Three Essays on Internal Migration and Nutrition in Tanzania

Kalle Hirvonen

Thesis Submitted for the Degree of Doctor of Philosophy

April 2014

Department of Economics

University of Sussex

Declaration

I hereby declare that this thesis has not been and will not be, submitted in whole or in part to another University for the award of any other degree.

Signature:

Kalle Hirvonen

UNIVERSITY OF SUSSEX

Kalle Hirvonen

DEGREE OF DOCTOR OF PHILOSOPHY

Three Essays on Internal Migration and Nutrition in Tanzania

SUMMARY

This thesis is formed of three separate essays. The essays are empirical in nature and use the Kagera Health and Development Survey from Tanzania. The survey spans a 19-year period offering a unique opportunity to study many long-run dynamic processes of development in rural Africa.

In the first essay, a version of which was co-authored with Joachim De Weerdt, we use these data to shed light on how mass internal migration changes the nature of informal risk-sharing. By quantifying how shocks and consumption co-move across linked households, our analysis shows that migrants unilaterally insure their extended family members who remain at home. This finding contradicts risk-sharing models based on reciprocity, but is consistent with assistance driven by social norms. Migrants sacrifice three to five per cent of their consumption growth to provide this insurance, which seems too trivial to have a stifling effect on their growth through migration.

The second essay studies the role of exogenous income shocks on long-term migration decisions. The results reveal that temperature shocks cause large fluctuations in household consumption and inhibit long-term migration among men. These findings suggest that liquidity constraints are binding and prevent potential migrants from tapping into the opportunities brought about by internal migration.

The final essay focuses on child nutrition and examines whether under-nourished children are able to recover the height losses later in life. The essay questions the methods used in the existing empirical literature and challenges the conventional view that recovery is nearly impossible after five years of age. The empirical part of the essay documents how puberty offers an opportunity window for recovery in the case of children in Kagera.

Acknowledgements

I want to dedicate this work to my mother Sirpa Tynkkynen and my sister Kati Pettersson. Without their love and support I would have never managed to fulfil this dream. Thank you, kiitos.

I have been lucky to have been guided by three excellent supervisors. Barry Reilly is the reason I wanted to do my PhD at Sussex. His attention to detail and dedication to research are some of the qualities I hope to take with me. The support from Andy McKay has been vital throughout these years. He has tirelessly reviewed my drafts and improved them with a number of sharp comments and suggestions. I thank Alex Moradi for opening a door to a new exciting research path and also for his support and impeccable advice.

In late 2008 I sent a hopeful email to Joachim De Weerdt to ask whether they needed any help with the looming KHDS field work. This email – and the response I got after waiting for about 15 minutes completely changed the course of my PhD experience. As an applied economist I should of course be careful in making claims about parallel universes but it has truly been a tremendous learning experience working with Joachim. The field work in Kagera was a great experience and fulfilled a long lasting dream. I wish to thank all my colleagues and friends at E.D.I., especially Aris Mgohamwende and Leonard Kyaruzi. I would also like to acknowledge the help I got from Mama Anna and Benjamin Kamukulu. I thank Respichius Mitti and Thaddeus Rweyemamu for sharing their wisdom regarding the art of field work and local customs and culture in Kagera. Neil Chalmers for teaching me everything I now know about data management. Brian Dillon for being a great house mate. I enjoyed working with Timothy Kyessy and the field supervisors; George J. Musikula, Bernard M. Matungwa, Allan Katemana and Mwenge Godlaid.

Without the excellent team of interviewers and data processors none of this would have been possible. Finally, I would like to thank those 6,353 individuals and their families for taking part in KHDS. Wakora, Asante.

Throughout this journey I have received support and advice from many great economists, especially from Kathleen Beegle, Helene Bie Lilleor, Jean-Francois Maystadt, Stefan Dercon, Sofya Krutikova, Richard Tol, Alan Winters, Shqiponja Telhaj and Torfinn Harding. I thank the staff at UNU-WIDER for an enjoyable and highly productive research visit back in spring 2012. Joy Blake has assisted me with numerous admin related issues.

I am grateful for the support from dear friends back home and beyond, especially Jaska ja Olga, Petri ja Jonna, Maija, Stikson, Aino ja Artturi, Sinkka ja Masa, Jere, Tiera, Chema y Maria, Carl och Karin, Christophe and Vojta. I thank the doctors, teachers and lumberjacks of Kenttura, Enonkoski, Finland for their guidance, advice and support – and helping me to keep my feet firmly on the ground.

I have been fortunate to gain true friends at Sussex. The PhD community has been a great resource during this journey. Special thanks go to Edgar Cooke who has not only been a great friend but has also helped me with my research problems. I also thank Janani Ramasubramanian for memorable conference trips and Rashaad Shabab for sharing countless of painful kilometres in the woods of South Downs. Thanks to Miniva, Amrita, Ani, Elsa, Hector, Giulia, Marinella, Marta, Paola, and other friends for happy times.

I am grateful to the Yrjö Jahnsson foundation, Finnish Cultural Foundation and Economic and Social Research Council for funding my PhD studies.

Table of Contents

List of Figures

List of Tables

List of Abbreviations and Acronyms

2SLS	Two-stage Least Squares
AEU	Adult Equivalent Unit
Baseline	1991-94 survey rounds of KHDS
BCS70	1970 British Cohort Study
CAPI	Computer-assisted personal interviewing
CDC	Centers for Disease Control and Prevention
DHS	Demographic and Health Surveys
EDI	Economic Development Initiatives
FE	Fixed Effects
GDP	Gross Domestic Product
GPS	Global Positioning System
GPRS	General packet radio service
HAZ	Height-for-Age Z-score
HH	Household
IV	Instrumental Variables
KHDS	Kagera Health and Development Survey
LPM	Linear Probability Model
LSMS	Living Standards Measurement Study
NASA	The National Aeronautics and Space Administration
NCHS	National Center for Health Statistics
NFE	Network Fixed Effects
OLS	Ordinary Least Squares
PHHM	Previous Household Member (Panel Respondent)
Tsh	Tanzanian shilling
URT	United Republic of Tanzania
WHO	World Health Organization

1 Introduction

> Economic development is a dynamic process with a time scale of more than a few years. The scarcity of household-survey data covering long periods of time has meant that most empirical evidence on hypotheses about growth and development has come from cross-country aggregate time-series data. (Rosenzweig 2003p. 112)

This thesis in part is a paean to longitudinal surveys of households and individuals. Administering long panel surveys is costly but the benefits come, as pointed out above by Rosenzweig (2003), in the ability to document long-run development patterns and to investigate many of the most exciting questions in Development Economics. Or as Dercon, Krishnan, and Krutikova (2013)p. 16 put it:

> Village studies and long-term micro-data sets are not quite the current flavour in development economics research, but without them, we will not be able to document carefully long-term changes and create the new stylised facts for the next generation of development research.

In this thesis, I employ a unique panel data from the Kagera region in Tanzania. At 19-years, the Kagera Health and Development Survey (KHDS) is one of the longest running panel surveys in Sub-Saharan Africa. These data are described in detail in the next chapter with a focus on the latest round of the survey in 2010 that I participated in collecting. Attrition is a particular concern in any panel survey. Attrition is unlikely to be random and as such is likely to induce considerable biases in any estimates of welfare (Alderman et al. 2001, McKay and Lawson 2003). For these reasons, the KHDS survey has paid particular attention to tracking individuals over time. The survey has maintained an exceptional tracking rate of panel respondents over its life providing unique research opportunities for examining long-run poverty and welfare trends in one part of Sub-Saharan Africa. Using these data I contribute to two central themes within Development Economics: internal migration and nutrition.

Historically, internal migration has moved in step with development and poverty reduction in both the rich developed and fast growing developing countries (Chenery and Syrquin 1975, Collier and Dercon 2009). Early models of structural formation saw internal migration assuming central stage when a country modernizes and moves away from subsistence agriculture (Lewis 1954, Ranis and Fei 1961). The demographic flows that accompany such modernisation processes will affect traditional rural communities profoundly. For migrants, geographical mobility is associated with large income gains and presents one of the main routes out of poverty (Beegle, De Weerdt, and Dercon 2011, Christiaensen, De Weerdt, and Todo 2013, Dercon, Krishnan, and Krutikova 2013).

These internal migration flows are captured in the Kagera survey. By 2010, after excluding those who were deceased, more than half of the panel respondents had migrated away from their baseline village where they were first interviewed in 1991-94. More than 95 per cent of these migrants are internal migrants and relocated either within the region or elsewhere in Tanzania. Less than four per cent migrated to another East-African country (mostly Uganda) and only 0.5 per cent moved outside East-Africa.

In the first essay, joint work with Joachim De Weerdt, an attempt is made to understand how internal migration, a core part of the structural process, interacts with a traditional institution like informal risk-sharing to shape economic mobility and vulnerability. We find that migrants rapidly progress in life, while the living standards of those who decided to stay grew at a much slower pace. By quantifying how shocks and consumption co-move across linked households, we show how migrants unilaterally insure their extended family members at home. This finding contradicts risk-sharing models based on reciprocity, but is consistent with assistance driven by social norms. An average migrant sacrifices 2.9 to 5.0 percentage points out of a total growth of 108 per cent to insure his or her relatives. This estimate is equivalent to a 'tax rate' of between 2.7 and 4.6 per cent. We regard this tax rate as too trivial to exert any constraining effect on migrants. Our findings challenge the popular notion of the entrepreneurial African migrant who is weighed down by demands for assistance from those living at home.

In the second essay I study whether financial constraints in the rural areas inhibit out-migration. I use temperature changes as a measure of income shocks. Previous studies based on aggregated longitudinal cross-country data find that annual temperature changes cause large fluctuations in economic outcomes in Sub-Saharan Africa. In this essay, I present, to my knowledge, the first attempt to measure the impact of temperature changes using micro-level data from Sub-Saharan Africa. The results reveal how household per capita consumption co-moves with temperature in the previous growing season. Migration is associated with various costs that are borne up front. In the absence of well-functioning credit markets, migration decisions should then respond to these negative income shocks induced by temperature changes. The empirical analysis based on discrete time-event techniques (Allison 1982, Jenkins 1995) confirms these predictions. I find that a one standard deviation increase in the previous year's average monthly growing season temperature reduces the overall male migration rate by about 13 per cent. However, female migration is not affected by these shocks possibly due the fact that, in this Tanzanian context, such migration appears to be largely motivated by marriage and family.

The final essay of the thesis moves away from internal migration and focuses on childhood nutrition. Prolonged and severe malnutrition in utero and early childhood hinders neural and physical development. A conventional view in the literature is that early childhood conditions largely determine adult outcomes: stunted children become shorter and less educated adults compared to their healthy and well-nourished peers. This implies that policy interventions after the first five years of life cannot be effective. The empirical evidence supporting this view, however, typically originates from data sets that follow children only over a short period of time providing an incomplete portrait of children's long-term development patterns.

In the final essay I study the extent to which short and stunted children are able to recover the height losses experienced in early childhood. I use simple statistical theory and the 1970 British Cohort study to demonstrate methodological flaws in the existing empirical literature measuring such catch-up growth. I then go on to document how puberty offers an opportunity window for the Kagera children to recover from past nutritional shocks. This essay is forthcoming in *Annals of Human Biology*.

The structure of this thesis can now be outlined. The next chapter begins with a short overview of the macro-economic situation in Tanzania over the survey period. It then provides an in-depth description of the Kagera Health and Development survey with a particular focus on the latest round of the survey in 2010. The following three chapters comprise the essays. Chapter three presents the first essay titled "Risk Sharing and Internal Migration" while chapter four contains the second titled "Temperature Shocks, Household Consumption and Internal Migration: Evidence from rural Tanzania". The fifth chapter provides the third essay titled "Measuring catch-up growth in malnourished populations". Chapter six contains some concluding remarks.

2 Data and Context

2.1 Tanzania 1990-2010

Similar to many other African countries (see Radelet 2010, Young 2012, McKay 2013), Tanzania recorded impressive growth rates over the period of study in this thesis. Figure 2.1 shows that after a period of sluggish growth in early 1990s, the country's GDP per capita increased by roughly 50 per cent over the past two decades.

Figure 2.1: GDP per capita in Tanzania, 1990-2010 (1990 = base year)

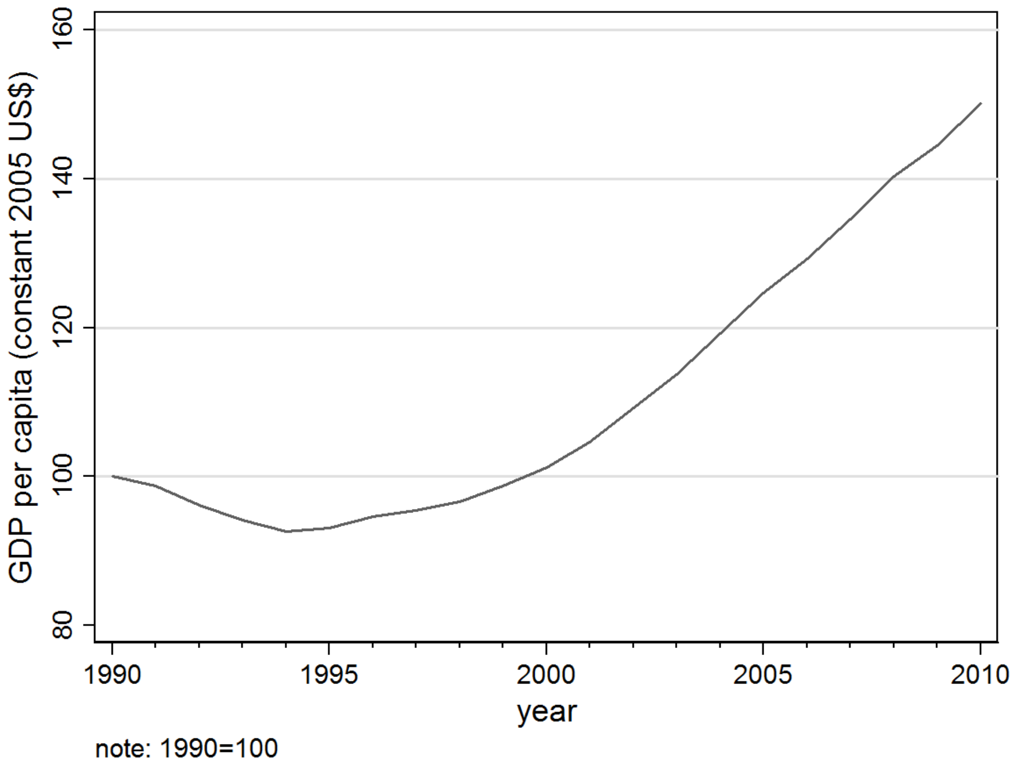

Source: own calculations from the World Bank World Development Indicators data base (World Bank 2013)

Although more than two thirds of the Tanzanian working population engage in agricultural production (URT 2009), at the macro-level the economy has become less and less dependent on agriculture (Lokina et al. 2011). This can be seen in Figure 2.2 that shows each sector's contribution to the GDP in 1990, 2000 and 2010. By the latter years, the service sector, owing largely to the growth in tourism, leapfrogged agriculture as the largest sector constituting nearly half of Tanzanian GDP in 2010. The share of the industrial sector, mainly driven by mining and construction, has nearly doubled in these two decades.

Figure 2.2: Value added as a share of GDP by sector

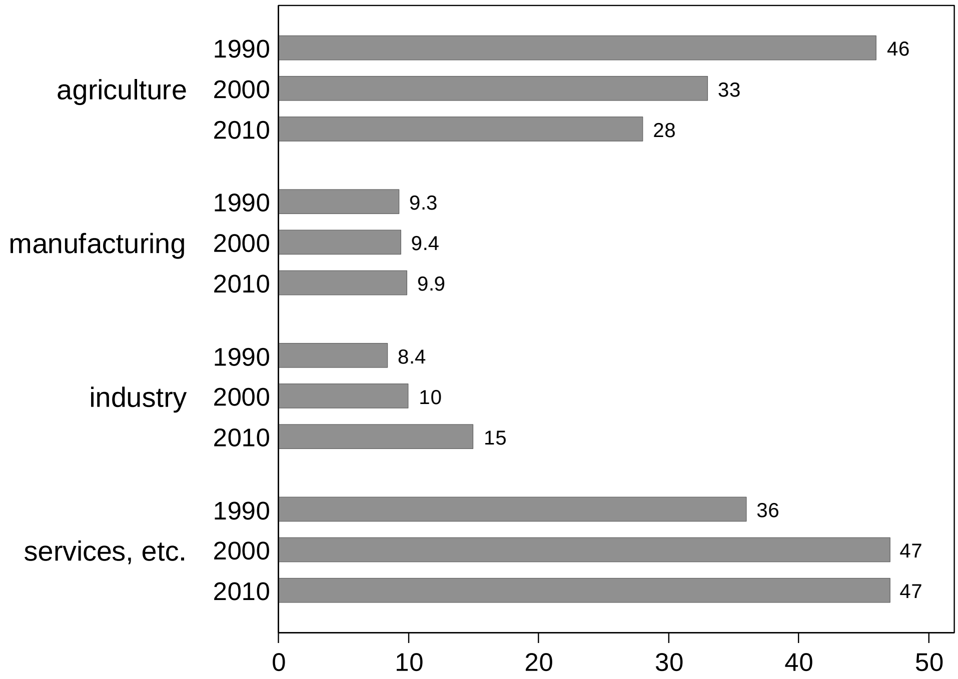

Source: own calculations from the World Bank World Development Indicators data base (World Bank 2013)

The decline of the agricultural sector is, however, only relative, not absolute. Figure 2.3 shows that the total value added in agriculture more than doubled in 1999-2010 but the other sectors recorded even more impressive growth rates. The value added in the industrial sector more than tripled over the 20 years.

Figure 2.3: Growth of the value added by sector

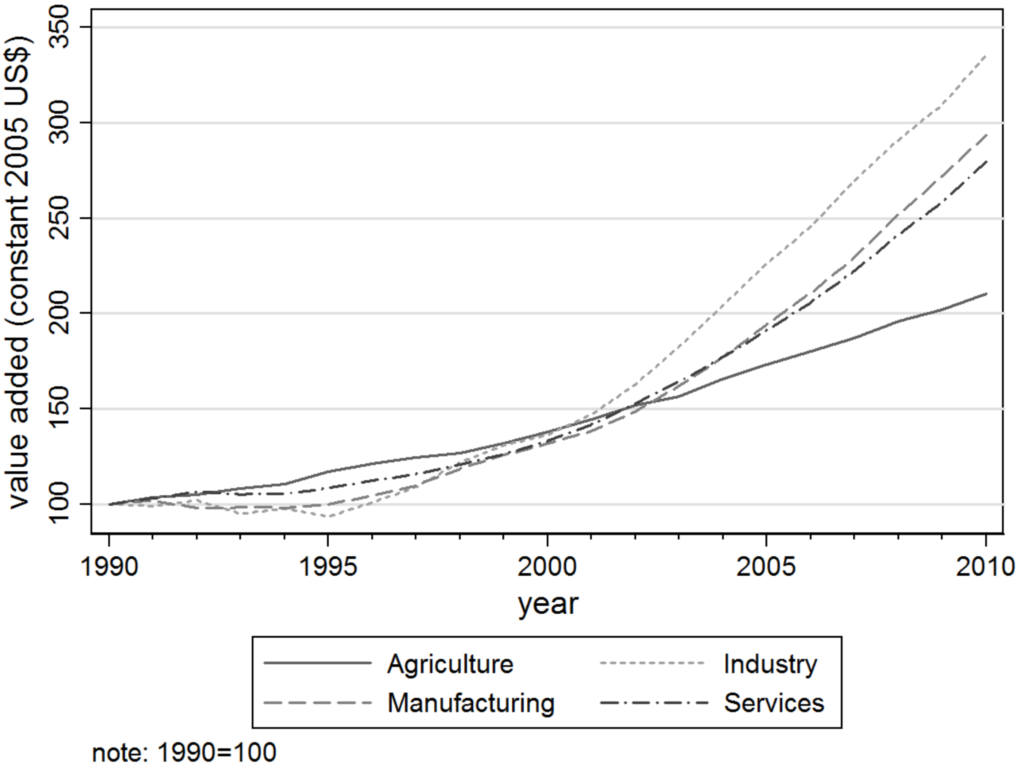

Note: these values are *not* expressed in per capita terms.

Source: own calculations from the World Bank World Development Indicators data base (World Bank 2013).

These impressive growth figures have not, however, been translated into significant poverty reduction. The data from the Household Budget Surveys indicate that although the median expenditure per capita doubled between 2000-2007, the population living under the basic needs poverty line declined by only three percentage points over the same period (URT 2009).[1] It seems then that the Tanzanian 'growth miracle' has not been pro-poor. However, this discrepancy may also be explained by the conceptual differences, including the use of different price deflators, between the national accounts and household survey data (Atkinson and Lugo 2010).

[1] According to the 2007 Household Budget Survey, the percentage of the population living under the basic needs poverty line was 39 in 1991/92, 36 in 2000/1 and 33 in 2007 (URT 2009).

The quality of African national accounts data has recently been called into question (Jerven 2013). [2] However, some non-monetary indicators also provide support to the growth story. Figure 2.4 shows that the infant mortality rate has decreased from 100 per 1,000 live births in 1990 to 42 in 2010.

Figure 2.4: Infant mortality in 1990-2010

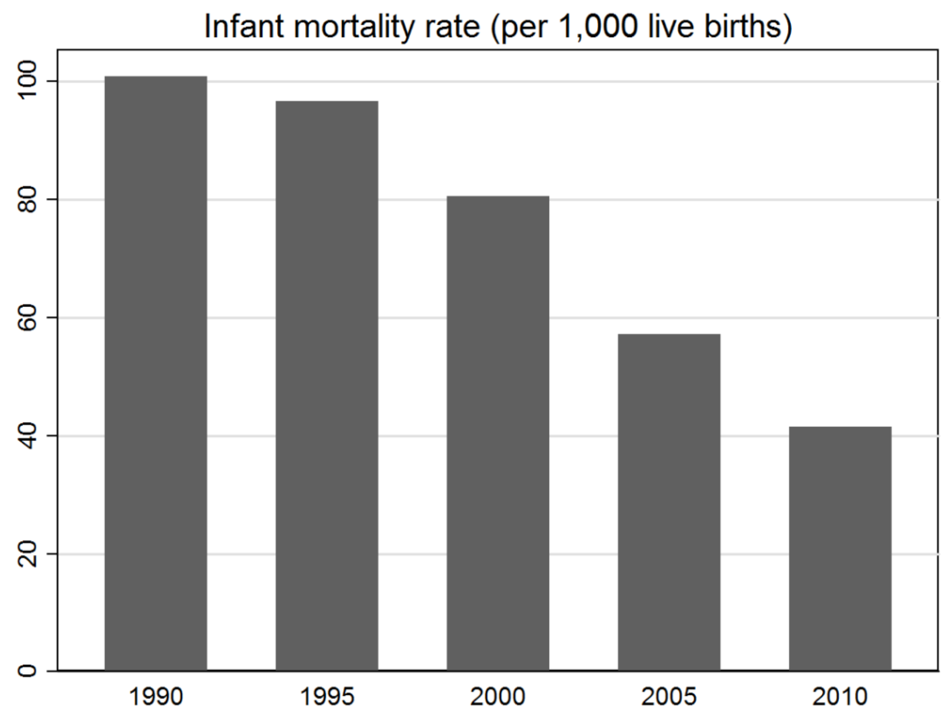

Source: own calculations from the World Bank World Development Indicators data base (World Bank 2013)

Finally, these high GDP growth rates accompanied by a declining share of agriculture follow the predictions of the early models of structural transformation (Lewis 1954, Ranis and Fei 1961) where a country develops by moving away from subsistence agriculture to other, more productive, activities. The national accounts data certainly suggest, as also

[2] Lipton (2013) provides an insightful (and entertaining) review of Jerven's book and of the 'Africa's National-accounts Mess'.

pointed out by Moyo et al. (2012), that the structural transformation is already underway in Tanzania. This, of course, should then take place together with increasing urbanization. Unfortunately, at the time of writing, the National Bureau of Statistics has not yet released the 2012 census information about urbanization rates. The current figures are therefore largely based on projections from the growth rates between the previous censuses administered in 1988 and 2002. Using these data, the United Nations (2012) estimate that the share of the urban population in Tanzania grew from 19 to 26 per cent in 1990-2010.

2.2 Kagera

Kagera is a region in the north-western part of Tanzania (see Figure 2.5). A large part of Lake Victoria is contained within this region and it shares a border with Burundi, Rwanda, and Uganda. Out of all the regions of Tanzania, Kagera is the farthest from Dar es Salaam – the commercial capital of the country. [3] The land area consists of roughly 29,000 km^2. The region is overwhelmingly rural and agricultural production is the most important source of income, with more than 80 per cent of the region's economically active population engaged in it (URT 2012). Bananas, beans, maize, and cassava comprise the main food crops while coffee, tea, and cotton are important cash crops. Recent years have witnessed a rise in improved banana varieties and sugar for use as cash crops. Only two per cent of agricultural households reported having access to irrigation in 2007/08 (URT 2012). Finally, according to the 2012 census, the region has a population of roughly 2.5 million people (URT 2013).

[3] The 'as the crow flies' distance from the region's capital (Bukoba) to Dar es Salaam is more than 1,000 km.

Figure 2.5: Kagera on a map

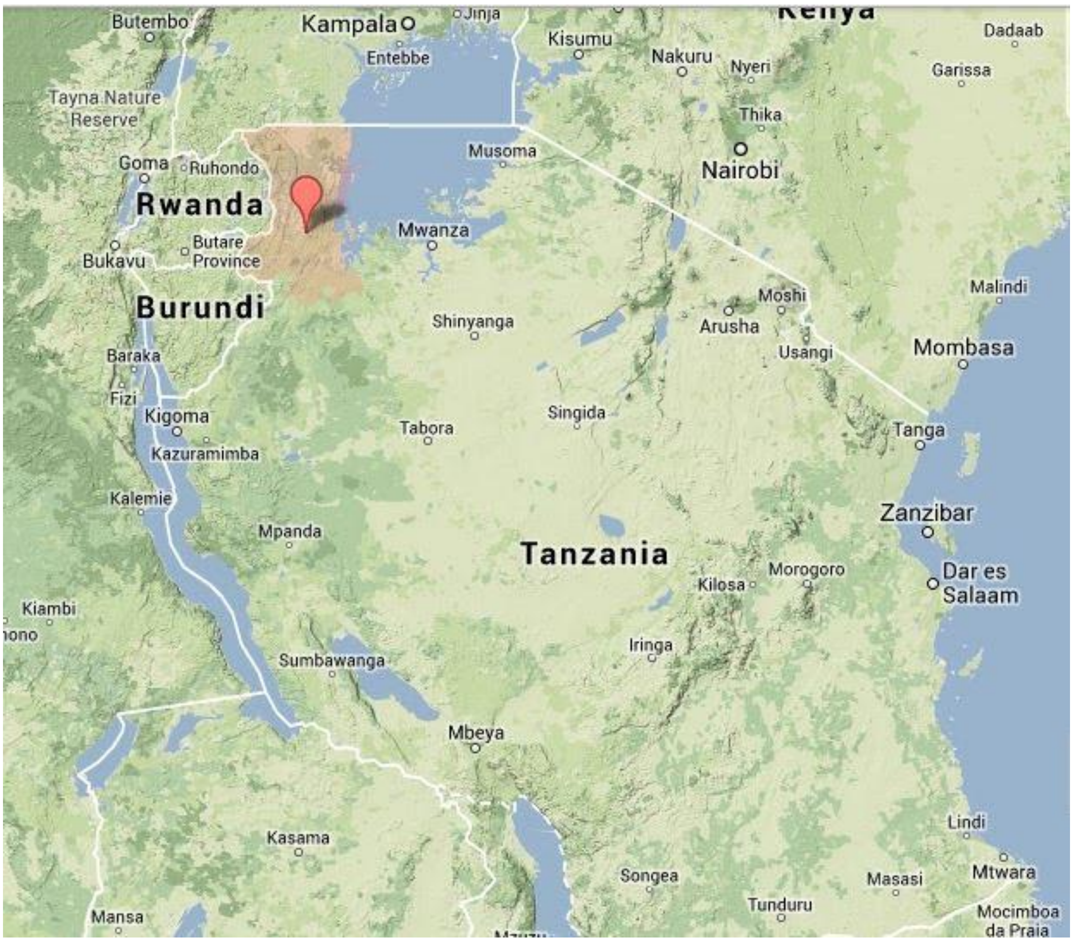

source: Google Maps (2013)

2.3 Kagera Health and Development Survey [4]

The Kagera Health and Development Survey (KHDS) was originally designed to study the economic impact of adult deaths on surviving household members. This focus was motivated by the fact that Kagera is historically a region with a high HIV/AIDS incidence

[4] This section builds on and contains passages from a previously published working paper: De Weerdt, Joachim, Kathleen Beegle, Helene Bie Lilleør, Stefan Dercon, Kalle Hirvonen, Martina Kirchberger, and Sofya Krutikova. 2012. "Kagera Health and Development Survey 2010: Basic Information Document." Rockwool Foundation Working Paper Series no. 46.

(for details, see Iliffe 2006). The first rounds of the survey were designed and implemented by the World Bank and the Muhimbili University College of Health Sciences. The survey consisted of 915 households from 51 villages that were interviewed up to four times from autumn 1991 to January 1994. Figure 2.6 shows the baseline villages on a map. The sample was drawn from a random sample stratified by adult mortality rates in the communities and the agro-climatic zones in the region. [5] The World Bank (2004) provides a comprehensive description of the baseline survey.

Figure 2.6: KHDS baseline villages

source: own construction using GPS Visualizer (2013)

[5] Comparisons of various welfare indicators with the 1991/92 Tanzanian Household Budget Survey suggest that KHDS provides a representative sample for the region during this period (Beegle, De Weerdt and Dercon, 2011).

The KHDS-2004 survey aimed to re-interview all individuals that were ever interviewed in the baseline survey and were alive in 2004. This effectively meant that the original household panel survey turned into a panel of individuals. A full household questionnaire was administered in a household where a panel respondent was found residing. Due to household dynamics, the sample size increased to more than 2,700 households. Excluding households for which all previous members were deceased (17 households and 27 respondents), the KHDS 2004 field team managed to re-contact 93 per cent of the baseline households. Beegle, De Weerdt, and Dercon (2006) provide a description of this survey round.

2.3.1 Field work for the 2010 round

The KHDS 2010 was designed to provide data to understand changes in living standards of the sample of individuals originally interviewed 16-19 years ago. Similarly to KHDS-2004, the 2010 round attempted to re-interview all respondents ever interviewed in the KHDS 91-94 – irrespective of whether the respondent moved out of the original village, region, or country, or was residing in a new household.

The KHDS 2010 was primarily funded by the Rockwool Foundation and the World Bank, with additional funds provided by the William and Flora Hewlett Foundation through the Agence Inter-établissements de Recherche pour le Développement (AIRD). The project was initiated and designed by Kathleen Beegle (Development Research Group, World Bank), Helene Bie Lilleør (Rockwool Foundation), Joachim De Weerdt (EDI) and Sonya Krutikov (University of Oxford), with crucial inputs from Stefan Dercon (University of Oxford), Martina Kirchberger (University of Oxford) and me. Neil Chalmers programmed and designed the Electronic Survey (CAPI) application. The field work was implemented in 2010 by Economic Development Initiatives (EDI) under the direction of

Joachim De Weerdt, with Respichius Mitti and Leonard Kyaruzi forming the coordination team and Thaddeus Rweyemamu in charge of data processing. I was advising on survey technical issues and questionnaire content during the field work. My other responsibilities included ensuring that the data collected met the quality standards required for rigorous statistical and econometric analysis.

Electronic survey methods

A novel feature of the KHDS 2010 compared to the earlier waves is that data were collected using electronic survey questionnaires administered on handheld computers.

Apart from saving paper, the key advantage of electronic questionnaires is the identification of errors and missing fields during the interview.[6] In KHDS 2010, most of these checks were programmed into the questionnaires from the beginning but many were added as the survey progressed based on the feedback from the field and from the data processing team. Towards the end of the survey, there were more than 1,400 in-built consistency checks in the programme. Before leaving the household, interviewers ran the final validation check and resolved any problems directly with the respondent while still in the household. This resulted in extremely clean data with very few coding errors. Furthermore, the data were recorded and stored into a relational database with referential integrity. As a result, the end user of the data will be able to merge data files with ease and there will not be any duplicate entries in the data files. Another advantage of electronic questionnaires is the immediate availability of the data. In KHDS 2010, data were uploaded from the field to a secure server using GPRS-enabled mobile phone networks. A third advantage is that data from previous waves were carried forward to the

[6] For a formal analysis of the benefits of computer-assisted personal interviewing (CAPI), using similar software, see Caeyers, Chalmers and De Weerdt (2012).

questionnaires providing the opportunity to resolve possible inconsistencies between the waves. Finally, the ability to upload data daily provided a range of opportunities to improve tracking practices involved in finding panel respondents. In particular, data files could be sent across teams located in different areas based on the updated tracking information. Furthermore, if the information was not sufficient to track the migrant, the data file was sent back to the teams still located in the baseline village or to other location for additional tracking information.

Organization of field work

The KHDS 2010 field work was managed by Economic Development Initiatives (EDI). The project headquarters were the EDI office in Bukoba, Kagera. The Bukoba office was responsible for all the contractual and legal matters related to the survey. In addition, logistical arrangements as well as the maintenance and storage of the field equipment were managed from the office. The next sections provide an overview of the recruitment procedures, main field work and tracking.

Recruitment

The recruitment for the interviewers and data processing officers started in January 2010. The four KHDS supervisors were EDI employees with past survey experience. The supervisors were involved in the development and piloting of the survey instrument. They also took part in the field work planning and the preparation of the field work manuals. The training of 26 interviewers began in February 2010. The training consisted of three weeks of class and 10 days of field training. In June, five more interviewers were trained and they joined the teams in July. These interviewers had previous survey experience and therefore field training consisted of observing other interviewers during real household visits. In total, 31 interviewers were employed for the field work. The project also had three full time data processing officers based in the headquarters.

Main Field Work

The main field work started in April 2010. It consisted of four teams of 6-8 interviewers and each team was led by one supervisor. The main field work was split into two phases. In the first phase, the four teams visited all original clusters in districts close to Bukoba. The teams visited the households that were still located in the original baseline community and also tracked individuals who had moved to nearby villages.

In the second phase, the teams split up. Two teams visited baseline communities in Northern and Southern districts of the region and tracked migrants in these areas. The other two teams were sent out to locations outside Kagera to track migrants. The next section provides details about the tracking in KHDS 2010.

Before starting the work in a community, the supervisor arranged a meeting with the village chairman and presented the necessary documents such as the research permits. The supervisor also arranged accommodation for the field team from the village. Households were grouped by the sub-villages and each interviewer was allocated questionnaires based on these locations.

The village chairman and local guides assisted in finding the respondents. The interviewers usually agreed on the date and time of the visit with the household beforehand. The interviewers completed the household and other questionnaires where applicable in one or two visits depending on the size of the household and the number of sections and questionnaires that applied to the household.

Questionnaire checks were done in four stages:

- First, before leaving the household, the interviewer ran the final validation check. The check went through all pre-written validation rules and alerted the interviewer for any errors or missing fields in the questionnaire.
- Second, in the first months of the field work, the supervisor went through the questionnaire with the interviewer and discussed any problematic issues. They also provided feedback to the headquarters if question specific instructions or new consistency checks were needed.
- Third, after the supervisor had uploaded the data to the headquarters, the data processing team went through each questionnaire. At best, the data was reviewed by the data processing team after a day or two of the interview. If the mistakes were discovered at this stage, the questionnaire was sent back to the supervisor. The supervisor then, depending on the type of problem, consulted the interviewer or arranged a revisit to the household.
- Fourth, the field management checked the overall quality of the data using statistical packages and provided feedback to the field teams.

The quality control was managed by the Data Processing Team. This included direct observations during interviews and re-visits to households. During re-visits, the data processing officers asked a set of questions from the Household Questionnaire to verify the validity of the data. These re-visits also provided an opportunity to correct questionnaires and to improve tracking information.

Tracking

The tracking phase started in August 2010 when a team of seven interviewers were sent to the neighbouring region of Mwanza, nine interviewers were sent to Dar es Salaam and one interviewer travelled to Uganda. During the tracking phase, one team of 11 interviewers remained in Kagera to visit remaining original baseline communities and track individuals based in Kagera. Each team was led by one supervisor.

In the beginning of October, most of the interviewers were set up to work independently. This facilitated fast turnover of the data files and rapid dispatching of interviewers to conduct interviews in remote parts of Tanzania. These interviewers were responsible for their logistical planning and they transferred data files directly to the headquarters. Supervisors were in close communication with these interviewers.

The main tracking phase ended at the end of October 2010. After that, the field work continued for eight weeks with four full time interviewers in Kagera and one part-time interviewer in Dar es Salaam. Supervision during this phase was organised from the headquarters. This small team was responsible for tracking remaining panel respondents and re-visiting households for questionnaire corrections.

Re-contacting rates

The re-contact rates in the KHDS 2010 are in line with the ones achieved in KHDS 2004. Table 2.1 reports the KHDS 2010 re-contacting rates in terms of the baseline households. Excluding the households in which all panel respondents were deceased, 92 per cent of the households were re-contacted.

Although KHDS is a panel of individuals and the definition of a household loses meaning after 10-19 years, it is common in panel surveys to consider re-contact rates in terms of households. Taking into account the long, 10 or 16 year periods between surveys, the attrition rates in KHDS-2004 and KHDS-2010 are extremely low by the standards of such panels (Alderman et al. 2001).

As in KHDS 2004, households that were interviewed four times at the baseline were more likely to be found in 2010. Excluding the households in which all members had died, 95 per cent of these households were re-interviewed in 2010.

Table 2.1: KHDS 2010 re-interview rates in terms of households

Number of interviews during 1991-1994	Re-interviewed	Deceased	Untraced	Total
1	22 *54%*	4 *10%*	15 *37%*	**41**
2	36 *78%*	2 *4%*	8 *17%*	**46**
3	54 *78%*	2 *3%*	13 *19%*	**69**
4	706 *93%*	18 *2%*	35 *5%*	**759**
Overall	818 *89%*	26 *3%*	71 *8%*	**915**

Notes: "Re-interviewed" means that at least one member of the baseline household was re-interviewed. "Deceased" means that all Previous Household Members are reported to be dead. "Untraced" means that no Previous Household Member was re-interviewed.

The KHDS 2010 re-contact rates in terms of panel respondents are provided in Table 2.2. The older respondents, if alive, were much more likely to be re-contacted than younger respondents. In the oldest age category, 60 years and older at the baseline, the interview teams managed to re-contact almost 98 per cent of all survivors. The length of the KHDS survey is evident in this age category, however, as almost three quarters of the respondents had passed away by 2010.

Table 2.3 provides the KHDS 2010 re-contact rates by location. More than 50 per cent of the re-interviewed panel respondents were located in the same community as in KHDS 91-94. Nearly 14 per cent of the re-contacted respondents were found in regions other than Kagera. The survey team also tracked panel respondents in Uganda where one per cent of the interviewed panel respondents were located. The location of the untraced respondents is based on the tracking data. More than half of the untraced respondents are reported to be living in Kagera.

Table 2.2: KHDS 2010 re-interview rates in terms of individuals

Age at baseline 1991-1994	Re-interviewed	Deceased	Untraced	Re-interview rate among survivors
<10 years	1403 *76%*	173 *9%*	276 *15%*	84%
10-19 years	1523 *77%*	176 *9%*	284 *14%*	84%
20-39 years	891 *63%*	369 *26%*	149 *11%*	86%
40-59 years	397 *63%*	205 *32%*	30 *5%*	93%
60+ years	122 *26%*	352 *74%*	3 *1%*	98%
Overall	4336 *68%*	1275 *20%*	742 *12%*	85%

Notes: Sample of individuals interviewed in KHDS 91-94. Age categories are based on age at first interview. "Re-interviewed" means that the person was found and was re-interviewed. "Untraced" means that the person was not found or refused to be re-interviewed.

Table 2.3: KHDS 2010 re-contact rates by location

	Number	Location	%
Baseline sample	6,353		
Re-interviewed	4,336		
		Same community	52
		Nearby village	9
		Elsewhere in Kagera	24
		Other region	14
		Uganda[a]	1
Untraced	742		
		Kagera	53
		Dar es Salaam	9
		Mwanza	9
		Other region	10
		Other country[b]	8
		Not known	11
Deceased	1,275		

Notes: Location for untraced respondents is reported by other household members from the baseline survey who were successfully located, interviewed, and able to provide location information on the respondent. In some cases, this information comes from other relatives or neighbours residing in the baseline communities.

a. KHDS 2010 tracked international migrants in Uganda only.
b. Countries to which the 58 untraced respondents had moved are: Burundi, Denmark, Kenya, Norway, Rwanda, South-Africa, Sweden, UK and USA.

2.4 What makes these data special?

In terms of length, there are only a few other comparable African household panels. The Ethiopian Rural Household Survey covers a 20-year period spanning 1989-2009. Initially, in 1989, the study included six rural villages in Central and Southern Ethiopia. In 1994, the survey added nine more villages from other parts of the country. This increased the sample size from 450 to 1,500 households. The latest round in 2009 also included a tracking survey that followed migrants from the 2004 round. Another long panel comes from Zimbabwe where Dr Bill Kinsey with his colleagues has been following 400 households over a 17-year period. Surveys were conducted in 1983, 1987 and on an annual basis from 1992 to 2000. The last round of the survey in 2000 included a tracking mode with the aim of tracing children who were under three years of age during the 1982-84 droughts. Unfortunately, this data set is not publicly available. Finally, the KwaZulu-Natal Income Dynamics Study (KIDS) from South-Africa covers a 10-year period (1993-2004). [7]

Apart from the nearly unparalleled length of the survey in Sub-Saharan African context, KHDS has a number of qualities that makes it particularly suited for investigating the research questions set out in this thesis.

The data set is ideal for migration analysis. The tracking feature of the survey means that we have rich information about the migrants, including information about their living standards. In Chapter 3, with Joachim De Weerdt, we document how migrants have grown much richer than those who remained in villages were they were first interviewed during the baseline survey in 1991-94. This implies, as discussed in Beegle, De Weerdt, and Dercon (2011), that had we not tracked migrants we would have seriously underestimated improvements in living standards using this panel sample. In Chapter 3, we exploit the fact that household dynamics over this long period has created extended family networks that have been dispersed into different parts of Tanzania and Uganda. This analysis would not have been possible without the great effort we put into tracking the panel respondents in 2010.

The 2010 round collected extensive information about the migration histories since the baseline survey in 1991-94. Among other things, these data provide the year when the migrant first left his/her baseline village. I exploit this information to develop a time-event history model to study how migration decisions respond to temperature changes.

[7] There are also a few exciting cohort studies. For example, the Young Lives Study tracks two cohorts of children in Ethiopia for a 15-year period (2002-2016).

Furthermore, the baseline data allow me to control for various household and individual level characteristics prior to migration.

We also recorded the GPS coordinates for each interviewed household. An external GPS device, linked to the CAPI application through a Bluetooth connection, took approximately 10 GPS readings per household. The average of these readings was computed and linked to the household level file. These steps ensured that the GPS data are unlikely to contain errors. In Chapter 4, I use these high-quality GPS data to link baseline villages to historical temperature and rainfall records from external data sources.

Finally, collecting anthropometric data from all adults and children older than five years is surprisingly rare (see e.g. Moradi 2010). For example, the Demographic and Health Surveys (DHS) collect anthropometric data only from children under five and from women aged 15-49. In the third essay, this feature of the survey performs a central role as I document how children in puberty experience considerable catch-up growth in height.

3 Essay 1: Risk Sharing and Internal Migration

Joint work with Joachim De Weerdt

3.1 Introduction

If, in the next decades, Africa catches up with the rest of the world, then that will almost certainly coincide with intergenerational mobility out of rural into urban areas and out of agriculture into non-agricultural activities (Lewis 1954, Harris and Todaro 1970). Historically, in both rich developed countries and fast-growing developing countries, this type of migration has moved in lockstep with development and poverty reduction (Collier and Dercon 2009). Recently, China's urban population officially surpassed its rural one: of China's 1.35 billion people, 51 per cent lived in urban areas at the end of 2011, rising from less than 20 per cent in 1980 (UN, 2012). Furthermore, UNDP (2009) reports that of the one billion migrants worldwide, three-quarters are internal migrants. With international migration open to relatively few Africans, we should expect massive internal migration to form a core part of the development process.

The scale of this demographic process is captured in the data that form the empirical basis of this essay, further motivating our focus on internal migration. These data are part of an exceptional panel data set from the Kagera region in Tanzania, spanning nearly two decades of migration and development. The 2010 follow-up survey attempted to trace all 6,353 individuals listed on the baseline 1991/94 household rosters and re-interview them irrespective of their location. Once we exclude the 1,275 individuals who had died by

2010, we are left with 4,996 baseline individuals whose 2010 locations are known.[8] Of those, 45 per cent were found residing in the baseline village, 53 per cent had migrated within the country, two per cent to another East African country (primarily Uganda), and 0.3 per cent had moved outside of East Africa. This region – not atypical of remote rural Africa – is clearly on the move, with internal migration dwarfing international migration.

We attempt to understand how this powerful current of internal migration, which is part and parcel of the modernization process, interacts with a traditional institution like informal risk-sharing to shape economic mobility and vulnerability. This is a key question because, as Munshi and Rosenzweig (2006, p. 1230) put it

> [...] a complete understanding of the development process must not only take account of the initial conditions and the role of existing institutions in shaping the response to modernization and globalization, but must also consider how these traditional institutions are shaped in turn by the forces of change.

Geographical mobility in rural Tanzania is associated with large income gains. Our data show that despite only minor welfare differences during the 1991-94 baseline survey, those who moved out of the region to other parts of Tanzania have grown roughly twice as rich as those who did not by the time they were interviewed again nearly two decades later. As we are measuring consumption and not income, it is clear that the main beneficiaries of this migration-led growth were the migrants themselves and certainly not their relatives who remained at home.

---

[8] We lack location information on 82 individuals. Because this is after multiple attempts through various sources it is unlikely that these individuals have moved outside of East Africa. Information on such an important, low-occurrence event is unlikely to be hidden.

But did these migrants simply leave and never look back, or did they maintain links with the home community? We investigate this question by exploiting the fact that the 3,314 households interviewed in 2010 are grouped in 817 geographically dispersed extended family networks. Using techniques from the risk-sharing literature, we quantify how migrants' consumption responds to shocks experienced by others in their extended family network. Much of the migration literature has a very strong focus on dealing with the selectivity of the migration decision. Interestingly, in this essay, the endogeneity of migration turns out to be irrelevant for our most important contribution: the documentation of the long-run dynamics of risk-sharing arrangements among extended family members in a context of large internal migration flows. Whether or not migration is causally responsible for any of our findings is an interesting, but secondary question, which we will nevertheless attempt to address in Section 3.8.

Our analysis departs from a number of other studies in the migration literature by focusing on consumption instead of transfers. This choice of the outcome variable is motivated by the fact that risk sharing and other economic exchange could occur through a multitude of different mechanisms, of which transfers is just one. Other mechanisms could include looking for a job for someone, employing them directly, providing them with tips, advice or a network link, or providing migration opportunities (Munshi 2003). By analyzing consumption we focus on the joint and final effect of all such mechanisms.

The observed divergence between migrants and non-migrants in these data also persists within extended family networks (Beegle, De Weerdt, and Dercon 2011), which violates the full risk-sharing hypothesis (Townsend 1994), and does not support the notion that migration is the result of a household level maximization strategy (Stark and Bloom 1985, Rosenzweig and Stark 1989, Grimard 1997). It could, however, be consistent with other reciprocity-based models (e.g. limited commitment, moral hazard, or hidden income) that permit the co-existence of divergent consumption growth and risk sharing. In our empirical analysis, we find that migrants are affected by shocks to others in the network whereas non-migrants are not. Such unilateral insurance leads us to reject the reciprocity-based risk-sharing models.

One explanation for this observed lack of reciprocity could be that migrants insure non-migrants in exchange for other benefits (Lucas and Stark 1985, Hoddinott 1994). These benefits could accrue to the migrant later in life and outside the purview of our survey data. We consider, but reject a number of such longer-run transactional motives for the observed unilateral insurance. The results are, however, very much in line with findings from the diverse literature on social norms (Platteau 2000, Cox and Fafchamps 2007, Burke and Young 2011), where those who move ahead remain obligated to their extended family in the home community.

Our results speak to an emerging literature that worries about home communities imposing a stifling 'kin-tax' on the upwardly mobile. Baland, Guirkinger, and Mali (2011) show how people take out costly loans in order to conceal their income, while Platteau (2012) sees migration as a means to escape the implied prying eyes and incessant demands of the kinship group. The kinship poverty trap model of Hoff and Sen (2006) predicts possible resistance from the home communities as they feel threatened by productive forces leaving and severing links with home to escape taxing demands for assistance. Anticipating this, the home community may set up subtle exit barriers, which could lead to below-optimal levels of migration. Jakiela and Ozier (2012) report laboratory evidence from Kenya that women feel obliged to share four to eight per cent of the income gains realized in the experiment. In our sample, Tanzanian migrants sacrifice 2.9 to 5.0 percentage points out of a total growth of 108 per cent to insure their relatives. This estimate is equivalent to a ‘tax’ of 2.7 to 4.6 per cent. We regard this tax rate as too trivial to exert any constraining effect on migrants.

After describing the model, the data and the estimation strategy in the next three sections, we discuss the results in Section 3.5. Section 3.6 tackles the issue of longer-run transactional motives. In Section 3.7 we calculate the cost of this insurance provision. Section 3.8 discusses the endogeneity of migration and Section 3.9 contains some further robustness checks. Section 3.10 provides a concluding discussion.

3.2 Risk sharing in theory [9]

The full risk-sharing hypothesis is based on the idea that the network acts as if it was a single household that maximized utility subject to a joint budget constraint. The model predicts that incomes are completely pooled (according to predetermined weights) and all idiosyncratic income shocks are smoothed through the network (e.g. Altonji, Hayashi, and Kotlikoff 1992, Townsend 1994).

In a simple two-household extended family network, both households derive utility from consumption: $v(c)$. Insurance and credit markets are missing, and income (y_s) is uncertain and depends on the state of the world (s). [10] We assume that households live infinitely. [11]

[9] We thank Stefan Dercon who alerted us to the attractiveness of contrasting full and partial insurance concepts through Equations (3.1) and (3.4).
[10] To simplify notation, we abstract away savings. This does not affect the main predictions of the model (see, for example, Ligon, 1998 for a characterisation of the full risk sharing model with savings). However, the ability to save may exacerbate the efficiency problems if the key assumptions listed below do not hold (see Ligon, 1998; Chandrasekhar, Kinnan and Larreguy, 2012).
[11] If the time frame is finite, in the absence of altruism, households would not have any incentive for risk sharing in the final period, and as result in T-1, T-2, etc. The assumption of an infinite time frame holds if the new household head inherits from the previous head and maintains the risk sharing contract with same households. See Fafchamps (1992) for an alternative justification for this assumption.

Assuming that households maximize a well behaving utility function [12], the standard utility maximization problem yields a following first order condition:

$$\frac{u'[c_1(y)]}{u'[c_2(y)]} = \lambda = \frac{\omega_2}{\omega_1}, \tag{3.1}$$

where λ, the Lagrange multiplier, is the marginal utility of income. According to Equation (3.1), households equate their marginal utilities of consumption in all states of the world. The allocation depends on the Pareto weights ω_1 and ω_2 that are determined by the extended family.

If utility functions follow a constant relative risk aversion function: $u(c) = \frac{c^{1-\psi}}{1-\psi}$, where ψ is a measure of risk aversion [13], the first order conditions for household i at time t become: $\omega_i c_{it}(y)^{-\psi} - \lambda = 0$. Equating these conditions for the two households, taking logarithms and re-arranging yields:

$$\Delta \ln c_1\,(y) = \Delta \ln c_2\,(y). \tag{3.2}$$

Equation (3.2) implies that if full risk sharing takes place, we should not expect to see households within the same extended family growing at different rates. Furthermore, assuming that there are no frictions between the households in the extended family, the model predicts that all idiosyncratic shocks experienced by households are completely smoothed through the extended family. These two predictions form the basis for our test – and rejection of the full risk-sharing model. First, the descriptive statistics in Section

[12] The utility function is inter-temporally separable, strictly increasing but concave ($v' > 0$ & $v'' < 0$).
[13] We assume that the risk preferences within the networks are identical. The implications of this assumption are discussed in Section 3.4.

3.3 confirm highly unequal consumption growth between migrants and non-migrants. Second, in Section 3.5 we show that after controlling for extended family fixed effects, household consumption growth remains responsive to idiosyncratic income shocks.

The rejection of full risk sharing is neither novel nor surprising and emerged as an empirically established stylized fact early on within this strand of literature, being valid across a variety of different contexts (e.g. Altonji, Hayashi, and Kotlikoff 1992, Townsend 1994, Grimard 1997). Most studies, however, find that at least some degree of insurance takes place and explain this theoretically by adding additional constraints (relating to the failure of assumptions regarding perfect information and full commitment) to the full risk-sharing model. We will discuss each of these constraints in turn, but point out that an important common feature across all these augmented models is that, if the risk-sharing contract survives, the ratios of marginal utilities become state contingent and are, therefore, no longer constant over time. This could allow the share of some members (migrants in our case) to increase over time.

In the presence of enforcement problems, the better-off households have an incentive to leave the arrangement and live in autarky. The limited commitment model (e.g. Coate and Ravallion 1993, Attanasio and Ríos-Rull 2000, Ligon, Thomas, and Worrall 2002, Kinnan 2012) augments the full risk-sharing model with participation constraints (one for each household):

$$\sum_{t=1}^{\infty} \beta^t \sum_{s=1}^{S} \pi(y_s)\{v_1[c_{1t}(y_s)]\} \geq \mathrm{u}_A \,, \tag{3.3}$$

where u_A is the expected utility received in autarky, β is the discount rate and π is the probability attached to the state of the world s. Solving the augmented maximization problem yields a following first-order condition:

$$\frac{u'[c_1(y)]}{u'[c_2(y)]} = \frac{\omega_2 + \sum_{s=1}^{S} \mu_2(y_s)}{\omega_1 + \sum_{s=1}^{S} \mu_1(y_s)} \tag{3.4}$$

where μ_1 and μ_2 are the Lagrange multipliers attached to the participation constraints. Now, as can be seen from Equation (3.4), if the participation constraints bind, the ratio of marginal utilities becomes state contingent. In the context of migration, a growth premium has to be granted to the migrant whose autarky options have improved. As a result, risk sharing is no longer efficient: the impact of idiosyncratic income shocks is not equally shared within the extended family network.

The other frictions have similar analytical consequences. If households cannot monitor other network members, the problem of free riding emerges. In moral hazard models (Lim and Townsend 1998, Kinnan 2012), the full risk-sharing model is augmented with incentive-compatibility constraints. The *ex ante* information asymmetry leaves the extended family to balance effort and insurance. Migrants are motivated to exert effort by rewarding them with higher consumption. This comes with an efficiency cost: idiosyncratic shocks are not completely smoothed within the network. Finally, if there is imperfect information about the realized incomes, households may have an incentive to misreport their incomes to avoid payments or even claim transfers from other households. In hidden income models (Townsend 1982, Fafchamps 1992, Kinnan 2012), the maximization problem is augmented with truth-telling constraints that require that households will not gain from misreporting. To encourage truthful reporting, migrants are allowed to enjoy a larger share of the consumption cake. As a consequence, Pareto-efficient risk sharing is again sacrificed.

These frictions can have important implications for the degree of risk sharing. [14] Distinguishing which of the three models of constrained insurance explains our data best is beyond the scope of this essay. [15] One common feature, however, is that despite friction, reciprocity remains intact: households engage in reciprocal risk sharing but the degree of its efficiency varies. In Section 3.5, we study the existence of such reciprocal but partial risk-sharing arrangements by testing whether households are responsive to income shocks faced by other households in the same extended family network.

In this essay, we contrast these reciprocity-based models with those that take into account social norms. Redistributive values may have been instilled since childhood and carefully nurtured through oral transmission, rituals and ceremonies in which the importance of the kinship group is strongly emphasized (Lévi-Strauss 1969). Remittances and other forms of assistance may also buy social prestige, political power or serve to perpetuate subordination (Platteau 2012, Platteau and Sekeris 2010). In the risk-sharing literature, social norms have been seen as the glue that keeps the risk-sharing contract from breaking apart by alleviating enforcement and information problems (Stark and Lucas 1988, Fafchamps 1999, Foster and Rosenzweig 2001). Theoretically this can be modeled as subjective satisfaction that individuals receive from participation. [16] The satisfaction can stem from the fulfillment of obligations and the avoidance of social sanctions, such as

---

[14] For example, Chandrasekhar, Kinnan and Larreguy (2011, 2012), using field experiments from Southern India find that limited commitment reduces transfers by 10 per cent and hidden income by 40 per cent.
[15] See Kinnan (2012) for such an exercise with data from rural Thailand.
[16] In the context of limited commitment, we can re-write the right-hand side of Equation (3.3) as $\mathrm{u}_A - A$, where A captures such satisfaction (Fafchamps, 1999; Foster and Rosenzweig, 2001; De Weerdt and Fafchamps, 2011).

guilt, shame or ridicule, or fear of witchcraft. It can also include altruism, which we do not attempt to distinguish from social norms.

A recent empirical literature relying on experimental design highlights the importance of these forces. Chandrasekhar, Kinnan and Larreguy (2011, 2012) find that in the presence of hidden income and limited commitment, social proximity between the risk-sharing partners increases the amounts transferred. The field experiments of Leider et al. (2009) and Ligon and Schechter (2012) show that altruism is more important than repeated interaction in determining the size of the transfer. Furthermore, social norms could weaken the constraints to risk sharing to the extent that they never bind and allow for the existence of sustained, unreciprocated transfers. [17] Below we will find evidence of such unilateral relations and argue that this is consistent with risk sharing motivated by social norms.

3.3 Data and descriptive analysis

This essay exploits the fact that the 2010 round of the survey includes all tracked split-offs from the original household and contains particularly rich information on the current links between them. The 2010 sample contains 3,314 households, originating from 816 initial households. The average baseline household spawned 4.1 households by 2010, out of which 1.8 were non-migrant and 2.3 were migrant households. Approximately three per cent of the initial households (99 households) did not have any split-offs. In what follows we will refer to these networks as extended family networks.

Figure 3.1 provides an overview of migration patterns. By 2010 nearly 45 per cent of the households were still residing in their original baseline community. We define migrants

[17] Schechter and Yuskavage (2011) empirically document unreciprocated relations in Paraguay.

as households that in 2010 are not located in the original village but were found in a nearby village, elsewhere within Kagera or outside Kagera. [18]

Figure 3.1: KHDS-2010 – Re-contacting after 16+ years

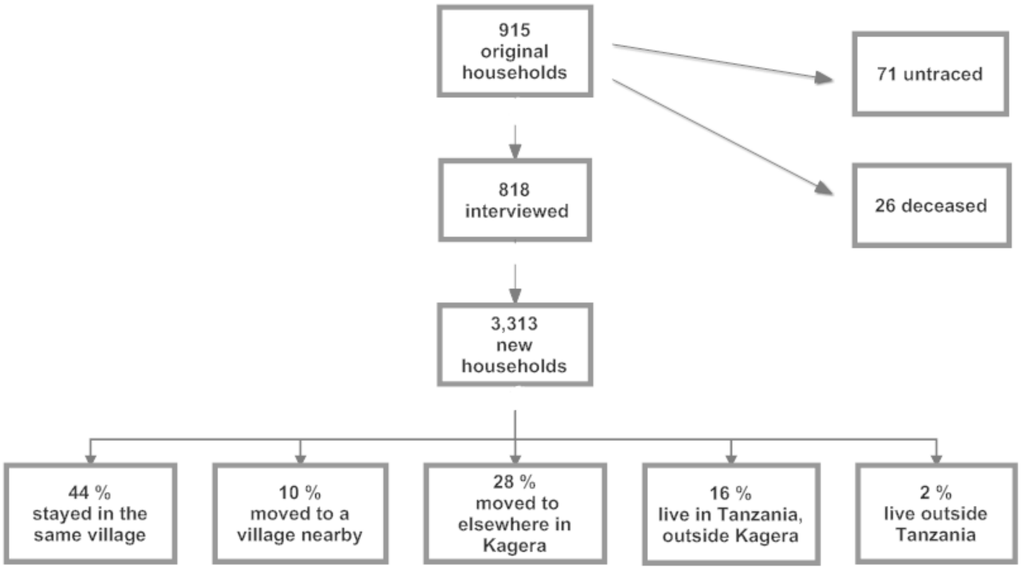

These internal migration flows are associated with structural transformation in our data.[19] Table 3.1 shows that out of the 1,850 migrant households, only one-third reported agriculture as their main-income generating activity. For the 1,460 non-migrant households this is nearly 65 per cent. More than 25 per cent of the migrant households engage in informal or formal wage employment and 11 per cent are self-employed (non-agriculture). These income generating activities differ across locations. Migrants who move farther from the baseline village are also less likely to engage in agriculture and more likely to be in wage employment. Finally, only a handful of migrants residing in urban areas (cities, in district or regional capitals) report agriculture as their main income generating activity. These differences in the income portfolios between migrants and non-

[18] Our results are robust to alternative migrant definitions, such as also defining households that moved to a nearby village as non-migrant households.
[19] This is also documented by Christiaensen, De Weerdt, and Todo (2013) who use the same data to study the role of urbanization and diversification in poverty reduction.

migrants provide scope for mutually beneficial risk sharing arrangement since there is less likely to be strong co-movement of incomes between the two parties. [20]

Table 3.1: Main income generating activity by migrant status

	non-migrant HHs	migrant HHs					
		All	nearby village	elsewhere in Kagera	outside Kagera	rural	urban
	%	%	%	%	%	%	%
agriculture	64.9	33.0	51.3	41.9	8.5	54.5	6.5
fishing	1.8	1.7	0.6	2.6	0.9	2.4	0.8
trading	11.7	17.2	17.2	15.1	20.7	14.7	20.4
wage employment	6.2	26.8	12.0	20.1	45.8	13.1	43.6
transfers & savings	1.2	2.5	1.8	1.4	4.6	0.5	5.0
self-employed	8.8	11.2	10.5	9.5	14.4	7.1	16.4
casual labour	5.5	7.6	6.7	9.5	5.3	7.8	7.4
number of HHs	**1,460**	**1,850**	**343**	**917**	**590**	**1,021**	**829**

Note: Agriculture category includes farming and livestock keeping, trading includes agriculture and non-agricultural trading. Wage employment can be either in formal or informal employment. Transfers include pensions, remittances and rental income. Self-employed category only considers self-employment outside agriculture. The information is missing for 2 non-migrant and 5 migrant households.

Remittances offer one possible medium for risk sharing between households Table 3.2 provides a summary of the average remittance flows over the past 12 months in 2010 between the migrant households and households living in or near their baseline villages. While non-migrant households were net receivers of remittances, Table 3.2 shows that transfers flow both ways. This could lead one to think – mistakenly as the analysis below reveals – that these are relationships of reciprocal risk sharing. The data in Table 3.2 are self-reported and it is interesting to note that migrants claim to send more home than non-

[20] Strong correlation of income shocks between migrants and non-migrants would be problematic to our econometric analysis. In Section 3.5 we note that the intra-class correlation of shocks in the networks is close to zero.

migrants acknowledge. A similar discrepancy does not exist in migrant-migrant or stayer-stayer dyads [21].

Table 3.3 provides an overview of the reasons for leaving the baseline village. More than one-third of the female respondents but none of the male respondents cited marriage as the reason for migrating, which is what one would expect in a culture with patrilocal marriages. Less than 15 per cent of the female respondents reported that they left because of work. In contrast, almost 45 per cent of the male migrants reported to have moved because they had found work or went looking for work. Despite these differences in migration motives across the two gender groups, we do not find any statistically significant differences in risk sharing between male and female migrants (see Appendix A).

Table 3.2: Reported remittances in and out between migrants and non-migrants

Dyad	gifts out	gifts in	net in
stayer-migrant	8,920	13,577	4,656
migrant-stayer	19,044	10,567	-8,477
stayer-stayer	8,310	10,476	2,167
migrant-migrant	17,096	14,816	-2,280
Total	**13,208**	**12,328**	**-880**

Note: The table gives the amount of transfers as reported by the first half of the dyad with respect to transfers to (first column) or from (second column) the second half of the dyad. This data are based on self-reported remittance flows in households in the past 12 months in 2010.

[21] By 'dyad' we refer to a pair of households.

Table 3.3: Reasons for leaving the baseline village

Reason	males (%)	females (%)
To look for work	29.8	7.5
Own schooling	16.0	10.3
Found work	15.1	6.7
To live in a healthier environment	10.4	11.7
Marriage	0.0	38.9
Other reason	28.7	24.9
Total	**100.0**	**100.0**

The consumption data originate from extensive food and non-food consumption modules in the survey, carefully designed to maintain comparability across survey rounds and controlling for seasonality. The aggregates are temporally and spatially deflated using data from a price questionnaire included in the survey. Consumption is expressed in annual per capita terms using 2010 Tanzanian shillings. [22]

Table 3.4 provides the summary of the consumption and poverty developments of the panel respondents with respect to their 2010 location. On average, consumption levels in the sample almost doubled over 19 years. Individuals who stayed in their community saw their consumption increase by more than 40 per cent. Consumption growth for migrants was much higher: those who left Kagera saw their consumption nearly triple over the same two decades. The poverty statistics tell the same story: nearly all respondents who left the region managed to escape poverty, while poverty reduction among non-migrants was more modest. These descriptive statistics reinforce the results reported in Beegle, De Weerdt, and Dercon (2011): individuals who moved did considerably better than those who decided to stay.

[22] Using adult equivalent units as the denominator instead of household size produces almost identical results across all specifications.

Table 3.4: Consumption and poverty movements of the panel respondents in 1991-2010 by 2010 location

	mean 91	mean 2010	Δ in means	N
Consumption per capita (Tsh) by 2010 location				
Within community	343,718	492,398	148,680***	2,224
Nearby community	364,099	569,438	205,339***	382
Elsewhere in Kagera	357,930	695,951	338,021***	1,007
Out of Kagera	389,379	1,110,827	721,449***	658
Full Sample	355,926	642,558	286,632***	4,271
Consumption Poverty Head Count (%) by 2010 location				
Within community	31	19	-13***	2,224
Nearby community	30	20	-10***	382
Elsewhere in Kagera	31	16	-15***	1,007
Out of Kagera	23	3	-21***	658
Full Sample	30	16	-14***	4,271

Note: All consumption values are in annual per capita terms and expressed in 2010 Tanzanian shillings. Significance of the difference in means using a paired t-test; *** p<0.01, ** p<0.05, * p<0.1.

After moving, migrants remain linked to extended family members at home: 90 per cent of the migrants report that they communicated with a non-migrant network member in the 12 months preceding the survey. Migrants who maintained some form of communication experienced an average consumption growth of 110 per cent, while those who did not grew by 88 per cent. [23] This difference is statistically significant at the one per cent level. The severing of the most basic links does not seem to be associated with higher consumption growth – if anything, the reverse is true.

We use data from shock modules administered in 2004 and 2010. During both of these rounds, the panel respondents were asked to consider each year between the survey rounds and indicate whether a particular year was, in economic terms, 'Very good', 'Good',

[23] The mean consumption growth among those who maintained contact was 394,679 Tsh and among those who severed links 286,991 Tsh.

'Normal', 'Bad', 'Very bad'. For each 'Very bad' response, the respondents were asked to provide the main reason for the hardship. We consider each 'Very bad' response as an economic shock. More than 60 per cent of the panel respondents reported experiencing at least one such shock between 1994 and 2009.

Table 3.5 provides an overview of the shocks experienced. Most frequently reported economic shocks were death of a family member, serious illness, and a poor harvest due to bad weather.

Table 3.5: Shocks reported by the panel respondents 1994-2009

Type of shock	**Freq.**	**Percentage**
Death of family member	797	26%
Poor harvest due to adverse weather	638	21%
Serious illness	577	19%
Loss in wage employment	219	7%
Loss of assets	205	7%
Eviction/resettlement	99	3%
Poor harvest due to pests or crop diseases	98	3%
Low crop prices	85	3%
Loss in off-farm employment	78	3%
Low income due to lower remittances	43	1%
Loss of livestock	6	0.2%
Loss of gifts and support by organizations	4	0.13%
Other reasons	172	6%
Total	**3,021**	**100%**

The shock data were collected at the individual level – in particular for each person on the 2010 roster who also appears on the original 1991/94 rosters. Since our focus is to examine the role of shocks on household consumption, the data had to be reformatted from the individual to the household level. [24] If at least one individual in the household

[24] We repeated the complete analysis of the following sections using individual level data and find it does not change the conclusions.

reported to have experienced a shock, we interpret it as a household level shock. We should also exclude shocks that occurred before the households split. Fortunately, we know the year in which the respondents moved to their 2010 location, allowing us to include only shocks that occurred at least one year after this move. [25]

Furthermore, some of the shock categories are problematic to our network analysis. Mortality shocks may trigger inheritance flows within extended families. As such, a negative shock in one household may actually be a positive income shock in another household. A similar problem arises with the loss of remittance shocks, if these capture the loss of transfers from a household within the same extended family. We therefore exclude these two shock categories from our final shock variable.

Another concern is that because of the self-reported nature of the shock variable uninsured shocks may go unreported or reports may differ along unobserved characteristics. To address this, in Section 3.7, we extend our analysis with an alternative shock measure based on historical rainfall data collected from 212 weather stations in Kagera and the migration destinations of our panel respondents.

Finally, there are 439 households that belong to a network that contains only non-migrants or only migrants. As our interest lies in the role of migration in risk sharing, we cannot use these households for empirical identification for risk sharing between migrant and non-migrant households. These households are therefore dropped from the final sample. Table 3.6 presents the summary statistics for the final sample of 2,349 households by 2010 migration status.

---

[25] This means that for households that remained in the baseline village we consider shocks that took place between 1994 and 2009. An alternative strategy would be to only use shocks that occurred after these household lived with *any* other network household member. Applying this strategy does not, however, change the conclusions of this chapter.

Table 3.6: Descriptive statistics

	Migrant households		Non-migrant households	
	mean	std dev	mean	std dev
1991 household per capita consumption	355,038	193,321	344,095	188,122
2010 household per capita consumption	739,033	634,925	488,830	358,197
per capita consumption growth in 1991-2010	383,995	643,235	144,736	352,247
natural log of per capita consumption growth in 1991-2010	0.5754	0.793	0.2944	0.618
own shock	0.2102	0.408	0.5320	0.499
# of hhs that reported a shock in the network	0.7937	1.120	1.4467	1.322
2010 household characteristics :				
age of oldest PHHM in the 2010 hh	31.376	11.132	44.405	18.288
a PHHM is head of this 2010 hh	0.4094	0.492	0.8100	0.392
a PHHM is spouse of this 2010 hh's head	0.4528	0.498	0.3058	0.461
a PHHM is child of this 2010 hh's head	0.0559	0.230	0.2067	0.405
divorced PHHM in 2010 hh	0.0433	0.204	0.0556	0.229
a widowed PHHM in 2010 hh	0.0417	0.200	0.1881	0.391
a married PHHM in 2010 hh	0.6614	0.473	0.6942	0.461
max yrs edu of PHHM in this 2010 hh	6.7811	3.083	6.1696	2.945
number of PHHMs in this 2010 hh	1.1024	0.419	1.5681	1.021
household size in 2010 hh	4.4732	2.450	4.8406	2.322
hh size in aeu in 2010 hh	3.5236	1.939	3.8314	1.878
Initial household characteristics:				
natural log value of assets	13.7058	1.100	13.7210	1.061
Educ of hh head	4.4189	3.139	4.2132	2.979
head was male	0.7646	0.424	0.7850	0.411
Age of hh head	48.9764	15.697	48.9296	15.599
age of head squared	2,645	1,596	2,637	1,575
Males 0-5 years	0.7622	0.896	0.7090	0.875
Males 6-15 years	1.3283	1.188	1.3040	1.123
Males 16-60 years	1.3756	1.022	1.4365	1.059
Males 61+ years	0.1913	0.394	0.2048	0.404
Females 0-5 years	0.8386	0.959	0.7609	0.877
Females 6-15 years	1.4591	1.340	1.3661	1.246
Females 16-60 years	1.8929	1.320	1.7822	1.186
Females 61+ years	0.2236	0.446	0.1937	0.407
hh had a non-earth floor in 1991	0.1811	0.385	0.1455	0.353
Observations	1,270		1,079	

Note: PHHM refers to previous household member (i.e. person interviewed at the baseline).

3.4 Econometric strategy

We begin the econometric analysis by testing the full risk-sharing hypothesis for those extended family networks that contain both migrant and non-migrant households. The difference in logged per capita consumption between 2010 and the baseline ($\Delta \ln c_{ij}$) for household i in extended family j is formally modeled as:

$$\Delta \ln c_{ij} = \beta s_{ij} + x_{ij}' \gamma + \alpha_j + \varepsilon_{ij} \,, \tag{3.5}$$

where s_{ij} has a value 1 if the household experienced a shock in 1994-2009, or if a migrant household, after migrating to its current location. The term x_{ij} represents a vector of household characteristics in 2010 capturing the characteristics of the previous household members [26] such as the number of previous household members in the 2010 household, the age of the oldest and the education (in years) of the most educated previous household members in the household. We also include dummies capturing their relationship to the 2010 household head and their marital status. [27] The term α_j represents the network fixed effect and ε_{ij} is the error term. The inclusion of the network fixed effects means that we compare the impact of shocks between the households originating from the same initial household. As such, the full risk-sharing model presented earlier requires that $\beta = 0$.

The rejection of the full risk-sharing model using Equation (3.5) implies either that the risk-sharing arrangement is not efficient – or that the network does not engage in risk sharing at all. The rejection may also stem from the violation of the assumption that the risk preferences are identical within the network (Schulhofer-Wohl 2011, Chiappori et al.

[26] Previous household member refers to a person interviewed at the baseline in 1991/94.

[27] To address concerns about some of these 2010 household characteristics variables being potentially endogenous, we run all main regressions again, but drop each of these control variables in turn. We find the shock and network shock coefficients remain stable across all such re-specifications.

2011, Mazzocco and Saini 2012). [28] To explore the existence of reciprocal risk sharing, we assess whether household per capita consumption growth is responsive to shocks experienced by other households in the same extended family. This test builds on Equation (3.5). We drop the network fixed effects and replace them with baseline village fixed effects (θ_v) and network characteristics (w_j) comprising the number of migrant and non-migrant households in the network and variables capturing characteristics of the initial household (such as its demographic composition, the household head's characteristics, including education, gender, age and the quadratic of age). We also include (logged) per capita consumption at the baseline ($\ln c_{j,1991}$). The network shock variable, z_{ij}, measures the number of households affected by an income shock. The household's own shocks are excluded from this variable. The partial risk-sharing specification is formulated as:

$$\Delta \ln c_{ij} = \beta s_{ij} + \delta z_{ij} + x'_{ij}\gamma + w'_j\vartheta + \gamma \ln c_{j,1991} + \theta_v + \varepsilon_{ij}. \tag{3.6}$$

A negative and statistically significant δ would imply that some risk sharing takes place within the extended families.

We will assess the impact of these network shocks separately for migrant and non-migrant households and fully acknowledge that selection into migration is unlikely to be random. The differences in the observed level of risk sharing may be caused by migration or by some unobserved characteristics that differ between migrant and non-migrant households, or by some combination of both. As a result, these regressions do not allow us to say

[28] In a context of heterogenous risk preferences, Pareto-efficient contract allocates more aggregate risk to less risk-averse households. As demonstrated by Schulhofer-Wohl (2011), Chiappori et al (2011) and Mazzocco and Saini (2012) this would lead to an upward bias in β in Equation (3.5). The standard full risk sharing test is then biased against the null-hypothesis of full-risk sharing.

whether migration is causally responsible for the migrant taking on the role of insuring sedentary extended family network members, or whether the effect is driven by unobservables. In particular, we cannot make any statements about what would have happened if migrants had stayed home or the home-stayers had migrated. It is possible that in this parallel universe roles would have switched (migration is causally responsible) or not (it is driven by the unobserved differences between migrants and non-migrants). Our primary contribution lies in documenting the fact that migrants provide unilateral insurance to non-migrants, while at the same time shooting ahead of them in consumption terms. Nonetheless, we will dedicate Section 3.9 to shedding light on how selection fits into this analysis and will utilize information from additional survey rounds to exclude certain possible types of endogeneities.

Finally, the baseline per capita consumption variable in Equation (3.6) raises a concern about endogeneity. The error term ε_{ij} could be correlated, for example due to measurement error, with the lagged consumption variable on the right hand side of Equation (3.6). This would then bias the estimate measuring the impact of the lagged consumption but it may also affect other coefficients. Fortunately, we can think of a credible instrument that allows us to assess this possibility. Rainfall is one of the main inputs in agricultural production in Kagera and poor rainfall (i.e. droughts) can have serious consequences for incomes. Excess rains are less of a problem due to the focus of the production on tree crops and also because the terrain is relatively undulating. The region has two rainy seasons, a long rainy season usually between March and May and a short rainy season usually between October and December. The agricultural production takes place during these seasons. Therefore, we employ average monthly z-score deviations of rainfall during the two rainy seasons preceding the interview and truncate the positive rainfall deviations to zero. [29] Rainfall during the agricultural production is expected to influence consumption through income fluctuations but is unlikely to be correlated with the potential measurement error in the per capita consumption variable. The baseline village fixed effects (θ_v) in Equation (3.6) wipe out the level effects of rainfall in the first stage regression. Therefore, exploiting the fact that rainfall shocks will affect different types of households in different ways, we interact the rainfall variable with head's gender, age and education yielding a total of three instruments.

3.5 Results

We begin by testing the full risk-sharing model described above. Column 1 in Table 3.7 provides the results for the base specification of Equation (3.5) with network fixed effects (NFE). The control variables capture the characteristics of the previous household members, including their position within the 2010 household. The signs of the control variables are *a priori* correct. For example, education has a positive impact on consumption growth, while households with widowed or divorced previous household members experience lower consumption growth than others within the same extended family network.

The statistical significance of the shock coefficient, despite the inclusion of NFE, reveals that shocks are not insured within extended families. Households that experienced a shock had 14 percentage points lower consumption growth, on average and *ceteris paribus*, than households from the same extended family who did not experience a shock. The emergence of this wedge in the face of a shock implies a clear rejection of the full risk-sharing model in the extended family networks in this study.

[29] Beegle, De Weerdt and Dercon (2008) employ a similar instrumental variable approach for their lagged consumption variable in assessing the long-term impact of adult deaths on consumption growth in Kagera.

Table 3.7: The effect of shocks on consumption growth

Dependent variable: (logged) per capita consumption growth	1 OLS, NFE	2 OLS	3 2SLS
Own shock	-0.141***	-0.144***	-0.140***
	(0.026)	(0.024)	(0.024)
2010 household characteristics:			
Age of oldest PHHM in the household	-0.001	0.000	-0.000
	(0.001)	(0.001)	(0.001)
A PHHM is head of the household	0.164***	0.152***	0.151***
	(0.047)	(0.034)	(0.034)
A PHHM is spouse of the household head	0.110**	0.077*	0.071*
	(0.048)	(0.039)	(0.042)
A PHHM is child of the household head	-0.196***	-0.184***	-0.182***
	(0.051)	(0.037)	(0.036)
A divorced PHHM in the household	-0.342***	-0.306***	-0.300***
	(0.068)	(0.076)	(0.073)
A widowed PHHM in the household	-0.333***	-0.332***	-0.328***
	(0.056)	(0.052)	(0.051)
A married PHHM in the household	-0.483***	-0.454***	-0.449***
	(0.040)	(0.038)	(0.039)
Max years of education of PHHM in the hh	0.058***	0.063***	0.060***
	(0.006)	(0.005)	(0.006)
Number of PHHMs in the household	-0.008	-0.018	-0.013
	(0.023)	(0.020)	(0.019)
Network characteristics:			
Number of split-off households stayed		-0.056***	-0.060***
		(0.009)	(0.010)
Number of split-off households moved		-0.008	-0.011
		(0.010)	(0.011)
Household characteristics at the baseline:			
Natural log value of assets in 1991		0.006	-0.009
		(0.017)	(0.027)
Education of 1991 household head		0.003	0.001
		(0.006)	(0.007)
Head was male in 1991		-0.064*	-0.082
		(0.038)	(0.053)
Age of household head in 1991		0.010**	0.009*
		(0.005)	(0.005)
Age of head squared		-0.000**	-0.000**
		(0.000)	(0.000)

Table 3.7: The effect of shocks on consumption growth

Dependent variable: (logged) per capita consumption growth	1 **OLS, NFE**	2 **OLS**	3 **2SLS**
Num of males 0-5 years in the hh		0.005	0.016
		(0.019)	(0.025)
Num of males 6-15 years in the hh		0.049***	0.056***
		(0.012)	(0.016)
Num of males 16-60 years in the hh		0.006	0.002
		(0.016)	(0.017)
Num of males 61+ years in the hh		0.164***	0.199***
		(0.049)	(0.070)
Num of females 0-5 years in the hh		0.012	0.018
		(0.017)	(0.020)
Num of females 6-15 years in the hh		0.022*	0.027*
		(0.012)	(0.015)
Num of females 16-60 years in the hh		0.008	0.011
		(0.013)	(0.013)
Number of females 61+ years in the hh		0.026	0.033
		(0.033)	(0.036)
Household had a non-earth floor in 1991		-0.008	-0.054
		(0.047)	(0.095)
(logged) hh per capita consumption in 1991		-0.911***	-0.732**
		(0.042)	(0.290)
Number of observations	2,349	2,349	2,349
R^2	0.202	0.421	0.412
Adjusted R^2	0.199	0.414	0.392

note: *** p<0.01, ** p<0.05, * p<0.1. Cluster-robust standard errors by baseline village are in parenthesis. Regressions in column 1 includes NFE, regressions in columns 2 and 3 include baseline village fixed effects. PHHM refers to previous household member (i.e. person interviewed at the baseline).

In column 2 we drop the NFE and replace with network characteristics, such as the number of migrant and non-migrant network members (which together control for network size and composition) and the wealth and demographics of the baseline household from which the network is formed. We also include baseline village fixed effects. The size of the shock coefficient is nearly identical to the one obtained with NFE, giving confidence in the network level controls we use in a later analysis of reciprocal risk sharing.

Finally, column 3 provides the Two-Stage Least Squares results that address the potential endogeneity problem arising from the inclusion of the initial logged per capita consumption variable. The first stage regression results and the standard IV-diagnostic tests are presented in Appendix B. The included instruments show how households headed by older and more educated males enjoy higher baseline consumption. The excluded instruments are zero-truncated negative z-score deviations of rainfall interacted with the household head's age, education and gender. They show that the positive level effects of each of these three household head characteristics are attenuated with the inclusion of negative rainfall shocks. The Cragg and Donald (1993) test yields 19.1 indicating that our instruments are relevant. Comparison with the critical values provided in Stock and Yogo (2005) suggests that the bias of our IV-estimate is less than five per cent of the OLS estimate. The Hansen (1982) J-test provides a p-value of 0.570. Thus, the null hypothesis of zero correlation between the instrument and the error term is upheld at conventional levels. The shock coefficient and the standard error from the 2SLS estimates are almost identical to those from OLS, indicating that the potential endogeneity of the logged per capita baseline consumption has a negligible influence on the shock variable. [30] In the light of this, we use the more efficient OLS method to make inferences in the remainder of the text.

Next we test whether any risk sharing takes place in these networks. As discussed earlier, we replace the NFE with network characteristics and baseline village fixed effects and augment the specification with the network shock variable. The first column in Table 3.8 reports the results for the migrant households and the second column for the non-migrant households. For migrants, the network shock coefficient is negative and highly

[30] A C-test (see e.g., Hayashi 2000p. 220-221) with one degree of freedom yields $\chi 2$-test statistic of 0.104 (p=0.748). Thus, we cannot reject the null hypothesis that the lagged consumption variable is exogenous.

significant. These network shocks have a sizeable impact on migrant household consumption: on average, a shock in one household in the network resulted in a drop of five percentage points in consumption growth. As shocks are not correlated within the extended family networks (the intra-class correlation coefficient equals 0.017 with a standard error of 0.016), this finding reveals that migrants insure other households in their extended families. Non-migrant households, on the other hand, do not appear to be affected by the network shocks. The point estimate is nearly zero and insignificant. These results suggest that the risk-sharing arrangement is not reciprocal.

In order to investigate this further, we decompose the network shock variable into shocks in non-migrant and migrant households. The first variable measures the number of non-migrant households that experienced a shock in the extended family. The second network shock variable measures the number of migrant households affected by shocks. As before, the household's own shocks have been excluded from these variables.

Table 3.9 presents the regression results. Migrants are susceptible to shocks affecting other migrant and non-migrant households within their extended family network, while non-migrants are sensitive to neither. On average, a shock in one non-migrant household in the network leads to a drop of 5.5 percentage points in migrant household's consumption growth. Shocks in other migrant households have a negative effect of similar magnitude on a migrant's consumption than shocks experienced in stayer households but this coefficient is not statistically significant at a conventional level (p=0.127).

Given the patrilocal migration patterns where women migrate largely because of marriage (see Table 3.3), it is reasonable to ask to what extent they remain linked to their original family. In Tables A.1 and A.2 of Appendix A we explore the gender dimension in risk sharing by interacting the shock variables with the male dummy. The results show that the findings in Table 3.8 and 3.9 are not driven by male migrants suggesting that also female migrants engage in insuring their family members back home.

Table 3.8: The effect of network shocks on consumption growth

Dependent variable: (logged) per capita consumption growth	**Migrant households** **1** **OLS**	**Non-migrant households** **2** **OLS**
Number of households that experienced a shock in the network	-0.050***	0.008
	(0.018)	(0.015)
Own shock	-0.094**	-0.060*
	(0.043)	(0.032)
Number of split-off hhs stayed	-0.031*	-0.032**
	(0.016)	(0.013)
Number of split-off hhs moved	-0.008	-0.026**
	(0.013)	(0.013)
Age of oldest PHHM in the 2010 hh	0.002	-0.000
	(0.002)	(0.001)
A PHHM is head of this 2010 hh	0.160***	0.220***
	(0.056)	(0.054)
A PHHM is spouse of this 2010 hh's head	-0.042	0.185***
	(0.059)	(0.057)
A PHHM is child of this 2010 hh's head	-0.356***	-0.008
	(0.099)	(0.044)
A Divorced PHHM in 2010 hh	-0.363***	-0.165*
	(0.111)	(0.088)
A widowed PHHM in 2010 hh	-0.412***	-0.124**
	(0.125)	(0.052)
A married PHHM in 2010 hh	-0.431***	-0.246***
	(0.052)	(0.051)
Max years of education of PHHM in this 2010 hh	0.072***	0.030***
	(0.007)	(0.007)
Number of PHHMs in this 2010 hh	0.050	-0.063***
	(0.047)	(0.023)
(logged) hh per capita consumption in 1991	-0.978***	-0.848***
	(0.049)	(0.053)
Number of observations	1,270	1,079
R^2	0.462	0.390
Adjusted R^2	0.450	0.374

note: *** p<0.01, ** p<0.05, * p<0.1. Cluster-robust standard errors by baseline village are in parenthesis. Regressions include baseline village fixed effects and variables controlling for household characteristics at the baseline. PHHM refers to previous household member (i.e. person interviewed at the baseline).

Table 3.9: Network shocks in migrant and non-migrant households

Dependent variable: (logged) per capita consumption growth	**Migrant households**	**Non-migrant households**
	1	**2**
	OLS	**OLS**
Number of non-migrant hhs that experienced a shock in the network	-0.055**	0.013
	(0.028)	(0.022)
Number of migrant hhs that experienced a shock in the network	-0.043	0.002
	(0.028)	(0.023)
Own shock	-0.093**	-0.059*
	(0.043)	(0.032)
Number of split-off hhs stayed	-0.030*	-0.034**
	(0.018)	(0.015)
Number of split-off hhs moved	-0.009	-0.024*
	(0.013)	(0.014)
Number of observations	1,270	1,079
R^2	0.463	0.390
Adjusted R^2	0.450	0.374

note: *** p<0.01, ** p<0.05, * p<0.1. Cluster-robust standard errors by baseline village are in parenthesis. Regressions include baseline village fixed effects, 2010 household level variables capturing characteristics of the previous household members and variables controlling for household characteristics at the baseline.

We conclude that migrant households are partially and unilaterally insuring households that stay behind. This lack of reciprocity violates the predictions of the reciprocity-based models (without a social norms term). Because, on average, migrants are nearly twice as rich as those who remained at home, these findings are consistent with reciprocity-based models augmented with a social norms term, which attenuates the participation, truth-telling or incentive compatibility constraints.

3.6 Other transactional insurance motives

An alternative explanation to the observed lack of reciprocity could be that migrants insure non-migrants in exchange for other benefits. By concentrating on consumption differences we have considered only current pay-offs from any risk-sharing arrangement.

It is quite possible that the benefits are still to accrue to the migrant in the more distant future. Lucas and Stark (1985) mention that there could be exchange motives for insurance provision relating to the desire for non-migrants to look after local assets, the intention to return home and the aspiration to inherit. In a context that lacks technology to allow future income to be consumed now, we could confuse unilateral insurance with postponed reciprocity. Fortunately, the KHDS questionnaire is particularly rich and we are thus able to explore some of these issues.

The questionnaire asks each migrant about asset holdings in the baseline village. As our outcome variable is consumption growth we cannot use these asset holdings as explanatory variables: current wealth is surely endogenous to growth in wealth. We attempt to circumvent this problem by looking at the share of assets in the current portfolio that are located in the village. While it remains possible that portfolio composition is endogenous to consumption growth, we believe the results are informative enough to report.

About 28 per cent of migrants have assets in the baseline village and 25 per cent of migrants own land in the baseline village. For land we have exact area measurements, but not monetary values. If migrants engage in risk sharing with those who remain at home for the purpose of maintaining land and ensuring their continued entitlement to the land (which is important in a country with few formal land deeds), then we would expect more responsiveness to network shocks from people with a larger share of their land holdings in the baseline village. The first column of Table 3.9 explores this. As before, the dependent variable is logged per capita consumption growth. We interact the non-migrant network shocks with a variable measuring the share of the land in the baseline village. The coefficient on this interacted variable turns out insignificant implying that the share of land in the baseline village neither increases nor decreases the insurance provision. Land ownership and the eagerness to hold onto it may differ across the gender lines. To test whether the lack of statistically significant interaction effect is due to pooling male and female migrants together, we disaggregated the sample by gender and re-ran the first column of Table 3.9. The results in Table A.3 of Appendix A show that the finding holds for both sub-samples: the insurance arrangement is not driven by land ownership in the baseline village.

In the second column in Table 3.10 we interact the non-migrant network shock variable with the length of the migration spell. Following Dustmann and Mestres (2010), we argue this to be a measure of the permanence and success of the move and an inverse measure of the return likelihood. We find that the duration of the migration spell does not have any impact on migrant's insurance provision. This also holds when we use non-linear versions of the migration duration in the form of a piecewise linear spline. Due to the patrilocal migration patterns, it is plausible that the return likelihood of women is lower and in turn might correlate differently with the extent of their insurance provision. As before, we explored this by re-running the test with gender disaggregated samples but the original outcome holds for both sub-samples (see Table A.3 of Appendix A).

The third column in Table 3.10 investigates whether the expectation to inherit is a plausible motive for unilateral insurance. Since traditional law excludes women from inheriting land, a household can only inherit land if it has male members whose parents own clan land in the baseline village. [31] We construct a variable that captures these households and also control for cases where the parents no longer live in the baseline village in order to isolate the inheritance story from the effect of having parents in the

[31] Land in this context normally belongs to the clan and the purpose of the traditional law is to keep in the clan (see De Weerdt 2010).

baseline village. Nearly 42 percent of the migrant households have parental clan land holdings waiting for them in the baseline village. By interacting the non-migrant network shocks with a parental clan land holdings dummy, we find that that these households are no more (or less) engaged in insurance provision than households that do not expect to inherit land.[32]

A final transactional motive that could be consistent with the regression results is that non-migrants pay insurance premiums to migrants in return for their continued insurance provision. This does not seem consistent with the findings of Table 3.2, where we noted that migrants are net senders of transfers.

[32] Since this test already exploits the gender dimension in the inheritance law we are not able, nor feel necessary, to conduct this test for the two gender groups.

Table 3.10: Other transactional insurance motives

	Migrant households		
Dependent variable: (logged) per capita consumption growth	**1**	**2**	**3**
Number of non-migrant hhs that experienced a shock in the network	-0.042	-0.050	-0.048
	(0.028)	(0.063)	(0.045)
--- Interacted with:			
* Share of land in BLV in total land portfolio	-0.055		
	(0.054)		
* Number of years since the last PHHM migrated into this hh		-0.001	
		(0.005)	
* Hh has inheritable land in the baseline village			0.019
			(0.055)
* Hh member's parent lives in BLV			-0.032
			(0.063)
Own shock	-0.085**	-0.097**	-0.090**
	(0.043)	(0.045)	(0.044)
Share of land in BLV in total land portfolio	0.300***		
	(0.064)		
Household does not own land	0.262***		0.171***
	(0.047)		(0.043)
Number of years since the last PHHM migrated into this hh		-0.001	
		(0.004)	
Hh has inheritable land in the baseline village			-0.016
			(0.066)
Hh member's parent lives in BLV			0.054
			(0.063)
Number of split-off hhs stayed	-0.021	-0.028	-0.026
	(0.017)	(0.018)	(0.018)
Number of split-off hhs moved	-0.015	-0.014	-0.015
	(0.013)	(0.013)	(0.013)
Number of observations	1,270	1,270	1,270
R^2	0.487	0.462	0.470
Adjusted R^2	0.474	0.449	0.456

note: *** p<0.01, ** p<0.05, * p<0.1. Cluster-robust standard errors by baseline village are in parenthesis. Regressions include baseline village fixed effects, 2010 household level variables capturing characteristics of the previous household members and variables controlling for household characteristics at the baseline.

3.7 Is there a kin tax?

Does the migrant incur a significant cost for providing this unilateral insurance? From Table 3.9 we observe that for each shock to the extended family network at home there is a drop of 5.5 percentage points in the migrant's consumption, which appears to be a permanent deviation from the growth curve. The average migrant has 0.53 network shocks of non-migrants, resulting in an implied overall consumption growth penalty of 2.9 percentage points, on average, over the 19-year period. Over this same period, the average consumption growth among migrants was 108 per cent, implying that insurance constituted an average annual growth penalty of around 0.077 of one percentage point (reducing average annual growth roughly from 3.93 per cent to 3.86 per cent). [33] Put another way, migrants share about 2.7 per cent of their very substantial growth by insuring family members at their original location. [34]

This is a lower-bound estimate because we cannot exclude the possibility that we are only measuring a subset of relevant shocks: if shocks are self-reported then respondents may fail to mention those that were effectively insured. Fortunately, the survey provides an alternative shock measure, which is not self-reported. We have historical rainfall data from the Tanzanian Meteorological Agency for gauges in 212 weather stations in Kagera and the migration destinations in our sample. In a first step, each household is linked to all rainfall stations within a 100 km radius. Next, a monthly rainfall figure is calculated for each household by weighing each monthly rainfall reading with the inverse of the distance of the rainfall station where it was recorded to the household in question.

[33] We use geometric (rather than arithmetic) means to calculate the average annual growth rates.
[34] The 95%-confidence interval for the annual growth penalty is [0.0002, 0.1515] and for the 'kin-tax' [0.01, 5.40].

The mean distance to the nearest rainfall station is 10 km and the median is 23 km. For each household we can calculate average monthly z-score deviations of rainfall during the two rainy seasons, in relation to the 30 year average (1980-2010) for that village. Rainfall shocks are then constructed by truncating the positive yearly average rainfall deviations to zero. We calculate a non-migrant household's own shock as the most negative shock in the 1994-2009 period. The first row in Table 3.11 shows that rainfall shocks are important in determining consumption growth, with every standard deviation decrease in (negative) rainfall deviation causing consumption growth to decline by 12 percentage points for migrants and 18 percentage points for non-migrants.[35]

Table 3.11: Re-calculating the kin-tax through rainfall data

	Migrant households		**Non-migrant households**	
	mean	**1**	**mean**	**2**
max rain shock in own location [a)]	-0.70 [0.52]	0.115** (0.045)	-1.17 [0.30]	0.179** (0.076)
max rain shock in deviation in baseline village [b)]	-0.63 [0.56]	0.079* (0.046)		
max rain shock in deviation in migrant locations [b)]			-1.21 [0.29]	-0.014 (0.086)
Number of observations	1,270		1,079	
R^2	n/a	0.471	n/a	0.400
Adjusted R^2	n/a	0.458	n/a	0.382
baseline district FE?	n/a	yes	n/a	yes

note: *** p<0.01, ** p<0.05, * p<0.1.
[a)] For migrants this is after they moved to their 2010 location, for non-migrants this refers to 1994-2009.
[b)] After the migrant moved to their 2010 location.
Standard deviations in brackets. Cluster-robust standard errors by baseline village are in parenthesis. Regressions include 2010 household level variables capturing characteristics of the previous household members and variables controlling for household characteristics at the baseline.

[35] Out of the 1,270 migrant households 77 per cent report to derive (at least some) income from agricultural production and 73 per cent own land. The results reported in Columns 1 and 2 of Table 3.11 are robust to restricting the migrant sample to these households.

Knowing that rainfall shocks drive the incomes of both stayer and migrant households, we can use them as an alternative network shock indicator. We replace the network shock variable with the baseline village rainfall shock variable in Equation (3.6). For migrant households, this rainfall shock is constructed as the most negative rainfall deviation in the baseline village after the migrant left. For stayer households, we take the most negative rainfall deviation among the migrant household locations, after the migrant arrived to the current location. Due to the covariate nature of the rainfall shocks we cannot use the baseline village fixed effects in the stayer household specification. Similarly, there is little variation in the baseline village rainfall shock variable among migrants that originate from the same village. We therefore replace the baseline village fixed effects with baseline district fixed effects. [36] Column 1 reports the results for the migrant households. We see that after the migrants leave their consumption remains responsive to rainfall shocks at the baseline village. Once we use rainfall shocks the kin tax goes up to 5.0 percentage points. [37] We consider the parsimonious rainfall shock specification from Table 3.11 as the upper bound effect, binding the estimate between 2.7 and 4.6 per cent.

Column 2 reports the corresponding results for the non-migrant households. Consistent with the results presented in Section 3.5, we see that non-migrants are not affected by rainfall shocks that take place in migrant households.

It is interesting, if slightly misleading (see below), to point out that our implicit tax rate of three to five per cent is of the same order of magnitude as that found in two other studies. Jakiela and Ozier (2012) estimate that women in a laboratory setting in Kenya

[36] Households group into 51 baseline village and 6 baseline districts. For migrants, the baseline village fixed effects exploit the variation arising from the fact that migrants leave at different times. This yields almost an identical coefficient but due to the limited variation remaining in the data, the coefficient is not significant at a conventional level (p=0.155).
[37] The 95 per cent confidence interval for the kin tax ranges between 0.7 and 10.7 percentage points.

acted as if they were expecting to be pressured to share four per cent of their experiment winnings with relatives that were not present at the experiment. Ambler (2012) reports that El Salvadorian migrants living around Washington DC remit five per cent more of a windfall income if they know the potential recipients at home will be informed about it. However, we consider these striking similarities in magnitude to be slightly misleading. First, the above studies focus on the effect, on sharing, of providing full information on a windfall income, while our figures reflect the effect of shocks on sharing within real-world belief sets. Second, these experimental studies capture short-run dynamics while our study focuses on consumption growth over a more protracted 19-year period.

3.8 Endogeneity of migration

Much of the research on migration is concerned with establishing its causal effects. The primary goal of this essay, however, is quite different. Here, we aim to document what happens to a traditional institution, like informal insurance, in a society that modernizes and is characterized by massive internal migration. The strength of the essay lies in its ability to describe this process within linked extended family networks, over a protracted two decade time period and exploiting all the richness of the real world. Even then, causality remains important in that it relates to the external validity of our results. If migration is *not* causally responsible for the empirical patterns described in the previous sections then that should increase one's degree of skepticism about these same patterns applying in other settings where the selection process into migration is different.

Our data are not experimental and their real-world richness comes at the cost of not being able to provide iron clad proof of causality. Fortunately, we are able to exploit more survey rounds in order to speak to the causality issue and exclude certain forms of unobserved heterogeneity as explanations for the results. The purpose of this section is to be very specific about which remaining types of endogeneity could compete with causality to explain the results.

Thus far the essay has used waves one and six of the survey. This section will exploit the four waves that lie between. In what follows it is useful to bear in mind that consumption was measured comparably within waves 1, 5 and 6 and within waves 2, 3 and 4 – but not between these two groups of waves.[38]

As our left hand side variable is consumption growth, a first set of concerns relates to whether migrants start out from similar consumption positions at baseline (compared to non-migrants) and whether, prior to migration, they are following a similar growth path. We can test this by regressing baseline consumption on future migration status ($M_{i,2010}$) and use baseline village fixed effects (θ_v) and control for the initial household characteristics (w_j) used in Equation (3.6) (and listed in Table 3.6):

$$Y_i = \propto M_{i,2010} + w_j'\vartheta + \theta_v + \varepsilon_i. \tag{3.7}$$

In the first column ofTable 3.12, Y_i is the natural logarithm of the consumption level at baseline. The insignificance of the future migration coefficient shows that our extensive controls manage to capture the heterogeneity with respect to baseline start-off levels: migrants started off at the same position as non-migrants, *ceteris paribus*.

We then replace the right-hand side of Equation (3.7) with (logged) consumption at the baseline. The second column of Table 3.12 exploits information on consumption growth between rounds 2, 3 and 4 (1992-1994) to ask whether growth across those three rounds

[38] The primary concern is the difference in recall period. Beegle et al. (2012) look at how differential recall periods affect consumption aggregates, using data from a survey experiment conducted in Tanzania.

can be explained by future migration status. We find that it cannot: migrants and non-migrants were on the same growth curve prior to migration, given the controls we use.

Table 3.12: The effect of future migration on baseline consumption levels and short-run growth (1992-94)

Dependent variable:	**logged per capita consumption per capita in wave 1**	**consumption growth 1991-4 (waves 2-4): Δ ln(conspc)**
	1	**2**
migrant in 2010	-0.017	0.017
	(0.038)	(0.045)
(logged) household per capita consumption in 1991		-0.064
		(0.041)
Number of observations	803	782
R^2	0.198	0.032
Adjusted R^2	0.183	0.012

note: *** p<0.01, ** p<0.05, * p<0.1. Cluster-robust standard errors by baseline village are in parenthesis. Regressions include baseline village fixed effects (θ_v) and control for the initial household characteristics (w_j) used in Equation (3.6) (and listed in Table 3.6)

Finally, we preclude the possibility that there are any time invariant traits of migrants and non-migrants, such as risk preferences, that are jointly determining their migration decision and the respective roles they take on in the insurance arrangement. This test exploits the fact that some migrants migrated after we observe them in 2004. We restrict the sample to 1,146 households that had not migrated by 2004 and re-estimate Equation (3.6) using growth in (logged) consumption from 1991 to 2004 as the dependent variable.[39] We also include a dummy that captures those 151 households that will migrate between 2004 and 2010.

[39] The term x_{ij} now refers to household characteristics in 2004 with respect to its previous household members.

In column 1 of Table 3.13, we see that the future migration status is not significant indicating that, *ceteris paribus*, migrants and non-migrants were on the same long-run growth path. Households are negatively affected by own shocks and network shocks, although the latter is marginally insignificant at p=0.112. In the reduced sample baseline village fixed effects take a high toll on the degrees of freedom, with only an average of three migrant households per village, compared to 25 migrant households per village in the main regressions. In column 3, we replace the baseline village fixed effects with baseline district fixed effects. The coefficient on the migration dummy remains insignificant and both own and network shocks yield a significant and negative effect on consumption growth. Taken together this reveals that the sample of 1,146 non-migrant households was sharing risk in the period prior to their migration.

In columns 2 and 4 we interact the network shock variable with 2010 migration status. This interaction is not significant, irrespective of using baseline village or district fixed effects. This shows that migrant and non-migrant households are both responsive to each other's shocks prior to the move and the insurance relationship becomes unilateral only after the move. Similar to a difference-in-difference estimator, this excludes the possibility that time invariant characteristics of either party are driving the results.

Table 3.13: The effect of future migration on long-run growth and pre-migration insurance contract type

Dependent variable: (logged) per capita consumption growth 1991-2004	**1**	**2**	**3**	**4**
migrant in 2010	0.058	0.065	0.047	0.059
	(0.051)	(0.066)	(0.053)	(0.066)
own shock	-0.064*	-0.064*	-0.069**	-0.069**
	(0.037)	(0.037)	(0.033)	(0.033)
number of (other) stayer households affected by shock	-0.047	-0.046	-0.047*	-0.045*
	(0.029)	(0.029)	(0.027)	(0.026)
--- ***Interacted with:***				
* (migrant in 2010)		-0.009		-0.017
		(0.049)		(0.048)
Number of observations	1,146	1,146	1,146	1,146
R^2	0.405	0.405	0.412	0.424
Adjusted R^2	0.390	0.389	0.397	0.405
baseline village FE?	yes	yes	no	no
baseline district FE?	no	no	yes	yes

note: *** p<0.01, ** p<0.05, * p<0.1. Cluster-robust standard errors by baseline village are in parenthesis. Regressions include 2004 household level variables capturing characteristics of the previous household members and variables controlling for household characteristics at the baseline.

Taken together, the results in Table 3.12 and Table 3.13 tell us that there is no unobserved heterogeneity between migrants and non-migrants with respect to the starting position and slopes of their short-run (1992-1994) and long-run (1991-2004) growth paths prior to migration. Results from Table 3.13 further suggest that the insurance contract was characterized by reciprocity before the move and only became unilateral after the move. We can therefore be confident that the effect of migration on informal insurance is either causal, or it is driven by the occurrence of a time variant event (like a shock pulling or pushing someone into migration), or change in individual characteristic (like coming of age, achieving higher levels of education or winning a lottery), which causes one to both migrate and assume the special role in the insurance network.

3.9 Robustness

We conducted an array of robustness checks to validate our findings.[40] First, we find that the results are robust to an alternative migrant definition where also households that moved to a nearby village are defined as non-migrants (see footnote 18).

Second, the results are not driven by the configuration of the data. The shock data were initially defined at individual level while our outcome variable is measured at household level. Conducting the empirical analysis at individual level does not affect our main findings (see footnote 24).

Third, defining household consumption per adult equivalent instead of per household member yields close to identical results in all specifications (see footnote 22).

Fourth, changing the way we isolate the shocks that occurred before the households split does not change our results either (see footnote 25).

Fifth, we also checked whether the potential endogeneity of some of our control variables is driving our results. As discussed in Section 3.5, instrumenting the lagged consumption variable does not affect the shock coefficient. In addition, when the 2010 household level control variables are omitted one-by-one, the estimated shock coefficients remain stable across all specifications (see footnote 27).

3.10 Conclusions

In the next few decades internal migration is likely to drive the development process in Africa. This demographic process is evident in the data used in this essay. Starting from the household rosters of a representative household survey conducted nearly two decades ago in Kagera, we find that over half of the original household members had moved

[40] The results of these robustness checks are available upon request.

internally, while very few moved internationally. We document how this powerful current of migration, an integral part of development, interacts with a traditional institution like informal risk sharing to shape economic mobility and vulnerability.

We find that internal migrants provide unilateral insurance to those who remain at home, which seems to be driven by social norms rather than exchange motives. The total, final, long-run effect of this insurance provision on the migrant’s growth amounts to a 2.7 to 4.6 per cent sacrifice in consumption (2.9 to 5.0 percentage points off the 108 % total growth realized by the migrant).

While our study cannot conclusively say where migrants would be without their extended family networks back home, a 'tax rate' of 2.7 to 4.6 per cent seems too low to be an important brake on a migrant’s growth. We do know that migrants who have severed links with home perform worse than other migrants and one should not overlook the fact that, while starting out from similar baseline welfare levels, migrants realize a total consumption growth which is three times higher than that of non-migrants.

4 Essay 2: Temperature Shocks, Household Consumption and Internal Migration

4.1 Introduction

As noted in Chapter 3, internal migration typically takes a central stage when underdeveloped countries modernize and move away from subsistence agriculture (Lewis 1954, Ranis and Fei 1961, Chenery and Syrquin 1975, Collier and Dercon 2009). At the micro-level, an emerging empirical literature documents how internal migration is associated with large income gains and poverty reduction in Sub-Saharan Africa (Beegle, De Weerdt, and Dercon 2011, Christiaensen, De Weerdt, and Todo 2013, Young 2013).

In this essay, I investigate whether financial constraints in the rural areas inhibit migration. The essay views long-term migration as an investment. [41] Migration is associated with various costs (e.g., travel expenses, search costs in the destination area) that are borne up front. In the presence of liquidity and credit constraints, potential migrants can only finance migration by accumulating enough savings in the previous period. Approximately 97 per cent of African crop land is rainfed (Faurès and Santini 2008) rendering the livelihoods of rural households extremely vulnerable to weather variation. This strong correlation between weather outcomes and rural incomes imply that constrained households may not be able to finance migration costs in bad weather years. In other words, if liquidity constraints bind, migration rates decrease with negative weather shocks and increase with positive ones. Such correlation between weather and migration should not exist in the absence of binding liquidity constraints.

[41] This essay only considers long-term internal migration spells. For studies that analyse the impact of exogenous income shocks on temporary or seasonal internal migration, see (Jayachandran (2006), Bryan, Chowdhury, and Mobarak (2013)). For research focusing on international migration, see (Halliday (2006), Feng, Krueger, and Oppenheimer (2010), Angelucci (2013), Bazzi (2013)).

I again use the KHDS data to test this hypothesis. As noted in Chapter 2, the survey provides baseline data based on a comprehensive household questionnaire administered in 1991-1994. The latest round of the survey conducted in 2010 attempted to track all individuals, irrespective of whether they still resided in the same village or not. This follow-up round also collected detailed information on migration histories since 1994. These data are then linked to historical temperature and rainfall data in an attempt to shed light on how local weather shocks shape long-term migration decisions in a context of mass migration (see section 3.1). In Chapter 3, we found that migrants maintain insurance links to those who remain (for other studies documenting similar evidence, see Rosenzweig and Stark 1989, Yang and Choi 2007). However, the results reject reciprocal risk sharing between migrants and their non-migrant family members implying that, in this context, migration is not the result of a household level maximization strategy. Section 4.6 of this essay returns to this issue providing evidence that supports this finding.

In section 4.3, I show how household welfare co-moves with local weather events. Using consumption as the welfare metric, the results show that households are highly vulnerable to temperature changes during the growing season. Controlling for rainfall, household fixed effects and various time-varying factors, I find that a one standard deviation increase in the mean monthly growing season temperature decreases household per capita consumption by 4.9 per cent, on average and *ceteris paribus.* This part of the analysis speaks to the emerging literature that documents how temperature shocks cause large fluctuations in economic outcomes in Sub-Saharan Africa (Dell, Jones, and Olken 2012, Schlenker and Lobell 2010a, Burke et al. 2011). The use of national level output data, GDP (Dell, Jones, and Olken 2012) or country-level yields (Burke et al. 2011, Schlenker and Lobell 2010a), raises the question of whether these estimated impacts hold at the micro level.[42] This essay presents, to my knowledge, the first attempt to quantify the impact of temperature changes exploiting household level data from Sub-Saharan Africa. [43] The results presented here broadly agree with this macro literature.

Having established that household incomes co-move with weather outcomes, I then run reduced form regressions to model migration decisions. More specifically, I use discrete time event history models (Allison 1982, Jenkins 1995) to test whether these local weather shocks affect long-term migration decisions in Tanzania. Consistent with the hypothesis above, positive temperature deviations reduce long-term migration among men. A one standard deviation increase in the previous year's average monthly growing season temperature reduces the overall male migration rate by about 13 per cent. This finding provides evidence of binding liquidity constraints. Female migration is not affected by these shocks possibly due the fact that it is largely motivated by marriage and family in this context as already noted in the previous chapter.

A small but emerging body of literature provides mixed evidence on the impact of weather variation on internal migration.[44] Gray and Mueller (2012b) find that flooding had little effect on migration patterns in Bangladesh while crop shocks (mainly due to the lack of rainfall) tended to increase migration. In Ethiopia, droughts were found to increase men's labour related migration but decrease female marriage related migration (Gray and Mueller 2012a). There are two recent studies that examine the link between temperature and internal migration. Dillon, Mueller, and Salau (2011) find that men are more likely to migrate in Northern Nigeria if the temperatures were below or above the long-term

[42] Such aggregation problems are typical in Economics (see Blundell and Stoker 2005).
[43] Skoufias and Vinha (2013) study the impact of temperature and rainfall shocks on household consumption in rural Mexico. Mueller, Gray, and Kosec (2014) measure the impact of temperature and rainfall changes on household income in Pakistan.
[44] For recent literature reviews, see Lilleør and Van den Broeck (2011) and Marchiori, Maystadt, and Schumacher (2013).

mean. Mueller, Gray, and Kosec (2014) study the impact of temperature and rainfall using data from Pakistan. The authors find that positive temperature deviations increase migration rates both among men and women. Furthermore, recent evidence from cross-country studies suggests that weather anomalies have increased urbanization in Sub-Saharan Africa (Marchiori, Maystadt, and Schumacher 2012, Barrios, Bertinelli, and Strobl 2006). I contribute to this literature by providing a theoretical framework in which to rationalise the results that depart from the predictions of the conventional models of migration (namely, Harris and Todaro 1970). Finally, I explicitly show how weather shocks cause large fluctuations in household welfare. Surprisingly, this step of establishing a channel through which weather variation affects migration decisions is largely absent in the existing literature on this topic (see, Lilleør and Van den Broeck 2011). [45]

The structure of this essay can now be outlined. The next section provides a theoretical framework that links migration decisions to temporary income shocks. Section 4.2 describes the context and the data used in the analysis. Section 4.3 shows how household consumption co-moves with temperature in Kagera. Section 4.4 first describes the econometric model used to analyse migration patterns in Kagera and then presents the results. Section 4.5 discusses the robustness of the empirical results. Section 4.6 looks into the link between long-term weather variability and migration. Section 4.7 concludes.

[45] Recent study by Mueller, Gray, and Kosec (2014) is an exception. Similar to this essay, the authors study the impact of weather variation on household welfare and migration rates using long panel data from Pakistan.

4.2 Data

4.2.1 Context

As discussed in Chapter 3, the region has two rainy seasons in which agricultural production takes place. The long rainy season (*Masika*) occurs usually between March and May and the short rainy season (*Vuli*) between October and December. Figure 4.1 shows the average monthly rainfall patterns over a 30-year period. Figure 4.2 reveals that the average monthly temperatures in the region range between 21 and 23.5 °C.

Figure 4.1: Average monthly rainfall patterns in 1980-2009 in Kagera

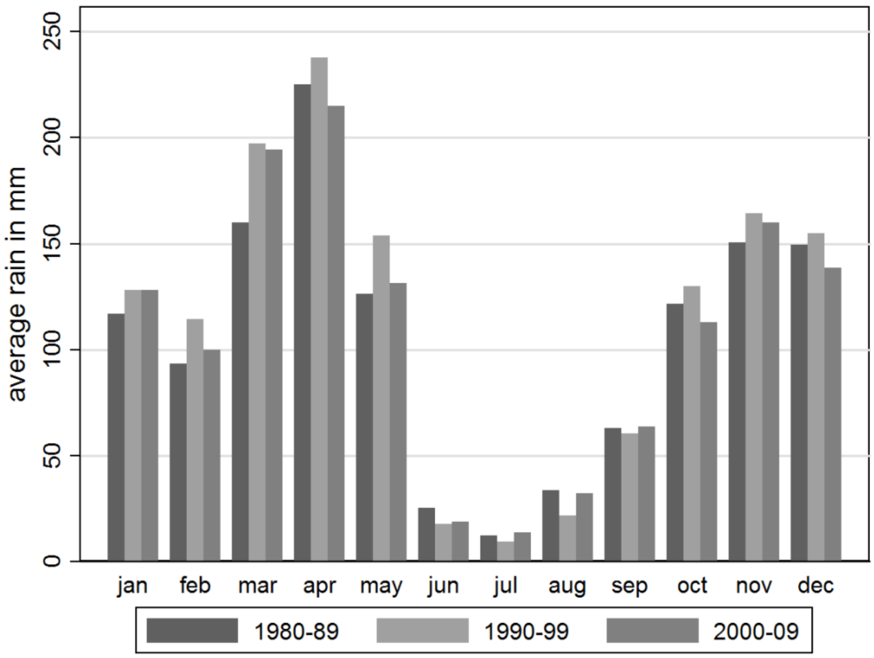

Source: own calculations from rainfall data from the Tanzanian Meteorological agency.
Note: The long rainy season months are March, April and May and the short rainy season October, November and December.

Figure 4.2: Average monthly temperatures in 1981-2009 in Kagera

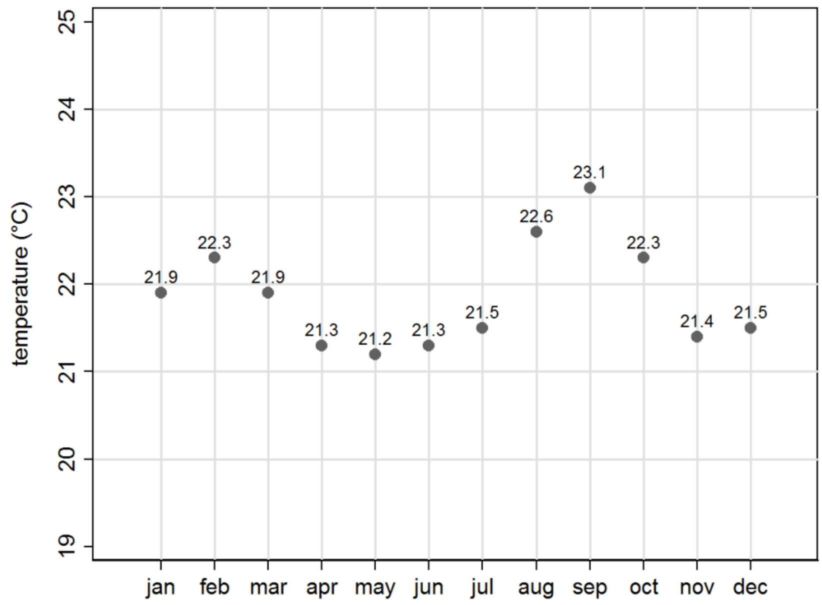

Source: own calculations from NASA's Modern-Era Retrospective analysis for Research and Applications (MERRA) data.
Note: The long rainy season months are March, April and May and the short rainy season months are October, November and December.

4.2.2 Migration

Earlier work using these data report large income gains from internal migration. Using the 1991-2004 panel of the survey, Beegle, De Weerdt, and Dercon (2011) find that migrants enjoy a 36 percentage point growth premium over non-migrants – even after controlling for selection. As documented in Table 3.4 of Chapter 3, the 2010 data show that despite only minor differences in welfare in early 1990s, those who migrated out of the region became twice as rich as those who decided to stay.

The 2010 questionnaire collected detailed information about migration histories. The migration module focused on long-term migration spells of at least six months. Each panel respondent was asked how many times they had migrated and the year they first left the baseline village. For return migrants, the year of return to the baseline village is also known. The essay focuses on long-term migration decisions. This focus is motivated by the data: the mean – right censored – migration spell is 8.4 years (median: 8).

The recall nature of the questionnaire is less of a concern here. First, as we tracked individuals who migrated, these questions were asked directly from the migrants and not, as often is the case, from those who stayed behind. Second, the questionnaire focused on long-term migration spells. Psychological and survey methodological literature suggests that individuals are (more) likely to remember salient events (Dex 1995), such as weddings or long-term migration spells. This is further confirmed by Smith and Thomas (2003) who employ matched migration history data from Malaysia to model recall bias associated with migration. They find that the length of the migration spell was positively associated with the likelihood that the respondent reported to have migrated and recommend that migration history modules should focus on moves that lasted at least six months.

As described in Table 3.3 of Chapter 3, more than one-third of the female respondents but none of the male respondents cited marriage as the reason for migrating. Typically, the future husband and his household are responsible for the bride price. This eases the financial burden associated with migration at the origin household. Therefore, it is anticipated that female migration patterns are less affected by temperature shocks. Less than 15 per cent of the female respondents reported that they left because of work. In contrast, almost 45 per cent of the male migrants reported to have moved because they had found work or went looking for work. Given these differences in out-migration motives, in Section 4.4, I analyse the impact of weather shocks on migration patterns separately for men and women.

Age is one of the main determinants of migration decisions. This can be seen in Figure 4.3 and Figure 4.4 that depict the annual migration rates for each age group by gender. We see that the fraction of individuals leaving is low at about two per cent before the age of 15. After this age, many children finish schooling and leave because of work or marriage. The hazard rate peaks around the age of 20 and declines thereafter. The mobility after the age of 39 is low. This motivates restricting the migration analysis to 15 to 39 year individuals.[46]

Figure 4.3: Fraction of males migrating by age

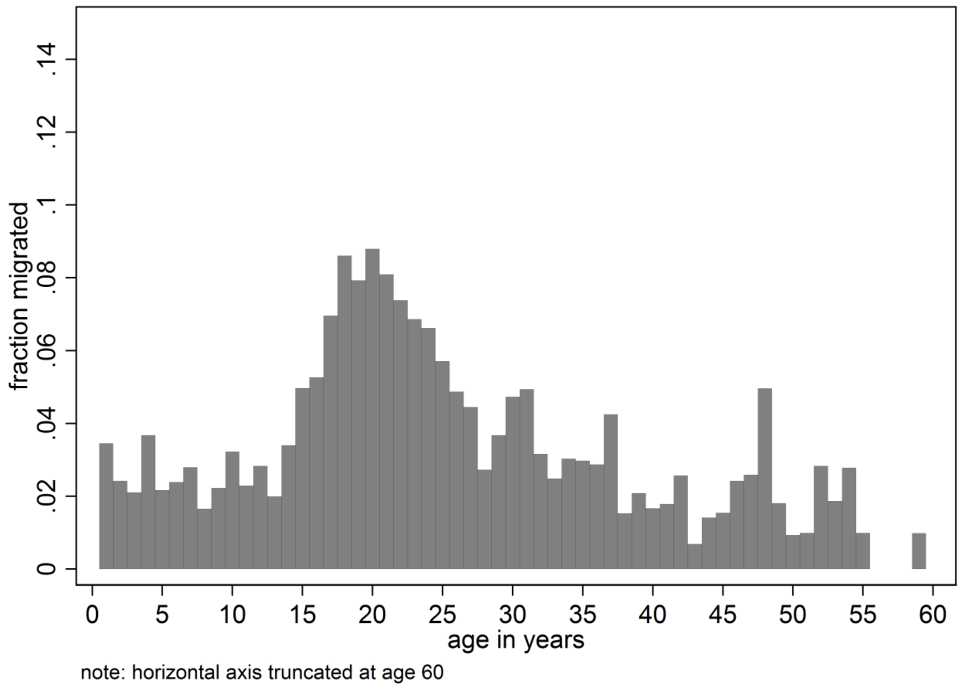

[46] Also previous studies have focused on this age group (e.g., Gray and Mueller 2012a, 2012b).

Figure 4.4: Fraction of females migrating by age

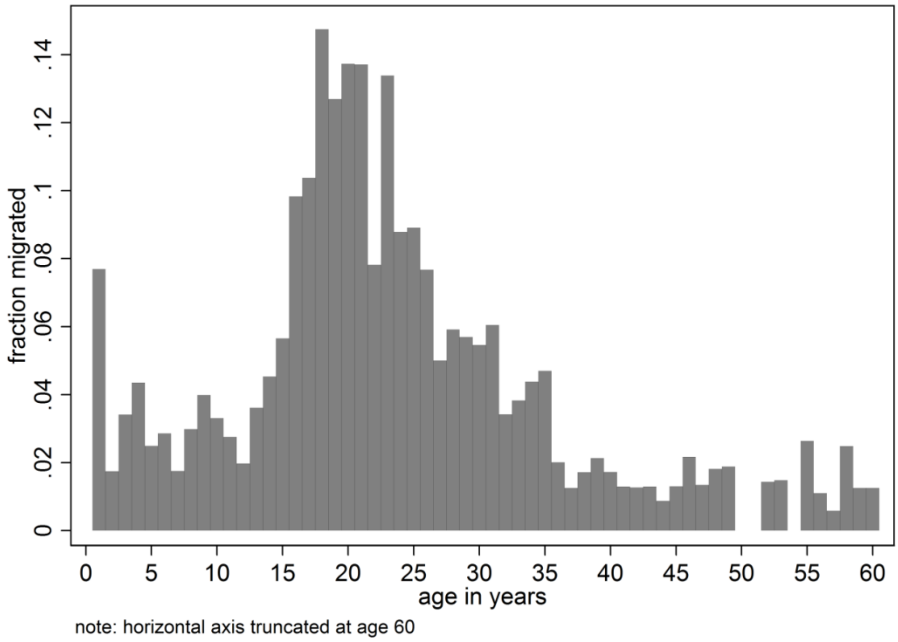

Education is another important determinant of out-migration. Figure 4.5 provides the annual migration rates by educational status for men and Figure 4.6 for women. The rates vary across educational categories with individuals with secondary degrees more likely to leave.

Figure 4.5: Annual hazard rates among men by level of education

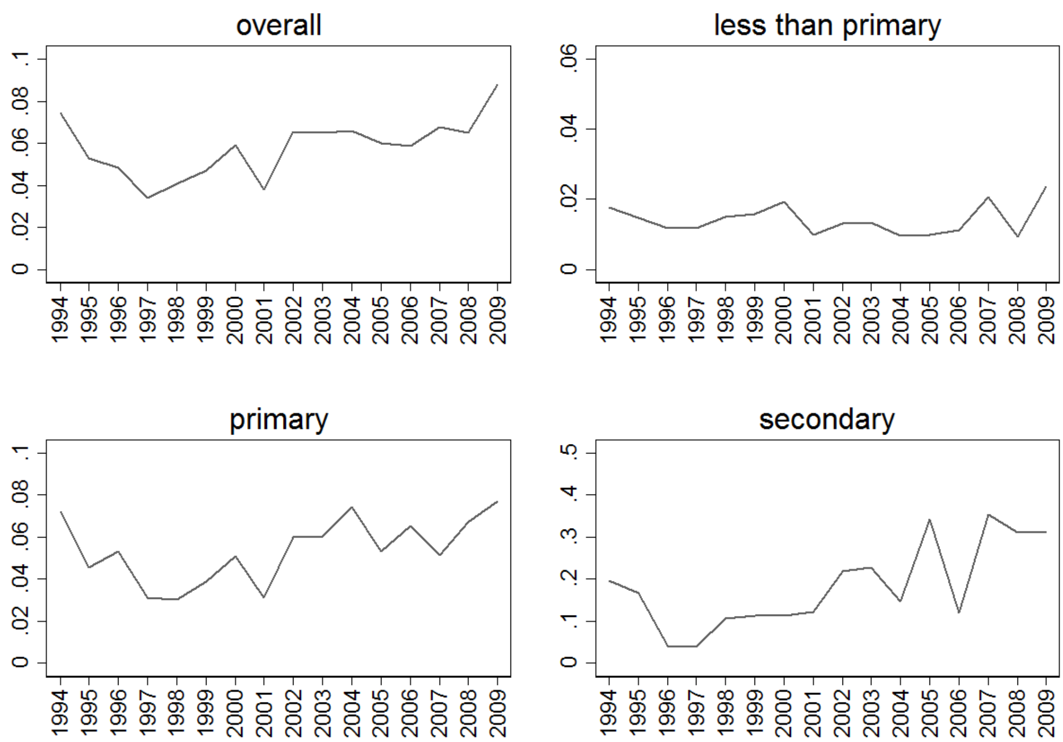

Note the changes in scale in the vertical axis

Mean annual hazard rates:

	mean	std. dev.
less than primary	0.014	0.004
primary	0.054	0.015
secondary	0.183	0.102
overall	**0.058**	**0.014**

Figure 4.6: Annual hazard rates among women by level of education

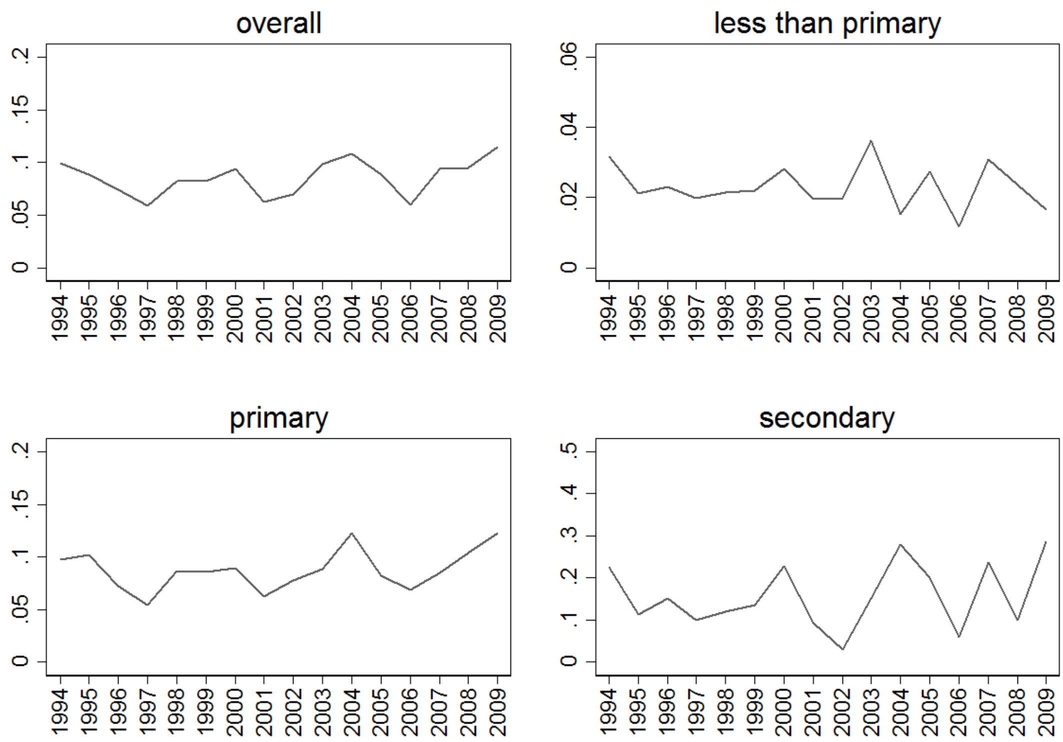

Note the changes in scale in the vertical axis

Mean annual hazard rates:

	mean	**std. dev.**
less than primary	0.023	0.006
primary	0.088	0.019
secondary	0.157	0.078
overall	**0.086**	**0.017**

4.2.3 Temperature and rainfall data

For a given geographical area, the quality of the weather data typically depends on the number of active weather stations. In poor countries, the collection of accurate weather data may not always be a priority. In Africa, like in many other continents, the number of active weather stations has been in steady decline over the past decades (Lorenz and Kunstmann 2012). Thus, the so called re-analysis data based on combining satellite and *in situ* observations provide a valuable resource for regions with a sparse station network.

The temperature data come from NASA's Modern-Era Retrospective analysis for Research and Applications (MERRA). MERRA is a global gridded data set based on retrospective analysis of historical weather data obtained from satellite images and weather stations (Rienecker et al. 2011). I use GPS coordinates to link the gridded temperature data to the baseline villages. The data set provides *daily* temperature measures aggregated into grids that are 1/2° in latitude × 2/3° in longitude (roughly 55 km × 75 km at the equator).

I also need data on rainfall. Temperature and rainfall are likely to be correlated and without the inclusion of rainfall the identification may suffer from the classic omitted variable bias (Auffhammer et al. 2013).[47] For rainfall I can exploit data based on a dense station network.[48] I have historical rainfall data from the Tanzanian Meteorological Agency for gauges in more than 200 weather stations in Tanzania. For Kagera, I take all weather stations within the 100 km radius from each village (17 stations in total) and link

[47] The mean within-village correlation coefficient between the growing season temperature and rainfall over 1994-2009 is -0.347 with a standard deviation of 0.189.

[48] The network of stations that provide temperature data for Tanzania is sparse. For example, a widely used global gridded data set provided by the Climate Research Unit (CRU), hosted by the University of East Anglia, uses only three weather stations from Tanzania (Rowhani et al. 2011) – none of them from Kagera. This motivates the reliance on re-analysis data.

the rainfall data to each village using an inverse distance weighing method.[49] The results are also robust to using the coarser rainfall data provided in the MERRA data set.

All weather data sets span the entire study period. I construct the mean monthly temperature and rainfall during the previous two rainy seasons. [50] [51] As only very few households have access to irrigation (see Section 2.2), I expect the weather conditions in the growing season to largely determine the contemporaneous incomes. Table 4.1 provides the summary statistics of the weather variables in the 51 villages in 1994-2009. The mean monthly growing season temperature across the villages is 23 °C. The standard deviation stands at 1.29 °C but this variation mainly originates from the differences between the villages. The within-village standard deviation is only 0.33 °C implying that temperatures remain fairly stable over time. The empirical strategies described in Sections 4.3 and 4.4 exploit this within-village variation in temperatures. Therefore, in what follows, I will interpret the impact of temperature shocks with respect to the within-village standard deviation (0.33 °C).

Table 4.1: Summary statistics of the weather variables in the 51 villages in 1994-2009

		std. dev.		
	mean	**overall**	**between villages**	**within villages**
temperature (°C)	22.70	1.29	1.26	0.33
growing degree days	2706.2	237.6	231.9	60.6
rainfall (mm)	140.93	68.16	47.8	49.0

[49] The method involves estimating rainfall for each village using all available data from the stations. A village specific monthly rainfall value is calculated by weighing each monthly rainfall reading with the inverse of the distance to the rainfall station where it was recorded.
[50] For temperature this means first calculating the mean daily temperature for each growing season month (March, April, May, October, November and December) and then taking the mean over these to construct an annual measure. For rainfall, I first calculate the total rainfall for each growing season month and then take a mean over these. The results are robust to using growing degree days instead of mean temperatures (see below).
[51] Note that within a village all temperature and rainfall readings at any one time are identical. As such, all variation at the village level is through time.

The agronomical literature often uses growing degree days (GDDs) to capture temperature's apparent non-linear relationship with plant growth (e.g. Ritchie and NeSmith 1991). Recent studies in economics and climate change research have also adopted this measure (Schlenker, Hanemann, and Fisher 2006, Schlenker and Lobell 2010a, Burke et al. 2011, Dillon, Mueller, and Salau 2011, Fisher et al. 2012). The construction of GDDs involves summing up the number of 'temperature days' over the growing seasons. [52] Table 3 shows that the mean number of growing degree days in the data is 2,706 with a within-village standard deviation of 60.6. This one standard deviation corresponds to a year that had 61 growing season days in which the temperature was 1 °C above its long-run mean.[53]

The data based on self-reports suggest that weather related shocks are common reasons for income fluctuations in Kagera. The 2004 and 2010 rounds collected extensive retrospective information about the income shocks experienced by the panel respondents. Nearly 60 per cent of the respondents reported to have experienced at least one shock in 1994-2009. As documented in Chapter 3, the most frequently reported shocks were the death of a family member (27% of all shocks), poor harvest due to bad weather (21%) and serious illness (19%).

4.3 Temperature shocks and household consumption

The econometric analysis begins by estimating the impact of temperature shocks on household consumption. I use the four waves of household panel survey administered in 1991-94. As described in Chapter 2, this baseline survey consisted of 915 households that

[52] The literature typically sets two thresholds, one at 8 °C and the other at 32 °C. A day with an average temperature of 9 °C counts as one degree day. Similarly, a temperature of 20 °C accounts for 12 degree days. Temperatures above 32 °C are considered harmful for plant growth and thus do not contribute to the degree days, so the maximum daily value is 24 °C. I adopt these thresholds and then sum the degree days over the two growing seasons.

[53] The growing season in Kagera has 184 days.

were interviewed up to four times. After dropping the households that were interviewed only once and a small number of households with missing consumption data, I am left with 3,277 observations from 912 households.

These early rounds of the survey also collected data on incomes. However, measuring income in a context where most households engage in self-employed agriculture is difficult and subject to a large margin of error (Deaton 1997). In such a setting, consumption has been viewed to provide a more reliable measure of welfare (e.g., Deaton and Grosh 2000). This motivates the use of consumption as the outcome variable to estimate the impact of weather shocks on household welfare. [54]

The consumption data originate from extensive food and non-food consumption modules in the survey. The consumption basket includes 97 food items (home produced and purchased) and 36 non-food items. The aggregates are temporally and spatially deflated using data from a price questionnaire included in the survey. Consumption is expressed in per capita terms using 1991 Tanzanian shillings (Tsh). The recall period was 12 months in the first wave (administered usually in 1991 or 1992). The waves 2, 3 and 4 were administered 6 to 7 months after the previous wave and therefore the recall period in the consumption module was six months. Following Bengtsson (2010) [55], the consumption values are normalised by dividing the annual value in wave 1 by two. The mean household per capita consumption for the pooled sample is 34,191 Tsh with a standard deviation of

[54] The OLS estimate for temperature appears statistically significant and negative when household per capita income or agricultural income is used as the outcome variable. However, the estimated impacts are implausibly large. The median quantile regression produces more similar (and statistically significant) temperature estimates compared to the consumption models. This suggests, as discussed in Bengtsson (2010) and confirmed by graphical analysis done by the author (not reported), that the income data are characterised by numerous outliers. Therefore the OLS results based on the income data should be interpreted with extreme caution. Results of these income regressions are available from the author upon request.

[55] Bengtsson (2010) employs the first four waves of the survey to show how child and adult body weights respond to temporary changes in household income.

26,914 Tsh. The econometric strategy described below controls for any questionnaire specific traits and seasonality aspects of consumption through including controls for wave and the quarter of year when the interview took place.

I use a fixed-effects approach (Burke et al. 2011, Dell, Jones, and Olken 2012, Fisher et al. 2012, Dell, Jones, and Olken 2013) to identify the effect of temperature changes on household consumption. Using household and year fixed effects, the impact of the temperature changes is identified from household (or community) specific deviations in temperature while controlling for annual shocks common to all villages.

Building on the work in Bengtsson (2010), the natural log of per capita consumption for household h in wave t ($\ln c_{ht}$) is modelled as:

$$\ln c_{ht} = \beta_1 T_{h,t-1} + \beta_2 R_{h,t-1} + \gamma_h + q_{ht} + \omega_t + \varepsilon_{ht}, \tag{4.1}$$

where t = 1,…,4. The term $T_{h,t-1}$ captures the temperature and $R_{h,t-1}$ the rainfall during the rainy season months in the 12 months preceding the interview.[56] The term q_{ht} is a set of dummies controlling for the quarter of year when the interview took place. Wave dummies are captured by ω_t. The term γ_h represents household fixed effects capturing all time-invariant characteristics of the households. The computed standard errors are clustered at the village level.

Besides the wave fixed effects, the survey design permits the further control for unobserved time-variant effects (Bengtsson 2010). First, the wave does not fully

[56] I did not find evidence of a non-linear (quadratic) relationship between the weather variables and consumption. This motivates the focus on the linear specification here. Moreover, the *within-transformation* makes the interpretation of the quadratic terms cumbersome in the fixed effect models (see McIntosh and Schlenker 2006).

correspond to the month or even to the calendar year of the interview. This was because the survey team visited the baseline villages at different times of the year and sometimes even in different calendar years within waves. [57] For example, in the first wave, 47 per cent of the households were interviewed in 1991, 50 per cent in 1992 and 3 per cent in 1993.[58] Furthermore, the weather variables are constructed so that they capture the conditions in the growing season months in the 12 months prior to the interview month. These two features induce additional random cross-sectional variation to the data. The first additional specification includes a set of year dummies (θ_t) that capture any macro-level shocks specific to the calendar year. I then replace these with district by year interaction terms ($d_h * \theta_t$). [59] While these interaction terms capture district specific trends, such as district level price effects, they come with a cost of absorbing a large amount of the variation in the temperature data.

Table 4.2 reports the results. The first three columns contain the estimates for equation (4.1) using mean monthly growing season temperature in the past 12 months before the interview date. Columns 4-6 use growing degree days (GDDs). Furthermore, columns 1 and 4 control for time effects only through wave dummies. Columns 2 and 5 add the year dummies and columns 3 and 6 replace these with district by year interaction terms. The temperature variables appear with negative coefficients and are well determined in all specifications. The impact of one standard deviation increase (0.33 °C) in average monthly growing season temperature decreases household per capita consumption by 4.9 to 5.5 per cent (on average and *ceteris paribus*) depending on the specification. In the most preferred specification, containing the fixed effects models with year dummies

[57] The visits across villages were randomised at the first round of interviews (World Bank 2004).
[58] If a household dropped out from the survey it was replaced with another household. This explains why few households were interviewed in 1993.
[59] The 51 villages group into six districts.

(column 2), a one standard deviation increase in temperature decreases household per capita consumption by 4.9 per cent. The 95 per cent confidence interval for this negative effect is [0.9; 8.8]. The specifications using GDDs report similar impacts. In column 5 we see that a one standard deviation (60.6) increase in the previous year's growing degree days causes a 4.8 per cent fall in per capita consumption (on average and *ceteris paribus*).

Columns 3 and 6 include the district-by-year interaction terms. The coefficients on the temperature variables (mean temperature and GDDs) are of similar magnitude as in other columns, although, as expected, somewhat less precisely estimated ($p<0.10$). This implies that the channel through which temperature affects consumption is not mediated through the local price effects. Furthermore, following Harari and La Ferrara (2012), I also estimated Equation (4.1) using weather variables that capture the dry season (or off-season) temperature and rainfall. These variables did not exert independent effects on household consumption further supporting the notion that temperature affects consumption through the agricultural yield channel. [60] This conjecture is in line with two recent studies from Tanzania. Ahmed et al. (2011) and Rowhani et al. (2011) use regional level data on agricultural yields and document how positive temperature deviations have a large negative impact on yields. [61]

Rainfall appears with a significant, and economically meaningful, coefficient only when logged and when temperature is excluded from the specification. In addition, using natural logs instead of raw values of mean temperature and rainfall provides qualitatively similar results. [62]

[60] The results are not reported but are available upon request.
[61] The production functions in these two studies also include rainfall but similar to the results in this essay, it exerts only a negligible independent impact on yields.
[62] The regression results of rainfall only and log transformed weather variables are not reported but available upon request.

Table 4.2: Impact of temperature changes on household consumption

dependent variable: (logged) household per capita consumption	**1**	**2**	**3**	**4**	**5**	**6**
	mean temperature			**growing degree days (GDD) a)**		
temperature (°C or GDD)	-0.148**	-0.147**	-0.168*	-0.008**	-0.008**	-0.009*
	(0.060)	(0.060)	(0.095)	(0.003)	(0.003)	(0.005)
rainfall (mm)	-0.000	-0.001	0.001	-0.000	-0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
household fixed effects *χ2 (869) = 5681.2*	yes	yes	yes	yes	yes	yes
year dummies *χ2 (2) = 4.87*	no	yes	no	no	yes	no
district by year interactions *χ2 (17) = 51.03*	no	no	yes	no	no	yes
wave dummies *χ2 (3) = 121.9*	yes	yes	yes	yes	yes	yes
quarter of year dummies *χ2 (3) = 2.82*	yes	yes	yes	yes	yes	yes
Number of observations	3,277	3,277	3,277	3,277	3,277	3,277
Adjusted R^2	0.175	0.176	0.190	0.175	0.177	0.190
R^2	0.177	0.179	0.195	0.177	0.179	0.195

a) The growing degree day variable has been divided by 10 to facilitate interpretation.

note: *** p<0.01, ** p<0.05, * p<0.1. Standard errors (in parenthesis) are clustered at the village level. Mean temperature (columns 1-3) is measured as average monthly growing season temperature (in °C) 12 months before the interview month. Growing degree days (columns 4-6) are measured as growing degree days 12 months before the interview month. Rainfall is measured as average monthly growing season rainfall (in mm) 12 months before the interview month. The Wald tests compare the sum of squared residuals between the restricted and unrestricted models. These tests are based on the model presented in Column 3 (except in the case of year dummies).

Finally, the OLS approach implicitly assumes that temperature exerts a similar impact across the entire consumption distribution. The validity of this assumption was tested by estimating a pooled model with village and year dummies using quantile regression techniques. Appendix C presents the results. According to the inter-quintile tests the differences between the coefficients at the lower and upper quintiles are not statistically different from zero. Furthermore, since the quantile regression method focuses only on the sign and not the size (as OLS) of the residual, it is less sensitive to outliers in data and departures from normality (e.g., Koenker and Hallock 2001). As the median and the corresponding OLS regression produce broadly similar temperature coefficients it is unlikely that outliers are driving the results reported in Table 4.2.

These estimated impacts are large but broadly comparable to what have been documented in the recent empirical literature based on aggregated macro-level data. In an influential paper, Dell, Jones, and Olken (2012) investigate the within-country impact of annual temperature deviations on economic growth. Employing a cross-country panel spanning 1950-2003 with country and region by year fixed effects, the authors find that a 1 °C increase in annual temperature reduces annual GDP growth rate on average by 1.3 percentage points in poor countries. [63] This negative impact is larger (1.9 percentage points) when the sample is constrained to Sub-Saharan African countries.

Schlenker and Lobell (2010a) use cross-country panel data based on Sub-Saharan African countries from 1961 to 2007 to study the impact of temperature fluctuations on five key staple crops. Their model estimates, based on country fixed effects and quadratic time trends, suggest that a one degree increase in the growing season temperature decreases average Maize yields (tons/ha) by 10 per cent, Sorghum yields by 6 per cent, Millet yields by 9 per cent, Groundnut yields by 8 per cent and Cassava yields by 7 per cent (Schlenker and Lobell 2010b). The estimated impacts were even larger for South-Africa and Zimbabwe, two countries characterised by higher fertilizer use.

Burke et al. (2011) employ an African cross-country panel from 1961 to 2008 to study the impact of annual temperature deviations on Maize yields. Using country and year

[63] Dell *et al.* define a country as poor if the PPP GDP per capita is below the median in the first year in the data.

fixed effects, the authors find that a 1 °C increase in average growing season temperature reduces maize yields on average by 10 to 30 per cent depending on the specification.

The foregoing has established that temperature shocks induce considerable fluctuations in household consumption. The next section analyses the impact of these income shocks on long-term migration decisions.

4.4 Migration and weather shocks

4.4.1 Econometric strategy

I employ event history analysis based on discrete time methods (Allison 1982, Jenkins 1995). The model allows for the inclusion of time-varying explanatory variables, such as temperature and rainfall. It also accommodates the apparent (independent) censoring of the data. The procedure involves the transformation of the data to person-year level. I consider mobility between 1994 and 2009.[64] Furthermore, given the age restriction, individuals enter the data set when they turn 15 and leave when they migrate or turn 40. This yields an unbalanced data set of 24,962 person-year observations. The total number of individuals is 3,240 (1,662 women and 1,578 men) with 1,480 right censored individuals, those who had not migrated by 2009. Because the migration motives are different for men and women, the analysis is separated by gender.[65]

The event history analysis typically adopts the logit model. The logit model constraints the probability to a [0,1] interval by assuming a cumulative density function (CDF) that follows a logistic distribution. Here the estimated model can be formulated as:

[64] The 2010 migration module asked about mobility between 1994 and 2010. The 2010 fieldwork began in April and the last interviews were held in December. Since 2010 was not a full year, I exclude it from the analysis.
[65] Also the data support the separation by gender. A likelihood ratio test akin to Chow (1960) with one degree of freedom yields a $\chi 2$ test statistic of 89.0 (p=0.000).

$$\ln\left[\frac{m_{ivt}}{1-m_{ivt}}\right] = x'_{ivt}\beta + \delta_t + \vartheta_v. \tag{4.2}$$

The out-migration event of an individual i is captured by the binary variable m_{ivt} that assumes a value 1 if the person migrates from village v at a year t, and zero otherwise. The term ϑ_v represents a set of village dummies and δ_t year dummies. As for the consumption specification, the standard errors are clustered at the village level.

The vector x_{ivt} captures individual and household level characteristics affecting an individual's probability of migrating. The time varying variables include dummies for education and mean monthly growing season temperature [66] and rainfall in the previous year. The changing risks of migrating associated with age are modelled using age dummies. [67] The time invariant variables are measured from responses in wave 1 administered before 1994. They include an individual's gender and relationship to the household head, head's age and gender, land size and the demographic composition of the household. Table 4.3 provides the summary statistics of these covariates. The reason for not controlling for initial consumption (or income) is that it is likely to be highly correlated with current income or consumption (which I do not observe). This would make the interpretation of the weather variables as income shocks cumbersome.[68]

---

[66] As for the consumption regressions, using growing degree days instead of mean temperature produces qualitatively similar results. Results are not reported but available from the author.

[67] The results are robust to using an age specific spline variable instead of dummies. The dummy variable approach is preferred here to facilitate the interpretation of the age effects.

[68] Nevertheless, the results are not sensitive to the inclusion of the initial consumption variable in the model.

Table 4.3: Summary statistics

	mean	std. dev.
Time-variant variables:		
temperature (°C)	22.61	1.290
rainfall (mm)	139.6	67.78
age in years	23.90	6.735
less than primary education (reference)	0.297	0.457
primary education	0.656	0.475
secondary education	0.046	0.210
Time-invariant variables: *		
male	0.530	0.499
head or spouse	0.127	0.332
child of head	0.583	0.493
grandchild of head	0.115	0.319
other relation to the head (reference)	0.175	0.380
male head	0.786	0.410
age of the head	47.95	16.10
land size (acres)	5.086	5.017
number of males 0-5 years	0.715	0.877
number of males 6-15 years	1.301	1.158
number of males 16-60 years	1.426	1.020
number of males 61+ years	0.187	0.390
number of females 0-5 years	0.737	0.916
number of females 6-15 years	1.315	1.292
number of females 16-60 years	1.769	1.257
number of females 61+ years	0.184	0.403
Number of person-years	**24,962**	

* measured in wave 1 (usually administered in 1991 or 1992)

The identification strategy here is similar to the fixed effects approach employed in Section 4.3. Village dummies absorb the spatial variation in these data, so only the within-village variation in temperatures over time is exploited. The year dummies control for any annual macro-level shocks common to all villages.[69]

[69] The results are robust to using the spell at risk variable instead of the year dummies (see Section 0).

4.4.2 Results

Table 4.4 provides the results for men. The odd columns show the log-odd ratios and the even columns the marginal effects. The first two columns provide the results without village dummies where the weather variables are constructed as deviations from village mean. The remaining four columns show the results with the village dummies included. The estimates are broadly similar across the models with (Columns 3-6) and without village dummies (Columns 1-2) suggesting that the estimated impacts on variables capturing the individual level characteristics are not influenced by the time invariant village characteristics. Focusing first on the control variables, we see that, as anticipated, education increases the risk of migrating. The coefficients on the age dummies reflect the age pattern observed in Figure 4.3. Furthermore, the ties to the household head are important predictors of migration. Other things equal, the closer the familial ties are to the head the less likely the individual is to migrate. This finding is line with previous research in Tanzania (Beegle, De Weerdt, and Dercon 2011) and Malawi (Beegle and Poulin 2013).

Table 4.4: Impact of temperature changes on migration among males

	1	2	3	4	5	6
	log-odds	**mfx**	**log-odds**	**mfx**	**log-odds**	**mfx**
temperature §	-0.487***	-0.026***	-0.448***	-0.023***	-0.434***	-0.020***
	(0.141)	(0.008)	(0.138)	(0.007)	(0.136)	(0.006)
rainfall §	-0.001	-0.000	-0.000	-0.000	-0.001	-0.000
	(0.001)	(0.000)	(0.001)	(0.000)	(0.001)	(0.000)
secondary education	0.315***	0.017***	0.338***	0.018***	0.366***	0.017***
	(0.110)	(0.006)	(0.111)	(0.006)	(0.126)	(0.006)
primary education	1.624***	0.085***	1.674***	0.087***	1.597***	0.074***
	(0.150)	(0.007)	(0.160)	(0.008)	(0.170)	(0.008)
less than primary educ	reference	reference	reference	reference	reference	reference
age: 15-19	reference	reference	reference	reference	reference	reference
age: 20-24	-0.064	-0.003	-0.029	-0.002	0.123	0.006
	(0.125)	(0.007)	(0.124)	(0.006)	(0.136)	(0.006)
age: 25-29	-0.641***	-0.034***	-0.592***	-0.031***	-0.364**	-0.017**
	(0.154)	(0.008)	(0.160)	(0.008)	(0.163)	(0.008)
age: 30-34	-0.775***	-0.041***	-0.748***	-0.039***	-0.541***	-0.025***
	(0.156)	(0.009)	(0.159)	(0.008)	(0.174)	(0.008)
age: 35-39	-1.020***	-0.054***	-0.992***	-0.052***	-0.778***	-0.036***
	(0.223)	(0.012)	(0.224)	(0.012)	(0.224)	(0.010)
head or spouse	-0.684***	-0.036***	-0.684***	-0.036***	-0.701***	-0.032***
	(0.215)	(0.011)	(0.219)	(0.011)	(0.219)	(0.010)
child of head	-0.701***	-0.037***	-0.635***	-0.033***	-0.707***	-0.033***
	(0.101)	(0.005)	(0.104)	(0.005)	(0.107)	(0.005)
grandchild of head	-0.308**	-0.016**	-0.316**	-0.016**	-0.415***	-0.019***
	(0.139)	(0.007)	(0.142)	(0.007)	(0.150)	(0.007)
other relation to the head	reference	reference	reference	reference	reference	reference
household level controls	Yes		yes		yes	
village dummies	No		yes		yes	
year dummies	Yes		yes		yes	
Sample	full		full		educ migrants dropped	
Observations	13,225		13,225		12,699	
Pseudo-Log likelihood	-2,746.613		-2,685.366		-2,358.522	

note: *** p<0.01, ** p<0.05, * p<0.1.
§ = constructed as deviation from the mean in columns 1 and 2.
Standard errors (in parenthesis) clustered at the village level. The standard errors for the marginal effects (mfx) are calculated using the delta-method.
Temperature is measured as average monthly growing season temperature (in °C) in the previous year.
Rainfall is measured as average monthly growing season rainfall (in mm) in the previous year.
The household level controls are measured in wave (i.e. before 1994) and include head's characteristics (gender, age and education), household composition (number of males and females of certain age) and land ownership (acres).

The attention now shifts to the weather variables. As predicted by the theoretical model, we see that an increase in the mean temperature in the previous year's growing season decreases the probability of migrating. In column 4, a one standard deviation (0.33 °C) increase in the previous year's average monthly growing season temperature decreases the probability of migrating by 0.76 of a percentage point, on average and *ceteris paribus*. This corresponds to a reduction in the overall male migration rate of about 13 per cent. [70] The estimate is well determined. Columns 5 and 6 show the results when the individuals who reported migrating because of education are dropped from the sample. The coefficient on the temperature variable is of similar magnitude and remains significant at the one per cent level. This implies that individuals migrating because of schooling are not driving the results obtained. Rainfall in the previous year's growing season does not exert an independent impact on migration decisions.

Table 4.5 provides the results for women. Similar to men, the likelihood of leaving increases with education. Also the age patterns correspond to the age specific hazard rates presented in Figure 4.4. The temperature variable appears with a negative sign in all columns but the effect is not statistically significant at conventional levels. This confirms the established prior that female migration is less responsive to income shocks induced by weather due to the fact that it is largely motivated by marriage and family. As for men, the results are not sensitive to dropping the education migrants from the sample.

[70] The overall male migration rate is 5.8 per cent (see Figure 4.5).

Table 4.5: Impact of temperature changes on migration among females

	1 log-odds	2 mfx	3 log-odds	4 mfx	5 log-odds	6 mfx
temperature §	-0.014	-0.001	-0.015	-0.001	-0.027	-0.002
	(0.189)	(0.014)	(0.188)	(0.014)	(0.195)	(0.014)
rainfall §	-0.001*	-0.000*	-0.001	-0.000	-0.001	-0.000
	(0.001)	(0.000)	(0.001)	(0.000)	(0.001)	(0.000)
secondary education	0.264***	0.020***	0.295***	0.022***	0.288***	0.021***
	(0.073)	(0.006)	(0.083)	(0.006)	(0.083)	(0.006)
primary education	0.843***	0.064***	0.804***	0.061***	0.580***	0.042***
	(0.128)	(0.009)	(0.149)	(0.011)	(0.184)	(0.013)
less than primary educ	reference	reference	Reference	reference	reference	reference
age: 15-19	reference	reference	Reference	reference	reference	reference
age: 20-24	0.097	0.007	0.138	0.010	0.115	0.008
	(0.097)	(0.007)	(0.104)	(0.008)	(0.113)	(0.008)
age: 25-29	-0.398***	-0.030***	-0.328**	-0.025**	-0.301**	-0.022**
	(0.133)	(0.010)	(0.140)	(0.011)	(0.139)	(0.010)
age: 30-34	-0.676***	-0.051***	-0.603***	-0.045***	-0.534***	-0.038***
	(0.134)	(0.010)	(0.145)	(0.011)	(0.145)	(0.010)
age: 35-39	-1.175***	-0.089***	-1.109***	-0.083***	-1.031***	-0.074***
	(0.215)	(0.016)	(0.219)	(0.017)	(0.220)	(0.016)
head or spouse	-0.781***	-0.059***	-0.905***	-0.068***	-0.943***	-0.068***
	(0.182)	(0.014)	(0.180)	(0.013)	(0.183)	(0.013)
child of head	-0.016	-0.001	-0.017	-0.001	-0.051	-0.004
	(0.094)	(0.007)	(0.101)	(0.008)	(0.105)	(0.008)
grandchild of head	0.079	0.006	0.072	0.005	0.028	0.002
	(0.133)	(0.010)	(0.145)	(0.011)	(0.148)	(0.011)
other relation to the head	reference	reference	Reference	reference	reference	reference
household level controls	yes		yes		yes	
village dummies	no		yes		yes	
year dummies	yes		yes		yes	
Sample	full		full		educ migrants dropped	
Observations	11,737		11,737		11,377	
Pseudo-Log likelihood	-3,261.807		-3,225.630		-3,023.916	

note: *** p<0.01, ** p<0.05, * p<0.1.
§ = constructed as deviation from the mean in columns 1 and 2.
Standard errors (in parenthesis) clustered at the village level. The standard errors for the marginal effects (mfx) are calculated using the delta-method.
Temperature is measured as average monthly growing season temperature (in °C) in the previous year.
Rainfall is measured as average monthly growing season rainfall (in mm) in the previous year.
The household level controls are measured in wave (i.e. before 1994) and include head's characteristics (gender, age and education), household composition (number of males and females of certain age) and land ownership (acres).

4.5 Robustness

The village and year fixed effects absorb a considerable amount of the variation in the weather data. The village fixed effects are vital so that the effect of temperature is identified from the village-specific deviations rather than from differences in mean temperature levels between villages. The duration modelling framework, however, offers an alternative method to model time to migration. The spell at risk captures the number of elapsed years since the individual entered the data at time t.[71] Since individuals enter the data set at different calendar years, this variable is less correlated with the annual weather variables. Appendix D reports estimates for the migration models using the spell at risk variable instead of calendar year dummies. First, the likelihood ratio tests show that the year dummies provide the best fit to the data. Among the spell at risk specifications, the higher order polynomials provide the best fit and yield similar temperature coefficients to the ones presented in Table 4.4 and Table 4.5.

A concern in the heavy use of village fixed effects is that perhaps some of these observed effects are originating from a single village that does not conform to the general pattern in the region. I addressed this issue by re-running the consumption and migration regressions by omitting each village in turn. Appendix E shows how the temperature coefficients remain stable across these 51 regressions.

Finally, the migration regressions presented in Section 4.4 include controls at the individual level but do not fully control for (time-invariant) individual heterogeneity. To address this concern, the odd columns in Table F.1 of Appendix F repeat Column 4 in Table 4.4 and Table 4.5 but using the linear probability model (LPM). The use of this model allows the introduction of individual-level fixed effects. First, columns 1 and 3

[71] The spell year variable records a value of 1 in the 1st year, and 2 in the 2nd year and so forth.

show that the LPM produces identical marginal effect estimates to those reported for the logit in Table 4.4 and Table 4.5. The results are thus not sensitive to the choice of the underlying statistical model (logit or LPM). The even columns replace the village fixed effects with individual level fixed effects. As expected, the use of individual level fixed effects inflates the standard errors but otherwise the coefficients on the weather variables appear broadly similar. This implies that unobserved time-invariant heterogeneity does not appear to be driving the results obtained.

4.6 What about the long-run weather variability?

The primary focus of this essay is to analyse the impact of *ex post* temperature shocks on long-run migration decisions. The New Economics of Labour Migration (Stark and Bloom 1985, Rosenzweig and Stark 1989) posits that in an environment characterised by large weather uncertainty households engage in strategic migration. More specifically, households send migrants to areas where the variability of income is not correlated with the variability of income at the origin. The purpose of this section is to study whether this hypothesis holds in the Tanzanian context.

The econometric strategy in this essay has modelled migration decisions using village fixed effects. In these models, the long-run variability of temperature and rainfall are absorbed in the fixed effect. Next I retrieve the fixed effect (ϑ_v) from Equation (4.2) and decompose it to an unknown part (unobserved time-invariant heterogeneity) and to a part associated with long-run temperature and rainfall variability.[72] In the first stage, I use the LPM and re-run regression based on Equation (4.2) without the time-varying temperature and rainfall variables. In the second stage, I regress the retrieved village level fixed on the long run variability of temperature and rainfall while controlling for the elevation of

[72] Earlier research has also used this type of approach to study the role of time-invariant variables in fixed effect models (see e.g., Reilly and Witt 1996, Falco et al. 2011).

the village. If households engage in strategic migration we should expect to see higher out-migration rates in villages characterised by larger long-run weather variability.

I measure the long-run variability using the coefficient of variation (CV) based on temperature and rainfall data spanning 1981-2010. [73] Table 4.6 presents the results. The odd columns provide the results for OLS regressions. In the even columns I use the standard errors from the first stage as weights. This Weighted Least Squares approach gives more weight to villages that had a more precisely estimated coefficient in the first stage. The first two columns in Table 4.6 provides the results for men and the last two columns for women. We see that long-run weather variability is not associated with higher or lower out-migration rates in any of the columns. These results suggest, similar to the findings in Chapter 3, that migration is not a strategic household level decision to spread risk in this context. [74]

[73] CV is constructed as the standard deviation divided by the mean.

[74] This finding is robust to 1) omitting the elevation variable from the second stage, 2) restricting the right hand side variables in the first stage only to village and year dummies and 3) using mean and standard deviation instead of CV. These results are not reported but available from the author.

Table 4.6: Decomposing the Fixed Effect

	means	Males		females	
		OLS	**WLS**	**OLS**	**WLS**
CV of temperature	0.300	-0.860	-0.852	0.216	0.194
	[0.128]	(0.925)	(0.926)	(0.965)	(0.983)
		((1.113))	((1.136))	((0.821))	((0.837))
CV of rainfall	0.019	0.007	0.007	0.010	0.012
	[0.004]	(0.027)	(0.027)	(0.028)	(0.028)
		((0.020))	((0.020))	((0.030))	((0.030))
elevation in km	1.332	-0.031	-0.030	-0.001	0.001
	[0.163]	(0.021)	(0.021)	(0.022)	(0.022)
		((0.017))	((0.017))	((0.022))	((0.022))
intercept		0.111***	0.109***	0.100***	0.098***
		(0.033)	(0.033)	(0.034)	(0.035)
		((0.029))	((0.030))	((0.036))	((0.037))
Observations		51	51	51	51
R^2		0.067	0.063	0.003	0.004

note: *** p<0.01, ** p<0.05, * p<0.1.
Dependent variable: Fixed effect (ϑ_v) from a LPM regression based on Equation (4.2) but without rainfall and temperature variables.
Standard deviations are in square brackets, *standard errors* in parenthesis and *robust standard errors* in double-parenthesis.
WLS: Weighted Least Squares model using the standard errors in the first stage model as weights. CV: coefficient of variation in 1981-2010.

4.7 Conclusions

A rapidly emerging literature documents how economic outcomes in African countries co-move with temperature. Previous studies use panel data aggregated to a country level (e.g., Dell, Jones, and Olken 2012, Burke et al. 2011, Schlenker and Lobell 2010a). This study is the first one that takes the analysis to a micro-level using household level panel data from rural Tanzania. The fixed effects model confirms the findings from the macro-level studies. Household consumption is highly responsive to temperature changes in Kagera with a one standard deviation increase in the mean monthly growing season temperature decreasing household per capita consumption by 4.9 per cent.

Due to the lack of reliable data on agricultural yields, I am not able to conclusively isolate the mechanism in which temperature changes affect consumption outcomes. The fact that off-season temperatures do not exert an impact on consumption suggests that the mechanism is through the agricultural yield channel. This conjecture is in line with two recent studies that use regional level data from Tanzania and report how agricultural yields co-move strongly with temperatures (Rowhani et al. 2011, Ahmed et al. 2011).

Temperature shocks also inhibit internal migration. Each year, five to eight per cent of potential migrants leave their village to look for work or for other reasons. This type of geographical mobility is associated with large income gains and presents one of the main routes out of poverty in rural Tanzania (Beegle, De Weerdt, and Dercon 2011, Christiaensen, De Weerdt, and Todo 2013). A one standard deviation increase in the growing season temperature causes a 13 per cent reduction in the overall male migration rate, *ceteris paribus* and on average. This result implies that liquidity constraints bind in this Tanzanian context. Female migration is unaffected by temperature changes possibly due to the fact that it is largely motivated by marriage and family.

These findings highlight the importance of temperature changes on welfare both at the individual and household level in Sub-Saharan Africa. The uninsured income shocks inhibit young adults from tapping into the opportunities brought about by geographical mobility.

5 Essay 3: Measuring catch-up growth in malnourished populations

This essay is forthcoming in *Annals of Human Biology*
(DOI: 10.3109/03014460.2013.827239)

5.1 Introduction

Growth faltering in developing countries typically begins in the first three months of life and persists until the age of two or three (Eveleth and Tanner 1976, Shrimpton et al. 2001, Victora et al. 2010). This is usually caused by insufficient or poor nutrition or by infectious diseases (Golden 1994). Since children are growing rapidly during these years, even a short retarded growth spell in this period quickly leads children to fall behind from their fast growing peers (Shrimpton et al. 2001, Victora et al. 2010). To bounce back to their original growth curves, short and stunted children need to experience higher growth rates than their healthy and well-nourished peers.

In the clinical and epidemiological literature, such catch-up growth is defined as height velocity that is above the expected for the child's age and occurs after a period of growth retardation (Tanner 1981, Williams 1981, Ashworth and Millward 1986, Boersma and Wit 1997). A complete catch-up takes place if the original, pre-retardation, growth curve is attained (Williams 1981, Ashworth and Millward 1986). Height-for-age z-score (HAZ) measures the distance in height to the median child of a healthy and well-nourished population providing the exact empirical counterpart for this definition. An increase in HAZ means that height velocity is above what is expected for child's age and gender.

Recent literature has studied the existence of catch-up growth by regressing adult or pre-adolescent height on early childhood height (e.g. Martorell et al. 1992, Hoddinott and

Kinsey 2001, Fedorov and Sahn 2005, Alderman, Hoddinott, and Kinsey 2006, Victora et al. 2008, Mani 2012, Outes and Porter 2013). This essay demonstrates that these tests misinterpret catch-up growth with within-population convergence. Such convergence occurs if short and stunted children experience faster growth than others or taller children grow at a slower pace than others – or because of some combination of both. By regressing current height on early childhood height researchers cannot tell whether children are catching-up or faltering. Moreover, by focusing on the movement *within* the population these regressions fail to capture catch-up growth at the population level. The interpretation of these regression coefficients is further hampered by regression-to-the-mean (Friedman 1992, Quah 1993).

The evidence on catch-up growth is further limited due to the scarcity of longitudinal data sets that follow children over their entire childhood (Eckhardt, Gordon-Larsen, and Adair 2005). The empirical part of this essay provides new evidence for Sub-Saharan Africa. Using a 19-year longitudinal data set for the Kagera region in Tanzania, I find considerable catch-up growth. The mean HAZ-score in the cohort improves from -1.86 in early childhood to -1.20 in adulthood. Without catch-up growth, the average girl in the sample would have been nearly 5 centimetres shorter in adulthood. For boys, the difference between the predicted adult height in early childhood and the actual adult height is around 4.5 centimetres. Graphical analysis shows that this catch-up growth takes place in puberty. Recently, Coly et al. (2006) and Prentice et al. (2013) documented similar extensive pubertal catch-up growth among Senegalese and Gambian children.

The findings from these three African cohorts challenge the conventional view in the literature where catch-up growth after early childhood is difficult and seldom observed (e.g. Martorell, Rivera, and Kaplowitz 1990, Grantham-McGregor et al. 2007, Schroeder

2008). The results in this essay add to the emerging evidence that puberty offers an opportunity window for catch-up growth (Haas and Campirano 2006, Moradi 2010, Prentice et al. 2013) and prompt re-thinking the policy recommendation where nutritional interventions are *only* targeted to young children (for a similar point, see Prentice et al. 2013). Furthermore, a number of studies use adult height as a proxy for health and wealth in early childhood (Almond and Currie 2011). The findings in this essay imply that height in adulthood may not be such a good indicator of childhood conditions.

Finally, the pubertal catch-up growth may provide a missing piece to the *African height puzzle*. In a seminal essay, Deaton (2007) shows that, contrary to most other countries in the world, disease environment and national income are not good predictors of female adult height in African countries. African women are taller than their economic and disease environment suggests. Deaton (2007), (Bozzoli, Deaton, and Quintana-Domeque (2009)) and Gørgens, Meng, and Vaithianathan (2012) explain this puzzle with a selection effect: childhood mortality is concentrated on the genetically short children thus shifting the mean adult height right. Moradi (2010; 2012) proposes an alternative explanation that African children experience a considerable catch-up growth in puberty. The results presented here provide support to this hypothesis.

The structure of this essay can now be outlined. The next section demonstrates the methodological flaws in the recent empirical literature. Section 5.3 provides a brief survey of studies employing longitudinal surveys and defining catch-up growth as the change in height-for-age z-score. Section 5.4 presents the data used in the empirical analysis. Section 5.5 provides the results and section 5.6 analyses the timing of the catch-up growth in the Kagera cohort. Section 5.7 concludes.

5.2 How *not* to measure catch-up growth

Recent empirical literature (Martorell et al. 1992, Hoddinott and Kinsey 2001, Fedorov and Sahn 2005, Alderman, Hoddinott, and Kinsey 2006, Victora et al. 2008, Mani 2012, Outes and Porter 2013) subscribes to the definition of catch-up growth used in the medical literature but employs regression analysis, usually using the following type of specification:

$$h_{it} = \alpha + \beta h_{i,t-1} + e_{it}\ , \tag{5.1}$$

where α is the intercept and e_{it} the error term. The term h_{it} is height (or height-for-age z-score) of individual i in period t and $h_{i,t-1}$ height (or height-for-age z-score) in a previous period. In studies that have an access to a sufficiently long data set, t usually refers to height measured in adulthood and *t-1* to height measured in early childhood. The β coefficient is then interpreted as the measure of catch-up growth. A zero β coefficient on the lagged height is taken as a complete catch-up: initial height does not predict adult height. A coefficient equal to one is interpreted as evidence that no catch-up growth takes place: short or stunted children remain locked into their lower growth trajectory. Finally, few studies (e.g. Hoddinott and Kinsey 2001) use change in height as the dependent variable. This growth specification changes the interpretation of β but leads to qualitatively identical conclusions about the extent of catch-up growth (see Fedorov and Sahn 2005).

There are a number of problems with the regression approach. Most importantly, these tests confuse catch-up growth with within-population convergence. To illustrate this, we can express the β coefficient as the ratio of covariance between the adult height and initial height, $cov(y,x)$, and the variance of the initial height, $var(x)$, where y refers to $h_{i,t}$ and x to $h_{i,t-1}$:

$$\beta = \frac{cov(y,x)}{var(x)}. \tag{5.2}$$

Abstracting from the case when the variance of childhood height approaches infinity, β approaches zero if the covariance between adult height and childhood height approaches zero. In other words, β approaches zero when childhood height is not a good predictor of adult height. Such an outcome is possible if short and stunted children catch-up others – or if the growth of the taller children in the sample stagnates. As such, β is only picking up movements within the HAZ distribution over time but researchers cannot tell whether this is due to improvements in HAZ or not.

Second, as the β coefficient focuses on the changes *within* the distribution it fails to capture a widespread catch-up growth in the population. Such uniform movement of the distribution (as for example in Figure 5.3 presented later) only affects the estimated constant (α), not the β coefficient. This is a particular concern if the population under scrutiny is largely malnourished.

Another important caveat is introduced by regression-to-the-mean (even in the absence of measurement error). Finding that β<1 may imply convergence (reduction in the dispersion of the height distribution over time) or simply reflect a natural variation in growth rates (height distribution remains intact). Quah (1993) demonstrates this algebraically using the Cauchy-Schwarz inequality:

(5.3)
$$cov(y,x) \leq \sqrt{var(x)}\sqrt{var(y)}.$$

Now, if there is no convergence or dispersion it follows that:

(5.4)
$$var(y) = var(x).$$

Using this and the Cauchy-Schwarz inequality in Equation (5.2) gives:

(5.5)
$$\beta \leq \frac{\sqrt{var(x)}\sqrt{var(x)}}{var(x)} = 1.$$

As such, even when imposing constant variance over time, β is smaller than one (the equality in the Cauchy-Schwarz holds only in the degenerate case when x and y are linearly dependent). This is why β<1 does not necessarily imply convergence – let alone catch-up growth. It is of no surprise then that most of these studies document partial convergence (i.e. 0<β<1). The β estimates vary between 0.2 (Fedorov and Sahn 2005) to 0.7 (Martorell et al. 1992).

I further demonstrate these concerns using the 1970 British Cohort Study (Elliott and Shepherd 2006). The cohort represents a healthy and well-nourished population. By default, it does not contain malnourished children and therefore I should not find any catch-up growth. Table F.1 of Appendix G contains the summary statistics for the final sample of 9,635 individuals. Figure 5.1 provides the distribution of the children's HAZ scores in early childhood (solid line) and adulthood (dashed line). As expected, the means lie close to zero in both periods and the distributions are virtually on top of each other.

Figure 5.1: The evolution of the height-for-age z-scores (HAZ) distribution over time in British cohort study (BCS70)

I then estimate Equation (5.1) using height-for-age z-scores (see Appendix H on why using raw measures of height is not appropriate here). Table 5.1 presents the regression results. The β coefficient is estimated as 0.601. The 99 per cent confidence interval for this point estimate is [0.578; 0.623]. This finding seriously questions the validity and the interpretation of findings in the previous literature using the regression approach to study catch-up growth. First, it is inconceivable to find catch-up growth in a healthy and well-nourished population. The descriptive and graphical analyses also show that there is negligible movement in the distribution over time. Yet, the current literature would interpret this regression result as partial catch-up growth. Second, a more sensible interpretation of the result is convergence in the height distribution (i.e. the dispersion diminished over time). This is also not supported by the graphical analysis.

Table 5.1: Estimating convergence in the British cohort (BCS70)

(Dependent variable: HAZ in adulthood)

	BCS70
HAZ in early childhood	0.601***
	(0.009)
intercept	0.164***
	(0.008)
Number of observations	9,635
R^2	0.383

note: *** p<0.01. White (1980) adjusted standard errors are in parenthesis.

Could this be explained by measurement error? It is true that measurement error in initial height ($h_{i,t-1}$) causes a downward bias, potentially leading to a false inference on convergence. Measurement error in $h_{i,t}$ is less harmful as it only leads to inflated standard errors. These statements hold if we make the plausible assumption that measurement error is random (i.e. not correlated with children's characteristics) with zero mean. In Appendix I, I show that the magnitude of the potential measurement error is too small to drive the results in Table 5.1.

To further analyse what is going on, in Figure 5.2, I plot HAZ in adulthood on HAZ in early childhood. If the nutritional status in early childhood perfectly predicted adult outcomes (i.e. children remain locked into their growth trajectories), the regression line (dashed line in the figure) would lie on the 45-degree line (solid line). Instead, the regression line is flatter. From the figure it becomes clear that this is an outcome of regression-to-the-mean: those initially unusually short (the dots lying on the middle-left in the figure) and tall (the dots lying on the middle-right) reverted towards the mean by adulthood.

Figure 5.2: HAZ in adulthood on HAZ in early childhood, the British cohort (BCS70)

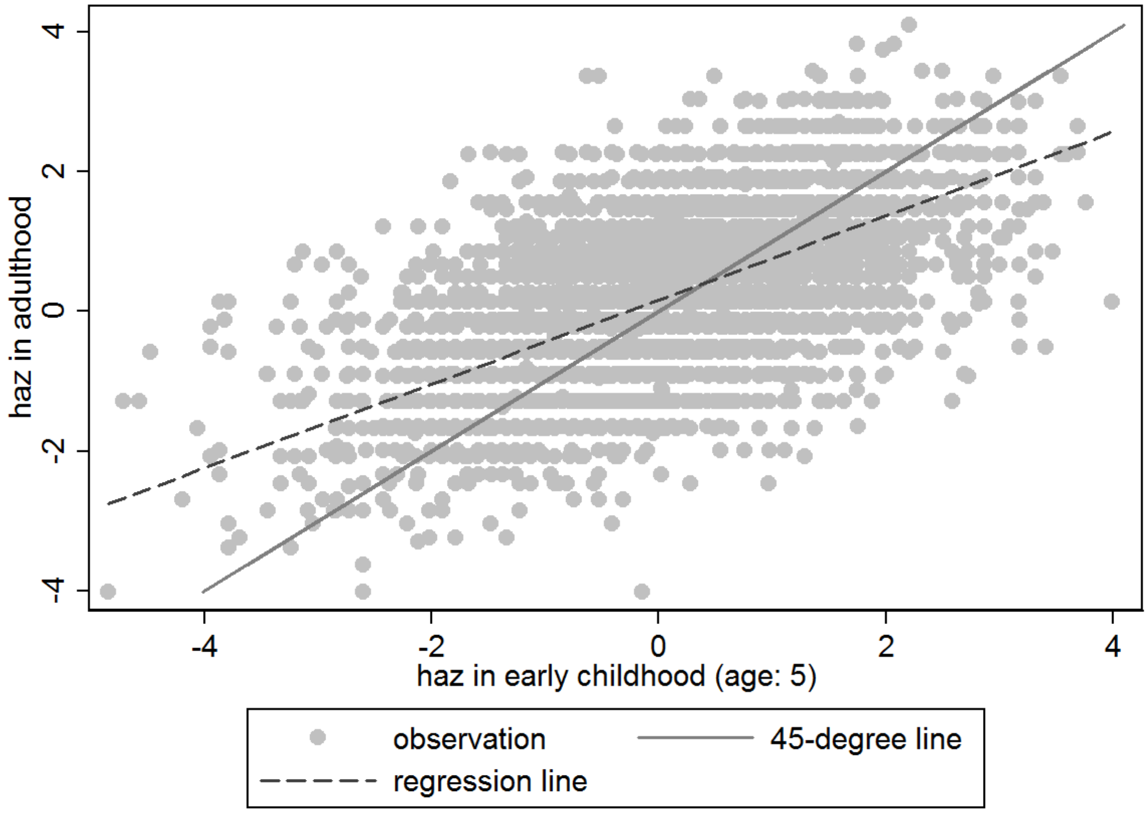

The economics literature generally augments Equation (5.1) with control variables, such as child, household and community characteristics. This shifts the focus from *unconditional* to *conditional* convergence but does not address any of the criticisms raised in this essay. Furthermore, the same literature is concerned with establishing the causal effect of initial height on adult height. In the presence of time-invariant unobserved characteristics, such as a child's innate healthiness or parental preferences, the lagged height variable ($h_{i,t-1}$) is, by construction, correlated with the error term (e_{it}). As a result, β captures both the relationship of adult height and early childhood height (i.e. convergence) as well as the relationship between adult height and the omitted or unobserved variables. The usual strategy is to use instrumental variable (IV) techniques to eliminate the role of unobserved omitted characteristics (Hoddinott and Kinsey 2001, Fedorov and Sahn 2005, Alderman, Hoddinott, and Kinsey 2006, Mani 2012, Outes and

Porter 2013). The approach starts with a quest for variables that are correlated with the initial height but not with adult height (other than through initial height). In the first stage, these *instrumental* variables are used to predict the initial height. In the second stage, the predictions are plugged into the original equation in order to purge the correlation between the error term and the initial height. The IV-approach is not, however, a panacea. These IV-estimates continue to measure within-sample convergence. Moreover, the approach does not solve the regression-to-the-mean issue (see Quah 1993). Finally, researchers often remove the time-invariant unobserved variables from Equation (5.1) by taking first differences and then use an instrument to purge the remaining correlation with the lagged height variable and the error term (e.g. Mani 2012, Fedorov and Sahn 2005). This is unfortunate as it is precisely the impact of these variables, rather than the β coefficient, that would provide useful information to the health practitioners and policy makers (see Durlauf and Quah 1998 for a similar point in the context of economic growth analysis).

5.3 A review of studies measuring catch-up growth as the change in HAZ

In an influential study, Adair (1999) defines catch-up growth as a recovery from stunting. Using the Cebu Longitudinal Health and Nutrition Survey from the Philippines, she finds that the proportion of the stunted children (HAZ<-2) fell from 63 per cent at the age of two years to 50 per cent by the age of 12. Using the NCHS/WHO 1977 growth reference to convert heights into HAZ scores, she finds that the mean HAZ scores improved from -2.41 at the age of two to -1.94 by the age of 12. As she does not observe the adult heights, these catch-up growth estimates provide only tentative evidence of the final catch-up growth in the cohort. Later analyses by Eckhardt et al (2005a,b) show that growth faltering took place during puberty, especially among females.

A small number of studies group children according to their degree of stunting in early childhood and compare the height increments from early childhood to adulthood between groups and to an external healthy and well-nourished reference population. Martorell, Rivera, and Kaplowitz (1990) use data from Guatemala where more than half of the children in the sample were stunted at the age of five. These children did not experience larger height gains than their American peers if anything, the opposite was true. In addition, the authors did not find any marked differences in height increments between shorter and taller Guatemalan children. Satyanarayana, Nadamuni Naidu, and Narasinga Rao (1980) and Satyanarayana et al. (1981) study growth increments using Indian cohorts of boys and girls respectively. They group children according to their severity of stunting and compare the height increments between the age of five and 17 across the groups and to a longitudinal study from Boston. These Indian boys were not found to experience any catch-up growth and also the height gains between the groups were similar (Satyanarayana, Nadamuni Naidu, and Narasinga Rao 1980). The initially most nutritionally deprived girls (HAZ < -4) were found growing considerably faster than the Boston girls and other Indian girls in the same sample (Satyanarayana et al. 1981). This result may however be an outcome of measurement error in initial height leading to an incorrect grouping of the children.

Among the very few longitudinal studies from sub-Saharan Africa, Coly et al. (2006) follow more than 2,800 Senegalese children for nearly two decades and compare their growth rates to the NCHS/WHO 1977 reference. The mean HAZ in early childhood among girls -1.3 and among boys was -1.4. By adulthood these means reduced to -0.41 for girls and -0.58 for boys implying nearly complete catch-up. This widely neglected study also makes a clear distinction between catch-up growth relative to a healthy reference population (global catch-up growth) and relative to other children in the population (local catch-up growth). More recently, Prentice et al. (2013) provided similar evidence from Gambia. Exploiting longitudinal data on 160 children whose heights were measured multiple times between the ages of 8 and 24 and measuring heights against the UK-1990 reference population, they also document nearly complete catch-up growth. The mean HAZ among boys improved from -1.25 to -0.5 and from -1.1 to -0.2 among girls.

5.4 Data

I use KHDS for the empirical part of the essay. As discussed in Chapter 1, the sample was drawn from a random sample stratified by adult mortality rates in the communities and the agro-climatic zones in the region. Statistical tests (not reported but available from the author) did not reveal statistical differences in HAZ scores between children residing in high or low adult mortality communities.

Anthropometric measurements were taken from all respondents who were present at the time of the household interview. In 1991-94, trained anthropometrists were responsible for measuring all household members. In 2004 and 2010, enumerators, carefully trained by a qualified nurse, took the measurements usually after the household questionnaire was administered in the household. In all survey rounds, heights of children less than three years old were measured using a length board with a sliding foot piece. The heights of adults and children older than three years were measured using a height board with a sliding head piece. All heights were recorded to the nearest millimetre.

Height-for-age Z-scores were calculated using the *zanthro* command in Stata 11.2 (see Vidmar et al. 2004). I use the US 2000 NCHS/CDC as the reference population (see Kuczmarski et al. 2002). The 2006 WHO Child Growth Standards (see WHO 2006) in conjunction with the WHO Reference 2007 for 5-19 years (see de Onis, Onyango, et al. 2007) resulted in more catch-up growth than NCHS/CDC. This is driven by the differences between the two references. In early childhood (ages 0-5), the median child in the WHO Child Growth Standards is taller than in the NCHS/CDC reference population (de Onis, Garza, et al. 2007). In adulthood (at the age of 19), however, the height difference is negligible when comparing the WHO Reference 2007 for 5-19 years and NCHS/CDC. This highlights the difficulty in comparing studies that employ different growth standards or reference populations. I prefer the NCHS/CDC growth reference as it allows the calculation of HAZ scores from birth to 20 years of age, and is constructed using the same reference throughout the entire growth period.

The sample for this study is constructed from children who were between 12 and 59 months old at the time of the four waves of the baseline survey in 1991-94 and who are at least 18 years old in 2010. This cohort of 884 children is followed from the 1991-94 rounds through the 2010 round. In 2010, 559 of these children were interviewed and measured, 69 had died and 256 were not found or their heights were not measured. I drop all children whose height was not measured in 2010 or whose date of birth is not known. After dropping the few children with implausible height measurements (HAZ< -5 or ΔHAZ> ±4) the final sample contains 540 children (269 girls and 271 boys) from 365 households. If the child was measured more than once when she was 12 to 59 months old during the four interview rounds at the baseline, I took the last measurement. An alternative strategy would have been to take the mean over these observations to address the potential measurement error in these data. This yields nearly identical mean HAZ score suggesting that the potential measurement error either has a zero mean or is absent in these data. Finally, I use the difference between the date of the interview and the child's date-of-birth to calculate the ages.

The sample attrition poses a concern to the analysis. Attrition due to death is less of a problem as catch-up growth analysis is focused on the height developments of the surviving. Attrition due to other causes is more problematic. If such attrition is positively correlated with health, then studies are likely to under-report catch-up growth in the sample. To address this, in Table 5.2, I compare children's height-for-age scores by attrition status. Children who were not represented in the final sample have slightly higher HAZ scores than those who did. A two-tailed t-test shows, however, that this difference is not statistically significant. Further examination, presented in Table J.1 of Appendix J, reveals that children who deceased after the first round had lower HAZ-scores than those who survived and comprise the final sample. Children who were not traced or present at the time of the measurement have slightly better HAZ scores than children in the final sample. However, according to a two-tailed t-test, neither of these observed differences is statistically significant. Sample attrition does not seem to be associated with higher or lower HAZ scores.

Table 5.2: Attrition in KHDS: initial HAZ-scores by sample category

	KHDS	
final sample	no	yes
observations	344	540
HAZ:		
mean	-1.804	-1.864
std. dev.	1.262	1.155
difference	0.06	
t-test	0.71	

note: t-test based on Welch t-test on the difference in means between the two groups

Finally, Table J.2 of Appendix J, provides an overview of the HAZ-scores in KHDS by migration status for each child in the sample. By 2010, I find that half of the sample had migrated. Had we not tracked individuals, I would have lost half of the sample. Surprisingly, however, migration does not seem to be correlated with nutrition status.

According to a two-tailed t-test, the difference in the adult HAZ-scores between children who remained in the baseline village and those who migrated by 2010 is not statistically different from zero.

5.5 Catch-up growth in the Kagera cohort

Next, I analyse the extent of catch-up growth in the Kagera cohort. As discussed above, I measure catch-up growth as the change in HAZ and focus on analysing the data using descriptive and graphical methods. Table 5.3 shows the summary statistics for the cohort. The mean HAZ scores in 1991-94, when the children are less than five years old, is 1.86 standard deviations below the median of the US-reference group. Approximately 44 per cent of the children are stunted (HAZ < -2) and 16 per cent are severely stunted (HAZ < -3). These percentages agree with the statistics reported in the 1991/92 Tanzanian Demographic and Health Survey for the same region: 44 per cent of the children under 5 years old in Kagera were found stunted and 19.5 per cent severely stunted (Ngallaba et al. 1993).

Table 5.3: Evolution of HAZ scores in the Kagera cohort

	mean	std. dev.
age in months (t=0)	41.08	14.686
age in years (t=1)	20.25	1.591
height (t=0)	89.64	10.269
height (t=1)	161.7	8.267
HAZ_1 (t=0)	-1.86	1.155
HAZ_2 (t=1)	-1.20	1.002
difference in HAZ:	0.66	1.07
t-test: $HAZ_{t=0} = HAZ_{t=1}$	10.01	
percentage: $HAZ_{t=0} < -2$	44%	
percentage: $HAZ_{t=1} < -2$	20%	
percentage: $HAZ_{t=0} < -3$	16%	
percentage: $HAZ_{t=1} < -3$	3%	
Observations	540	

note: t=0 refers to early childhood, t=1 to adulthood. HAZ scores calculated using the US 2000 NCHS/CDC as the reference population.

Interestingly, in 2010, the cohort has been able to catch-up with the reference group: the mean height-for-age z-score in the sample is now -1.20. In 2010, 20 per cent are stunted and only three per cent severely stunted. There is also a gender difference. The mean HAZ-score in the female sample increases from -1.71 to -0.98 whereas for males the catch-up growth is somewhat more modest: mean HAZ increases from -2.01 to -1.42. These statistics show that the Kagera children are able to catch-up the growth losses incurred in early childhood. Figure 5.3 shows the full distributions of the HAZ scores in both periods. The figure reinforces the summary statistics: there is considerable catch-up growth in the sample. This can be seen as the adult height distribution shifting right relative to the early childhood distribution.

Figure 5.3: The evolution of the HAZ distribution over time in Kagera cohort

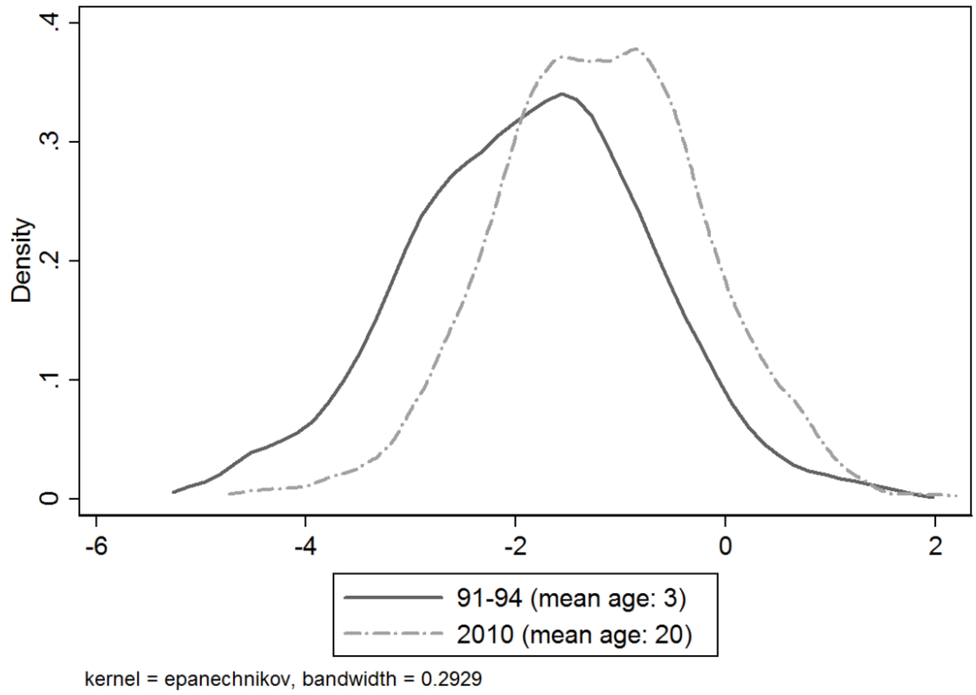

Figure 5.4 offers another cut of the data. Here I plot the change in HAZ on HAZ in early childhood. The horizontal line goes through zero. The dots above this line are children whose HAZ score improved between the two periods. The dots below this line belong to children whose HAZ score worsened. The figure also contains a vertical line that goes through zero. The dots on the left of this line belong to children who had a negative HAZ score in early childhood. The top left corner (marked with A) contains the children who experienced catch-up growth: after initial growth retardation ($HAZ_{t-1} < 0$), they experienced growth that was above what was expected for their age ($\Delta HAZ > 0$). This corner corresponds directly to the definition of catch-up growth used in the clinical and epidemiological literature.

Figure 5.4 shows that 95 per cent of Kagera children had negative HAZ scores when they were first measured. More than 74 per cent of the sample experienced catch-up growth. The mean improvement in HAZ scores among these children is 1.1 units of standard deviation (median = 0.97). Nearly 21 per cent of the children had negative HAZ-scores in early childhood and fell further behind later in life. Only five per cent of the children had initially positive HAZ-scores. Most of these children experienced slower growth than was expected.

Figure 5.4: Change in HAZ between early childhood and adulthood on HAZ in early childhood the Kagera cohort (KHDS)

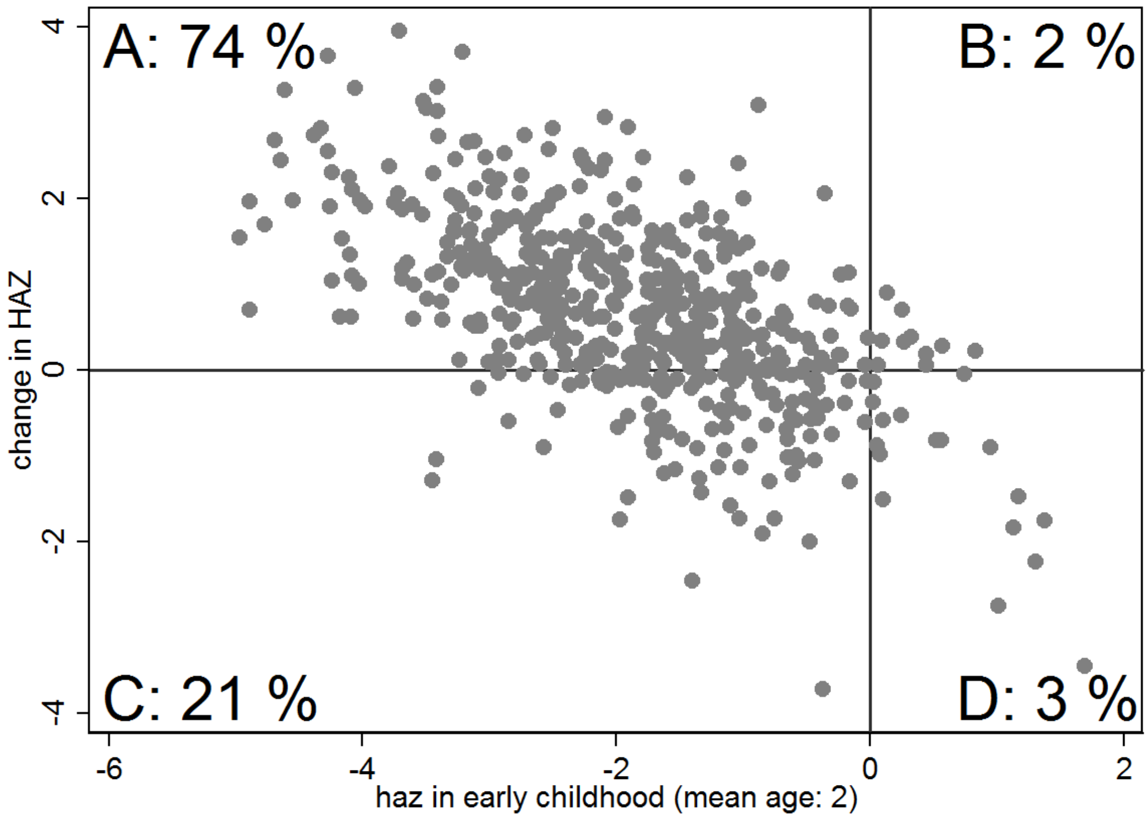

Note:

A: under-nourished and growing B: well-nourished and growing
C: under-nourished and faltering D: well-nourished and faltering

5.6 Timing of the catch-up growth in the Kagera cohort

The analysis in the previous section finds considerable catch-up growth in the KHDS cohort. Coly et al. (2006) find that the near complete catch-up growth documented in the Senegalese cohort takes place in puberty. In a recent study combining longitudinal and cross-sectional data from rural Gambia, Prentice et al. (2013) show how children experience substantial catch-up growth during puberty.

To analyse the timing of the catch-up growth in the Kagera cohort, I would need to compare growth rates at different points in childhood. Unfortunately, assessing children's annual growth patterns is not feasible using the longitudinal data as I only have two or three data points for each child. Fortunately, I can use the cross-sectional data to mimic children's growth patterns (see Moradi, 2010, for a similar exercise using data for Cote d'Ivoire and Ghana). Using the baseline data in 1991-94, I constructed mean HAZ-scores at each age until the age of 23. To account for concerns that the observed catch-up growth in these cross-sectional data is an artefact of selective mortality (see Bozzoli, Deaton, and Quintana-Domeque 2009, Rouanet 2011), I drop all children who did not survive to their 18th birthday. This mortality information originates from the 2010 survey when the panel respondents are adults. Figure 5.5 shows the growth patterns for both gender groups (solid lines). Similar to the evidence presented in Shrimpton et al. (2001), growth retardation begins immediately after birth and continues until 2-3 years of life. After four years of age, the HAZ-scores remain relatively stable until the age of 10 to 11. At this age, the median child in the US reference group enters the adolescent growth spurt. The HAZ-scores fall rapidly at this point suggesting that puberty is delayed for the Kagera children. For boys, the HAZ-scores continue to fall until the age of 15 at which point the HAZ-curve shoots up. The growth ceases at the age of 19 for girls and at the age of 22 for boys. By now boys have caught-up the height losses incurred during puberty but HAZ-scores have also improved further to nearly restore their early childhood levels. Girls begin their adolescent growth spurt earlier and are able to completely restore their initial growth curves.

The shortcoming of using cross-sectional data to construct growth curves is that I cannot be sure whether the observed growth patterns are driven by age or by birth cohort effects. That is, whether the observed differences between the age cohorts are arising because they are observed at different ages (age effect) or because they were born into different environmental conditions (birth cohort effect). To circumvent this, I exploit the longitudinal feature of the data and plot the growth curves for the same children whose pre-adult HAZ-scores I have in the later rounds of the survey (i.e. 2004 and 2010 rounds). These are the two dashed lines in Figure 5.5 starting at the age of 11. These curves lie very close to the original growth curves lending support to the age effects story. The three longitudinal data observations further confirm these patterns. At the average age of three, the mean HAZ score for girls is -1.71 and for boys is -2.02. In puberty (average age of 14), girls' mean HAZ score has improved to -1.25 and boys' to -1.89. In adulthood, the girls' mean HAZ score stands at -0.98 and boys' at -1.42.

This evidence shows that puberty is an important opportunity window for the Kagera children to catch-up their healthy and well-nourished peers. Moreover, this growth pattern does not seem to be cohort specific.

The finding is in tune with recently emerging evidence from studies employing cross-sectional data sets suggesting that puberty is a sensitive period for children's growth. Interestingly, most of this evidence is from studies analysing populations originating from Sub-Saharan Africa. Moradi (2010), using multivariate regression analysis and a sample of more than 200,000 women from Sub-Saharan Africa, finds that economic growth has a positive effect on adult height during early childhood but also during puberty. Akresh et al. (2012) provide indirect evidence on this by comparing female adult heights between cohorts exposed to the Nigerian civil war in 1967-70. Strikingly, they find that the most negatively affected children were the ones aged 13 to 16 during the war. Economic historians have documented intense pubertal catch-up growth among 19th century African Americans. In a famous study, Steckel (1987) documents a remarkable catch-up growth of African American slaves during puberty. Using manifests of 50,000 slaves transported in 1820-1860, he shows that children remain highly malnourished until puberty during which the height deficit to an external reference population halved. This catch-up is credited to improved diets when the slaves entered the labour force between eight and 12 years (Steckel, 1987; 2000) . Komlos (1992) documents similar catch-up growth among free African Americans in Antebellum Maryland.

Figure 5.5: HAZ scores by age (cross-sectional data) in Kagera (KHDS)

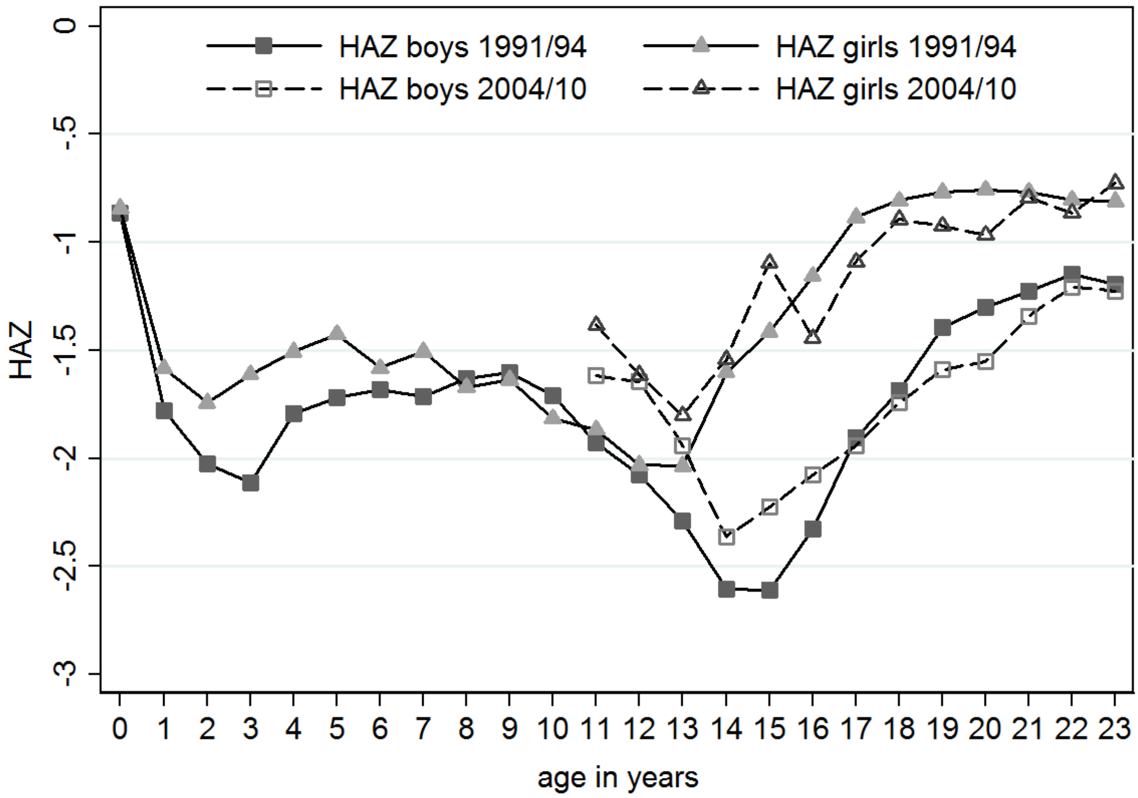

Note: The dashed curves are less smooth because of fewer observations.

Among studies based on non-African populations, Stein (1975) shows that children exposed to the Dutch Winter Famine in 1944-1945 were able to completely catch-up other, non-exposed, cohorts. Finally, Haas and Campirano (2006), using cross-country data, plot the pubertal growth rate on height just before the onset of puberty. They find that children from populations that have lower pre-pubertal heights seem to experience greatest growth during puberty.

5.7 Concluding discussion

The empirical analysis of catch-up growth requires longitudinal data or cohort studies that span the entire growth period. Catch-up growth is defined as growth in height that is above the expected for the child's age and occurs after a period of growth retardation. Height-for-age z-score measures the height distance in standard deviations to a healthy and well-nourished reference population. The evolution of HAZ scores over time therefore provides the exact counterpart of the definition used in the clinical and epidemiological literature.

Part of the existing evidence is plagued by studies that confuse catch-up growth with convergence. By regressing HAZ in adulthood on childhood HAZ researchers cannot tell whether children are catching-up or faltering. In addition, these tests focus on movement *within* the cohort and are therefore not able to detect widespread catch-up growth in populations. The interpretations of these regression estimates are further hampered by regression-to-the-mean.

Aside from the methodological contribution, I also provide new evidence for sub-Saharan Africa. Using longitudinal data for the Kagera region in Tanzania, I document considerable catch-up growth. Nearly 75 per cent of the children experienced catch-up growth: after a period of growth retardation (HAZ < 0) in early childhood they experienced growth that was above the expected (ΔHAZ > 0). The mean HAZ-score in the cohort improves from -1.86 in early childhood to -1.20 by adulthood. The difference between the predicted adult height in early childhood and the actual adult height is around 4.5 centimetres for boys and 5 centimetres for girls. The graphical analysis shows that most of this observed catch-up growth takes place in puberty.

The existing literature emphasises the importance of early childhood conditions on adult outcomes. This essay shows that not everything is set in stone after early-life: short and stunted children do not necessarily end up short and stunted adults. The observed pubertal catch-up growth may also have implications for economic outcomes in adulthood. Several studies have documented strong correlation between adult height and earnings in both developed and developing countries (for recent reviews, see Currie and Vogl 2013, Steckel 2009). Studies from developing countries find that taller people are more productive, and therefore earn more, in tasks that require physical strength (Haddad and Bouis 1991, Schultz 2002, Dinda et al. 2006, Thomas and Strauss 1997). In rich developed countries, the existence of the height premium is more puzzling as occupations increasingly require more brain than brawn. Studies using data for the US and the UK have attributed the height premium to the correlation between adult height and non-cognitive skills (Persico, Postlewaite, and Silverman 2004) or cognitive development in early childhood (Case and Paxson 2008), or both (Schick and Steckel 2010). Recently, Lundborg, Nystedt, and Rooth (2013) employing rich data on Swedish conscripts, find that taller individual earn more because adult height is positively correlated with having a better family background and better cognitive and non-cognitive skills. Somewhat surprisingly, also physical capacity, measured as muscular strength, plays a role in explaining the height premium in Sweden. Due to data limitations, these types of studies are rare in developing country contexts. Therefore, whether the origins of the height premium in poor countries lie in the early-life conditions (through early childhood nutrition affecting cognitive development) or in adult outcomes (e.g. physical strength) remains an unresolved question (Vogl 2012, Currie and Vogl 2013).

6 Conclusions

This thesis has employed truly exceptional panel data to examine two central themes within development economics. The length of the panel offers unprecedented opportunities to study long-term development processes in a Sub-Saharan African context. Careful tracking of those who moved makes the data set an ideal for analysing various aspects of internal migration. It also minimizes (potentially non-random) attrition when studying long-run health and poverty dynamics at the individual or household level.

The three essays presented in this thesis provide important contributions within each of their respective topics. However, any claim that this thesis has created "[…] *new stylised facts for the next generation of development research*" (Dercon, Krishnan, and Krutikova 2013p. 16), would certainly be far too extravagant. The purpose of this concluding chapter is to discuss the relative strengths and weaknesses of each essay, and suggest avenues for future research.

The main objective of the first essay is to document what happens to a traditional institution, like informal insurance, in a society undergoing a modernization process accompanied by massive internal migration. The economics literature often treats institutions as fixed in time. In this essay we document how a traditional institution adapts to and interacts with the transformation process. The strength of the essay lies in its ability to describe this process within linked extended family networks, over a long two decade time period and with 'real world' data. The drawback of not having experimental data is that we cannot provide iron-clad proof of causality. Although we rule out certain sources of endogeneity, we cannot make the claim that our findings are generalizable to other settings where the selection into migration is different. Finally, one of the initial motivators for this study was to understand whether the social norms impose a significant cost to migration. In the kinship poverty trap model (see Hoff and Sen 2006) costly actions imposed by the home-community on the migrants could lead some migrants to sever links with home. Knowing that some of its productive members may leave and never look back, the kinship network can set up endogenous exit barriers (e.g. through values and norms about migration), preventing some of its most productive members to leave in the first place. Such exit barriers would then provide an explanation to the puzzle why people do not migrate despite the high returns (causally) attached to it (Beegle, De Weerdt, and Dercon 2011). Our findings do not support this notion. First, we find that 90 per cent of the migrants maintain links to their home community. Second, the estimated 'tax rate' of three to five per cent seems too low to be a significant brake on a migrant's growth. It appears then that the missing piece to this migration puzzle lies elsewhere. The search therefore continues and constitutes an exciting path in my future research.

In the second essay, I examine how local weather shocks shape short-run consumption outcomes and long-run migration patterns in Kagera. This essay makes a number of important contributions to the literature. First, it is the first to quantify the welfare impacts of temperature shocks using micro-level data from Sub-Saharan Africa. Second, I study the impact of these temperature shocks on long-term migration decisions over a 15-year period. Finally, the simple theoretical model provided in this essay helps to interpret the findings that go against the more traditional view of migration (e.g. Harris and Todaro 1970). The limitations of this study relate to external validity. By focusing on a fairly small geographical area, I cannot be certain whether these results generalise to other regions of Tanzania, let alone other parts of Africa. Although the data requirements are high, the way forward is to undertake similar research in other parts of Africa. Finally, the structural transformation process (e.g. Lewis 1954) appears to be underway in emerging Africa. In order to understand the role of weather and climate over the full-cycle of this process, it would be interesting to exploit historical data on migration patterns and temperatures, for example for 19^{th}-20^{th} century England, or for another modern-day rich country. This constitutes another path in my future research.

The third essay makes an important methodological and empirical contribution in the field of Economics and Human Biology. First, it demonstrates that the recent literature measures catch-up growth incorrectly by confusing catch-up growth with within-population convergence. In addition, the empirical part of the essay concludes that not everything is set in stone after early childhood: short and stunted children do not necessarily become short and stunted adults. The strength of the empirical analysis lies in the length of the survey permitting me to observe the children's entire growth process. Previous studies have largely employed panel data sets of shorter length and, therefore, have not been able to analyse the growth process over the pubertal period. Finally, this study opens up a number of exciting avenues for future research. Understanding the causes for this pubertal catch-up growth is the first step. Moradi (2010) puts forward a number of interesting propositions. First, the disease environment is less harsh in puberty than in early childhood. As such, given that children grow rapidly also in puberty, the elasticity with which calories translate into growth in height is likely to be larger. Second, adolescents contribute to the household income and this may improve their intra-household bargaining power leading to increased food shares within the household. Investigating these channels forms an interesting path for future research. The second step is to understand how this catch-up growth shapes economic outcomes in adulthood. As discussed in Chapter 5, adult height is associated with higher earnings in a wide variety of contexts. In contrast to the rich developed world, a large part of the African working population is still engaged in agricultural production. As such, the returns to brawn may still be considerable. If height is then correlated with strength, then this pubertal catch-up growth is likely to improve living standards in adulthood. Attempting to shed light on this issue comprises part of my future research agenda.

Bibliography

Adair, Linda S. 1999. "Filipino children exhibit catch-up growth from age 2 to 12 years." *The Journal of Nutrition* no. 129 (6):1140-8.

Ahmed, S. A., N. S. Diffenbaugh, T. W. Hertel, D. B. Lobell, N. Ramankutty, A. R. Rios, and P. Rowhani. 2011. "Climate volatility and poverty vulnerability in Tanzania." *Global Environmental Change-Human and Policy Dimensions* no. 21 (1):46-55. doi: DOI 10.1016/j.gloenvcha.2010.10.003.

Akresh, Richard, Sonia Bhalotra, Marinella Leone, and Una Okonkwo Osili. 2012. "War and Stature: Growing Up during the Nigerian Civil War." *American Economic Review: Papers & Proceedings* no. 102 (3):273-77.

Alderman, Harold, Jere Behrman, Hans-Peter Kohler, John Maluccio, A., and Susan Watkins. 2001. "Attrition in Longitudinal Household Survey Data." *Demographic Research* no. 5 (4):79-124.

Alderman, Harold, J. Hoddinott, and B. Kinsey. 2006. "Long term consequences of early childhood malnutrition." *Oxford Economic Papers-New Series* no. 58 (3):450-474.

Allison, P.D. 1982. "Discrete-time methods for the analysis of event histories." *Sociological methodology* no. 13 (1):61-98.

Almond, Douglas, and Janet Currie. 2011. "Human capital development before age five." In *Handbook of Labor Economics*, edited by David Card and Orley Ashenfelter, 1315-1486. Amsterdam, The Netherlands: Elsevier.

Altonji, Joseph G., Fumio Hayashi, and Laurence J. Kotlikoff. 1992. "Is the Extended Family Altruistically Linked - Direct Tests Using Micro Data." *American Economic Review* no. 82 (5):1177-1198.

Ambler, Catherine. 2012. Don’t Tell on Me: Experimental Evidence of Asymmetric Information in Transnational Households. Mimeo, University of Michigan.

Angelucci, Manuela. 2013. "Migration and Financial Constraints: Evidence from Mexico." *The Review of Economics and Statistics* (forthcoming).

Ashworth, A., and D. J. Millward. 1986. "Catch-up Growth in Children." *Nutrition Reviews* no. 44 (5):157-163.

Atkinson, Anthony B, and Maria Ana Lugo. 2010. Growth, poverty and distribution in Tanzania. *International Growth Centre Working Paper Series* 10/0831.

Attanasio, Orazio P., and José-Víctor Ríos-Rull. 2000. "Consumption smoothing in island economies: Can public insurance reduce welfare?" *European Economic Review* no. 44 (7):1225-1258.

Auffhammer, Maximilian, Solomon M Hsiang, Wolfram Schlenker, and Adam Sobel. 2013. "Using Weather Data and Climate Model Output in Economic Analyses of Climate Change." *Review of Environmental Economics and Policy* no. 7 (2):181-198.

Baland, Jean-Marie, Cathrine Guirkinger, and Charlotte Mali. 2011. "Pretending to Be Poor: Borrowing to Escape Forced Solidarity in Cameroon." *Economic Development and Cultural Change* no. 60 (1):1-16.

Barrios, S., L. Bertinelli, and E. Strobl. 2006. "Climatic change and rural-urban migration: The case of sub-Saharan Africa." *Journal of Urban Economics* no. 60 (3):357-371. doi: DOI 10.1016/j.jue.2006.04.005.

Bazzi, Samuel. 2013. Wealth heterogeneity, income shocks, and international migration: theory and evidence from Indonesia. Mimeo, University of California, San Diego.

Beegle, Kathleen, Joachim De Weerdt, and Stefan Dercon. 2006. Kagera Health and Development Survey, 2004: Basic Information Document. Mimeo, The World Bank.

Beegle, Kathleen, Joachim De Weerdt, and Stefan Dercon. 2008. "Adult mortality and consumption growth in the age of HIV/AIDS." *Economic Development and Cultural Change* no. 56 (2):299-326.

Beegle, Kathleen, Joachim De Weerdt, and Stefan Dercon. 2011. "Migration and Economic Mobility in Tanzania: Evidence from a Tracking Survey." *Review of Economics and Statistics* no. 93 (3):1010-1033.

Beegle, Kathleen, Joachim De Weerdt, Jed Friedman, and John Gibson. 2012. "Methods of household consumption measurement through surveys: Experimental results from Tanzania." *Journal of Development Economics* no. 98 (1):3-18.

Beegle, Kathleen, and Michelle Poulin. 2013. "Migration and the transition to adulthood in contemporary Malawi." *The ANNALS of the American Academy of Political and Social Science* no. 648 (1):38-51.

Bengtsson, Niklas. 2010. "How responsive is body weight to transitory income changes? Evidence from rural Tanzania." *Journal of Development Economics* no. 92 (1):53-61.

Blundell, Richard, and Thomas M Stoker. 2005. "Heterogeneity and aggregation." *Journal of Economic Literature* no. 43 (2):347-391.

Boersma, Bart, and Jan Maarten Wit. 1997. "Catch-up growth." *Endocrine Reviews* no. 18 (5):646-61.

Bozzoli, Carlos, Angus Deaton, and Climent Quintana-Domeque. 2009. "Adult height and childhood disease." *Demography* no. 46 (4):647-69.

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mobarak. 2013. Seasonal migration and risk aversion. *mimeo, London School of Economics*.

Burke, Marshall, John Dykema, David Lobell, Edward Miguel, and Shanker Satyanath. 2011. Incorporating climate uncertainty into estimates of climate change impacts, with applications to US and African agriculture. *National Bureau of Economic Research (NBER) Working Paper* No. 17092.

Burke, Mary A., and H. Peyton Young. 2011. "Chapter 8 - Social Norms." In *Handbook of Social Economics*, edited by Jess Benhabib, Alberto Bisin and Matthew O. Jackson, 311-338. San Diego, CA: North-Holland.

Caeyers, Bet, Neil Chalmers, and Joachim De Weerdt. 2012. "Improving consumption measurement and other survey data through CAPI: Evidence from a randomized experiment." *Journal of Development Economics* no. 98 (1):19-33.

Cameron, N., M. A. Preece, and T. J. Cole. 2005. "Catch-up growth or regression to the mean? Recovery from stunting revisited." *American Journal of Human Biology* no. 17 (4):412-7.

Case, A., and C. Paxson. 2008. "Stature and status: Height, ability, and labor market outcomes." *Journal of Political Economy* no. 116 (3):499-532.

Chandrasekhar, Arun G, Cynthia Kinnan, and Horacio Larreguy. 2011. Information, networks and informal insurance: evidence from a lab experiment in the field. Mimeo, Massachusetts Institute of Technology and Northwestern University.

Chandrasekhar, Arun G, Cynthia Kinnan, and Horacio Larreguy. 2012. Informal Insurance, Social Networks, and Savings Access: Evidence from a lab experiment in the field. Mimeo, Massachusetts Institute of Technology and Northwestern University.

Chenery, Hollis B, and Moises Syrquin. 1975. *Patterns of development, 1950-70*. New York, NY: Oxford University Press.

Chiappori, Pierre-André, Krislert Samphantharak, Sam Schulhofer-Wohl, and Robert M Townsend. 2011. "Heterogeneity and risk sharing in village economies." *NBER Working Paper* no. 16696.

Chow, Gregory C. 1960. "Tests of equality between sets of coefficients in two linear regressions." *Econometrica: Journal of the Econometric Society*:591-605.

Christiaensen, Luc, Joachim De Weerdt, and Yasuyuki Todo. 2013. "Urbanization and Poverty Reduction–The Role of Rural Diversification and Secondary Towns." *Agricultural Economics* no. 44 (4-5):435–447.

Coate, Stephen, and Martin Ravallion. 1993. "Reciprocity without commitment: Characterization and performance of informal insurance arrangements." *Journal of Development Economics* no. 40 (1):1-24.

Collier, Paul, and Stefan Dercon. 2009. African Agriculture in 50 Years: Smallholders in A Rapidly Changing World? Paper presented at the 2009 Food and Agriculture Organization Expert Meeting on ‘How to Feed the World in 2050. Available at ftp://ftp.fao.org/docrep/fao/012/ak983e/ak983e00.pdf.

Coly, A. N., J. Milet, A. Diallo, T. Ndiaye, E. Benefice, F. Simondon, S. Wade, and K. B. Simondon. 2006. "Preschool stunting, adolescent migration, catch-up growth, and adult height in young senegalese men and women of rural origin." *The Journal of Nutrition* no. 136 (9):2412-20.

Cox, Donald, and Marcel Fafchamps. 2007. "Extended family and kinship networks: economic insights and evolutionary directions." *Handbook of Development Economics* no. 4:3711-3784.

Cragg, John G., and Stephen G. Donald. 1993. "Testing identifiability and specification in instrumental variable models." *Econometric Theory* no. 9:222-222.

Currie, Janet, and Tom Vogl. 2013. "Early-Life Health and Adult Circumstance in Developing Countries." *Annual Review of Economics* no. 5:1-36.

de Onis, M., C. Garza, A. W. Onyango, and E. Borghi. 2007. "Comparison of the WHO child growth standards and the CDC 2000 growth charts." *The Journal of Nutrition* no. 137 (1):144-8.

de Onis, M., A. W. Onyango, E. Borghi, A. Siyam, C. Nishida, and J. Siekmann. 2007. "Development of a WHO growth reference for school-aged children and adolescents." *Bulletin of the World Health Organization* no. 85 (9):660-7.

De Weerdt, J. 2010. "Moving out of Poverty in Tanzania: Evidence from Kagera." *Journal of Development Studies* no. 46 (2):331-349. doi: 10.1080/00220380902974393.

De Weerdt, Joachim, Kathleen Beegle, Helene Bie Lilleør, Stefan Dercon, Kalle Hirvonen, Martina Kirchberger, and Sofya Krutikova. 2012. "Kagera Health and Development Survey 2010: Basic Information Document." *Rockwool Foundation Working Paper Series* no. 46.

De Weerdt, Joachim, and Marcel Fafchamps. 2011. "Social Identity and the Formation of Health Insurance Networks." *Journal of Development Studies* no. 47 (8):1152-1177. doi: Doi 10.1080/00220388.2010.527952.

Deaton, Angus. 1997. *The analysis of household surveys: a microeconometric approach to development policy*. Baltimore, MD: Published for the World Bank by Johns Hopkins University Press.

Deaton, Angus. 2007. "Height, health, and development." *Proceedings of the National Academy of Sciences of the United States of America* no. 104 (33):13232-7.

Deaton, Angus, and Margaret Grosh. 2000. "Consumption." In *Designing Household Survey Questionaires for Developing Countries: Lessons from 15 Years of Living Standards Measurement Study*, edited by Margaret Grosh and Paul Glewwe, 91-133. Washington, DC: The World Bank.

Dell, M., B. F. Jones, and B. A. Olken. 2012. "Temperature Shocks and Economic Growth: Evidence from the Last Half Century." *American Economic Journal-Macroeconomics* no. 4 (3):66-95.

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken. 2013. What Do We Learn from the Weather? The New Climate-Economy Literature. National Bureau of Economic Research.

Dercon, Stefan, Pramila Krishnan, and Sofya Krutikova. 2013. "Changing Living Standards in Southern Indian Villages 1975–2006: Revisiting the ICRISAT Village Level Studies." *The Journal of Development Studies* (forthcoming).

Dex, Shirley. 1995. "The reliability of recall data: a literature review." *Bulletin de Methodologie Sociologique* no. 49 (1):58-89.

Dillon, A., V. Mueller, and S. Salau. 2011. "Migratory Responses to Agricultural Risk in Northern Nigeria." *American Journal of Agricultural Economics* no. 93 (4):1048-1061.

Dinda, Soumyananda, PK Gangopadhyay, BP Chattopadhyay, HN Saiyed, M Pal, and P Bharati. 2006. "Height, weight and earnings among coalminers in India." *Economics & Human Biology* no. 4 (3):342-350.

Durlauf, Steven N., and Danny T. Quah. 1998. The New Empirics of Economic Growth. *National Bureau of Economic Research (NBER) working paper series* no. 6422.

Dustmann, Christian, and Josep Mestres. 2010. "Remittances and temporary migration." *Journal of Development Economics* no. 92 (1):62-70.

Eckhardt, C. L., P. Gordon-Larsen, and L. S. Adair. 2005. "Growth patterns of Filipino children indicate potential compensatory growth." *Annals of Human Biology* no. 32 (1):3-14.

Elliott, J., and P. Shepherd. 2006. "Cohort profile: 1970 British Birth Cohort (BCS70)." *International Journal of Epidemiology* no. 35 (4):836-43.

Eveleth, Phyllis B., and J. M. Tanner. 1976. *Worldwide variation in human growth, International biological programme*. Cambridge: Cambridge University Press.

Fafchamps, Marcel. 1992. "Solidarity networks in preindustrial societies: Rational peasants with a moral economy." *Economic Development and Cultural Change* no. 41 (1):147-174.

Fafchamps, Marcel. 1999. "Risk sharing and quasi-credit." *Journal of International Trade & Economic Development* no. 8 (3):257-278.

Falco, Paolo, Andrew Kerr, Neil Rankin, Justin Sandefur, and Francis Teal. 2011. "The returns to formality and informality in urban Africa." *Labour Economics* no. 18:S23-S31.

Faurès, Jean-Marc, and Guido Santini. 2008. *Water and the rural poor: interventions for improving livelihoods in sub-Saharan Africa*. Rome, Italy: Food and Agriculture Organization (FAO).

Fedorov, Leonid, and David E Sahn. 2005. "Socioeconomic Determinants of Children's Health in Russia: A Longitudinal Study." *Economic Development and Cultural Change* no. 53 (2):479-500.

Feng, Shuaizhang, Alan B Krueger, and Michael Oppenheimer. 2010. "Linkages among climate change, crop yields and Mexico–US cross-border migration." *Proceedings of the National Academy of Sciences* no. 107 (32):14257-14262.

Fisher, Anthony C, W Michael Hanemann, Michael J Roberts, and Wolfram Schlenker. 2012. "The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather: comment." *The American Economic Review* no. 102 (7):3749-3760.

Foster, Andrew D., and Mark R. Rosenzweig. 2001. "Imperfect commitment, altruism, and the family: Evidence from transfer behavior in low-income rural areas." *Review of Economics and Statistics* no. 83 (3):389-407.

Friedman, M. 1992. "Do old fallacies ever die?" *Journal of Economic Literature* no. 30 (4):2129-2132.

Golden, M. H. 1994. "Is complete catch-up possible for stunted malnourished children?" *European Journal of Clinical Nutrition* no. 48 Suppl 1:S58-70.

Google Maps. *Kagera (S 1 52.041 and E 31 19.774)* 2013.

Gørgens, Tue, Xin Meng, and Rhema Vaithianathan. 2012. "Stunting and selection effects of famine: a case study of the great Chinese famine." *Journal of Development Economics* no. 97:99-111.

GPS Visualizer. *Kagera (S 2 04.362E 31 11.070)* 2013.

Grantham-McGregor, S., Y. B. Cheung, S. Cueto, P. Glewwe, L. Richter, B. Strupp, and International Child Development Steering Group. 2007. "Developmental potential in the first 5 years for children in developing countries." *Lancet* no. 369 (9555):60-70.

Gray, Clark L, and Valerie Mueller. 2012a. "Drought and Population Mobility in Rural Ethiopia." *World Development* no. 40 (1):134-145.

Gray, Clark L, and Valerie Mueller. 2012b. "Natural disasters and population mobility in Bangladesh." *Proceedings of the National Academy of Sciences* no. 109 (16):6000-6005.

Grimard, Franque. 1997. "Household consumption smoothing through ethnic ties: evidence from Cote d'Ivoire." *Journal of Development Economics* no. 53 (2):391-422.

Haas, J. D., and F. Campirano. 2006. "Interpopulation variation in height among children 7 to 18 years of age." *Food and Nutrition Bulletin* no. 27 (4):S212-S223.

Haddad, L. J., and H. E. Bouis. 1991. "The Impact of Nutritional-Status on Agricultural Productivity - Wage Evidence from the Philippines." *Oxford Bulletin of Economics and Statistics* no. 53 (1):45-68.

Halliday, Timothy. 2006. "Migration, risk, and liquidity constraints in El Salvador." *Economic Development and Cultural Change* no. 54 (4):893-925.

Hansen, Lars Peter. 1982. "Large Sample Properties of Generalized-Method of Moments Estimators." *Econometrica* no. 50 (4):1029-1054.

Harari, Mariaflavia, and Eliana La Ferrara. 2012. Conflict, Climate and Cells: A disaggregated analysis. *University of Bocconi Working Paper Series* 461.

Harris, John R., and Michael P. Todaro. 1970. "Migration, Unemployment and Development - A Two-Sector Analysis." *American Economic Review* no. 60 (1):126-142.

Hayashi, Fumio. 2000. *Econometrics*. Princeton: Princeton University Press.

Hoddinott, John. 1994. "A Model of Migration and Remittances Applied to Western Kenya." *Oxford Economic Papers-New Series* no. 46 (3):459-476.

Hoddinott, John, and Bill Kinsey. 2001. "Child Growth in the Time of Drought." *Oxford Bulletin of Economics and Statistics* no. 63 (4):409-36.

Hoff, Karla, and Arjit Sen. 2006. "The Kin System as a Poverty Trap?" In *Poverty Traps*, edited by Samuel Bowles, Steven Durlauf and Karla Hoff, 91-115. Princeton University Press.

Iliffe, John. 2006. *The African AIDS Epidemic: A History*: Athens, Ohio: Ohio University Press.

Jakiela, Pamela, and Owen Ozier. 2012. "Does africa need a rotten kin theorem? experimental evidence from village economies." *World Bank Policy Research Working Paper* (6085).

Jayachandran, Seema. 2006. "Selling labor low: Wage responses to productivity shocks in developing countries." *Journal of Political Economy* no. 114 (3):538-575.

Jenkins, S.P. 1995. "Easy estimation methods for discrete-time duration models." *Oxford Bulletin of Economics and Statistics* no. 57 (1):129-136.

Jerven, Morten. 2013. *Poor Numbers: How We Are Misled by African Development Statistics and What to Do about It*. Ithaca, NY: Cornell University Press.

Kinnan, Cynthia. 2012. Distinguishing barriers to insurance in Thai villages. Mimeo, Northwestern University.

Koenker, Roger, and Kevin Hallock. 2001. "Quantile regression: An introduction." *Journal of Economic Perspectives* no. 15 (4):43-56.

Komlos, John. 1992. "Toward an anthropometric history of African-Americans: the case of the free blacks in antebellum Maryland." In *Strategic Factors in Nineteenth Century American Economic History: A Volume to Honor Robert W. Fogel*, 297-329. Chicago, IL: University of Chicago Press.

Kuczmarski, R. J., C. L. Ogden, S. S. Guo, L. M. Grummer-Strawn, K. M. Flegal, Z. Mei, R. Wei, L. R. Curtin, A. F. Roche, and C. L. Johnson. 2002. "2000 CDC Growth Charts for the United States: methods and development." *Vital and Health Statistics, Series 11, Data from the national health survey* (246):1-190.

Leider, Stephen, Markus M. Möbius, Tanya Rosenblat, and Quoc-Anh Do. 2009. "Directed Altruism and Enforced Reciprocity in Social Networks." *Quarterly Journal of Economics* no. 124 (4):1815-1851.

Lévi-Strauss, Claude. 1969. *Elementary structures of kinship*. London: Social Science Paperbacks (Revised edition).

Lewis, W. Arthur. 1954. "Economic Development with Unlimited Supplies of Labour." *The Manchester School* no. 22 (2):139-191.

Ligon, Ethan. 1998. "Risk sharing and information in village economies." *Review of Economic Studies* no. 65 (4):847-864.

Ligon, Ethan, and Laura Schechter. 2012. "Motives for sharing in social networks." *Journal of Development Economics* no. 99 (1):13-26.

Ligon, Ethan, Jonathan P. Thomas, and Tim Worrall. 2002. "Informal insurance arrangements with limited commitment: Theory and evidence from village economies." *Review of Economic Studies* no. 69 (1):209-244.

Lilleør, H. B., and K. Van den Broeck. 2011. "Economic drivers of migration and climate change in LDCs." *Global Environmental Change-Human and Policy Dimensions* no. 21:S70-S81.

Lim, Youngjae, and Robert M Townsend. 1998. "General equilibrium models of financial systems: Theory and measurement in village economies." *Review of Economic Dynamics* no. 1 (1):59-118.

Lokina, Razack, Måns Nerman, Justin Sandefur, and Sierra Leone. 2011. "Poverty and Productivity: Small-Scale Farming in Tanzania, 1991-2007." *Preliminary results for International Growth Centre commissioned study*.

Lorenz, Christof, and Harald Kunstmann. 2012. "The Hydrological Cycle in Three State-of-the-Art Reanalyses: Intercomparison and Performance Analysis." *Journal of Hydrometeorology* no. 13 (5):1397-1420.

Lucas, Robert EB, and Oded Stark. 1985. "Motivations to remit: Evidence from Botswana." *The Journal of Political Economy*:901-918.

Lundborg, Petter, Paul Nystedt, and Dan-Olof Rooth. 2013. "Height and Earnings: The Role of Cognitive and Non-Cognitive Skills." *Journal of Human Resources* no. (forthcoming).

Mani, Subha. 2012. "Is there Complete, Partial, or No Recovery from Childhood Malnutrition? – Empirical Evidence from Indonesia." *Oxford Bulletin of Economics and Statistics* no. 74 (5):691–715.

Marchiori, Luca, Jean-Francois Maystadt, and Ingmar Schumacher. 2012. "The impact of weather anomalies on migration in sub-Saharan Africa." *Journal of Environmental Economics and Management* no. 63 (3):355-374. doi: DOI 10.1016/j.jeem.2012.02.001.

Marchiori, Luca, Jean-Francois Maystadt, and Ingmar Schumacher. 2013. Is environmentally-induced income variability a driver of migration? A macroeconomic perspective. In *IPAG Business School Working Paper Series, no. 2013-17*.

Martorell, Reynaldo, Juan Rivera, and Haley Kaplowitz. 1990. "Consequences of Stunting in Early Childhood for Adult Body Size in Rural Guatemala." *Annales Nestlé* no. 48:85-92.

Martorell, Reynaldo, Juan Rivera, Haley Kaplowitz, and Ernesto Pollitt. 1992. "Long-term consequences of growth retardation during early childhood." In *Human growth: basic and clinical aspects*, edited by M Hernandez & J Argente, 143-149. Amsterdam: Elsevier Science Publishers.

Mazzocco, Maurizio, and Shiv Saini. 2012. "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences." *The American Economic Review* no. 102 (1):428-468.

McIntosh, Craig T, and Wolfram Schlenker. 2006. Identifying non-linearities in fixed effects models. Mimeo, University of California, San Diego.

McKay, Andrew, and David Lawson. 2003. "Assessing the extent and nature of chronic poverty in low income countries: issues and evidence." *World Development* no. 31 (3):425-439.

McKay, Andy. 2013. "Growth and Poverty Reduction in Africa in the Last Two Decades: Evidence from an AERC Growth-Poverty Project and Beyond." *Journal of African Economies* no. 22 (suppl 1):i49-i76.

Moradi, Alexander. 2010. "Nutritional status and economic development in sub-Saharan Africa, 1950-1980." *Economics and Human Biology* no. 8 (1):16-29.

Moradi, Alexander. 2012. Selective Mortality or Growth after Childhood? What really is Key to Understand the Puzzlingly Tall Adult Heights in Sub-Saharan Africa. Mimeo, University of Sussex.

Moyo, Mujobu, Rebecca Simson, Jacob Arun, and François-Xavier de Mevius. 2012. Attaining Middle Income Status. Tanzania. Growth and Structural Transformation Required to Reach Middle Income Status by 2025. *International Growth Centre Working Paper Series* 11/1019.

Mueller, V., C. Gray, and K. Kosec. 2014. "Heat stress increases long-term human migration in rural Pakistan." *Nature Climate Change* no. forthcoming.

Munshi, Kaivan. 2003. "Networks in the modern economy: Mexican migrants in the US labor market." *The Quarterly Journal of Economics* no. 118 (2):549-599.

Munshi, Kaivan, and Mark Rosenzweig. 2006. "Traditional Institutions Meet the Modern World: Caste, Gender, and Schooling Choice in a Globalizing Economy." *The American Economic Review*:1225-1252.

Ngallaba, Sylvester, Saidi Hussein Kapiga, Ireneus Ruyobya, and J. Ties Boerma. 1993. Tanzania Demographic and Health Survey 1991/1992. Dar es Salaam, Tanzania: Bureau of Statistics, Planning Commission (Tanzania) / Macro International, Inc. (Maryland, USA).

Outes, Ingo, and Catherine Porter. 2013. "Catching up from early nutritional deficits? Evidence from rural Ethiopia." *Economics and Human Biology* no. 11 (0):148-163.

Persico, Nicola, Andrew Postlewaite, and Dan Silverman. 2004. "The Effect of Adolescent Experience on Labor Market Outcomes: The Case of Height." *Journal of Political Economy* no. 112 (5):1019-1053.

Platteau, Jean-Philippe. 2000. *Institutions, Social Norms, and Economic Development*: Harwood Academic Publishers.

Platteau, Jean-Philippe. 2012. "Redistributive pressures in Sub-Saharan Africa: Causes, consequences, and coping strategies." *The Historical Roots of Africa's Underdevelopment. Cambridge University Press, Cambridge*.

Platteau, Jean-Philippe, and Petros G Sekeris. 2010. "On the feasibility of power and status ranking in traditional setups." *Journal of Comparative Economics* no. 38 (3):267-282.

Prentice, Andrew M, Kate A Ward, Gail R Goldberg, Landing M Jarjou, Sophie E Moore, Anthony J Fulford, and Ann Prentice. 2013. "Critical windows for nutritional interventions against stunting." *American Journal of Clinical Nutrition* no. 97 (5):911-918.

Quah, Danny. 1993. "Galton's fallacy and tests of the convergence hypothesis." *The Scandinavian Journal of Economics* no. 95 (4):427-43.

Radelet, Steven C. 2010. *Emerging Africa: how seventeen countries are leading the way*. Washington, D.C.: Center for Global Development Books.

Ranis, Gustav, and John CH Fei. 1961. "A theory of economic development." *The American Economic Review* no. 51 (4):533-565.

Reilly, Barry, and Robert Witt. 1996. "Crime, Deterrence and Unemployment in England and Wales: an empirical analysis." *Bulletin of Economic Research* no. 48 (2):137-159.

Rienecker, Michele M, Max J Suarez, Ronald Gelaro, Ricardo Todling, Julio Bacmeister, Emily Liu, Michael G Bosilovich, Siegfried D Schubert, Lawrence Takacs, and

Gi-Kong Kim. 2011. "MERRA: NASA's modern-era retrospective analysis for research and applications." *Journal of Climate* no. 24 (14):3624-3648.

Ritchie, J.T., and D.S. NeSmith. 1991. "Temperature and crop development." In *Modeling plant and soil systems*, edited by John Hanks and J.T. Ritchie, 5-29. Madison,WI: American Society of Agronomy. Monograph 31.

Rosenzweig, Mark R. 2003. "Payoffs from panels in low-income countries: Economic development and economic mobility." *The American Economic Review* no. 93 (2):112-117.

Rosenzweig, Mark R, and Oded Stark. 1989. "Consumption Smoothing, Migration, and Marriage - Evidence from Rural India." *Journal of Political Economy* no. 97 (4):905-926.

Rouanet, L. 2011. The Double African Paradox: What does selective mortality tell us? Paper read at Proceedings of the German Development Economics Conference, Berlin 2011.

Rowhani, P., D. B. Lobell, M. Linderman, and N. Ramankutty. 2011. "Climate variability and crop production in Tanzania." *Agricultural and Forest Meteorology* no. 151 (4):449-460.

Satyanarayana, K., A. Nadamuni Naidu, and B. S. Narasinga Rao. 1980. "Adolescent growth spurt among rural Indian boys in relation to their nutritional status in early childhood." *Annals of Human Biology* no. 7 (4):359-65.

Satyanarayana, K., A. Nadamuni Naidu, M. C. Swaminathan, and B. S. Narasinga Rao. 1981. "Effect of nutritional deprivation in early childhood on later growth--a community study without intervention." *The American Journal of Clinical Nutrition* no. 34 (8):1636-7.

Schechter, Laura, and Alex Yuskavage. 2011. Reciprocated versus unreciprocated sharing in social networks. Mimeo, University of Wisconsin and U.S. Treasury.

Schick, Andreas, and Richard H. Steckel. 2010. Height as a proxy for cognitive and non-cognitive ability. *National Bureau of Economic Research (NBER) Working Paper* No. 16570.

Schlenker, W., W. M. Hanemann, and A. C. Fisher. 2006. "The impact of global warming on US agriculture: An econometric analysis of optimal growing conditions." *Review of Economics and Statistics* no. 88 (1):113-125.

Schlenker, Wolfram, and David B Lobell. 2010a. "Robust negative impacts of climate change on African agriculture." *Environmental Research Letters* no. 5 (1):014010.

Schlenker, Wolfram, and David B Lobell. 2010b. "Robust negative impacts of climate change on African agriculture: Supplementary Data." *Environmental Research Letters* no. 5 (1):014010.

Schroeder, Dirk G. 2008. "Malnutrition." In *Nutrition and health in developing countries*, edited by Richard D. Semba and Martin W. Bloem, 341-376. Totowa, NJ.: Humana Press Inc.

Schulhofer-Wohl, Sam. 2011. "Heterogeneity and Tests of Risk Sharing." *Journal of Political Economy* no. 119 (5):925-958.

Schultz, T Paul. 2002. "Wage Gains Associated with Height as a Form of Health Human Capital." *American Economic Review* no. 92 (2):349-353.

Shrimpton, R., C. G. Victoria, M. de Onis, R. C. Lima, M. Blossner, and G. Clugston. 2001. "Worldwide timing of growth faltering: Implications for nutritional interventions." *Pediatrics* no. 107 (5):art. no.-e75.

Skoufias, Emmanuel, and Katja Vinha. 2013. "The impacts of climate variability on household welfare in rural Mexico." *Population and Environment* no. 34 (3):370-399.

Smith, James P, and Duncan Thomas. 2003. "Remembrances of things past: test–retest reliability of retrospective migration histories." *Journal of the Royal Statistical Society: Series A (Statistics in Society)* no. 166 (1):23-49.

Stark, Oded, and David E Bloom. 1985. "The New Economics of Labor Migration." *American Economic Review* no. 75 (2):173-178.

Stark, Oded, and Robert EB Lucas. 1988. "Migration, remittances, and the family." *Economic Development and Cultural Change*:465-481.

Steckel, Richard H. 1987. "Growth depression and recovery: the remarkable case of American slaves." *Ann Hum Biol* no. 14 (2):111-132.

Steckel, Richard H. 2000. "Diets Versus Diseases in the Anthropometrics of Slave Children: A Reply." *The Journal of Economic History* no. 60 (1):247-259.

Steckel, Richard H. 2009. "Heights and human welfare: Recent developments and new directions." *Explorations in Economic History* no. 46 (1):1-23.

Stein, Zena. 1975. *Famine and human development: the Dutch hunger winter of 1944-1945, Oxford medical publications*. New York ; London: Oxford University Press.

Stock, James, and Motohiro Yogo. 2005. "Testing for weak instruments in linear IV regression." In *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, edited by Donald Andrews and James Stock. Cambridge University Press.

Tanner, J. M. 1981. "Catch-up Growth in Man." *British Medical Bulletin* no. 37 (3):233-238.

Thomas, Duncan, and John Strauss. 1997. "Health and wages: Evidence on men and women in urban Brazil." *Journal of Econometrics* no. 77 (1):159-185.

Townsend, Robert M. 1982. "Optimal multiperiod contracts and the gain from enduring relationships under private information." *The Journal of Political Economy*:1166-1186.

Townsend, Robert M. 1994. "Risk and Insurance in Village India." *Econometrica* no. 62 (3):539-591.

United Nations, Department of Economic and Social Affairs, Population Division 2012. World Urbanization Prospects: The 2009 Revision.CD-ROM Edition, 2010. Accessed at http://esa.un.org/unpd/wup/index.htm on 19 January 2012.

United Nations Development Programme. 2009. Human Development Report 2009. Overcoming barriers: Human mobility and development. United Nations, New York.

URT. 2009. Household Budget Survey 2007 – Tanzania Mainland. The United Republic of Tanzania (URT).

URT. 2012. National Sample Census Of Agriculture 2007/2008: Regional Report – Kagera Region (volume V) The United Republic of Tanzania (URT).

URT. 2013. Population and Housing Census: Population Distribution by Administrative Units; Key Findings. National Bureau of Statistics (NBS) and Office of Chief Government Statistician (OCGS), Zanzibar. United Republic of Tanzania (URT).

Victora, C. G., L. Adair, C. Fall, P. C. Hallal, R. Martorell, L. Richter, H. S. Sachdev, Maternal, and Group Child Undernutrition Study. 2008. "Maternal and child undernutrition: consequences for adult health and human capital." *Lancet* no. 371 (9609):340-57.

Victora, C. G., M. de Onis, P. C. Hallal, M. Blossner, and R. Shrimpton. 2010. "Worldwide timing of growth faltering: revisiting implications for interventions." *Pediatrics* no. 125 (3):e473-80.

Vidmar, S., J. Carlin, K. Hesketh, and T. Cole. 2004. "Standardizing anthropometric measures in children and adolescents with new functions for egen." *Stata Journal* no. 4 (1):50-55.

Vogl, Tom. 2012. Height, skills, and labor market outcomes in Mexico. *National Bureau of Economic Research (NBER) Working paper* No. 18318.

White, H. 1980. "A Heteroskedasticity-Consistent Covariance-Matrix Estimator and a Direct Test for Heteroskedasticity." *Econometrica* no. 48 (4):817-838.

WHO. 2006. "WHO Child Growth Standards based on length/height, weight and age." *Acta paediatrica* no. Suppl 450:76-85.

Williams, J. P. G. 1981. "Catch-up Growth." *Journal of Embryology and Experimental Morphology* no. 65:89-101.

World Bank. 2004. User's Guide to the Kagera Health and Development Survey Datasets. Mimeo, The World Bank.

World Bank. 2013. World Development Indicators (WDI) database. Data retrieved: 20 November 2013.

Yang, Dean, and HwaJung Choi. 2007. "Are remittances insurance? Evidence from rainfall shocks in the Philippines." *The World Bank Economic Review* no. 21 (2):219-248.

Young, Alwyn. 2012. "The African Growth Miracle." *Journal of Political Economy* no. 120 (4):696-739.

Young, Alwyn. 2013. "Inequality, the Urban-Rural Gap, and Migration." *The Quarterly Journal of Economics* no. 128 (4):1727-1785.

Appendices

Appendices: Essay 1

Appendix A: Is there a gender dimension in risk sharing?

Table A.1: replicating Column 1 in Table 3.8 using gender interactions

Dependent variable: (logged) per capita consumption growth	
Number of households that experienced a shock in the network	-0.050**
	(0.024)
--- Interacted with:	
* male migrant	0.001
	(0.034)
Own shock	-0.078
	(0.047)
--- Interacted with:	
* male migrant	-0.053
	(0.105)
male migrant	0.146**
	(0.067)
Number of observations	1,270
R^2	0.465
Adjusted R^2	0.452

Note: a household was categorized as a male migrant household if there was at least one male PHHM residing in the household. Out of the 1,270 migrant households, 63 per cent are categorised as female migrant households and 37 per cent as male migrant households. For other notes see Table 3.8.

Table A.2: replicating Column 1 in Table 3.9 using gender interactions

Dependent variable: (logged) per capita consumption growth	
Number of non-migrant hhs that experienced a shock in the network	-0.063** (0.032)
--- Interacted with:	
*male migrant	0.025 (0.050)
Number of migrant hhs that experienced a shock in the network	-0.026 (0.038)
--- Interacted with:	
*male migrant	-0.035 (0.071)
Own shock	-0.076 (0.047)
--- Interacted with:	
*male migrant	-0.053 (0.105)
male migrant	0.146** (0.067)
Number of observations	1,270
R^2	0.466
Adjusted R^2	0.451

Note: a household was categorized as a male migrant household if there was at least one male PHHM residing in the household. Out of the 1,270 migrant households, 63 per cent are categorised as female migrant households and 37 per cent as male migrant households. For other notes see Table 3.9.

Table A.3: Other transactional insurance motive tests in Table 3.10 by gender

Dependent variable: (logged) per capita consumption growth	Column 1		Column 2	
	Males	**Females**	**Males**	**Females**
Number of non-migrant hhs that experienced a shock in the network	-0.018	-0.062*	0.007	-0.129
	(0.047)	(0.034)	(0.077)	(0.096)
--- Interacted with:				
* Share of land in BLV in total land portfolio	-0.001	-0.075		
	(0.097)	(0.065)		
* Number of years since the last PHHM migrated into this hh			-0.005	0.005
			(0.007)	(0.008)
Own shock	-0.106	-0.093*	-0.153*	-0.092*
	(0.089)	(0.049)	(0.092)	(0.051)
Share of land in BLV in total land portfolio	0.373***	0.211***		
	(0.101)	(0.071)		
Household does not own land	0.194**	0.315***		
	(0.084)	(0.065)		
Number of years since the last PHHM migrated into this hh			0.004	-0.005
			(0.006)	(0.005)
Number of observations	474	796	474	796
R^2	0.533	0.449	0.502	0.423
Adjusted R^2	0.500	0.426	0.469	0.401

For notes, see Table 3.10.

Appendix B: First-stage regression results of Column 3 in Table 3.7

Dependent variable: (logged) hh per capita consumption in 1991	
Included instruments:	
Own shock	-0.019
	(0.018)
Number of split-off hhs stayed	0.018
	(0.019)
Number of split-off hhs moved	0.019
	(0.014)
Age of oldest PHHM in the 2010 hh	0.001**
	(0.001)
A PHHM is head of this 2010 hh	0.005
	(0.023)
A PHHM is spouse of this 2010 hh's head	0.036
	(0.025)
A PHHM is child of this 2010 hh's head	-0.008
	(0.034)
A divorced PHHM in 2010 hh	-0.037
	(0.038)
A widowed PHHM in 2010 hh	-0.015
	(0.033)
A married PHHM in 2010 hh	-0.032
	(0.023)
Max years of education of PHHM in this 2010 hh	0.015***
	(0.004)
Number of PHHMs in this 2010 hh	-0.024*
	(0.012)
Natural log value of assets in 1991	0.081***
	(0.027)
Education of hh head in 1991	0.024**
	(0.011)
Head was male in 1991	0.212***
	(0.075)
Age of hh head in 1991	0.006
	(0.006)
Age of head squared	-0.000
	(0.000)
Males 0-5 years in 1991	-0.064***
	(0.023)
Males 6-15 years in 1991	-0.045***
	(0.015)
Males 16-60 years in 1991	0.022
	(0.018)
Males 61+ years in 1991	-0.176**
	(0.079)

Females 0-5 years in 1991	-0.035
	(0.025)
Females 6-15 years in 1991	-0.034**
	(0.017)
Females 16-60 years in 1991	-0.017
	(0.015)
Females 61+ years in 1991	-0.030
	(0.042)
Hh had a non-earth floor in 1991	0.268***
	(0.059)
Excluded instruments:	
(Negative rainfall deviation) * (Age of hh head in 1991)	0.002
	(0.004)
(Negative rainfall deviation) * (Education of hh head in 1991)	0.048*
	(0.027)
(Negative rainfall deviation) * (Head was male in 1991)	0.355**
	(0.159)
Number of observations	2,349
R^2	0.228
Adjusted R^2	0.201
Under-identification test:	
Kleibergen-Paap rk LM statistic	8.780
p-value	0.032
Weak identification tests:	
Cragg-Donald Wald F Statistic	19.12
Kleibergen-Paap rk Wald F statistic	5.419
Over-identification test:	
Hansen-J statistic	1.124
p-value	0.570

note: *** p<0.01, ** p<0.05, * p<0.1. Cluster-robust standard errors by baseline village are in parenthesis. Regression includes baseline village fixed effects. PHHM refers to previous household member (i.e. person interviewed at the baseline).

Appendices: Essay 2

Appendix C: Quantile regression results

C.1: Impact of temperature variation on household consumption

dependent variable: (logged) household per capita consumption	1 OLS	2 q10	3 q25	4 q50	5 q75	6 q90	7 q25 vs. q75	8 q10 vs. q90
temperature (°C)	-0.149**	-0.169	-0.283***	-0.180**	-0.147	-0.139	0.136	0.030
	(0.060)	(0.125)	(0.105)	(0.089)	(0.096)	(0.131)	(0.123)	(0.174)
rainfall (mm)	-0.001	-0.001	-0.001	-0.001	-0.000	0.000	0.001	0.002
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
household level controls	yes	yes	yes	yes	yes	yes	yes	yes
village dummies	yes	yes	yes	yes	yes	yes	yes	yes
year dummies	yes	yes	yes	yes	yes	yes	yes	yes
wave dummies	yes	yes	yes	yes	yes	yes	yes	yes
quarter of year dummies	yes	yes	yes	yes	yes	yes	yes	yes
Number of observations	3,277	3,277	3,277	3,277	3,277	3,277	3,277	3,277
R^2/ Pseudo-R^2	0.452	0.247	0.244	0.265	0.295	0.330	n/a	n/a

note: *** p<0.01, ** p<0.05, * p<0.1
Column 1 repeats Column 2 in Table 4.2 without household fixed effects and uses village dummies and household level controls instead.
Columns 2-6 presents the results based on quantile regression and columns 7-8 inter-quantile regression techniques.
Village clustered (Column 1) or bootstrapped (Columns 2-8) standard errors (200 repetitions) are in parenthesis.
Household level controls include household assets (land size and type of floor), head's characteristics (gender, age, education) and variables capturing household composition and size.

Appendix D: Alternative specifications for the migration models

Table D.1: Impact of temperature variation on migration among males

	1 mfx	2 mfx	3 mfx	4 mfx	5 mfx
temperature	-0.023***	0.002	-0.006	-0.015**	-0.017***
	(0.007)	(0.008)	(0.007)	(0.006)	(0.006)
spell year			0.003***	-0.004**	-0.018***
			(0.001)	(0.002)	(0.005)
(spell year) ^{2}				0.000***	0.002***
				(0.000)	(0.001)
(spell year) ^{3}					-0.000***
					(0.000)
year dummies	yes	no	no	no	no
Likelihood-ratio test	$\chi 2(15)$ = 57.25	n/a	$\chi 2$ (1) = 22.18	$\chi 2$ (2) = 36.28	$\chi 2$ (3) = 42.16
Observations	13,225	13,225	13,225	13,225	13,225
Pseudo-Log likelihood	-2685.366	-2721.045	-2706.245	-2699.604	-2695.605

Note: Column 2 is the restricted model in the Likelihood ratio tests in Columns 1, 3-5.
The covariates and other notes are the same as in Table 4.4.

Table D.2: Impact of temperature variation on migration among females

	1 mfx	2 mfx	3 mfx	4 mfx	5 mfx
temperature	-0.001	0.007	0.005	-0.004	-0.005
	(0.014)	(0.009)	(0.010)	(0.010)	(0.010)
spell year			0.001*	-0.006***	-0.013**
			(0.001)	(0.002)	(0.006)
(spell year) ^{2}				0.000***	0.002*
				(0.000)	(0.001)
(spell year) ^{3}					-0.000
					(0.000)
year dummies	yes	no	no	no	no
Likelihood-ratio test	$\chi 2(15)$ = 44.53	n/a	$\chi 2$ (1) = 8.22	$\chi 2$ (2) = 11.98	$\chi 2$ (3) = 13.72
Observations	11,737	11,737	11,737	11,737	11,737
Pseudo-Log likelihood	-3,225.630	-3,247.281	-3,246.080	-3,242.010	-3,241.139

Note: Column 2 is the restricted model in the Likelihood ratio tests in Columns 1, 3-5. The covariates and other notes are the same as in Table 4.5.

Appendix E: Testing the stability of the temperature estimates

Figure E.1: temperature coefficients for males, omitting each village in turn

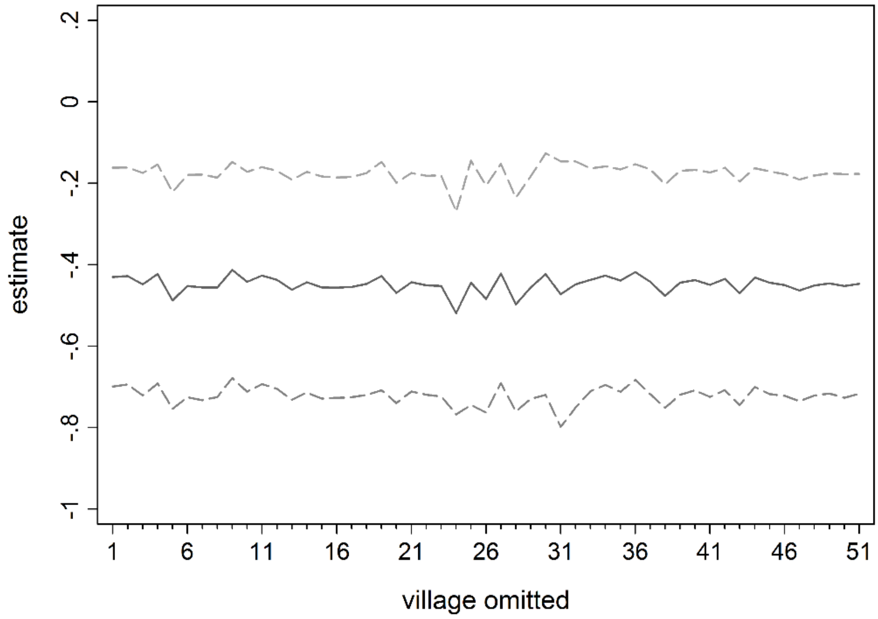

Note: dashed lines represent the 95% confidence intervals

Figure E.2: temperature coefficients for females, omitting each village in turn

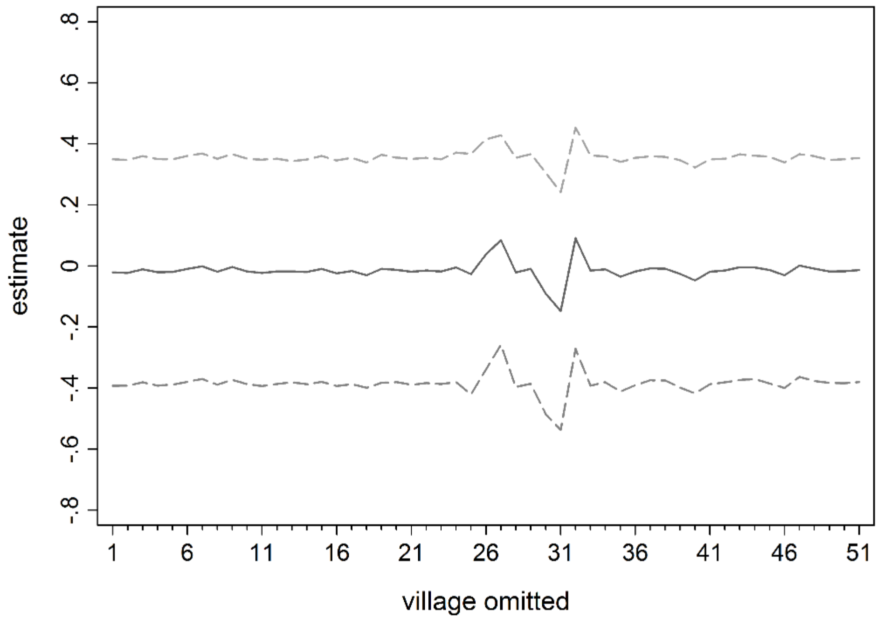

Note: dashed lines represent the 95% confidence intervals

Appendix F: Linear probability models

Table F.1: linear probability models with village and individual level fixed effects

	males		females	
	1	**2**	**3**	**4**
temperature	-0.025***	-0.032*	-0.000	0.002
	(0.007)	(0.017)	(0.015)	(0.013)
rainfall	-0.000	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)
individual level FE	no	yes	no	yes
village FE	yes	no	yes	no
household level controls	yes	no	yes	no
year dummies	yes	yes	yes	yes
Number of observations	13,225	13,225	11,737	11,737
Adjusted R2	0.025	0.069	0.019	0.090
R2	0.028	0.071	0.022	0.092

Standard errors (in parenthesis) clustered at the village level. The odd columns contain the same set of control variables as in Table 4.4 and Table 4.5.

Appendices: Essay 3

Appendix G: Descriptive statistics for the British Cohort Study

Table G.1: Evolution of HAZ scores in the British cohort (BCS70)

	BCS70	
	mean	**std. dev.**
age in months (t=0)	60.84	1.307
height (t=0)	108.8	5.035
height (t=1)	171.0	9.881
HAZ_1 (t=0)	0.01	1.056
HAZ_2 (t=1)	0.17	1.026
difference in HAZ:	0.16	0.91
t-test: $HAZ_{t=0} = HAZ_{t=1}$	10.56	
percentage: $HAZ_{t=0} < -2$	3%	
percentage: $HAZ_{t=1} < -2$	1%	
percentage: $HAZ_{t=0} < -3$	1%	
percentage: $HAZ_{t=1} < -3$	0%	
sample size	9,635	

note: t=0 refers to early childhood, t=1 to adulthood. HAZ scores calculated using the US 2000 NCHS/CDC as the reference population.
Adult heights are based on self-reports. Sensitivity tests (available from the author) were conducted to see if the main findings are affected by this. They were not.

Appendix H: Why estimating Equation (5.1) using raw measures of height is not appropriate?

Estimating Eq. (5.1) using raw measures of height artificially inflates the β coefficient. This can be demonstrated if we transform Equation (5.2) to $= r_{y,x}\,\frac{\theta_y}{\theta_x}$, where $r_{y,x}$ is the correlation coefficient and $\frac{\theta_y}{\theta_x}$ is the ratio of standard deviations of the two variables. The dispersion in the height distribution varies with age and increases as children grow older. This inflates the ratio of standard deviations and leads to a larger β coefficient (other things being equal). The use of height-for-age z-scores circumvents this problem as they measure the distance to the normal growth curve at each age, making the standard deviations less dependent of the age when the child was measured. Cameron, Preece, and Cole (2005) make a similar point discussing the relationship between catch-up growth and regression-to-the-mean.

Appendix I: The impact of measurement error on β

In a simple linear bivariate regression such as Equation (5.1), the magnitude of the downward bias can be calculated using the following formula (see, for example, Deaton 1997, 99):

H.1 $$\beta \frac{{\theta_x}^2}{{\theta_e}^2 + {\theta_x}^2} = \beta\lambda,$$

where θ_x is the true standard deviation of the correctly measured height (unobserved), and θ_e is the standard deviation of the measurement error. The measurement error inflates the denominator causing a downward bias in the convergence estimate. The term λ is then the reliability ratio measuring the magnitude of the downward bias.

The standard deviation in the early childhood height observed in the British cohort is 5.034. To get a sense of the potential bias, I calculate the impact of various level of zero-mean measurement error on β. As can be seen in Table below, a small level of measurement error leads to a negligible downward bias in the β coefficient. Even measurement error that has a standard deviation of 1 cm (i.e. 32 % of the height measurements contain more than 1 cm of measurement error), biases the estimate downward only by four percentage points and cannot explain the convergence finding in Table 5.1. It is difficult to imagine measurement error of this magnitude in any carefully constructed survey. In this light, adjustments, such as instrumental variables techniques, to address measurement error seem unnecessary.

Table I.1: the impact of measurement error (ME) on $\boldsymbol{\beta}$

std. dev. of ME	reliability ratio
0.01	0.999996
0.1	0.999606
0.25	0.997541
0.5	0.990235
0.75	0.978293
1	0.962051
1.25	0.941944
1.5	0.918482

Note: standard deviation of height used in this example is 5.034, based on my own calculations from the BCS70 data

Appendix J: Additional attrition tests for the Kagera Health and Development Survey

Table J.1: KHDS Attrition tests: initial HAZ scores by sample category

	N	mean	sd	t-stat*
in the final sample	540	-1.86	1.155	n/a
deceased after 1991-94	69	-2.10	1.396	1.37
not measured in 2010	257	-1.74	1.220	-1.32
missing date of birth	18	-1.51	1.198	-1.22

* Welch t-test testing the difference in means against the first category (in the final sample)

Table J.2: KHDS adult HAZ scores by migration status in 2010

	2010			
	N	mean	sd	t-stat*
non-migrant	258	-1.21	1.012	n/a
migrant	282	-1.20	0.995	-0.0299

* Welch t-test testing the difference in means against the first category (non-migrant)