

# Agricultural child labour in response to climate change\*

Author: Mathias Weidinger (i6113552)<sup>†</sup>

Supervisor: Prof. Dr. Pierre Mohnen<sup>‡</sup>

September 10, 2021

## Abstract

Anthropogenic climate change severely impacts agriculture in developing countries and with it the millions of children actively involved in that sector. Economic theory can be used to characterise the effects that climate change might have on agricultural child labour. Merging two strands of theoretical literature, from Environmental Economics and Development Economics, this paper presents a simple household decision model of child labour supply under climate change. To take the model to the data, I combine a household panel from Nigeria with geo-coded weather records and estimate the dose-response functions for temperature and precipitation. Higher mean-temperatures are estimated to increase climate change, whereas the effects of extreme weather events are mixed and highly non-linear.

*JEL*: D11, D12, D13, J22, J43, I32, Q54

*Keywords*: child labour, climate change, agriculture, global warming

---

\*This thesis constitutes the final deliverable for the Masters degree in Public Policy and Human Development, with a specialization in Social Protection Policy, at UNU-MERIT and Maastricht University.

<sup>†</sup>Student in the M.Sc. Public Policy and Human Development programme, cohort September 2018, Maastricht University and UNU-MERIT, [m.weidinger@maastrichtuniversity.nl](mailto:m.weidinger@maastrichtuniversity.nl).

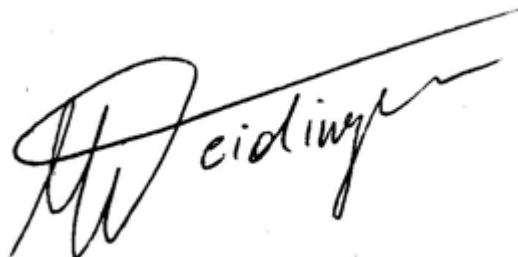
<sup>‡</sup>Professor of Microeconometrics at Maastricht University, Professorial Fellow at UNU-MERIT.

## Declaration of Academic Integrity

I, Mathias Weidinger, hereby declare with relation to my master's thesis *Agricultural child labour in response to climate change* that:

- I am aware of and have understood the rules and regulations stipulated in the Education and Examination Regulations (EER) regarding fraud and plagiarism;
- I am aware of the possible consequences and disciplinary measures in the case of fraud and plagiarism in my master's thesis;
- I have conducted myself in accordance with the Thesis Guidelines, Education and Examination Regulations and generally established standards of academic integrity in writing my master's thesis;
- I have carefully marked and referenced all direct quotes and references, and indirect quotes included in my master's thesis;
- my master's thesis is an original result of my own work and does not include the work of others except in the case of direct and indirect quotes that are recognizable as such.

Signed at Seewalchen am Attersee on September 10, 2021.

A handwritten signature in black ink, appearing to read 'M Weidinger', with a long horizontal stroke extending to the right.

# Dedication

I would like to express my heartfelt gratitude to my peers, colleagues, and mentors. Without your input this thesis would not have seen the light of day.

To Professor Pierre Mohnen for his patience and guidance.

To Patrick and Nina for sparking my interest in environmental economics, and to Oral for reminding me not to believe all stories economists might tell me.

A los jóvenes de la Feria Libre, a la Fundación PACES, a VoloB, y a mis hermanos de la comunidad Yanuncay - por plantar en mí el deseo de seguir su ejemplo de solidaridad vivida y humanidad sincera. Este labor no existiera si no fuera por la inspiración que me brindaron Ustedes.

To my friends through space and time - life, like research, would be dull without you.

To Nanaya, for believing in me beyond reason.

To Mum and Dad for everything.

## List of abbreviations

**2SLS** Two-stage Least Squares; **ARE** Agricultural, Resource and Environmental Economics; **CDD** Correlation Decay Distance; **CLD** Cloud Cover, CRU-TS; **CRE** Correlated Random Effects; **CRU-TS** Climatic Research Unit gridded Time Series; **DTR** Diurnal Temperature Range, CRU-TS; **EA** Enumeration Area, GHG; **EIA** Environmental Impact Assessment; **EKC** Environmental Kuznets Curve; **FAO** Food and Agricultural Organization of the United Nations; **FE** Fixed Effects; **FRS** Freezing Days, CRU-TS; **GCM** Global Climate Model; **GDP** Gross Domestic Product; **GHG** Green House Gas; **GHS** General Household Survey, Nigeria; **HIC** High Income Country; **HYV** High Yield Variety; **ILO** International Labour Organization; **IMF** International Monetary Fund; **IPCC** International Panel on Climate Change; **IPEC** International Programme on the Elimination of Child Labour, ILO; **LGA** Legislative Government Area, Nigeria; **LMIC** Low- and Middle Income Country; **LSMS** World Bank's Living Standards Measurement Surveys; **MPI** Multi-dimensional Poverty Index; **NAERLS** National Agricultural Extension and Research Liaison Services, Nigeria; **NBS** National Bureau of Statistics, Nigeria; **OLS** Ordinary Least Squares; **OPHI** Oxford Poverty and Human Development Initiative; **PET** Potential Evapotranspiration, CRU-TS; **PRE** Precipitation, CRU-TS; **SDG** Sustainable Development Goal; **SSA** Sub-Saharan Africa; **TMN** Minimum Temperature, CRU-TS; **TMP** Mean Temperature, CRU-TS; **TMX** Maximum Temperature, CRU-TS; **UNCRC** United Nations Convention on the Rights of the Child; **UNDP** United Nations Development Program; **UNICEF** United Nations International Children's Emergency Fund; **USD** United States Dollar; **VAP** Vapour Pressure, CRU-TS; **WET** Wet Days, CRU-TS;

# Contents

<b>1</b>	<b>Introduction</b>	<b>6</b>
<b>2</b>	<b>Literature</b>	<b>9</b>
2.1	Child labour . . . . .	10
2.2	Agriculture and labour supply . . . . .	15
2.3	Economic effects of climate change . . . . .	17
<b>3</b>	<b>Model</b>	<b>21</b>
<b>4</b>	<b>Case selection and data processing</b>	<b>27</b>
4.1	Climate and agriculture in Nigeria . . . . .	28
4.2	Data sources . . . . .	30
4.3	Variables . . . . .	31
4.3.1	GHS variables . . . . .	32
4.3.2	Climate variables . . . . .	34
4.4	Point value extraction from the 2D grid . . . . .	36
<b>5</b>	<b>Econometric strategy</b>	<b>38</b>
5.1	Recovering local and instantaneous non-linear effects . . . . .	39
5.2	Correcting for selection bias . . . . .	43
<b>6</b>	<b>Findings</b>	<b>47</b>
<b>7</b>	<b>Conclusion</b>	<b>52</b>
	<b>References</b>	<b>55</b>
<b>A</b>	<b>Annex: Extensions to the theoretical model</b>	<b>68</b>
A.1	Easing assumptions on time allocation . . . . .	68
A.2	Children’s productivity and adult supervision . . . . .	68

# 1 Introduction

Globally, nearly one in ten children are subjected to child labour, predominantly in low and middle income countries (LMICs). This fraction doubles when focusing on Africa alone. On any given day in 2016, 152 million African children aged five to seventeen were in some form of child labour; half of them worked under hazardous conditions ([ILO, 2017](#)). The consequences of this are well-understood:

Child labour can result in extreme bodily and mental harm, and even death. It can lead to slavery and sexual or economic exploitation. And in nearly every case, it cuts children off from schooling and health care, restricting their fundamental rights and threatening their futures ([UNICEF, 2020](#)).<sup>1</sup>

Unfortunately, Child labour is notoriously difficult to regulate. This is in part due to its informality: Most child labour is unpaid and takes place far off the formal labour market, on family farms or in family owned enterprises. Consequently, the compliance costs of anti-child labour legislation are so high that even where such laws are in place, enforcement tends to be lax, if not entirely absent.<sup>2</sup> Despite sustained international efforts to eliminate it, child labour decreased by only one percent in 2008-2012 and the [ILO \(2017\)](#) estimates that at the current rate of progress, 121 million children will still be working in 2025.<sup>3</sup>

Disaggregating the data reveals that 59 per cent of working children in Africa are between 5 and 11 years old and that almost nine in ten of them work in agriculture - 61.4 million children in absolute terms. The uniquely high concentration of child labour in agriculture - particularly in Sub-Saharan Africa (SSA) - begs a question: What happens to these children if agriculture becomes an increasingly unstable sector? Do they work

---

<sup>1</sup>For a recent review of the adverse health outcomes of child labour, see [Ibrahim et al. \(2018\)](#).

<sup>2</sup>Note that this is not necessarily due to negligence by the government. Most countries with high rates of child labour also suffer from limited fiscal capacity and, consequently, prohibitively small budgets.

<sup>3</sup>Article 32 of the 1989 UN Convention on the Rights of the Child (UNCRC) recognises the right of every child to be protected from economic exploitation and from performing work that is hazardous or harmful to their health and development or that interferes with their education. The ILO conventions on minimum age for admission to employment (no. 138, ratified by 173 countries) and on the the elimination of the worst forms of child labour (no. 182, ratified by 187 countries) are binding international agreements to this effect. Targets 8.7 and 16.2 of the United Nation's Sustainable Development Goals additionally target child labour, pledging to end it in all its forms by 2025.

more or less hours on average? Does their work become more or less hazardous? How is their overall welfare affected? These concerns bear unprecedented relevance in the wake of anthropogenic climate change.

There is mounting evidence that the human-induced accumulation of Green House Gases (GHG) in the earth's atmosphere has caused the global climate to change and will continue to do so in the coming centuries (Pachauri et al., 2015). The occurrence of many new record highs in recent decades suggest that the earth's mean temperature is rising. Munasinghe et al. (2012), for instance, show that the frequency of extremely high temperatures across the global landmass increased tenfold between the beginning of the twentieth century and 1999–2008. The high frequency of new record lows over the same period suggests a simultaneous rise in the variance of temperature (Auffhammer and Schlenker, 2014). As a consequence, climate change persistently increases the probabilities of extreme weather events such as droughts, floods, snow storms, heat waves, cyclones, and hurricanes (Pachauri et al., 2015). These changes in the natural environment profoundly alter the setting for human economic activity on planet earth.

It is of course true that climate change affects different regions in different ways. This heterogeneity has inspired some to speculate, whether the negative impacts experienced in some areas could be offset by positive effects elsewhere<sup>4</sup>, or be it in strictly economic terms. Tol, p. 34, for instance, notes that “although the world population is concentrated in the tropics, where the initial effects of climate change are probably negative, the relatively smaller size of the economy in these areas means that gains for the high-income areas of the world exceed losses in the low-income areas”. He later qualifies this statement, adding that, in the long term, warming above 1-2 degrees will likely have negative total effects. But even for the short-term, arguments like Tol's ought to be read with some scepticism. They are not inherently incorrect, but their accounting method entirely ignores the distributional

---

<sup>4</sup>For instance, while agricultural yields are likely to decrease in SSA, some conjecture that large parts of northern Eurasia, which are currently covered by permafrost, could become arable as temperatures rise. Moreover, an elevated concentration of atmospheric CO<sub>2</sub> reduces water stress in plants and may make them grow faster. This phenomenon, known as CO<sub>2</sub>-stimulation, and its beneficial effect on crop yields remain understudied and, therefore, unincorporated in most projections (e.g., in Schlenker and Lobell, 2010).

consequences of climate change (and indirect economic effects that stem from them).

The world's richest countries, which are predominantly situated in the global North, are responsible for almost eighty percent<sup>5</sup> of GHG emissions in 1850-2011 and high income countries continue to emit the most GHGs by far<sup>6</sup> (Ritchie and Roser, 2017). Due to their location, high income countries also experience more benign effects of climate change compared to their more climate-vulnerable counterparts further south, which face the harshest consequences. Intensified temperature extremes, precipitation anomalies, and natural disasters are all projected to disproportionately afflict Latin America, South Asia and the Pacific, and Sub-Saharan Africa. Together, their populations account for the vast majority of the world's population and, in particular, the world's poor. Thus, Climate Change exacerbates existing global inequalities and there is an argument to be made for distributional considerations beyond simple aggregate cost-benefit analysis. Distributional aspects are highly relevant for migration policy, international cooperation, and conflict-prevention efforts - not to mention global solidarity.

In SSA alone, the effects of climate change are responsible for at least 1,000 deaths, 13 million people seriously afflicted<sup>7</sup>, and 520 USD million in direct economic damages since 2000. One-third of the world's droughts occur in SSA, and the frequency of storms and floods is growing fastest in this region (IMF, 2020). Agricultural yields, meanwhile, are projected to fall significantly for five of the continent's most important food crops: by 22 percent for maize, 17 percent for Sorghum, Millet and Groundnut, and eight percent for cassava (Schlenker and Lobell, 2010). Similar negative impacts have been projected for major cash-crops like coffee (Craparo et al., 2015)<sup>8</sup> or cocoa (Boeckx et al., 2020). These findings place agriculture at the centre of climate vulnerability, and with it the millions of children working on farms throughout the world.

---

<sup>5</sup>The percentage is even higher when accounting for final consumption of goods, since great part of GHG emissions in low and middle income countries stem from the production of export goods, predominantly destined for high income countries.

<sup>6</sup>Note, however, that relative to 1990 most high income countries now have lower annual growth rate of emissions whereas LMICs have increased their emissions substantially.

<sup>7</sup>This includes individuals who were injured, left homeless, food insecure, or lacking water and sanitation.

<sup>8</sup>For contrasting findings, based on CO<sub>2</sub>-stimulation effects, see DaMatta et al. (2019).



Even so, there is virtually no literature investigating the implications of climate change for working children. The little empirical evidence that exists displays three common flaws: limited data, misspecification, and a lack of theoretical underpinning. These shortcomings and the policy-relevance and urgency<sup>9</sup> of the child labour debate provide a clear motivation for further research. This paper analyses the effects of climate change on the prevalence of child labour in agriculture, both from a theoretical and an empirical perspective: How does climate change - through persistent shocks to temperature and precipitation - affect children's labour supply?

After a detailed overview of the literature in section 2, section 3 extends the economic theory of child labour to account for a sustained shock to agricultural household production due to climate change. The resulting model produces a directional hypothesis that can be tested empirically: Negative climate shocks to agricultural production decrease children's work hours. Section 4 provides arguments for Nigeria as a case study and presents appropriate data for the empirical analysis. The data chosen mend the shortcomings noted in the literature; a longitudinal data set from Nigeria (2010-2016) is combined with geocoded weather data from suitable ground level stations, supplemented by satellite-based estimates. Section 5 outlines the empirical problem and describes the econometric framework to tackle it. The findings drawn from this analysis are discussed in section 6. The results offer novel and relevant input for weighing the costs and benefits of GHG emission abatement – not only monetarily but through the lens of child welfare. Section 7, addresses some caveats, provides ideas for further research, and concludes.

## 2 Literature

The topic at hand lies at the intersection of various strands of literature: Development and Labour Economics have provided some models of child labour. Environmental and Resource Economics feature a relatively recent and fast-growing literature on the economic

---

<sup>9</sup>Given the proximity of the 2025 deadline, scoped by SDG 8, to eliminate all forms of child labour, we might soon see the window of opportunity for comprehensive global policy action closing.

effects of climate change. Lastly, Agronomy has long studied the effects of weather shocks on crop yields. This section offers an overview of the first and the second of these fields, subsuming the third within the more recent climate studies, which draw heavily from it.

## 2.1 Child labour

The first economic texts to analyse child labour in detail date back to the late eighteenth and the early nineteenth centuries. Among those who discussed its prevalence in Europe during the Industrial Revolution were Smith, Malthus, Engels and Marx (Edmonds, 2007).<sup>10</sup> About a century later, the inception of human capital theory (Mincer, 1958; Schultz, 1961; Becker, 1964) fundamentally transformed the research agenda in labour economics and, as a corollary, gave some impetus to the study of child labour (see e.g. Rosenzweig and Evenson, 1977). The seminal contribution by Basu and Van (1998) eventually launched an enduring proliferation of the literature on child labour (Edmonds, 2007) in both theoretical and applied economics.<sup>11</sup>

Child labour is commonly modelled as a product of a constrained optimization exercise, undertaken as if each household were one collective decision maker. In practice, this implies that benevolent<sup>12</sup> household heads take decisions for the entire household in a nearly utilitarian fashion.<sup>13</sup> Basu and Van (1998) proved so influential because they formalized two long-held conjectures in axioms that became the bedrock of child labour theory.

The first of them, termed the ‘luxury axiom’, characterises parents’ preferences over child labour as lexicographic: child labour occurs if and only if families cannot cover their subsistence needs without it.<sup>14</sup> This characterization of preferences would imply a strictly

---

<sup>10</sup>With the exception of Malthus, the early economists viewed children as investment goods. Marx, in particular, remarked that “all family ties among the proletarians are torn asunder, and their children transformed into simple articles of commerce and instruments of labour” (Marx and Engels, 1848).

<sup>11</sup>For a recent overview, see the dedicated volume by Posso (2020).

<sup>12</sup>On the credibility and extent of such benevolence and its implications, see Bhalotra (2002).

<sup>13</sup>Nearly because child leisure is strictly preferred to child work and it is unclear how this preference affects the household’s aggregate utility.

<sup>14</sup>This axiom obtains its name from that fact that, from the household’s perspective, child leisure is a luxury good, more of which is consumed as income rises. An early proponent of this inverse relationship, Thomas Malthus noted that the prevalence of child labour in the late 18th century was proof that families were unable to meet their most basic needs (see Edmonds, 2007).

negative relationship between household income and the amount of child labour supplied by a household - a testable hypothesis for which [Basu and Van \(1998\)](#) have drawn substantial criticism (see e.g. [Edmonds and Schady, 2012](#)). Most notably, [Bhalotra and Heady \(2000, 2003\)](#) observe that children of land-rich families are often more likely to be in work than those of land-poor households.

Subsequent replications of this observation spurred numerous attempts at solving the ‘wealth paradox’ that had become apparent in child labour theory. [Bhalotra and Heady \(2003\)](#) attribute their findings to imperfections in the labour and credit markets as well as household size: If labour markets are imperfect, child labour is increasing in farm size and decreasing in household size, while access to credit spreads the effects out over time (viz., consumption smoothing). Others attempt to resolve the paradox by altering the original luxury axiom. [Basu et al. \(2010\)](#), for instance, propose that the relationship between land wealth and child labour is not monotone - positive or negative - but follows an inverted U-shape. Therefore, child labour will eventually decrease as households become wealthier. Initially, however, an increase in land wealth increases both income and the marginal benefit from child labour. It is only after a structural threshold, beyond which child labour’s marginal cost outweighs its benefits, that the relationship reverses direction and child labour decreases for good. This threshold could, for instance, be the income at which the household can afford to contract external labour. <sup>15</sup>

[Dwivedi and Marjit \(2017\)](#) propose another version of the luxury axiom by which child labour decreases in *relative* poverty. This allows for child labour to rise even as every household is made wealthier, as long as the relative distance in wealth between them increases. Their observation that Pareto-dominance may not fully characterise poverty mirrors an ongoing debate.<sup>16</sup> Absolute measures of poverty are prone to underestimation

---

<sup>15</sup>Note that this narrative - things need to worsen before they can get better - resembles the well-known tale of growth and inequality described by the Kuznets Curve ([Kuznets, 1955, 1963](#)) and that of growth and pollution described by the Environmental Kuznets Curve (EKC). Notwithstanding their popularity, these narratives remain empirically unsubstantiated. See, however, [Piketty \(2014\)](#) for evidence against the Kuznets curve, and [Mills and Waite \(2009\)](#) and [Özoku and Özdemir \(2017\)](#) for similar results regarding the EKC.

<sup>16</sup>For poverty measurement techniques of both absolute and relative poverty, as well as their distinction, see [Sen \(1976\)](#) and [Ravallion \(2020\)](#).

and run risk to ignore the amplification of human misery that springs from inequality. Relative poverty indicators, on the other hand, are only strictly relevant in conjunction with absolute measures - a millionaire surrounded by billionaires is not ‘poor’ in the common sense of the word.<sup>17</sup> While hybrid measures, like ‘weakly relative poverty’ (Ravallion and Chen, 2009) exist, I am not aware of attempts to characterise child labour responses to varying degrees of absolute poverty and income inequality in an integrated framework.<sup>18</sup> Contested to this day, the wealth paradox and the luxury axiom continue to inspire this type of research.<sup>19</sup>

As for their second premise, Basu and Van (1998) model child labour as a substitute for adult labour from the firm’s perspective - a notion captured in their ‘substitution axiom’. The degree of such substitutability has been debated as well. Due to their lack in experience and bodily development, children are commonly thought to be less suited for most jobs in terms of their productive potential. Bar and Basu (2009) model children to be productive (to some degree) only under adult supervision and entirely unproductive otherwise. While it extends the theory beyond the substitution axiom, their model is likely too restrictive. Toddlers aside, children - including those of very young ages - are capable of performing many simple tasks, making them productive labourers in their own right. This considered, the extensions proposed in annex A.2 allow for a more realistic characterisation of substitution with and without supervision.

As the literature continues to append Basu and Van’s model, it also outgrows it. Many of the more recent studies pay less attention to resolving the wealth paradox or characterising substitution between adult and child labour, but explore instead how child labour is affected through other channels. A natural extension is to explicitly model the allocation of children’s time between work and education in terms of opportunity cost. Dynamic models, owing much to human capital theory, describe how households weigh

---

<sup>17</sup>This is assuming that the currency used is not a highly inflated one (e.g. Euros rather than Venezuelan Bolívars) and that price levels are comparable to those in the real world.

<sup>18</sup>D’Alessandro and Fioroni (2016) show how political support for child labour regulation is shaped by inequality as it affects the welfare of low skilled adults differently than that of highly skilled ones.

<sup>19</sup>For more examples see e.g. Dumas (2007), Del Carpio (2008), Del Carpio and Loayza (2012), Edmonds and Schady (2012), Sarkar and Sarkar (2016), Oryoie et al. (2017), and Noack (2019).

the immediate short term benefits of income-generating child labour against children's future earning ability that increases in education (see e.g. [Bar and Basu, 2009](#); [Pal and Saha, 2012](#); [Dendir, 2014](#); [Edmonds and Shrestha, 2014](#); [Chakraborty and Chakraborty, 2018](#)). This crucially depends on the current extent of deprivation, on the rate by which parents discount the future, on whether they view their children's future earnings as a consumption smoothing mechanism for themselves (e.g. during old age), and also on their altruism towards their children.<sup>20</sup>

Another extension is to model the role of markets. [Baland and Robinson \(2000\)](#), [Bhalotra and Heady \(2003\)](#), and [Dumas \(2007\)](#) investigate how imperfections in the credit and (adult-)labour markets could cause households to deploy relatively more or less child labour. [Basu et al. \(2010\)](#) rely, in part, on the absence of functioning labour markets to arrive at their inverted U-shape. [Dumas \(2013, 2015, 2020\)](#) characterises a whole range of scenarios by selectively switching markets on and off one at a time and observing the effect on child labour in her model. To date, the literature largely confirms the moderating effects of labour markets and, less clearly so, the deferring effects of credit markets. Overall, child labour theory has evolved into a framework, capable of analysing child labour in relation to the wider micro-economic context that surrounds it.

Finally, the studies that are closest related to mine focus on the effects of external shocks on such a system. [Dumas \(2015, 2020\)](#) uses rainfall as a shock on household production in agriculture and, thus, indirectly on the households' child labour supply decision. She finds that, in the case of Tanzania, child labour increases in rainfall and that this increase is attenuated if the household has access to a well-functioning labour market. Moreover, her findings indicate that credit markets are less effective in smoothing such rainfall shocks. Manual labour provision can solve the household's problem contemporaneously (harvesting or sowing before the crops rot or the ground dries up), whereas consumption smoothing by credit only protracts costs into the future as long as the underlying problem remains unsolved. Of course, access to credit is usually quite limited in

---

<sup>20</sup>Knowing that they may well die before the fruits of child education can be reaped, purely self-interested parents may prefer sending their children to work and enjoy additional income presently.

LMICs and particularly so among the poorest. Furthermore, there is only so much credit a poor farming household can take up, assuming there is a functioning market, before it is worse off than before. Lastly, climate change is neither idiosyncratic nor transient, which makes it virtually impossible to insure against.

[Boutin \(2014\)](#) is the only paper to date which explicitly studies the link between climate change and child labour. Boutin’s approach is, however, less direct than that used in [Dumas’](#) rain-shock study. She measures climate vulnerability rather than effectuated climate change itself. Her independent variable is an aggregate of two indices that “take into account the multidimensionality of climate change” ([Boutin, 2014](#), p. 5); one index on biophysical vulnerability defined over climate volatility, variability, landscape typology and soil structure, and another index measuring a community’s adaptive capacities, which takes into account diversification strategies, financial capacities, and community-level amenities. One advantage of using such a composite index is that it overcomes potential endogeneity problems:

Households whose income and assets are most vulnerable to climate damage will benefit more from income diversification generated by child labour than other households. Consequently, they might be more likely to have a child working that can smooth consumption in the event of a severe climate shock. This would also be the case if risk-adverse households are more likely to have working children in order to diversify income sources and are also more likely to invest in self-protection mechanisms ([Boutin, 2014](#), p.5).

Endogeneity of this kind aside, concerns about measurement error as well as the cross-sectional nature of [Boutin’s](#) data still complicate causal inference. Her climate vulnerability index is certainly a valuable contribution, especially when it comes to designing forward looking adaptation policies. Much of the pertinent policy debate, however, centres around quantifying the impacts of exposures to adverse effects of climate change rather than the human response to an increase in vulnerability per se.<sup>21</sup> Therefore, this paper focuses on the effects of eventuated climate change - e.g., more frequent episodes of drought or flooding and persistent shifts in mean temperature and precipitation - on child labour.

---

<sup>21</sup>On the attribution of a single event to climate change, see [Hansen et al. \(2014\)](#).

Although [Boutin](#)'s metrics are different from the ones deployed here (vulnerability towards climate change versus actual, exogenously apportioned exposure to its negative effects), her findings provide a benchmark against which to compare the present study. She finds that climate vulnerability negatively affects child labour incidence and intensity, while it does not seem to have an impact on household chores. Thus, she concludes that child labour is an adjustment variable to local labour market conditions but not correlated with a given communities' joint climate resilience ([Boutin, 2014](#)).

The identification approach in this paper is much closer related to the rain induced shock used in [Dumas \(2020\)](#); I use multiple weather-component variables to shock the system and relate them back to changes in the climate. To guide the empirical strategy necessary, the next section provides a brief overview of the literature on climate change and its effects on the socio-economic sphere.

## 2.2 Agriculture and labour supply

In the field of Agricultural, Resource and Environmental Economics (ARE) there is a long history of using weather measures as explanatory variables in statistical models of agricultural productivity (e.g. [Fisher, 1925](#)). The reason for this is simple: Setting labour, machines, and fertilizer aside, the inputs to agricultural production are biophysical variables like soil nutrients, sunlight, rainfall, temperature, or pests. As such, a change in weather affects them much more immediately than it would affect most inputs in other sectors. But there are also considerable social and economic factors at play; Crop choice, the use of fertilizers or pesticides, capital and machinery used to till and harvest, and long-term planning of plot use rotation are all examples of human decision variables that interact with the aforementioned physical ones. The exact form of these human-nature interactions are subject to economic decision-making and depend on the farmers' preferences as well as on their constraints, environmental and otherwise. [Figure 1](#) schematically depicts such a production process from the farmer's perspective.

Labour supply depends in many ways on all the other factors to fall in place first: No

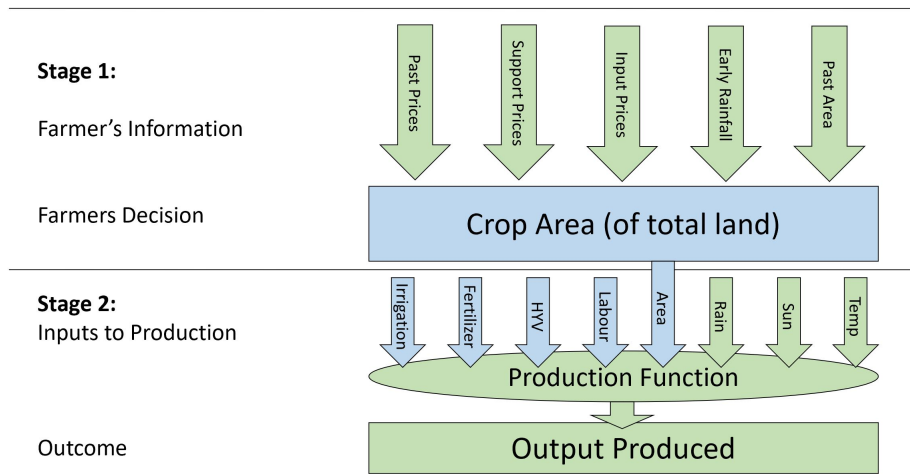


Figure 1: The farmer's problem.

*Note:* At stage 1, the farmer observes relevant information, the realization of which is exogenous. Based on this, the farmer apportions some land for planting a given crop. In stage 2, inputs enter the production function (HYV = high yield variety). Inputs depicted by blue arrows are choice variables for the farmer, green ones are not.

*Source:* Adapted from [Auffhammer \(2014\)](#).

additional harvesters are needed if, for example, a drought wipes out your crops before they are ripe. It is this connection between the physical conditions for agriculture and the human labour inputs to it, which may prove helpful for exploring the more specific characteristics of child labour supply in that sector.

One consequence of the complex nature-human interactions that characterise agricultural production is that the productivity of labour cannot be observed directly by the executive farmer, even if he were to monitor all workers constantly. Indirect measures of productivity, such as the number of tasks performed per worker, are also unavailable for most of the productive cycle: At harvest, one can judge the amount of produce harvested per worker and, at planting, the quantity of seeds sown or the area of land covered. Between planting and harvesting, however, labour productivity is not straight forward to measure indirectly as growth and crop quality are mostly determined by environmental factors. Since total production at the end of the harvest is a complex function of many worker's individual contributions and of the environment over time, individual workers'



inputs cannot be easily disaggregated or compared. Thus an informational asymmetry arises.

One result of this is the inverse relationship between a farm's size and its productivity, which was first noted in the case of India (Sen, 1962) and has since been evidenced in LMICs across the world. The literature shows quite convincingly that - holding inputs constant - small family owned farms are more productive than bigger enterprises (Sen, 1981a,b). There is disagreement, however, on whether this is due to the lesser extent of moral hazard (e.g. shirking) among family members vis a vis wage workers<sup>22</sup>, or due to the spatial dispersion of workers on bigger farms which drives up monitoring costs (Sen, 1981a). Either way, free-riding seems to be less of a problem on family-farms without labour market connections.

On the other hand, labour markets are crucial for insuring against crop failure. Kochar (1999) presents evidence suggesting that the smoothness of household consumption in the presence of farm-specific crop income shocks reflects the ability of households to smooth income directly, by increasing their market hours of work. This is an early pointer to the important role that labour markets play in household consumption smoothing and, consequently, its effects on child labour allocation discussed in Dumas (2015, 2020).

## 2.3 Economic effects of climate change

Since child labour is not prominently featured as a topic in the ARE, the importance of this literature for the present study lies predominantly in providing a framework that can link climate and weather to socio-economic phenomena *like* child labour.

Indeed, the relatively recent empirical literature on the economic impacts of climate change has turned the spotlight onto quantifying the effect of climate on many different socio-economic outcomes. Table 1 gives a non-exhaustive list of such studies. Most of them report highly non-linear relationships between climate and the outcomes of interest (e.g.

---

<sup>22</sup>Household members have a direct interest in maximising farm output because their consumption is directly dependent on it. Wage workers, on the other hand, can benefit from defecting behaviour as long as output remains high enough so that a decline can credibly be attributed to solely environmental factors and wages continue to be paid.

Table 1: Selected climate impact studies from ARE Economics.

	<b>Paper</b>	<b>Climate Variables</b>	<b>Response Variable</b>
<b>Agriculture</b>	<a href="#">Deschênes and Greenstone (2007)</a> <a href="#">Mendelsohn (2008)</a> <a href="#">Schlenker and Roberts (2009)</a> <a href="#">Schlenker and Lobell (2010)</a> <a href="#">Welch et al. (2010)</a> <a href="#">Lobell et al. (2011)</a> <a href="#">Hertel and Lobell (2014)</a>	various various various various various various various	yields (HIC) yields (LMIC) yields (HIC) yields (SSA) yields (LMIC) yields (globally) yields (LMIC)
<b>Macro/Trade</b>	<a href="#">Barrios et al. (2010)</a> <a href="#">Jones and Olken (2010)</a> <a href="#">Hsiang et al. (2015)</a> <a href="#">Costinot et al. (2016)</a> <a href="#">Deryugina and Hsiang (2017)</a> <a href="#">Dingel et al. (2019)</a> <a href="#">Burke and Tanutama (2019)</a> <a href="#">Schlenker and Taylor (2019)</a>	P T, P various various various various various various	growth (SSA) exports aggregate output trade advantage aggregate output trade inequality aggregate output market outlook
<b>Migration</b>	<a href="#">Feng et al. (2010)</a> <a href="#">Marchiori et al. (2012)</a> <a href="#">Cattaneo and Peri (2016)</a> <a href="#">Missirian and Schlenker (2017)</a>	various various T T	yields, migration migration (SSA) migration migration
<b>Health</b>	<a href="#">Deschênes (2014)</a> <a href="#">Isen et al. (2017)</a> <a href="#">Obradovich et al. (2017)</a> <a href="#">Burke et al. (2018)</a> <a href="#">Baylis (2020)</a>	various T (in utero) T T T	mortality rates adult well being sleep suicide rates temperament
<b>Warfare</b>	<a href="#">Hsiang et al. (2011)</a> <a href="#">Hsiang et al. (2015)</a>	various various	conflict conflict
<b>Energy</b>	<a href="#">Auffhammer and Mansur (2014)</a>	various	electricity demand

*Notes:* Studies by subject areas, denoting the type of climate variables used (P=precipitation, T=temperature, various) and the type of response variable for which the study controls. Location type (HIC = high income countries, LMIC = lower and middle income countries, SSA=Sub-Saharan Africa) provided where applicable. Note that climate variable types may indicate the use of more than one metric pertinent to that climate component (e.g. mean, maximum, minimum,...).

*Source:* Author's own compilation.

[Schlenker and Roberts, 2009](#); [Hsiang et al., 2015](#)), and warm temperatures seem to be a particularly relevant factor for many climate responses ([Auffhammer et al., 2013](#)).

While the details differ, most of these papers roughly follow a two step procedure: first weather data is merged with data on the outcome of interest as well as some control variables. From this data set, a function describing the climate effect - called dose-response function or damage function - is estimated. The relevant coefficients are often estimated numerically due to the highly non-linear nature of climate effects. Once obtained, the dose-response function is used to project future effects of climate change. To this end, climate projections from a global climate model (GCM) are passed to the function. The function's output are then interpreted as the projected future values of the dependent variable.<sup>23</sup>

The related methodological literature, in particular [Timmins and Schlenker \(2009\)](#); [Schlenker \(2010\)](#); [Auffhammer et al. \(2013\)](#); [Hansen et al. \(2014\)](#); [Auffhammer \(2018\)](#) and [Hsiang \(2016\)](#), has developed a range of econometric strategies for estimating dose response functions, forecasting, and drawing climate insights from weather data. My empirical analysis in section 5 uses the framework of [Hsiang \(2016\)](#) in building a first bridge between this empirically driven literature from ARE and the aforementioned literature on child labour. Conceptually, rainfall, temperature, and other weather components are modelled to be jointly drawn from an underlying distribution which is commonly called 'climate'.

In order to model climate change and its effects, one must first define what exactly climate is. We never directly observe climate. What people perceive in their immediate environment is weather and it is weather, not climate, that has an immediate effect on their lives. Hermeneutics complicate this distinction further: humans learn from repeated experience, develop foresight, and adapt. As a result, weather can have direct and indirect effects. Direct effects of rain, for instance, include getting wet or losing harvest due to flooding. Analogous indirect effects are wearing a rain coat next time and installing a drainage system to safeguard crops. Thus, direct and indirect effects of weather are increasingly convoluted over time and space.

This distinction also exists, as immediate versus down-stream effects, in the Environmental Impact Assessment (EIA) literature (see e.g., [Noble, 2015](#)). Furthermore, EIAs

---

<sup>23</sup>This approach by itself has been dubbed the "dumb farmer scenario" ([Auffhammer and Schlenker, 2014](#)), as it implicitly assumes that farmers continue business as usual despite climate change.

distinguish between environmental changes and effects. Assuming that there is always some level of change in the ecosystem, environmental effects are then those additional changes induced by external shock to the system - the difference in differences. This broad conceptualisation of all environmental effects, transitory or persistent, neatly encompasses climate change.

The notion of human foresight about the weather is crucial as it suggests an implicit characterisation of climate as a probability distribution. While we may not exactly know *ex ante* what the weather will be at our position in a given future moment, we might have accumulated prior expectations about the ‘usual’ extent to which precipitation, temperature, wind gusts, and other weather components tend to vary, about their typical transitory pace, and even about the approximate probabilities attached to each possible state conditional on the eventuated past and on the contemporaneous variance observed in other related weather components.

In developing such expectations, humans display an approximate understanding of the geophysical processes underlying the weather as they learn to appreciate the observable interdependence structure across weather components induced by these processes. Continuous observation of the weather is then equivalent to drawing independent random draws from an underlying distribution. As long as this distribution remains unaltered, the accuracy of one’s expectations increases with the number of draws. Precipitation, for instance, differs by temperature and we intuitively expect snow when the temperature falls below zero degrees centigrade, but rain above that threshold. As long seasonal variation of temperature stays constant, this enables long-term planning of all kinds of economic activity in humans, as well as survival strategies in animals (e.g., hibernation) and plants (e.g., shedding foliage) more broadly. Equivalent long-term strategies occur everywhere on planet earth, always specifically geared towards the local climate.

Recent evidence from the United States shows that humans are surprisingly fast in correcting their expectations to account for climate change. The subjective baseline against which temperature is evaluated appears to be dominated by recent experience. In consequence, “temperatures initially considered remarkable rapidly become unremarkable with

repeated exposure over a roughly five year timescale” (Moore et al., 2019, p. 4909). This rapid expectation adjustment relative to the pace of anthropogenic climate change has large implications for the notability of temperature anomalies as climate change progresses.<sup>24</sup> It also means that the two-fold effects of the static climate distribution - on effectuated weather and on people’s expectations about it respectively - may carry over to the non-static setting with climatic change quite seamlessly. The following model is geared towards giving a first intuition as to what happens to child labour as climate change progresses. It is followed by an empirical exploration.

### 3 Model

The model presented in this section draws from Jessoe et al. (2018). Their model, in turn, is a version of the standard agricultural household model (Singh et al., 1986) with weather as an additional production factor (as in Ravallion, 1988). In this class of models, the household solves the production and consumption sides simultaneously, resembling their dual role as producers and consumers.

Two additional extensions are necessary in order to adapt this framework to the analysis of child labour. First, the household’s aggregate utility function, which is left unspecified in Jessoe et al. (2018), should exhibit child labour aversion in line with the luxury axiom (Basu and Van, 1998). Secondly, the effective labour input  $L$  must be disaggregated into child and adult labour. Furthermore, the functional form of adult-equivalent child labour units needs to be specified in order to accommodate complementarity and appropriately scale total time input according to the differences in productive capacity between child and adult labour time units. Following some preliminaries, I implement these extensions and solve the household’s utility maximization problem.

---

<sup>24</sup>Notably, this attenuating presence-bias surrounding human perception of extreme weather may inhibit societal pressure for climate change mitigation efforts.

**Preliminaries:** The ILO’s International Programme on the Elimination of Child Labour (IPEC) defines child labour as

work that is mentally, physically, socially or morally dangerous and harmful to children; and that interferes with the children’s schooling by depriving them of the opportunity to attend school, either by obliging them to leave school prematurely, or by requiring them to attempt to combine school attendance with excessively long and heavy work (ILO, 2021).

Because the various components of this definition are often hard to establish in practice, age is commonly used as a proxy by which to distinguish benign work from child labour. The ILO’s Convention No. 138 stipulates the relevant ages that different countries use to define child labour. In applied work, therefore, the age threshold needs to be adjusted to the relevant country context, with the upper limit at eighteen years in any case. In formulating a theoretical model, such an adjustable age threshold can be implemented on any household’s age-ordering. Consider an agricultural household  $\mathcal{I} = \{1, 2, \dots, I\}$  whose  $I > 0$  members are ordered from youngest to oldest. Let  $a_i$  be the age of household member  $i \in \mathcal{I}$ . Moreover, let  $\underline{a}$  be the legal age at which an individual is no longer considered a child. Accordingly,  $i$  is a child if  $a_i < \underline{a}$  and an adult if  $a_i \geq \underline{a}$ . Let  $\hat{i}$  denote the oldest child in the household. The household can, accordingly, be divided into  $\hat{i}$  children and  $I - \hat{i}$  adults.

Every member of the household is endowed with time, normalized at 1. By assumption, the  $I - \hat{i}$  adults work all of the time, while the  $\hat{i}$  children’s time is divided between  $e \in [0, 1]$  time units of work effort on the household’s farm and  $(1 - e)$  time units of non-work.<sup>25</sup> For simplicity, the effort  $e$  is assumed to be equal for all children.<sup>26</sup> Note that aggregate time endowment up to a given individual is equal to the number of individuals it has been aggregated over. For example, the aggregate time endowment of adults in  $\mathcal{I}$  is  $(I - \hat{i})$ .

---

<sup>25</sup>This non-work may involve going to school as well as leisure. This difference is of no consequence in this static setting. Of course, in an intertemporal maximization exercise - as part of a dynamic model - the consequences of deciding between education and leisure will have a tangible effect on the accumulation of human capital and associated earning capabilities.

<sup>26</sup>While this assumption does not impact the model’s implications, it considerably decreases the complexity involved in solving it. Some possible extensions - allowing for  $e$  to vary across children - are discussed in annex A.1.

**Household preferences:** Household  $\mathcal{I}$  derives utility from the consumption of generic goods and services, which are denoted by  $c$ . Additionally, the household’s decision makers are averse to child labour, as long subsistence consumption levels can be reached without it.

The following Stone-Geary utility function<sup>27</sup> with an additional child-leisure factor captures this trade-off (cf. [Basu and Van, 1998](#)).

$$U(e, c) = (c - s)(1 - e), \tag{1}$$

where,  $s$  is the subsistence level of consumption and the second term on the right hand side captures children’s non-work. [Basu and Van \(1998\)](#) note that the specific form of this utility function implies “that the child’s leisure is a luxury good because doubling the household’s wealth (from non-child-labour sources) leads to a more than doubling of child labour” (p. 420, footnote 11). This is the implementation of their luxury-axiom.

**Labour equivalence:** Children tend to be less productive than adults. In agriculture, this stems from differences in physical strength, skill-refining experience, ability to manage equipment, and other productive traits - all of which favour adults.<sup>28</sup> It is common practice to multiply child labour with an adult-equivalence factor between 0 and 1, making it an imperfect substitute for adult labour ([Basu and Van, 1998](#); [Baland and Robinson, 2000](#); [Bhalotra and Heady, 2003](#); [Basu et al., 2010](#); [Dwivedi and Marjit, 2017](#); [Dumas, 2020](#)).<sup>29</sup>

---

<sup>27</sup>This functional form is commonly used to model subsistence levels into expenditure systems. [Stone \(1954\)](#) first estimated the linear expenditure system associated with it and [Geary \(1950\)](#) was the first to derive it in theory, in a comment on earlier work by [Klein and Rubin \(1947\)](#). Therefore, it is also known as Klein-Rubin utility function.

<sup>28</sup>The opposite is true when a productive activity requires specific characteristics that are more readily available in children. One particularly cruel example is mining in confined underground spaces, where children’s small size allows for easier access and faster extraction of precious ores while minimising drilling costs (see e.g. [O’driscoll, 2017](#)).

<sup>29</sup>Similar household equivalence scales are commonly used in estimating household consumption systems, with children (the elderly) consuming a fraction (multiple) of adult consumption. For reference, see [Pollak and Wales \(1979\)](#), [Lewbel \(1989\)](#), [Blundell et al. \(1994\)](#), [Lewbel \(1997\)](#), or [Deaton \(1999\)](#). While they are increasingly considered to be overly simplistic, equivalence scales are used in more involved methods, like the one presented by [Dunbar et al. \(2013\)](#).

In a similar fashion, let the full labour supply of household  $\mathcal{I}$  be

$$L = (I - \hat{i}) + \gamma \hat{i}e, \quad (2)$$

where the positive parameter  $\gamma \in [0, 1]$  defines the adult-equivalent productive capacity of child time allocated to farm work. That is, one child hour of farm work produces the same marginal returns as  $\gamma$  ( $< 1$ ) hours of adult farm work. To further limit the model's complexity, this adult-equivalence parameter is assumed to be the same across children. Some possible extensions that allow it to vary by children's individual characteristics (e.g. age, sex, or prior work experience) as well as in relation to parental supervision are suggested in annex [A.2](#).

**Agricultural production:** Agricultural goods are produced using labour  $L$  and quasi-fixed land and capital  $\bar{K}$ . The quantity produced is given by  $Q = f(L, \bar{K}, \theta)$ . As in [Ravallion \(1988\)](#), the factor  $\theta$  represents the realisation of weather, a higher value of which indicates weather that increases production.  $f(\cdot)$  is a continuous and twice differentiable production function for which the following four inequalities hold:

$$\frac{\partial f}{\partial L} > 0, \quad \frac{\partial f}{\partial \theta} > 0, \quad \frac{\partial^2 f}{\partial L^2} < 0, \quad \text{and} \quad \frac{\partial^2 f}{\partial L \partial \theta} > 0. \quad (3)$$

**Household problem:** Household  $\mathcal{I}$  is a price-taker in all markets. Formally, the household's problem is to maximise its utility, described in equation (1), subject to its budget constraint

$$(I - \hat{i})c + \hat{i}\beta c = f(L, \bar{K}, \theta), \quad (4)$$

where  $\beta \in [0, 1]$  is the adult-equivalent of a child's consumption.

Note that this formulation builds on one more, rather strong assumption; the household is assumed to live in autarky, consuming the entirety of what it produces and nothing more or less. There are no functioning labour markets, where household members could either sell their own labour for wages or contract external labourers, nor consumer good markets



for the household to trade its produce for money and supplement their consumption by buying externally provided goods and services.

While assuming this extreme form of economic isolation may be warranted, to some extent, for subsistence farms in remote rural areas, it clearly does not hold for many other environments where children commonly work. The implications of labour, credit, and goods markets on child labour have been explored thoroughly in the literature (see e.g., [Baland and Robinson, 2000](#); [Bhalotra and Heady, 2003](#); [Dumas, 2013, 2015, 2020](#)). I limit my analysis to the state of autarky and focus, instead, on the impacts of climate change.

Using Lagrange's method, and denoting the lagrangian multiplier as  $\lambda$ , the utility maximization problem yields the Lagrange function

$$\mathcal{L}(c, e, \lambda) = (c - s)(1 - e) - \lambda [(I - \hat{i})c + \hat{i}\beta c - f(L, \bar{K}, \theta)]. \quad (5)$$

The first order conditions, obtained by taking the partial derivative of  $\mathcal{L}$  with respect to the three choice variables and setting them equal to zero, yield an exactly identified system of three equations with three unknowns. Solving this system for  $c$  and  $e$  gives the following identities.<sup>30</sup>

$$c = \frac{f(L, \bar{K}, \theta)}{(I - \hat{i}) + \beta \hat{i}} \quad (6)$$

and

$$e = 1 + \frac{s((I - \hat{i}) + \beta \hat{i}) + f(L, \bar{K}, \theta)}{\gamma \hat{i} f'_L}. \quad (7)$$

In equation (7), the term  $f'_L$  is a short hand for

$$\frac{\partial f(L, \bar{K}, \theta)}{\partial L},$$

which is known to be strictly positive from the curvature assumptions formalised in equa-

---

<sup>30</sup>This system can be solved either with pen and paper or on a computer. A python script for this purpose is available at [www.github.com/mathiasweidinger/MPPTH/](http://www.github.com/mathiasweidinger/MPPTH/).

tion (3) above.

Equation (6) merely tells us that individual consumption is derived by dividing the households agricultural production by its member's adult-equivalent consumption needs. Equation (7), on the other hand, shows that children's time effort working is increasing in the household's combined subsistence levels and in the household's production function. The fact that household production, itself, increases in weather means that a weather shock directly affects child labour supply. The direction of this effect on  $e$  depends on the sign of the cross-derivative of output with respect to labor and weather.

In other words, child labour effort increases with weather that is "good" for production and decreases with weather that is "bad" for it. This means that no one variable's effect should be expected to be monotone; either consistently positive or negative. Rather, very low or high levels of precipitation or temperature will have negative effects on plant growth and, thus, on production while temperate weather should facilitate plant growth and enable production. For instance, rain is generally conducive to production but floods or droughts are not.

In terms of climate change, this would suggest that the direction of a sustained shift in weather patterns would also indicate the direction of the resulting shift in child labour. If, as most projections suggest, climate change has adverse effects on agricultural productivity, this result projects a decrease in child labour as a result of climate change.

Of course, the omission of labour markets strongly limits the real life relevance of this model. While negative weather shocks tend to lead to a lesser harvest and, thus, to less labour demand during that stage of production, it is not at all clear whether this will result in more or less child labour. Likely, a sustained decrease in production would have an income effect which, over time, could make it prohibitively expensive to contract external labour for harvesting, sowing, and other activities. If that is the case, the effect of climate change may well be more child labour.

The next section investigates empirically the effect of climate change on child labour in the case of Nigeria. Negative climate coefficients<sup>31</sup> would indicate that the model above

---

<sup>31</sup>In the sense that coefficients should be negative for those variables, which are expected to have a

is useful in explaining child labour supply in subsistence agriculture. In the opposite case, labour-market effects - which are omitted above - may be dominant.

## 4 Case selection and data processing

Nigeria is the most populous and the sixth-most densely populated country in Africa<sup>32</sup>, and its 2020 nominal GDP of USD 450 billion makes it the biggest economy on the continent ([World Bank, 2021](#)).<sup>33</sup> About half of the population live in urban areas, particularly the large metropolitan areas in the South and Southwest towards the country's coast on the Gulf of Guinea. Moving North, the population density decreases gradually, giving way to wide savannahs and steppes which are characterised predominantly by agriculture.

Nigeria exhibits some characteristics which make it a particularly well-suited case study in the context of this paper. First, while exact estimates are hard to come by, it is generally acknowledged that child labour is rampant. The [US Department of Labor \(2021\)](#) estimates that 13 million children work on a regular basis. In May 2019, the ILO's Country Office Director stated publicly that at least 43 percent of Nigerian children - 15 million - were trapped in forced, largely unpaid labour ([ILO, 2019](#)). This is double the SSA average, which itself is the highest regional average in the world. While child work is also common in other sectors, such as domestic servitude or gold mining, the vast majority of working children are employed on their families' farms.

Second, Nigeria exhibits high levels of (rural) poverty and socio-economic inequality. Although it is one of the fastest-growing countries in the world, almost half of its population still grapples with extreme poverty. With a Gini index of 35, income inequality lies close to the SSA average. This shows a highly uneven distribution of wealth between many poor and very few rich households. According to the most recent national report, 40 percent of the total population, or almost 83 million people, lived below the country's poverty line of 137,439 Naira (USD 381.75) per person per year. While most wealth is located in the

---

negative effects on the agricultural production function.

<sup>32</sup>With a population of 200 million, Nigeria is also the sixth-most populous country in the world.

<sup>33</sup>In terms of purchasing power parity, only Egypt narrowly outperforms Nigeria.

big metropolitan areas such as Lagos, Ibadan, and Abuja, poverty is both more prevalent and deeper in rural areas, where the headcount is above 50 percent and the poverty gap amounts to 17.4 percent of the poverty line on average in contrast to the four percent observed in urban areas (NBS, 2020).

The simultaneous concentration of (extreme) poverty and rampant child labour in rural areas closely resembles two key assumptions of the model derived in section 3: the luxury axiom, expressed by the subsistence terms in the Stone-Geary utility function, operationalises the assumption that poverty is the principle driver of child labour. The agricultural focus of most households' economic activity further suggests that a large enough subpopulation engages in *agricultural* child labour specifically. Taken together, these facts mean that Nigeria provides a highly relevant and suitable case to study empirically.

## 4.1 Climate and agriculture in Nigeria

Nigeria's climate can be divided into two main seasons: "dry winter from November to March, and rainy summer from April to October" (Shiru et al., 2019, p. 61). Beyond seasonality, Nigeria's climate is characterised by differences in humidity, which varies by latitude as shown in figure 2. The country's northernmost fifth is covered by arid steppe and desert (red and yellow), whereas the remaining southern portion displays a humid tropical climate (shades of blue); savannahs in the center and some monsoon and rainforest climates in the southernmost regions.

Partly on account of its geography, Nigeria produces a wide variety of crops. Data shows that most farmers "are involved in more than one agricultural activity, even though at varying degrees" (NAERLS, 2015, p. 99). At 99.1 percent of surveyed farmers, however, virtually every agricultural household grows at least some food crops for household consumption, whereas only 52 percent reported growing cash crops. This statistic emphasizes the role of subsistence agriculture in providing food security across the country.

Agriculture varies regionally, closely tracing the climatic differences outlined above. The most important food crops include Corn, Sorghum, Millet, Rice, Cassava, Yam, Co-




Figure 2: Climate zones of Nigeria.

*Notes:* Köppen-geiger climate classification at 1-km resolution. *Source:* Beck et al. (2018).

coyam, Groundnut, Cowpea, and Soybean. The drier northern areas mostly produce cereals and variety of crops increases southwards. While some crops are confined to particular climates, others are grown across the spectrum. Shiru et al. (2019) study the distribution of several crops across Nigeria’s climate zones and note that, while both Corn and Rice are grown throughout the country, their growing season in the North differs from that in the South. This is perhaps best explained by the seasonality of rainfall in the otherwise dry Northern regions. In total, Shiru et al. (2019) provide an approximate calendar of cropping seasons for five crops, which gives insight into the yearly agricultural cycle from sowing, thru growing, to harvesting (see figure 4 below). Yams total a growing season of up to eleven months, being planted as early as February thru April and harvested as late as October thru December. Millet, which only grows for five months (May to September) at most, has the shortest growing season, closely followed by corn. January is the only month during which the agricultural cycle lays completely dormant; the last harvests seize in December and planting starts in late February and becomes more frequent throughout March, April and May.

This brief overview of the climatic and agricultural characteristics of Nigeria provides the necessary context for the data that I introduce next. The extent of the growing seasons, for instance, informs the temporal units of the panel data used in the analysis and dictates

how to process relevant climate variables. The following section discusses the sources of this data and how they were selected.

## 4.2 Data sources

Analysing the prevalence of child labour as a function of climate change must involve at least two types of data: weather observations and individual-level work data for children. In order to control for possible confounders, other individual and household characteristics should be added. A last and absolutely vital requirement is, of course, that all these data are geo-coded in accordance with a coordinate reference system (usually degrees of latitude and longitude) so that they can be matched. This study relies on two data sources to meet these requirements.

First, the individual and household variables are provided by the Nigeria General Household Survey (GHS), which is part of the World Bank's Living Standards Measurement Surveys (LSMS) programme and is conducted by Nigeria's National Bureau of Statistics in collaboration with Nigeria's Federal Ministry of Agriculture and Rural Development, the National Food Reserve Agency, the Bill and Melinda Gates Foundation and the World Bank. GHS is an ongoing long-term project to create a growing panel of agricultural and household data in such a way as to allow the study of agriculture's role in household welfare over time. It collects information on household agricultural activities along with other information on the households like human capital, other economic activities, and access to services and resources.

The GHS-Panel lends itself extremely well to this paper's research question thanks to its succinctly agricultural focus. To date, the survey instrument has produced four waves (2010-2011, 2012-2013, 2015-2016, 2018-2019), the first three of which are used in this analysis. The decision to omit the most recent wave in this analysis was taken due to what seems to be non-random attrition of households after wave 3, concentrated in rural areas. Each wave of the GHS-Panel is a cross-section of approximately 22,000 individuals in approximately 5,000 households which are representative of the Nigerian population at

large. Each GHS wave consists of two visits, one at the end of the planting season, between September and November, and the other at the end of the harvest, February thru April of the following year. Thus, the final panel consists of six, approximately biannual cross sections.

The most relevant parts of the GHS panel are the individuals' labour information (e.g., hours worked, wages earned, sectoral information, and time spent on household chores), agricultural plot and production information by household, as well as the more general socio-economic standing of the households and their individual members. This includes education, which is of particular interest as it may compete with child labour for children's time allocation. Apart from its agricultural focus, what sets the GHS apart is the decision by its creators to include approximate spatial coordinates of the enumeration areas (EA) where the households are located.<sup>34</sup> This allows users to spatialise the dataset and combine it with other data sources according to the coordinates attached.

I spatialise the GHS panel dataset and combine it with climate data. Version 4.05 of CRU TS (Climatic Research Unit gridded Time Series) is a widely used global climate dataset. It is derived by the interpolation of monthly climate anomalies from extensive networks of weather station observations.<sup>35</sup> CRU TS 4.05 includes ten timeseries of geophysical variables that jointly depict the development of the world's climate at a monthly interval. This data, which is available from January 1901 to December 2020, is laid out on a 0.5° latitude by 0.5° longitude grid (approximately 50 kilometres by 50 kilometres cells) across all major land masses (except Antarctica). This high resolution allows capturing a great part of cross-community climate variation between EAs.

### 4.3 Variables

This section describes the main variables used in the analysis and how they were constructed.

---

<sup>34</sup>EAs are the smallest geographical partition present in the GHS, approximately equivalent to a community or a cluster of villages. Other, higher order units are equivalent to Nigeria's administrative divisions of Legislative Government Areas (LGA), provinces, and regions.

<sup>35</sup>For detail on how CRU TS was constructed, see the accompanying publication ([Harris et al., 2020](#)).

### 4.3.1 GHS variables

The survey recorded the “hours worked last week in a job”, as well as the “minutes spent on chores yesterday”. In order to streamline the time-frame of these two variables, the hours worked in a job are divided by seven and the minutes spent on chores are divided by sixty. Adding up the two results for every individual gives their daily work hours including chores.<sup>36</sup> Lastly, I filter the variable by age and set it to missing for individuals aged 18 or above. The final result is an approximate measure of daily hours spent in child labour including chores.<sup>37</sup> Other variables which are taken more or less unaltered from the GHS-panel<sup>38</sup> are household consumption, household size, sex, and age.

To assess the welfare of a given household, a multidimensional poverty index (MPI) is constructed, closely following the approach spearheaded by [Alkire and Foster \(2011\)](#) and most recently applied globally in [Alkire et al. \(2020\)](#). I try to retain the same dimensions, indicators, and relative weights used in [Alkire and Jahan \(2018\)](#); [Alkire et al. \(2020\)](#), subject to data availability in the GHS. Each of the indicators used is related explicitly to one or more of the SDGs.<sup>39</sup> Multidimensional indices are particularly attractive because they can be disaggregated ([Bourguignon and Chakravarty, 2003](#); [Foster et al., 1984](#)), in addition to fulfilling the properties usually required by poverty metrics (see [Sen, 1976](#)).<sup>40</sup>

The health dimension is comprised of two indicators. Nutrition, which is related to SDG 2 “Zero Hunger”, codes a household as deprived if there was a situation in the last 12 months when there was not enough food to feed the household members. Child mortality is measured by the number of male and female children, born to women between 12 and

---

<sup>36</sup>Note that this often yields non-integer values due to the conversion of the chore variables from minutes to hours. Moreover, this construction relies on the assumptions that (i) last week’s workload is representative for the overall workload, and (ii) that yesterday’s chore minutes are representative of every day’s time spent on chores.

<sup>37</sup>Alternative measures can be obtained, for instance by omitting the chores component or by filtering out observations outside the agricultural sector. All these alternative measures turn out to be virtually equivalent with correlations among them and the original variable ranging between 0.92 and 1. Since the additional constraints do, however, cost observations I proceed with the initially proposed variable.

<sup>38</sup>In some cases, they were recoded to subsume many categories into one or to facilitate interpretation of results.

<sup>39</sup>For a recent application to the SDGs, see the joint publication by [OPHI and UNDP \(2020\)](#).

<sup>40</sup>For a critical review of the MPI literature, see [Ravallion \(2011\)](#).



49, who died in the household. If this number is, on average across the household's eligible females, 1 or greater the household codes as deprived. Child mortality relates directly to SDG 3 "Health and Wellbeing". Both measures are equally weighted 1/6 each so that health accounts for a third of the MPI.

The education dimension also consists of two indicators. Years of schooling codes a household as deprived if no eligible household member has completed six years of schooling. Second, a household is deprived in terms of school attendance if any of its school-aged children is not attending school at the age-appropriate level up to class eight. Both education indicators are clearly related to SDG 4 "Quality Education". Again, each of the education indicators is weighted 1/6, giving education a total weight of 1/3.

The last dimension of poverty is living standards, measured by six variables which are each weighted 1/18. A household is deprived in cooking fuel, if it cooks using solid fuel such as dung, agricultural crop, shrubs, wood, charcoal, or coal. In terms of electricity, a household is deprived if it has no electricity in the home. Both cooking fuel and electricity are central to SDG 7 "Affordable and Clean Energy". Sanitation deprivation occurs if a household has no or only an unimproved sanitation facility, or if it has an improved facility but shared with another household. A household is coded as deprived of drinking water if its source of drinking water is not safe for human consumption or if safe drinking water is 30 minutes or longer away by foot (round trip). Both sanitation and drinking water are basic needs acknowledged in SDG 6 "Clean Water and Sanitation".

In relation to SDG 11 "Sustainable Cities and Communities", a household is deprived in housing if it has inadequate housing materials in any of the three components floor, roof, or walls. Lastly, a household is considered asset deprived if it does not own more than five of the these: sofa, chairs, table, mattress, bed, mat, sewing machine, gas cooker, electric stove, gas stove, kerosene stove, fridge, freezer, air conditioner, washing machine, electric clothes dryer, bicycle, motorbike, car or other vehicle, generator, fan, radio, cassette recorder, Hi-Fi sound system, microwave oven, iron, television set, computer, DVD player, satellite dish, and musical instrument. The indicator of asset deprivation is related to SDG 1 "No Poverty".

### 4.3.2 Climate variables

The ten different geophysical variables of CRU-TS can be further categorised as primary, secondary, and derived variables. Primary variables are those that are directly measured at ground station level and have no synthetic component. They are mean temperature over two months (TMP), diurnal temperature range over two months (DTR), and the precipitation rate (PRE).

Secondary variables differ from primary variables in that they have fewer direct observations available. The synthetic estimates are estimated using empirical relationships with the primary variables. Vapour pressure (VAP) is generated from TMP, and DTR as described in [Harris et al. \(2020\)](#). The count of wet days (WET) is defined as the number of days in a month on which  $PRE \geq 0.1$  millimetres - a metric used in diverse areas including evaluation of satellite observations and of potential evapotranspiration equations. It is compiled using the same interpolation algorithm as in [Harris et al. \(2014\)](#). Synthetic cloud cover observations (CLD) are computed from DTR as in [Harris et al. \(2014\)](#), using station-based values. CLD is the percentage of cloud coverage averaged over the month.

Derived variables are entirely synthetic as their values are imputed from the observed values using empirically validated formulae. Minimum and Maximum temperature over two months (TMN and TMX) are derived arithmetically from TMP and DTR as described in [Harris et al. \(2014\)](#). Both are useful metrics, for instance for monitoring droughts, agronomic production and river basin vegetation. Freezing days per month (FRS) is the number of days in a month that record  $TMN \leq 0$  degree Celsius. It is derived from TMN using an empirically determined function and is commonly used in dendroclimatology and health. Lastly, potential evapotranspiration (PET) is calculated from TMP, VAP, and CLD using the Penman-Monteith formula ([Allen et al., 1998](#)), the United Nations Food and Agriculture Organization's (FAO) standard method for modelling evapotranspiration, as explained in [Ekström et al. \(2007, pp. 1071 - 1072\)](#). PET is defined as the amount of evaporation from the ground into the atmosphere that would occur if a sufficient water source were available. It is an important variable with regards to agriculture as it directly

Table 2: CRU TS variables.

Variables	Code	Units	CDD(km)
<b>Primary</b>			
Mean 2m temperature	TMP	degrees Celsius	1200
Diurnal 2m temperature range	DTR	degrees Celsius	750
Precipitation rate	PRE	mm/month	450
<b>Secondary</b>			
Vapour pressure	VAP	hPA	1000
Wet days	WET	days/month	450
Cloud cover	CLD	percentage	600
<b>Derived</b>			
Frost days	FRS	days/month	750
Minimum 2m temperature	TMN	degrees Celsius	1200
Maximum 2m temperature	TMX	degrees Celsius	1200
Potential evapotranspiration	PET	mm/day	n/a

*Notes:* A wet day is one receiving  $\geq 0.1$  mm precipitation. Minimum and maximum temperatures are the monthly means of the individual daily minimum and maximum temperatures; they are not the overall minimum or maximum temperature recorded each month.

*Source:* Adapted from [Harris et al. \(2020\)](#).

affects plant growth. If PET exceeds PRE, a place is considered dryland.

Table 2 shows the variables of CRU TS, their units and their correlation decay distances (CDDs). A variable’s CDD is defined as the distance where the correlation between one station and all other stations decays below  $1/e$ . The search radius for the selection of stations for the interpolation in CRU TS is set equal to the CDD as this optimises accuracy ([Harris et al., 2020](#)). In other words, CDD is the distance between a station and a grid cell beyond which a station’s raw observations are not useful for interpolating the climate variable at the cell. Generally speaking, precipitation has a rather low CDD as it is highly spatially concentrated, whereas temperature and vapour pressure usually vary over greater distances.

## 4.4 Point value extraction from the 2D grid

The spatial scale of the CRU-TS data is laid out on a grid of 0.5 by 0.5 degrees according to the standard (viz. longitude and latitude) coordinate reference system. That means that CRU-TS has two-dimensional spatiality. The GHS data, on the other hand, is made up of observations at distinct (one-dimensional) points on the same standard coordinate reference system. These geo-code the location of the households present in GHS. For each of these locations, a one-dimensional value needs to be extracted from the two-dimensional CRU-TS grid.

For every EA in the GHS dataset, I compute the location-specific timeseries for every climate variable in CRU TS 4.05. To smoothen transitional differences between climate cells, every EA's climate variables are obtained by taking a weighted average of all the (two-dimensional) cells which fall within a circular buffer of, for instance, 20 kilometres around the (one-dimensional) EA's centroid coordinates. This computation increases the accuracy of the climate data further, mirroring the resolution of EAs by averaging over two or more cells. Therefore, depending on the length of the circle's radius and the resolution of the climate grid, the climate variable's value at the EA coordinates is the weighted mean of one or more grid cells' values.<sup>41</sup>

To consider this formally, let a circle with radius  $r_{it}$  be drawn around a household  $i$ 's position on each temporal layer grid  $t = \{1, 2, 3, \dots, T\}$ .<sup>42</sup> Let  $A_{it}$  be the total area enclosed by said circle. Let  $n_{it} = \{1, 2, \dots, N_{it}\}$  denote the  $N_{it}$  grid cells which fall, at least partly, within the circle's area  $A_{it}$ . Each  $n_{it}$ 's portion within the circle has an area of  $a(n_{it}) \leq A_{it}$ . Moreover, write the climate in grid cell  $n_{it}$  as  $c(n_{it})$ . Now the value of the climate variable at household  $i$  during time  $t$  is approximated by

$$c_{it} \approx \sum_{n_{it}=1}^{N_{it}} \frac{a(n_{it})}{A_{it}} c(n_{it}). \quad (8)$$

---

<sup>41</sup>It is not necessary to weigh this sum by CDD, because the radius of the buffer should be smaller than any variable's CDD by an order of magnitude.

<sup>42</sup>The midpoint circle algorithm is one of many digital circle generation algorithms that can be used to determine the points needed for rasterizing such a circular buffer (see, e.g., [Barrera et al., 2016](#)).

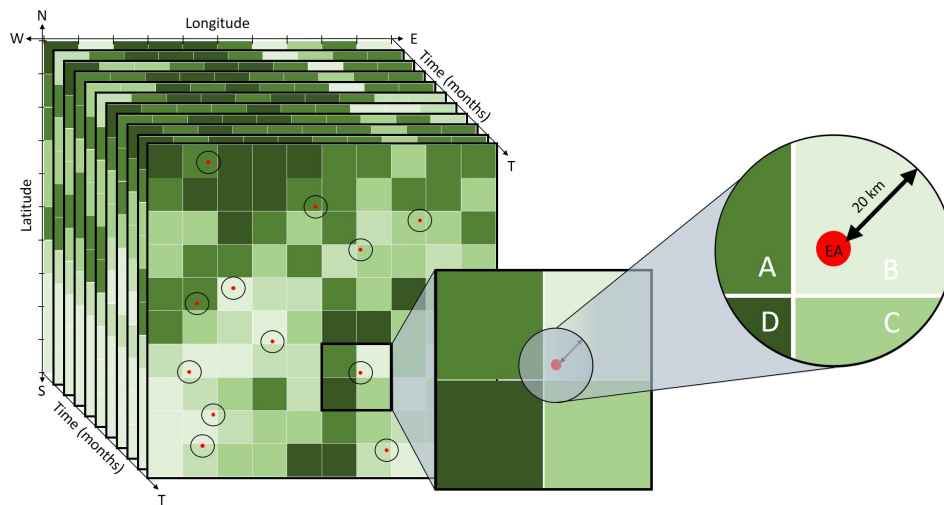


Figure 3: Smoothing climate variables using a circular buffer

*Note:* Time series of climate variables take the form of three dimensional “stack” of coordinate grids, one for each timepoint. Each of these grids is composed of cells. Every cell, in turn, holds the measured value of the variable within its spatial limits at the grid’s time point. The three dimensional grid stack on the left has the following axes: latitude from North (N) to South (S), longitude from West (W) to East (E), and discrete time in monthly intervals (T). Every grid’s cells are coloured in shades of green that indicate the value of the climate variable on a continuous scale (e.g. monthly mean temperature in degrees celsius). The red dots represent the EA locations from the spatialised GHS data. Around every EA centroid, a circular buffer with radius  $r = 20$  kilometres is drawn. The exact value at the EA’s location is then imputed by taking the weighted average of the cell segments (in order of size) B, A, C, and D that fall within the buffer, weighted by their respective surface areas.

*Source:* Author’s own compilation.

Figure 3 is a schematic representation of this procedure as applied to the data. I chose a radius of 20 kilometres for the circular buffer, which is far enough to allow for mixing of different grid cells, but lies well within the variables' CDDs. As there are no missing values in CRU TS, ten full monthly time series (one per variable), covering 1999-2020, can be created for each household.<sup>43</sup>

## 5 Econometric strategy

To identify the causal effect of climate on child labour hours, Hsiang (2016) prescribes that a panel data model with fixed effects (FE) is needed, as it accounts for unobserved, time-invariant variables. The regression in question can be represented by the usual linear, unobserved effects model (see e.g., Wooldridge, 2010): for any child  $i$  and time  $t$ ,

$$y_{it} = \alpha_i + \mathbf{c}_{it}\beta + \mathbf{x}_{it}\gamma + \varepsilon_{it}, \quad (9)$$

where  $\mathbf{c}_{it}$  is the vector of climate variables and  $\beta$  is the vector that includes the coefficients of interest. Moreover,  $\mathbf{x}_{it}$  includes control variables at the individual and household level, the effects of which are captured by the vector  $\gamma$ .  $\alpha_i$  denotes the unit-specific fixed effects and absorbs the effects of time-invariant factors, including unobservables. Lastly,  $\varepsilon_{it}$  is the idiosyncratic error term.

Huntington (1922) was perhaps first to note that this approach lends itself to estimating climate effects. In Equation 9, the climate variables are taken to be observed at the same time intervals  $t$  as child labour. However, the CRU-TS data is recorded monthly, whereas child labour (like all other GHS variables) occurs biannually. The next paragraphs closely trace Hsiang (2016) in recovering instantaneous effects from coarser data. They also extend the flexibility of the estimated climate effects.

---

<sup>43</sup>For implementing this in the R statistical programming language, see the `buffer` option of the `extract` command from the R-library `raster` (Hijmans, 2020) and its successor `terra` (Hijmans, 2021).

## 5.1 Recovering local and instantaneous non-linear effects

The existing literature suggests that there is some value in letting the climate variables enter the regression using a flexible functional form. This flexibility is justified for the following reasons: first, prior work estimating the effects of climate metrics has documented important non-linearities across a wide array of responses (Carleton and Hsiang, 2016) such as excess mortality from heat days (Deschênes and Greenstone, 2011) or crop yield damage from heat hours (Schlenker and Roberts, 2009); second, an appropriate flexible functional form should reveal the shape of the underlying response function, linear or otherwise (Hsiang, 2016); and third, it makes intuitive sense that plant-growth - the yield of which constitutes the physical output of household production - exhibits an optimal set of climate inputs, below or above which growth will be less than optimal. In fact, agricultural experiments show “that below (and above) certain thresholds, plants cannot absorb (additional) heat, while within the bounds of an upper and lower threshold heat absorption increases linearly with temperature” (Jesso et al., 2018, p. 239). Thus, agronomic studies suggest that accumulated exposure to heat over the growing season, rather than seasonal averages, determine crop growth. This also strongly suggests that responses are non-linear.

Following the approach presented in Hsiang (2016), let the instantaneous non-linear effect of climate for a household  $i$  in month  $\tau$  be  $f(\mathbf{c}_{i\tau})$ , and let this function be approximated as a linear combination of  $M$  simple non-linear functions, so that

$$f(\mathbf{c}_{i\tau}) \approx \beta_1 f_1(\mathbf{c}_{i\tau}) + \beta_2 f_2(\mathbf{c}_{i\tau}) + \dots + \beta_M f_M(\mathbf{c}_{i\tau}) = \sum_{m=1}^M \beta_m f_m(\mathbf{c}_{i\tau}), \quad (10)$$

where  $\beta_m$  are constant coefficients.

One common approach to model these non-linear responses is to fit  $M$  polynomials to the data. However, polynomials of degrees higher than three are not straight-forward to interpret.<sup>44</sup> An alternative to increasing the degree of the fitted polynomial to account

---

<sup>44</sup>A linear model represents a system with a constant rate of change, a quadratic model a system with a rate of change that increases or decreases at a fixed rate (the acceleration), and a cubic model represents a system in which the acceleration changes at a constant rate. Similar correspondences do not exist for higher order polynomials.

for curvature is to divide the range of the data into smaller pieces and apply polynomials of lesser degree to each of the subintervals, stipulating that the resulting function be continuous (see e.g. [Greene, 2018](#), pp. 275- 279). However, the use of such truncated polynomials can sometimes lead to numerical problems.<sup>45</sup> These problems can be avoided in practice by using a different set of functions in place of the truncated polynomials.

There exist several sets of functions which produce the same range of fitted curves, called equivalent bases ([Hastie and Tibshirani, 1987](#); [Hastie et al., 2009](#)). One commonly used variant are B-splines.<sup>46</sup> I use a natural cubic B-spline with knots at equally spaced interior points to achieve a flexible functional form. The spline divides the climate variables into five bins (i.e.  $M = 5$ ) according to the data's percentiles and fits the equivalent of a cubic function to each segment. This offers a non-linear alternative to using the climate data as is.<sup>47</sup>

To address the difference in frequency between the GHS and CRU-TS data, let  $\tau$  as used above denote the monthly interval of the CRU-TS data and write the biannual GHS frequencies as  $t$ . Assuming that the outcome of interest is a linear sum of  $f(\cdot)$  across points in time (i.e., temporal separability), the regression in Equation 9 can be rewritten as

$$y_{it} = \alpha_i + \left[ \sum_{\tau \in t} f(\mathbf{c}_{i\tau}) \right] + \mathbf{x}_{it}\gamma + \epsilon_{it}. \quad (11)$$

Substituting the approximation from Equation 10, this becomes

$$y_{it} \approx \alpha_i + \sum_{m=1}^M \beta_m \bar{f}_{mit} + \mathbf{x}_{it}\gamma + \epsilon_{it}, \quad (12)$$

---

<sup>45</sup>If two knots are very close then the associated truncated polynomial terms will be very similar for all observations, almost co-linear, and the fitting algorithm will become unstable. Additionally, the values in the design matrix can be very large which can lead to overflow errors (see e.g., [Hastie et al., 2009](#)).

<sup>46</sup>Although any function that can be fit with the truncated polynomials can also be fit with B-splines, and vice versa, the B-spline basis functions are not colinear making the computational algorithms for fitting splines much more stable.

<sup>47</sup>This and other approaches to model non-linearity are considered in more detail by [Hsiang \(2016, p. 18-19\)](#).



where  $\bar{f}_{mit}$  is the unit-specific sum of B-splines over time  $\tau \in t$ , that is

$$\bar{f}_{mit} = \sum_{\tau \in t} f_m(\mathbf{c}_{i\tau}). \quad (13)$$

Equation 12 can be estimated with a linear regression, using data at the monthly ( $t$ ) level. Doing so, thus, “recovers estimates for  $\beta_m$  describing the local and instantaneous function  $f(\cdot)$ , even though it uses coarser data” (Hsiang, 2016, p. 19).

The correspondence between months  $\tau$  and half years  $t$  can be inferred from figure 4, which overlays the cropping schedule presented in Shiru et al. (2019) with the surveying windows of the GHS, as described in the accompanying technical documents (NBS, 2020).

From figure 4 it is clearly visible that the agricultural growing season lasts approximately from March to December, leaving January and (most of) February outside the yearly crop cycle.

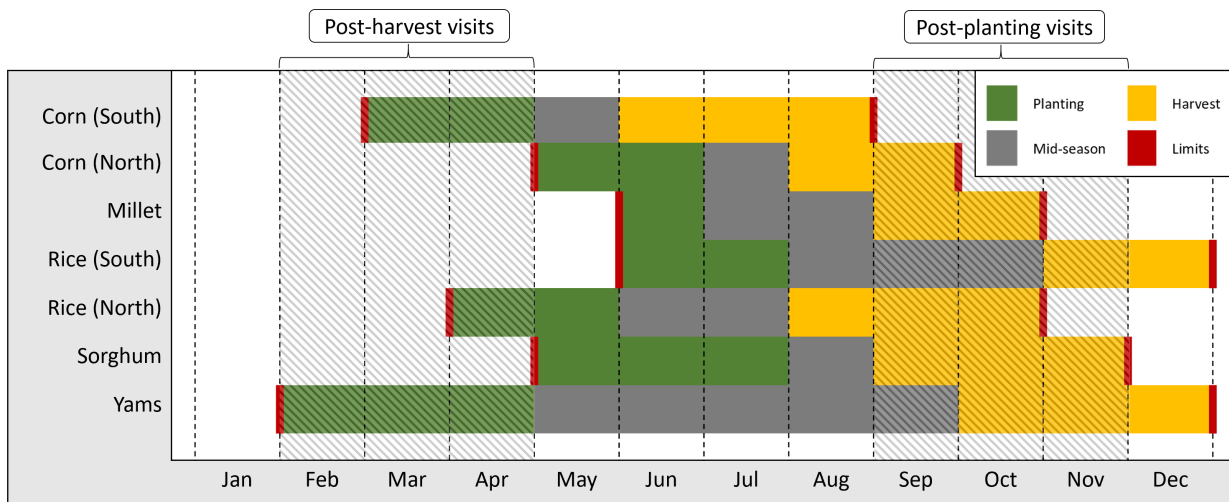


Figure 4: Cropping seasons for selected crops of Nigeria.

*Notes:* The cropping schedule published in Shiru et al. (2019, p. 62) is overlaid with the surveying schedule of the General Household Survey. The time windows during which post-harvest and post-planting visits took place are indicated by the crosshatched areas.

*Source:* Shiru et al. (2019, p. 62) and NBS (2020).

In the following empirical exercise, the climate variables PRE, TMN, TMX, and TMP

<sup>48</sup> are passed to the spline function in equation 13 for months  $\tau$  in the bi-annual time intervals  $t$  of the child labour outcome variable. Table 3 shows the correspondence of these two temporal scales for all six GHS visits (three waves) used. In fact, what NBS (2020) call “post-harvest” visits usually take place during the planting season, whereas “post-planting” visits fall within the harvest.

Table 3: Time correspondence of CRU-TS variables.

<b>wave</b>	<b>year</b>	<b>visit (<math>t</math>)</b>	<b>visit type</b>	<b>months (<math>\tau</math>)</b>
1	2010	1	post-planting	AUG - DEC
1	2011	2	post-harvest	MAR - JUL
2	2012	3	post-planting	AUG - DEC
2	2013	4	post-harvest	MAR - JUL
3	2015	5	post-planting	AUG - DEC
3	2016	6	post-harvest	MAR - JUL

*Notes:* Post-planting visits correspond to August thru December, whereas post-harvest visits are assumed to aggregate information from March thru July.

*Source:* author’s own compilation.

Based on this correspondence between  $t$  and  $\tau$ , I can now offer more specific insight into the variables that are included in the climate vector  $\mathbf{c}_{it}$ . Some of the climate variables are cumulative by nature and can, therefore, be aggregated by adding the monthly observations up over the five month periods associated with each visit. This concerns the variable PRE. Other variables - notably temperature - are mean values of daily observations.

Alas, the use of monthly average climate data attenuates much of the variation in daily weather and masks the importance of extreme temperatures to some extent. To not make matters worse, I do not average the monthly maximum and minimum temperature values further. Rather, the variables TMN and TMX are aggregated by choosing out of the five months associated to each visit the one that lies furthest from the mean temperature in either direction. TMP itself is, on the other hand, averaged over each visit’s time window.

---

<sup>48</sup>DTR is omitted as it would cause multicollinearity related to TMN and TMX. WET essentially measures the same phenomenon as PRE. VAP and PET are omitted as they have no clear theoretical connection to child labour and may cause problems with multicollinearity as well. FRS is omitted because there are no recorded days below freezing.

## 5.2 Correcting for selection bias

Plotting the dependent variable - children’s weekly working hours including household chores - in figure 5 shows a cluster at zero due to the fact that most children do in fact not work at all. In statistical terms, the data is censored at zero due to a selection process that selects children into two groups - those who work and those who do not.<sup>49</sup> As a result of this unobserved selection process, the effects estimated by standard FE may be biased (*viz.* selection bias). Thus, the approach presented above needs to be adjusted.

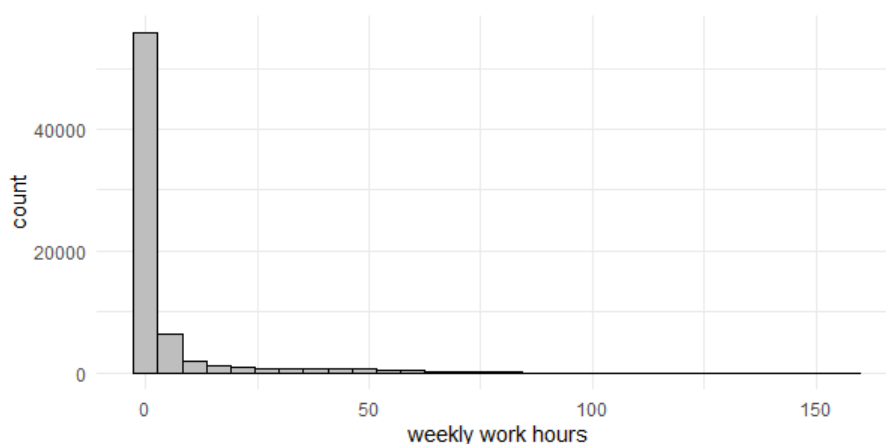


Figure 5: Child labour is censored at zero

*Note:* Observed child labour hours including household chores, which were constructed from the GHS, are censored at zero. The graph was drawn from the pooled data of all six cross sections. The individual cross sections show the same pattern. *Source:* Author’s own compilation.

In cross-sectional applications, sample selection is commonly corrected for by modelling the selection process and the observed process in two separate equations. The Tobit estimator (Tobin, 1958; Amemiya, 1984) and Heckman’s selection correction (Heckman, 1979) method are the most used estimators of this kind. Unfortunately, these frameworks do not readily extend to panel data models with FE.<sup>50</sup>

Wooldridge (1995), however, proposes an alternative method for the estimation of panel

<sup>49</sup>This is also known as the “extensive margin” of labour supply.

<sup>50</sup>The issue is that the first equation, which models the selection process, is a binary choice model that is commonly non-linear. Unlike in linear models, which remain consistent when the fixed effect is treated as a parameter to be estimated, non-linear models are generally not consistent in this setting (This is the well known “incidental parameters problem”).

data with unobserved, time-invariant effects and selection bias. [Semykina and Wooldridge \(2010\)](#) extend this estimator to handle the presence of arbitrary correlation between unobserved heterogeneity and explanatory variables. While I shall briefly introduce this estimator, interested readers are referred to the original papers for a rigorous presentation of the material.

To be precise, the estimator proposed by [Wooldridge \(1995\)](#) is a correlated random effects model. It is, however, of the fixed effects (FE) variety “in the sense that the unobserved component is allowed to be correlated with the observable explanatory variables, and we impose no distributional assumptions on the unobserved effect” (p.115). In addition, it does not require the idiosyncratic errors in the regression equation to have a known distribution, and they can have serial dependence of unspecified form.

Rewrite equation 9 as

$$y_{it} = x_{it}\theta + c_{i1} + u_{it1}, t = 1, \dots, T, \quad (14)$$

where, for ease of notation,  $x_{it}$  is a  $1 \times K$  vector of explanatory variables that includes both the adapted climate variables  $\bar{f}_{mit}$  and individual and household variables  $\mathbf{x}_{it}$  from equation 9.  $c_i$  is the unobserved effect, and  $u_{it1}$  is an idiosyncratic error.

Define a  $1 \times L > K$  vector of instruments,  $z_{it}$ , which is strictly exogenous on  $c_{i1}$ . Note that  $z_{it}$  contains all elements of  $x_{it}$  that are exogenous in equation 14 and that  $x_{it}$  and  $z_{it}$  both contain an intercept.

Next, [Semykina and Wooldridge \(2010\)](#) introduce selection by defining

$$s_{it}^* = z_{it}\delta_t + c_{i2} + u_{it2}, t = 1, \dots, T. \quad (15)$$

Here,  $c_{i2}$  is an unobserved effect and  $u_{it2}$  is an idiosyncratic error. Moreover, let

$$s_{it} = 1[s_{it}^* > 0] = 1[z_{it}\delta_t + c_{i2} + u_{it2} > 0], \quad (16)$$

where  $1[\cdot]$  is the indicator function. By definition,  $s_{it}$  is a selection indicator that equals

one if  $y_{it}$  is observed, and zero otherwise. Note that in the current application, the data clearly shows that the censored variable child labour  $y_{it} \equiv \max(0, s_{it}^*)$  is observed for all  $t$ ; we observe actual values for hours worked rather than a binary indicator only. [Wooldridge](#) calls this “partial observability of the selection variable” ([1995](#), p. 120).

Assume that the model summarised in equations [14](#), [15](#) and [16](#) holds. Moreover, let

$$c_{i1} = g(z_i) + a_{i1} = \bar{z}_i + a_{i1}, \quad (17)$$

where  $g(\cdot)$  is an arbitrary known function. Equation [17](#) is similar to the within transformation. It “produces fixed-effects slope estimators when using a balanced panel” ([Semykina and Wooldridge, 2010](#), p. 377; see also [Mundlak, 1978](#)). In an unbalanced panel, however, it is different from the within-transformation because it models  $c_{i1}$  as a function of  $\bar{z}_i$ , the elements of which are not distorted by selection. Therefore, “the model is ideologically similar to fixed effects, but is free of selection bias” ([Semykina and Wooldridge, 2010](#), p. 377). This makes it very attractive for the present case, since we are dealing with an unbalanced panel and sample selection.

For the case of an unbalanced panel, [Semykina and Wooldridge \(2010\)](#) find that plugging Equation [17](#) into Equation [14](#) yields

$$y_{it} = x_{it}\theta + \bar{z}_i\eta + \mathbb{E}(v_{it1}|z_i, s_{it}) + e_{it1}, \quad (18)$$

where, by construction,  $\mathbb{E}(e_{it1}|z_i, s_{it}) = 0$  for  $t = 1, \dots, T$ .<sup>51</sup> If  $\mathbb{E}(v_{it1}|z_i, s_{it})$  were known, Equation [18](#) could be consistently estimated by pooled Two Stage Least Squares (2SLS). They also derive that, for  $s_{it} = 1$ ,

$$y_{it} = x_{it}\theta + \bar{z}_i\eta + \mu\hat{\lambda}_{it} + e_{it1}, \quad (19)$$

---

<sup>51</sup>Note that it is not necessary that  $\mathbb{E}(e_{it1}|z_i, s_i) = 0$ . [Semykina and Wooldridge](#) emphasise that “in fact, generally  $e_{it1}$  will be correlated with selection indicators,  $s_{ir}$ , for  $r \neq t$ . This is a key benefit of the current approach: we can ignore selection in other time periods that might be correlated with  $v_{it1}$ ” ([2010](#), p. 380).

where  $\hat{\lambda}_{it}$  is an estimate of the inverse mills ratio, obtained by probit. This, in turn, leads to the final procedure to estimate selection-corrected estimates from (unbalanced) panel data with unobserved, time-invariant effects.

**Procedure 4.1.1. (Semykina and Wooldridge, 2010):**

1. For each  $t$ , estimate the selection equation by probit to get  $T$  inverse mills ratios  $\hat{\lambda}_{it}$ .
2. For  $s_{it} = 1$ , use pooled 2SLS to estimate Equation 19. Use  $z_{it1}$ ,  $\bar{z}_i$  and  $\hat{\lambda}_{it}$  as instruments.
3. “Panel-bootstrap” the asymptotic variance. “This involves resampling cross-sectional units (and all time periods for each unit sampled) and using the bootstrap sample to approximate the distribution of the parameter vector. The estimator is consistent for  $N \rightarrow \infty$  and  $T$  fixed” (Semykina and Wooldridge, 2010, p. 378).

□

Applying this estimator to the use case specified in Equation 9, I first run a Probit, regressing an indicator (1 if the child works, 0 otherwise) on the climate and household independent variables and retrieve the inverse mills ratios  $\hat{\lambda}_{it}$ . At the second stage, I estimate the equation

$$y_{it} = \sum_{m=1}^M \beta_m \bar{f}_{mit} + \mathbf{x}_{it}\gamma + \bar{z}_i\eta + \mu\hat{\lambda}_{it} + \epsilon_{it} \quad (20)$$

for all children who work non-zero hours by 2SLS. Household size is used as an identifying variable for selection.<sup>52</sup> Lastly, I bootstrap the asymptotic variance as suggested above. 200 bootstrap replications should be sufficient to obtain reliable standard errors (Efron and Tibshirani, 1994, p. 47).<sup>53</sup>

---

<sup>52</sup>Semykina and Wooldridge (2010) call this an “instrument”. The instrument should be strongly associated with the selection variable but unrelated to child labour hours. While household size is correlated more strongly with selection than with child labour hours, it remains a weak instrument, to be replaced whenever more suitable data is available.

<sup>53</sup>I use the XTHECKMAN module (Rios-Avila, 2021), which implements the estimator of Wooldridge (1995) and Semykina and Wooldridge (2010) in Stata.

## 6 Findings

Table 4 summarises the output from four different regression models, all of which estimate the effects of climate variables on child labour hours. The first two columns show the coefficients on the climate variables from a pooled Ordinary Least Squares (OLS, Model (1)) and a panel regression with unit Fixed Effects (FE, Model (2)). The third column presents the coefficients obtained from the Correlated Random Effects (CRE, Model (3)) model proposed by [Wooldridge \(1995\)](#) and [Semykina and Wooldridge \(2010\)](#). Unlike models (1) and (2), model (3) is a two stage model that corrects for selection bias.

Note that models (1), (2), and (3) use the climate variables in their level (and precipitation squared) rather than cubic b-splines. To further investigate nonlinearities of climate effects, model (4) combines CRE with cubic b-splines as discussed in the preceding section. The functional flexibility of b-splines comes, however, at the cost of interpretability which is reduced to directional, qualitative statements. Therefore, the most interesting coefficients in table 4 are perhaps those of model (3).

Table 4: Climate effects on weekly child labour: regression results.

Model	(1)	(2)	(3)	(4)				
Estimator	OLS	FE	CRE	CRE				
<i>cubic b-spline section</i>	-	-	-	1	2	3	4	5
Variable								
Precipitation (100 mls)	-0.673*** (0.20)	-1.023*** (0.22)	-4.035*** (0.44)	-17.644*** (2.35)	-18.669*** (1.74)	-15.521*** (2.15)	-24.795*** (1.85)	-24.219*** (2.97)
Precipitation squared	0.023* (0.01)	0.051*** (0.01)	0.241*** (0.03)					
Mean temperature (°C)	2.315*** (0.17)	1.008** (0.33)	4.118*** (0.55)	-4.961*** (1.16)	0.943 (1.03)	2.417* (1.05)	2.095 (1.22)	3.384* (1.38)
Maximum temperature (°C)	-1.24*** (0.08)	-0.947*** (0.09)	-1.701*** (0.19)	2.583 (1.79)	-3.890*** (1.10)	-3.611* (1.60)	-2.910* (1.23)	-6.238*** (1.59)
Minimum temperature (°C)	-0.712*** (0.06)	-0.259 (0.14)	-1.510*** (0.23)	6.079** (1.92)	10.585*** (1.34)	-2.524 (1.68)	1.556 (1.23)	0.774 (1.56)
Unit FE	no	yes	yes			yes		
Zone FE	yes	yes	no			no		
Time FE	yes	yes	yes			yes		
<i>N</i>	55,735	55,735	55,735			55,735		

Estimates of climate effects on weekly child labour hours from pooled OLS (1), Fixed Effects (2), Correlated Random Effects (Wooldridge, 1995) with level variables (3), and Correlated Random Effects with variables transformed through five-fold cubic b-splines to account for non-linearities (4).

Notes: \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ . Standard errors in brackets. Models (1) and (2) use standard errors that are clustered for the 18,735 individuals who reoccur in the panel. Models (3) and (4) use panel-bootstrapped standard errors calculated from 200 bootstrap replications, clustered the same way. All four models control for age, age squared, sex, MPI, log of household consumption, and sector (rural or urban). Models (3) and (4) use household size as an instrument for selection in the first stage. Precipitation squared is used in models (1), (2) and (3) to account for nonlinearities while using the level variables. Models (1), (2) and (3) can be interpreted the usual way, which is not true for model (4) due to the variables in b-splines. Source: Author's own compilation.



All models seem to agree that, on average, an increase in precipitation is associated with a decrease in child labour. This effect is not linear though since the coefficient on Precipitation squared is positive and significant in models (1) thru (3). In fact, plotting the predicted child labour hours from model (3) against precipitation in figure 6 shows that the negative effect is strongest at the lowest levels of precipitation and flattens for precipitation volumes of more than 400 ml. One interpretation of this fact is that households who already experienced a high load of precipitation need not send their children to water their crops. An additional increase in precipitation has less negative effects on child labour because the remaining tasks need to still be carried out. This may also imply that droughts affect the production function more negatively than normal and even high amounts of precipitation would. This negative shock to the production function would also decrease the demand for child labour in accordance with the theoretical perspective of section 3.

As for temperature, it seems that an increase in the mean temperature is, by itself, conducive to child labour. According to model (3), a one degree increase in mean temperature is associated with an additional 4 hours spent at work, conditional on that the child is working - this is a considerable increase, given working children's average of 8.4 hours weekly. The effect may not be as linear as this coefficient suggests, though. Child labour tends to decrease at temperatures in the low twenties and only picks up at temperatures of 27 degrees celsius or more - giving rise to the U-shape depicted in the upper panel of figure 6. These nonlinearities are not reliably evidenced by the b-splines of model (4), though, which means that the increasing right side of the U might also be the effect of other covariates.

As for the tails of the temperature distribution, it seems that both maximum and minimum temperatures are negatively associated with child labour hours. In fact, the coefficients for both maximum and minimum temperatures are consistently negative and highly significant with the exception of some spline sections in model (4).

Controlling for selection bias, model (3) suggests that a one degree increase in maximum temperature is associated with a decrease of child labour by almost two hours. The b-splines of model (4) show that this negative effect only takes shape for higher levels

of maximum temperature, whereas the effects of the first quintile is non-significant and, therefore, statistically not different from zero. Interestingly, plotting the predicted values of child labour against maximum temperature, an increasing line is obtained - see the lower left graph in the upper panel of figure 6. It is unclear why this is, but it is reasonable to think that the negative marginal effect of maximum temperature is being outweighed by, for instance, the positive effect of a mean-temperature increase.

One additional degree of minimum temperature decreases child labour by 1.5 hours on average. Once again, model (4) points to important nonlinearities in these effects. It suggests, that at the lower two quintiles, minimum temperature is actually associated with an *increase* in child labour. This leads to the inverted U-shape depicted in the lower right graph in the upper panel of figure 6. This is a reasonable result, considering that there are no freezing days in the data and that Nigeria, more generally, tends to suffer from drought rather than from freezing. Therefore, in our case study, lower minimum temperatures are unlikely to negatively affect plant growth, leaving the need for labour input unaffected at lower minimum temperature.

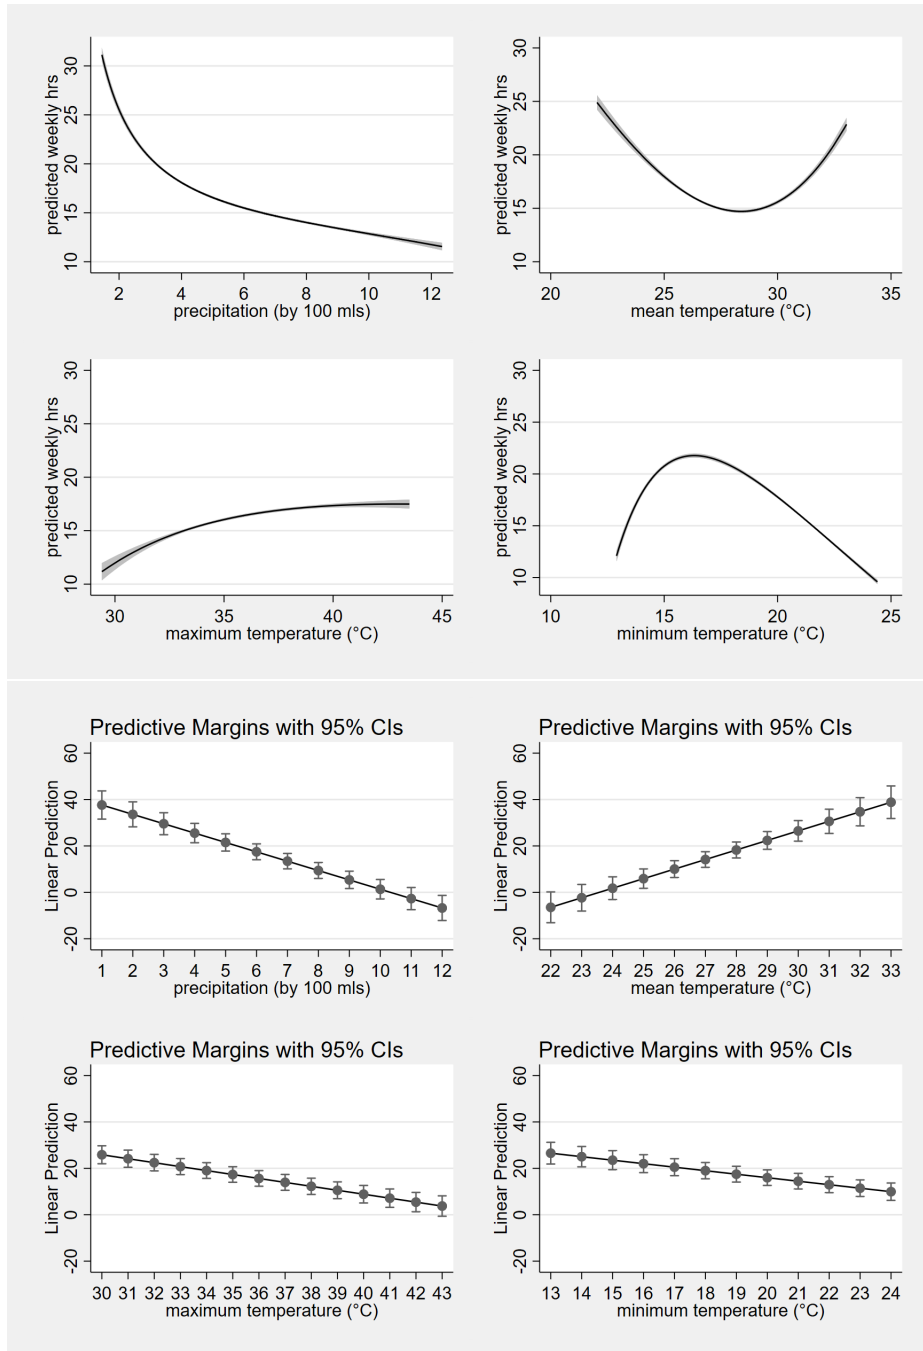


Figure 6: Predicted weekly child labour hours and marginal effects from Model (3).

*Note:* The upper four panels of this figure graph a fractional polynomial fit on the predicted child labour from the CRE-level model (3) in table 4 against the climate variables precipitation (by 100 mls), mean-, maximum, and minimum temperature (by °C). 95 percent confidence intervals in grey. The lower four panels plot the marginal effects of the same climate variables, also with 95 percent confidence intervals.  
*Source:* Author's own compilation.

In summary, the analysis suggests that a permanent increase in mean temperature due to global warming leads to an increase in child labour. This effect may be somewhat counter-acted by the negative influence of weather variability on crops, which is also to be expected as the climate continues changing (Pachauri et al., 2015). Higher maximum temperatures seem to attenuate child labour supply in Nigeria slightly, which might be explained by a decrease in labour demand when there is less crop to harvest during hot periods. Minimum temperatures display mixed effects, forming an inverted-U shape with its turning point at approximately 16 degrees celsius. Before this point, it seems that child work increases with minimum temperature. Above it, the effect changes direction. Precipitation decreases child labour, in particular at very low quantities which points to the importance of droughts in the region. This last finding that child labour decreases in precipitation stands in contrast to earlier findings for Tanzania (Dumas, 2020), which evidenced a positive relationship.

Some of the control variables, which have been omitted from table 4, also allow interesting insights. Somewhat surprisingly, the analysis finds no statistically significant effect of children’s gender on the amount of their work hours. The effect of total household consumption spending is that with an additional percentage point of household spending, children work 1.5 hours more. This is in line with the “wealth paradox” that was first described by Bhalotra and Heady (2003); children of wealthier families tend to work more rather than less on average. Interestingly, though, the same is not true when employing MPI as an alternative poverty metric. The coefficient on MPI is non-significant.

## 7 Conclusion

This paper sought to describe the effects of climatic change on agricultural child labour. To this end, section 3 developed a simple household decision model rooted in the child labour literature. The main hypothesis of this model is that, all things equal, child labour should decrease with ‘bad’ weather and increase with ‘good’ weather, where bad weather affects plant growth negatively and good weather aids plant growth. The essential idea is

that plant-growth is the one variable in the farm production function that is most closely linked to weather outcomes, which are in turn determined stochastically by the climate. As plants grow less, households' labour demand decreases and they react by decreasing the amount of child labour supplied, as they are assumed to be averse to child labour through the luxury axiom.

The hypothesis is tested in the empirical part of this study for the the case of Nigeria. The findings of this analysis are much more differentiated than the simple hypothesis posed. It is found that child labour tends to decrease, on average, with weather that is expected to hamper plant development. Extremely hot or cold temperatures and heavy rains are examples of this. However, the analysis also shows that child labour increases in mean temperature. This is counter-intuitive to the theoretical perspective of section 3. The use of b-splines in one of the regression models, moreover, suggest the existence of important non-linear effects for all four climate variables.

Thus, the evidence points out the limitations of my simple theoretical model and, consequently, of its hypothesis. One major criticism that ought to be made is that it does not at all account for markets, including the credit and labour markets which both have proven to be influential factors in other studies (e.g., [Bhalotra and Heady, 2003](#); [Dumas, 2013](#)). Their positive and negative effects (e.g., through selling labour or consumption smoothing) are not taken into consideration.

Moreover, there are some extensions that could have been made to the theoretical model's formulation of aggregate labour supply, of children's productivity in relation to that of adults, and the rather strict assumptions on time allocation. I chose not to make these extensions in the spirit of Occam's razor - instead, they are included in the annex. While these extensions might make the model's equations resemble real life a little better, they also would increase its complexity considerably. This would have been inefficient to do, given the limitations of the model in other domains. Rather than making the existing parts of the model more complex, future research should be geared towards integrating the roles of markets.

Considering only the empirical section of this study, another criticism remains. I

essentially employ what has been dubbed the “dumb farmer approach” by [Auffhammer and Mansur \(2014\)](#). My analysis does not take into consideration farmer’s climate adaptation strategies, like changing farming practices or migrating. Therefore, the findings I present here are only really useful for the short term. They may, however, retain their validity for longer when focusing on the type of farmers implied by this study. It is, perhaps, reasonable to assume that poor subsistence-farming households would be ill-equipped to adapt to climate change, compared to more affluent ones. The effects of climate adaptation on child labour remain an important field of study, not least because one strategy to save labour costs when times get rough is to employ one’s own children.

One more caveat concerns data availability. Due to the negative connotations surrounding the topic of child labour, parents may be inclined to under-report how much their children actually work. Therefore, the extent of child labour reflected in the data is likely lower than the real scope of the problem. Moreover, it is hard to come by high-definition climate data for the SSA region. While I use monthly climate data, it would be preferable to have daily climate variables at one’s disposal because monthly means already smoothen out a lot of variation from the data. The non-linearities reported here could be estimated more precisely with daily weather data.

This study merely represents a first attempt at gauging the complex implications of the ongoing climate crisis for working children. It shows that climate effects on child labour seem to be non-linear and, in particular, that the projected increase in mean temperatures may stall ongoing global attempts at eliminating child labour by 2025 along with SDG 8. In a sense, this study also relates to the literature on measuring the ‘cost of carbon’. If child labour is considered a global bad, appropriate resources need to be used in order to counter it. Thus, this paper provides yet another reason for taking urgent action, in accordance with the Paris Climate Accord and the latest climate science, to limit global temperature rise in the 21st century to 1.5°C above pre-industrial levels.

## References

- Alkire, S. and Foster, J. (2011). Counting and multidimensional poverty measurement. *Journal of Public Economics*, 95(7):476–487.
- Alkire, S. and Jahan, S. (2018). The New Global MPI: Aligning with the Sustainable Development Goals. OPHI Working Paper 121, Oxford Poverty and Human Development Initiative, Oxford.
- Alkire, S., Kanagaratnam, U., and Suppa, N. (2020). The Global Multidimensional Poverty Index (MPI) 2020. MPI Methodological Note 49, OPHI.
- Allen, R. G., Pereira, L. S., Raes, D., and Smith, M. (1998). Crop evapotranspiration - Guidelines for computing crop water requirements. Technical Report 56, Food and Agricultural Organization of the United Nations, Rome, Italy.
- Amemiya, T. (1984). Tobit models: A survey. *Journal of Econometrics*, 24(1):3–61.
- Auffhammer, M. (2014). Economic Impacts of Climate Change on Agriculture.
- Auffhammer, M. (2018). Quantifying Economic Damages from Climate Change. *Journal of Economic Perspectives*, 32(4):33–52.
- Auffhammer, M., Hsiang, S. M., Schlenker, W., and Sobel, A. (2013). Using Weather Data and Climate Model Output in Economic Analyses of Climate Change. *Review of Environmental Economics and Policy*, 7(2):181–198.
- Auffhammer, M. and Mansur, E. T. (2014). Measuring climatic impacts on energy consumption: A review of the empirical literature. *Energy Economics*, 46:522–530.
- Auffhammer, M. and Schlenker, W. (2014). Empirical studies on agricultural impacts and adaptation. *Energy Economics*, 46:555–561.
- Baland, J.-M. and Robinson, J. A. (2000). Is Child Labor Inefficient? *Journal of Political Economy*, 108(4).
- Bar, T. and Basu, K. (2009). Children, Education, Labor, and Land: In the Long Run and Short Run. *Journal of the European Economic Association*, 7(2-3):487–497.

- Barrera, T., Hast, A., and Bengtsson, E. (2016). A chronological and mathematical overview of digital circle generation algorithms – introducing efficient 4- and 8-connected circles. *International Journal of Computer Mathematics*, 93(8):1241–1253.
- Barrios, S., Bertinelli, L., and Strobl, E. (2010). Trends in Rainfall and Economic Growth in Africa: A neglected cause of the African growth tragedy. *The Review of Economics and Statistics*, 92(2):350–366.
- Basu, K., Das, S., and Dutta, B. (2010). Child labor and household wealth: Theory and empirical evidence of an inverted-U. *Journal of Development Economics*, 91:8–14.
- Basu, K. and Van, P. H. (1998). The Economics of Child Labor. *The American Economic Review*, 88(3):412–427.
- Baylis, P. (2020). Temperature and temperament: Evidence from Twitter. *Journal of Public Economics*, 184:104161.
- Beck, H. E., Zimmermann, N. E., McVicar, T. R., Vergopolan, N., Berg, A., and Wood, E. F. (2018). Present and future köppen-geiger climate classification maps at 1-km resolution. *Scientific Data*, 5(1):1–12.
- Becker, G. (1964). *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. Harvard University Press, Cambridge, MA.
- Bhalotra, S. (2002). Parent Altruism. Annual Conference, Royal Economic Society.
- Bhalotra, S. and Heady, C. (2000). Child Farm Labour: Theory and Evidence. Discussion Paper, Sunotory and Toyota International Centres for Economics and Related Disciplines at LSE, London.
- Bhalotra, S. and Heady, C. (2003). Child Farm Labor: The Wealth Paradox. *The World Bank Economic Review*, 17(2):197–227.
- Blundell, R., Preston, I., and Walker, I., editors (1994). *The Measurement of Household Welfare*. Cambridge University Press, Cambridge.



- Boeckx, P., Bauters, M., and Dewettinck, K. (2020). Poverty and climate change challenges for sustainable intensification of cocoa systems. *Current Opinion in Environmental Sustainability*, 47:106–111.
- Bourguignon, F. and Chakravarty, S. R. (2003). The Measurement of Multidimensional Poverty. *The Journal of Economic Inequality*, 1(1):25–49.
- Boutin, D. (2014). Climate Vulnerability, Communities’ Resilience and Child Labour. IZA Discussion Paper 8567, IZA, Bonn.
- Burke, M., González, F., Baylis, P., Heft-Neal, S., Baysan, C., Basu, S., and Hsiang, S. M. (2018). Higher temperatures increase suicide rates in the United States and Mexico. *Nature Climate Change*, 8:723–729.
- Burke, M. and Tanutama, V. (2019). Climatic Constraints on Aggregate Economic Output. NBER Working Paper 25779, National Bureau of Economic Research, Cambridge, MA.
- Carleton, T. A. and Hsiang, S. M. (2016). Social and economic impacts of climate. *Science*, 353(6304):1112–1127.
- Cattaneo, C. and Peri, G. (2016). The migration response to increasing temperatures. *Journal of Development Economics*, 122:127–146.
- Chakraborty, K. and Chakraborty, B. (2018). Low level equilibrium trap, unemployment, efficiency of education system, child labour and human capital formation. *Journal of Economics/Zeitschrift für Nationalökonomie*, 125(1):69–95.
- Costinot, A., Donaldson, D., and Smith, C. (2016). Evolving Comparative Advantage and the Impact of Climate Change in Agricultural Markets: Evidence from 1.7 Million Fields around the World. *Journal of Political Economy*, 124(1):205–248.
- Craparo, A. C., Van Asten, P. J., Läderach, P., Jassogne, L. T., and Grab, S. W. (2015). Coffea arabica yields decline in Tanzania due to climate change: Global implications. *Agricultural and Forest Meteorology*, 207:1–10.

- D'Alessandro, S. and Fioroni, T. (2016). Child labour and inequality. *Journal of Economic Inequality*, 14(1):63–79.
- DaMatta, F. M., Rahn, E., Läderach, P., Ghini, R., and Ramalho, J. C. (2019). Why could the coffee crop endure climate change and global warming to a greater extent than previously estimated? *Climatic Change*, 152(1):167–178.
- Deaton, A. (1999). *The Analysis of Surveys, a Microeconomic Approach to Development Policy*. John Hopkins University Press, Baltimore.
- Del Carpio, X. V. (2008). Does Child Labor Always Decrease With Income? An Evaluation In The Context Of A Development Program In Nicaragua. Policy Research Working Paper, World Bank, Washington, D.C.
- Del Carpio, X. V. and Loayza, N. V. (2012). The impact of wealth on the amount and quality of child labor. Policy Research Working Paper 5959, World Bank, Washington, D.C.
- Dendir, S. (2014). Children’s cognitive ability, schooling and work: Evidence from Ethiopia. *International Journal of Educational Development*, 38:22–36.
- Deryugina, T. and Hsiang, S. M. (2017). The Marginal Product of Climate. NBER Working Paper 24072, National Bureau of Economic Research.
- Deschênes, O. (2014). Temperature, human health, and adaptation: A review of the empirical literature. *Energy Economics*, 46:606–619.
- Deschênes, O. and Greenstone, M. (2007). The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather. *American Economic Review*, 97(1):354–35.
- Deschênes, O. and Greenstone, M. (2011). Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US. *American Economic Journal: Applied Economics*, 3(4):152–185.

- Dingel, J. I., Meng, K. C., and Hsiang, S. M. (2019). Spatial Correlation, Trade, and Inequality: Evidence from the Global Climate. NBER Working Paper 25447, National Bureau of Economic Research, Cambridge, MA.
- Dumas, C. (2007). Why do parents make their children work? A test of the poverty hypothesis in rural areas of Burkina Faso. *Oxford Economic Papers*, 59:301.
- Dumas, C. (2013). Market Imperfections and Child Labor. *World Development*, 42:127–142.
- Dumas, C. (2015). Shocks and child labor: The role of markets. Working Papers SES 458, Universität Freiburg / Université de Fribourg, Freiburg/Fribourg.
- Dumas, C. (2020). Productivity Shocks and Child Labor: The Role of Credit and Agricultural Labor Markets. *Economic Development and Cultural Change*, 68(3):763–812.
- Dunbar, G. R., Lewbel, A., and Pendakur, K. (2013). Children’s Resources in Collective Households: Identification, Estimation, and an Application to Child Poverty in Malawi. *The American Economic Review*, 103(1):438–471.
- Dwivedi, J. K. and Marjit, S. (2017). Relative Affluence and Child Labor - Explaining a Paradox. *Review of Development Economics*, 21(4):1178–1190.
- Edmonds, E. V. (2007). Child Labor. NBER Working Paper 12926, National Bureau of Economic Research, Cambridge, MA.
- Edmonds, E. V. and Schady, N. (2012). Poverty Alleviation and Child Labor. *American Economic Journal: Economic Policy*, 4(4):100–124.
- Edmonds, E. V. and Shrestha, M. (2014). You get what you pay for: Schooling incentives and child labor. *Journal of Development Economics*, 111:196–211.
- Efron, B. and Tibshirani, R. J. (1994). *An Introduction to the Bootstrap*. Chapman and Hall CRC, first edition.
- Ekström, M., Jones, P. D., Fowler, H. J., Lenderink, G., Buishand, T. A., and Conway, D. (2007). Regional climate model data used within the SWURVE project ? 1: Projected changes

- in seasonal patterns and estimation of PET. *Hydrology and Earth System Sciences Discussions*, 11(3):1069–1083.
- Feng, S., Krueger, A. B., and Oppenheimer, M. (2010). Linkages among climate change, crop yields and Mexico-US cross-border migration. *Proceedings of the National Academy of Sciences*, 107(32):14257–14262.
- Fisher, R. A. (1925). The influence of rainfall on the yield of wheat at Rothamsted. *Philosophical Transactions of the Royal Society of London. Series B, Containing Papers of a Biological Character*, 213(402-410):89–142.
- Foster, J., Greer, J., and Thorbecke, E. (1984). A Class of Decomposable Poverty Measures. *Econometrica*, 52(3):761–766.
- Geary, R. C. (1950). A Note on "A Constant-Utility Index of the Cost of Living". *The Review of Economic Studies*, 18(1):65–66.
- Greene, W. H. (2018). *Econometric Analysis*. Pearson Education Limited, Harlow, global edition ; 8th edition edition.
- Hansen, G., Auffhammer, M., and Solow, A. R. (2014). On the Attribution of a Single Event on Climate Change. *Journal of Climate*, 27(22):8297–8301.
- Harris, I., Jones, P. D., Osborn, T. J., and Lister, D. H. (2014). Updated high-resolution grids of monthly climatic observations – the CRU TS3.10 Dataset. *International Journal of Climatology*, 34(3):623–642.
- Harris, I., Osborn, T. J., Jones, P., and Lister, D. (2020). Version 4 of the CRU TS monthly high-resolution gridded multivariate climate dataset. *Scientific Data*, 7(1):109.
- Hastie, T. and Tibshirani, R. (1987). Generalized Additive Models: Some Applications. *Journal of the American Statistical Association*, 82(398):371–386.
- Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning*. Springer Series in Statistics. Springer New York, New York, NY.

- Heckman, J. J. (1979). Sample Selection Bias as a Specification Error. *Econometrica*, 47(1):153–161.
- Hertel, T. W. and Lobell, D. B. (2014). Agricultural adaptation to climate change in rich and poor countries: Current modeling practice and potential for empirical contributions. *Energy Economics*, 46:562–575.
- Hijmans, R. J. (2020). Package 'raster': Geographic Data Analysis and Modeling. <https://cran.r-project.org/web/packages/raster>.
- Hijmans, R. J. (2021). Package 'terra': Spatial Data Analysis. <https://cran.r-project.org/web/packages/terra>.
- Hsiang, S. M. (2016). Climate Econometrics. *Annual Review of Resource Economics*, 8(1):43–75.
- Hsiang, S. M., Burke, M., and Miguel, E. (2015). Climate and Conflict. *Annual Review of Economics*, 7(1):577–617.
- Hsiang, S. M., Meng, K. C., and Cane, M. A. (2011). Civil conflicts are associated with the global climate. *Nature*, 476(7361):438–441.
- Huntington, E. (1922). *Civilization and Climate*. Yales University Press, New Haven, CT.
- Ibrahim, A., Abdalla, S. M., Jafer, M., Abdelgadir, J., and De Vries, N. (2018). Child labor and health: A systematic literature review of the impacts of child labor on child's health in low-and middle-income countries. *Journal of Public Health*, 41(1):18–26.
- ILO (2017). Global Estimates of Child Labour: Results and trends 2012-2016. Technical report, International Labour Organization, Geneva.
- ILO (2019). ILO calls actors to join the movement against child labour in Africa. [http://www.ilo.org/africa/about-us/offices/abuja/WCMS\\_710120/lang-en/index.htm](http://www.ilo.org/africa/about-us/offices/abuja/WCMS_710120/lang-en/index.htm).
- ILO (2021). What is child labour? <https://www.ilo.org/ipec/facts/lang-en/index.htm>.
- IMF (2020). Regional Economic Outlook: Sub-Saharan Africa. Regional Economic Outlook, International Monetary Fund, Washington, D.C.

- Isen, A., Rossin-Slater, M., and Walker, R. (2017). Relationship between season of birth, temperature exposure, and later life wellbeing. *Proceedings of the National Academy of Sciences*, 114(51):13447–13452.
- Jessoe, K., Manning, D. T., and Taylor, J. E. (2018). Climate Change and Labour Allocation in Rural Mexico: Evidence from Annual Fluctuations in Weather. *The Economic Journal*, 128(608):230–261.
- Jones, B. F. and Olken, B. A. (2010). Climate Shocks and Exports. *American Economic Review: Papers and Proceedings*, 100:454–459.
- Klein, L. R. and Rubin, H. (1947). A Constant-Utility Index of the Cost of Living. *The Review of Economic Studies*, 15(2):84–87.
- Kochar, A. (1999). Smoothing consumption by smoothing income: Hours-of-work responses to idiosyncratic agricultural shocks in Rural India. *Review of Economics and Statistics*, 81(1):50–61.
- Kuznets, S. (1955). Economic Growth and Income Inequality. *The American Economic Review*, 45(1):1–28.
- Kuznets, S. (1963). Quantitative Aspects of the Economic Growth of Nations: VIII. Distributions of Income by Size. *Economic Development and Cultural Change*, 11(2):1–80.
- Lewbel, A. (1989). Household equivalence scales and welfare comparisons. *Journal of Public economics*, 39:377–391.
- Lewbel, A. (1997). Consumer demand systems and household equivalence scales. In Pesaran, H. and Schmidt, P., editors, *Handbook of Applied Econometrics (Microeconomics, Vol. II)*. Blackwell, Oxford.
- Lobell, D. B., Schlenker, W., and Costa-Roberts, J. (2011). Climate trends and global crop production since 1980. *Science*, 333(6042):616–620.

- Marchiori, L., Maystadt, J.-F. O., and Schumacher, I. (2012). The impact of weather anomalies on migration in sub-Saharan Africa. *Journal of Environmental Economics and Management*, 63:355–374.
- Marx, K. and Engels, F. (1848). *Manifest Der Kommunistischen Partei*. Bildungsgesellschaft für Arbeiter, London.
- Mendelsohn, R. (2008). The Impact of Climate Change on Agriculture in Developing Countries. *Journal of Natural Resources Policy Research*, 1(1):5–19.
- Mills, J. H. and Waite, T. A. (2009). Economic prosperity, biodiversity conservation, and the environmental Kuznets curve. *Ecological Economics*, 68(7):2087–2095.
- Mincer, J. (1958). Investment in Human Capital and Personal Income Distribution. *Journal of Political Economy*, 66(4):281–302.
- Missirian, A. and Schlenker, W. (2017). Asylum applications respond to temperature fluctuations. *Science*, 358(6370):1610–1614.
- Moore, F. C., Obradovich, N., Lehner, F., and Baylis, P. (2019). Rapidly declining remarkability of temperature anomalies may obscure public perception of climate change. *Proceedings of the National Academy of Sciences*, 116(11):4905–4910.
- Munasinghe, L., Jun, T., and Rind, D. H. (2012). Climate change: A new metric to measure changes in the frequency of extreme temperatures using record data. *Climatic Change*, 113(3-4):1001–1024.
- Mundlak, Y. (1978). On the Pooling of Time Series and Cross Section Data. *Econometrica*, 46(1):69–85.
- NAERLS (2015). Agricultural Performance Survey Report of 2015 Wet Season in Nigeria. Agricultural Performance Survey Report, National Agricultural Extension and Research Liaison Services, Zaria.
- NBS (2020). Poverty and Inequality in Nigeria 2019: Executive Summary. National Poverty and Inequality Report, National Bureau of Statistics, Abuja.

- Noack, F. (2019). Resource Abundance and Education. *Economic Development and Cultural Change*, 68(2):699–727.
- Noble, B. F. (2015). *Environmental Impact Assessment: A Guide to Principles and Practice*. Oxford University Press, Canada, Don Mills, Ontario, third edition.
- Obradovich, N., Migliorini, R., Mednick, S. C., and Fowler, J. H. (2017). Nighttime temperature and human sleep loss in a changing climate. *Science Advances*, 3(5):1–6.
- O’driscoll, D. (2017). Overview of child labour in the artisanal and small-scale mining sector in Asia and Africa. K4D Helpdesk Report, Institute of Development Studies, Brighton, UK.
- OPHI and UNDP (2020). Charting pathways out of multidimensional poverty: Achieving the SDGs.
- Oryoie, A. R., Alwang, J., and Tideman, N. (2017). Child Labor and Household Land Holding: Theory and Empirical Evidence from Zimbabwe. *World Development*, 100:45–58.
- Özokcu, S. and Özdemir, Ö. (2017). Economic growth, energy, and environmental Kuznets curve. *Renewable and Sustainable Energy Reviews*, 72:639–647.
- Pachauri, R. K., Allen, M. R., Barros, V. R., Broome, J., Cramer, W., Christ, R., Church, J. A., Clarke, L., Dahe, Q., Dasgupta, P., Dubash, N. K., Edenhofer, O., Elgizouli, I., Field, C. B., Forster, P., Friedlingstein, P., Fuglestvedt, J., Gomez-Echeverri, L., Hallegatte, S., Hegerl, G., Howden, M., Jiang, K., Jimenez Cisneros, B., Kattsov, V., Lee, H., Mach, K. J., Marotzke, J., Mastrandrea, M. D., Meyer, L., Minx, J., Mulugetta, Y., O’Brien, K., Oppenheimer, M., Pereira, J. J., Pichs-Madruga, R., Plattner, G.-K., Pörtner, H.-O., Power, S. B., Preston, B., Ravindranath, N., Reisinger, A., Riahi, K., Rusticucci, M., Scholes, R., Seyboth, K., Sokona, Y., Stavins, R., Stocker, T. F., Tschakert, P., van Vuuren, D., and van Ypersele, J.-P. (2015). Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II, and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Technical report, International Panel on Climate Change, Geneva.
- Pal, S. and Saha, B. (2012). Self-Employment and Child Labour: Theory and Evidence. 7th IZA/World Bank Conference: Employment and Development, World Bank, New Dheli.



- Piketty, T. (2014). *Capital in the Twenty-First Century*. Belknap Harvard, Cambridge, MA.
- Pollak, R. A. and Wales, T. J. (1979). Equity; the individual versus the family: Welfare comparisons and equivalence scales. *American Economic Review*, 69:216–221.
- Posso, A., editor (2020). *Child Labor in the Developing World: Theory, Practice, and Policy*. Springer Singapore, Singapore.
- Ravallion, M. (1988). Expected Poverty Under Risk-Induced Welfare Variability. *The Economic Journal*, 98(393):1171–1182.
- Ravallion, M. (2011). On multidimensional indices of poverty. *The Journal of Economic Inequality*, 9(2):235–248.
- Ravallion, M. (2020). On Measuring Global Poverty. *Annual Review of Economics*, 12(1):167–188.
- Ravallion, M. and Chen, S. (2009). Weakly Relative Poverty. Policy Research Working Paper 4844, World Bank, Washington, D.C.
- Rios-Avila, F. (2021). XTHECKMANFE: Stata module to fit panel data models in the presence of endogeneity and selection. Boston College Department of Economics.
- Ritchie, H. and Roser, M. (2017). CO2 and Greenhouse Gas Emissions. <https://ourworldindata.org/co2-and-other-greenhouse-gas-emissions>.
- Rosenzweig, M. R. and Evenson, R. (1977). Fertility, Schooling, and the Economic Contribution of Children of Rural India: An Econometric Analysis. *Econometrica*, 45(5):1065.
- Sarkar, J. and Sarkar, D. (2016). Why Does Child Labour Persist With Declining Poverty? *Economic Inquiry*, 54(1):139–158.
- Schlenker, W. (2010). Crop Responses to Climate and Weather: Cross-Section and Panel Models. In Lobell, D. and Burke, M., editors, *Advances in Global Change Research*, volume 37, pages 99–108. Springer, Dordrecht.

- Schlenker, W. and Lobell, D. B. (2010). Robust negative impacts of climate change on African agriculture. *Environmental Research Letters*, 5:1–8.
- Schlenker, W. and Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change. *Proceedings of the National Academy of Sciences*, 15(37):15594–15598.
- Schlenker, W. and Taylor, C. A. (2019). Market Expectations About Climate Change. NBER Working Paper 25554, National Bureau of Economic Research, Cambridge, MA.
- Schultz, T. W. (1961). Investment in Human Capital. *American Economic Review*, 51(1):1–17.
- Semykina, A. and Wooldridge, J. M. (2010). Estimating panel data models in the presence of endogeneity and selection. *Journal of Econometrics*, 157(2):375–380.
- Sen, A. (1981a). Market failure and control of labour power: Towards an explanation of 'structure' and change in Indian agriculture. Part 1. *Cambridge Journal of Economics*, 5(3):201–228.
- Sen, A. (1981b). Market failure and control of labour power: Towards an explanation of 'structure' and change in Indian agriculture. Part 2. *Cambridge Journal of Economics*, 5(4):327–350.
- Sen, A. K. (1962). An Aspect of Indian Agriculture. *The Economic Weekly*, (Annual Number 14):243–243.
- Sen, A. K. (1976). Poverty: An Ordinal Approach to Measurement. *Econometrica*, 44(2):219.
- Shiru, M. S., Shahid, S., Chung, E.-S., and Alias, N. (2019). Changing characteristics of meteorological droughts in Nigeria during 1901–2010. *Atmospheric Research*, 223:60–73.
- Singh, I., Squire, L., and Strauss, J., editors (1986). *Agricultural Household Models: Extensions, Applications, and Policy*. Johns Hopkins University Press, Baltimore and London.
- Stone, R. (1954). Linear Expenditure Systems and Demand Analysis: An Application to the Pattern of British Demand. *The Economic Journal*, 64(255):511–527.
- Timmins, C. and Schlenker, W. (2009). Reduced-Form Versus Structural Modeling in Environmental and Resource Economics. *Annual Review of Resource Economics*, 1(1):351–380.

- Tobin, J. (1958). Estimation of Relationships for Limited Dependent Variables. *Econometrica*, 26(1):24–36.
- Tol, R. S. J. (2009). The Economic Effects of Climate Change. *Journal of Economic Perspectives*, 23(2):29–51.
- UNICEF (2020). Child labour. <https://data.unicef.org/topic/child-protection/child-labour/>.
- US Department of Labor (2021). Child Labor in Nigeria. Child Labor Report, US Department of Labor, Washington, D.C.
- Welch, J. R., Vincent, J. R., Auffhammer, M., Moya, P. F., Dobermann, A., and Dawe, D. (2010). Rice yields in tropical/subtropical Asia exhibit large but opposing sensitivities to minimum and maximum temperatures. *Proceedings of the National Academy of Sciences*, 107(33):14562–14567.
- Wooldridge, J. M. (1995). Selection corrections for panel data models under conditional mean independence assumptions. *Journal of Econometrics*, 68(1):115–132.
- Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*. MIT Press, Cambridge, MA, second edition.
- World Bank (2021). World Bank Open Data - Nigeria. <https://data.worldbank.org/>.

## A Annex: Extensions to the theoretical model

The following pages present possible extensions to the theoretical model of section 3. A comprehensive repository of information and files related to this project, including additional detail about the empirical analysis and its findings and an outline of the computational implementation, is available online at [www.github.com/mathiasweidinger/MPPTH](http://www.github.com/mathiasweidinger/MPPTH).

### A.1 Easing assumptions on time allocation

The model in section 3 maintained the assumption that adults work all of their time, while children could set aside a share of their time  $e \geq 1$  to work effort. Alternatively, one could allow both children and adults to allocate their time between work and non-work. Keeping individual's time endowment standardised at 1, individual  $i$ 's work time share can be written as  $l_i \in [0, 1]$ . The household's total aggregate time labour hours are obtained by summing  $l_i$  across children:  $\sum_{i=1}^{\hat{i}} l_i = L_C$ . Conversely, all adult's labour input is captured by  $\sum_{i=I-\hat{i}}^I l_i = L_A$ . In fact, this setup would allow for work time to vary not only between adults and children, but also across individuals of the same type.

### A.2 Children's productivity and adult supervision

In the model presented in section 3, children were assigned an adult equivalence factor of  $\gamma \in [0, 1]$  that scaled their work units to those of adults, making child and adult work imperfect substitutes. Bar and Basu (2009) extend this substitution axiom (Basu and Van, 1998), contending that child labour must be supervised, using proportional amounts of adult labour, in order to be productive. In their model, children are employed when matched with an appropriate amount of adult supervision and left idle otherwise; child and adult labour are substitutes within a bound beyond which they become complements. Toddlers aside, this underlying assumption that unsupervised are completely unproductive might be too restrictive; Once instructed, even young children are able to replicate a multitude of simple tasks without constant monitoring. Rather than establishing child productivity, therefore, adult supervision can be thought to extend a low innate productive capacity. As a result, the household may make use of a given child's labour even when there is not sufficient adult time left to supervise it. This idea is developed further in

this possible extension to the simple model.

To consider this formally, assume that child labour is necessarily less productive than adult labour by a factor  $v(u)$ , which specifies the adult-equivalence of child labour if a fraction  $u \in (0, 1)$  of adult labour is used to supervise the fraction  $s \in (0, 1)$  of child labour. Supervision generally increases children's productivity, such that

$$0 < \underline{v} \leq v_0 = v(u = 0) \leq v(0 < u < 1) \leq v(u = 1) \leq \bar{v} \leq 1 \quad \forall u \in (0, 1],$$

where  $\underline{v}$  and  $\bar{v}$  are the absolute boundaries of children's productive potential. I call  $s$  the supervision rate,  $u$  the intensity of supervision,  $v(u)$  the technology of skill enhancement, and  $v_0$  the adult-equivalence factor of unsupervised child labour (or "raw skill"). In absence of supervision (viz.,  $s = 0$ ), the initial adult-equivalent of aggregate child labour input is  $v_0 L_C$  and the household's total labour input in this case is  $L_{\mathcal{I}} = L_A + v_0 L_C$ .

When supervision occurs, the technology of skill enhancement  $v(u)$  exhibits monotonicity and decreasing returns to scale from supervision, meaning that

$$\frac{\partial v(u)}{\partial u} > 0, \quad \text{and} \quad \frac{\partial^2 v(u)}{\partial u^2} < 0.$$

Moreover, it must hold that  $v(u) > u \quad \forall u \in (0, 1)$ , since otherwise it would not be worthwhile to invest in supervision. Any monotonically increasing and concave function can be used to represent  $v(u)$ , provided it meets these few requirements. Panel (a) of figure 7 illustrates the shape of  $v(u)$  by the example of the function  $v(u) = v_0 + (\bar{v} - v_0)u^{\frac{1}{3}}$ , using four different values of  $v_0$ .<sup>54</sup>

The adult-labour equivalents of any given combinations of adult and child labour, considering the complementary of supervision in addition to the substitution between  $L_a$  and  $L_c$ , can be captured by the household's effective labour supply:

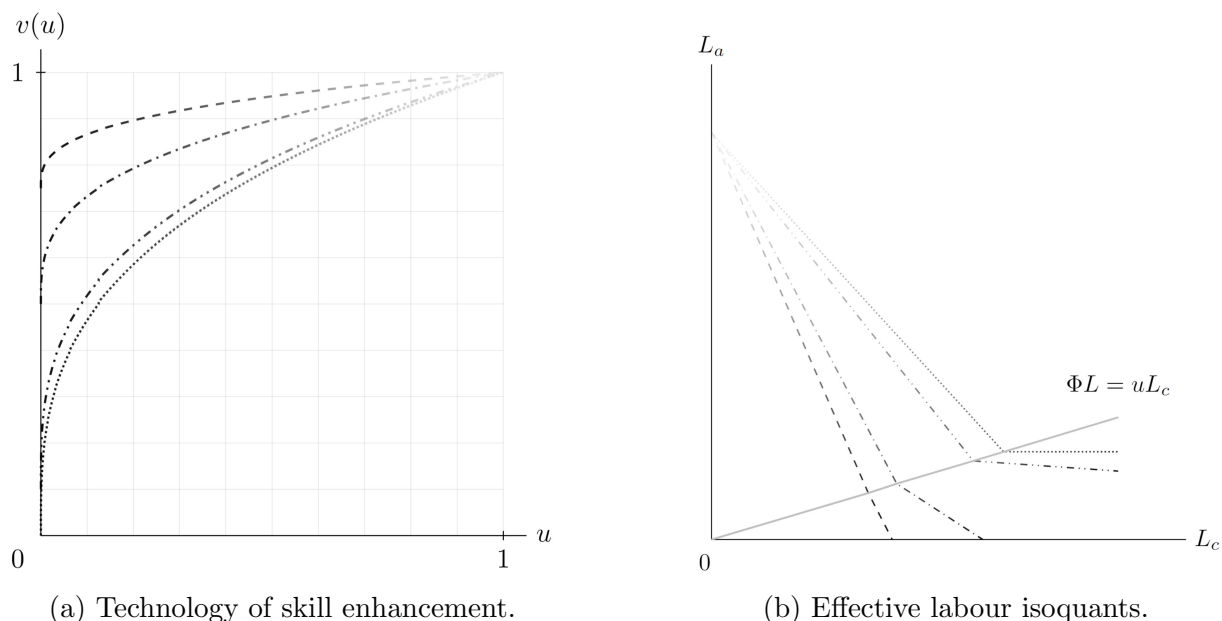
$$L_{\mathcal{I}} = \begin{cases} L_A - suL_C + sv(u)L_C + (1-s)v_0L_C & \text{if } suL_C < L_A \\ v(u)\frac{L_A}{u} + v_0\left(L_C - \frac{L_A}{u}\right) & \text{if } suL_C = L_A \end{cases}. \quad (21)$$

This specification allows for the household decision maker to choose flexibly not only the

---

<sup>54</sup>Note that  $\bar{v}$  is unambiguously defined as follows from the definition of  $v$  and  $u$ :  $v(u = 1) \leq \bar{v} \leq 1$  and  $v(u) \geq u \quad \forall u \in (0, 1) \iff v(u = 1) = \bar{v} = 1$ .

Figure 7: Supervision, child productivity, and labour supply.



*Note:* (a) The figure graphs four versions of the function  $v(u) = v_0 + (\bar{v} - v_0)u^{\frac{1}{3}}$  for different values of raw skill  $v_0$ . The dotted curve corresponds to  $v_0 = 0$  and the curves above it to  $v_0 = 0.1$ , to  $v_0 = 0.5$ , and to  $v_0 = 0.75$  respectively. (b) The dotted isoquant corresponds to  $v_0 = 0$ , whereas the isoquants to its left correspond to  $v_0 = 0.1$ , to  $v_0 = 0.5$ , and to  $v_0 = 0.75$  respectively. The upward sloping line in gray, which connects the isoquant kinks, is defined by its slope, the intensity of supervision  $u$ . *Source:* author's own compilations.

levels of adult and child labour supply, but also the fraction of children to work with and without supervision, as well as the intensity of such supervision; all this in view of optimising the household's welfare. Setting  $v_0 = 0$ ,  $s = 1$ ,  $u < 1$  and  $v(u) = 1 \quad \forall u$ , equation 21 simplifies to the labour supply function in Bar and Basu (2009). Panel (b) of figure 7 plots combinations of child and adult labour, using the same values for  $v_0$  as in Panel (a). The dotted curve on top corresponds to the Bar and Basu perspective that unsupervised child labour must be entirely unproductive (i.e.  $v_0 = 0$ ).<sup>55</sup> Note that along each of the piece-wise linear functions in panel (b), the effective labour input  $L$  stays the same, indicating that, holding other inputs constant, they are isoquants in the farm production.

<sup>55</sup>For all functions graphed in figure 7, it is assumed that  $u = 0.1$  and  $s = 1$ .