.055	0.097	0.081	0.013	0.014	-0.017
SE - WB	*0.132*	*0.106*	*0.081*	*0.069*	*0.08*	*0.134*	*0.129*	*0.066*	*0.08*	*0.085*
Year 1	0.043	0.119	0.075	0.054	0.073	0.02	0.098 *	0.104 *	.101 **	0.074
SE - WB	*0.132*	*0.147*	*0.093*	*0.039*	*0.091*	*0.121*	*0.159*	*0.103*	*(0.049)*****	*0.107*
Year 2	0.067	0.083	0.059	.055***	0.066	0.075	0.119 **	.08 *	.069 *	0.078
SE - WB	*0.13*	*0.124*	*0.089*	*0.03*	*0.084*	*0.124*	*0.12*	*0.099*	*(0.039)**	*0.09*
Settings										
n	2007	2007	2007	2007	2007	1960	1484	1408	1826	1117
HH controls	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
District controls	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Child controls	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
HH Income	N	N	N	N	N	N	N	N	N	N
SE method	Wild bootstrap	Wild bootstrap	Wild bootstrap	Wild bootstrap	Wild bootstrap	Wild bootstrap	Wild bootstrap	Wild bootstrap	Wild bootstrap	Wild bootstrap
R-squared	0.41	0.43	0.37	0.28	0.46					

Note:
This table contains regression point estimates with SE's calculated as specified in the SE method row. * = 90% significance, ** = 95% significance

Standard errors - different approaches -

Table 20: SE differences

							Completed all sections			
Clustering method	Variable	Motor	Early Literacy	Early Numeracy	Socio-emotional	Total IDELA	Motor	Early Literacy	Early Numeracy	Socio-emotional
Robust	In-utero	0.076	0.057	0.051	0.048	0.048	NA	NA	NA	NA
	Year 1	0.065	.048**	0.046	0.042	.041*	NA	NA	NA	NA
	Year 2	0.059	.043*	0.040	0.039	.038*	NA	NA	NA	NA
Wild Bootstrap	In-utero	0.132	0.106	0.081	0.069	0.08	NA	NA	NA	NA
	Year 1	0.132	0.147	0.093	0.039	0.091	NA	NA	NA	NA
	Year 2	0.130	0.124	0.089	.03 ***	0.084	NA	NA	NA	NA
Sub counties	In-utero	0.093	0.06	0.052	0.049	0.052	NA	NA	NA	NA
	Year 1	0.812	0.068*	0.063	0.044	0.057	NA	NA	NA	NA
	Year 2	0.066	0.054	0.041	0.036	0.04	NA	NA	NA	NA

Note:
This table contains SE's and corresponding significances calculated using different clustering methods. * = 90% significance, ** = 95% significance

B Appendix B - R Code to retrieve rainfall data

This R code snippet can be used to extract rainfall data from the ARC-2 data set using the rnoaa package. The variables for start and end dates, and lat-long boundaries need to be defined to restrict the amount of data downloaded. The connection tends to time out after a while, so it is advisable to download the data in chunks rather than in one pull.

```
library(tibble)
library(dplyr)
library(rnoaa)
## Loop to extract data: Crop first

### Retrieve data for several years
### Construct long term averages
### Short term aggregates (monthly?) for the periods of interest
start.date <- as.Date(c("1990-01-01"))
end.date <- as.Date(c("2019-01-01"))
end.date - start.date
# GPS min lat: -.52 max lat: 1.15 min long: 29.73 max long: 34.26
# Set GPS boundaries
start.lat <- -1
end.lat <- 1.5
start.long <- 29
end.long <- 35

rainfall.tib <- tibble(date = NA,
                       lat = NA,
                       lon = NA,
                       precip = NA)

# j <- 1
precip.df = tibble(date = as.Date(character(0)),
                lon = numeric(0),
                lat = numeric(0),
```

```
                precip = double(0))
for(i in start.date:end.date) {
temp <- arc2(as.Date(i,origin = "1970-01-01"))
temp <- temp[(temp$lat > start.lat) & (temp$lat < end.lat) &
               (temp$lon > start.long) & (temp$lon < end.long), ]
temptbl <- tibble(date = rep(as.Date(i,origin = "1970-01-01"),nrow(temp)),
                  lon = temp$lon,
                  lat = temp$lat,
                  precip = temp$precip)
precip.df <- bind_rows(precip.df,temptbl)

cat("Processing: ", i, "\n")
}

save(precip.df,file="precip.dftemp.Rda")
```

C Appendix C - IDELA questions

Socio-emotional skills

- Self-awareness: child's name, caretaker's name, village and country name
- Emotional awareness: Identification of situations and linking them with basic emotions
- Empathy: Ability to identify/recognise emotions in others and suggest solutions
- Conflict resolution: Identify potential solutions to conflicts
- Social skills: Ability to recall friend names

Motor skills

- Copy shapes
- Draw a person
- Fold a piece of paper
- Hopping

Early Literacy

- Expressive vocabulary: name common HH items and colours
- Print awareness: Recognise and identify text
- Letter recognition
- Phonemic awareness
- Oral comprehension: Ability to understand and recall a short statement
- Writing

Early Numeracy

- Compare sizes of circles and lengths of sticks
- Sort cards in ascending or descending order
- Identify basic geometric shapes
- Number recognition
- Addition and subtraction
- Solving a simple jigsaw puzzle