The London School of Economics
and Political Science

Essays on the Economic Consequences of Weather and Climate Change

Jonathan Colmer

March 2016

A thesis submitted to the London School of Economics for the degree of Doctor of Philosophy.

For my wife Naomi, whom I married just before the beginning of this project, and for my daughter Grace, who was born just before its end.

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work unless otherwise clearly indicated.

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that this thesis consists of approximately 80,000 words.

Statement of Conjoint Work

Two out of the four chapters that form this thesis involve conjoint work:

Chapter 2 was co-authored with Yonas Alem. Overall, my contribution amounts to 75% of the paper.

Chapter 4 was co-authored with Ralf Martin, Mirabelle Muûls, and Ulrich Wagner. Overall, my contribution amounts to 50% of the paper.

Abstract

This thesis seeks to advance our understanding of climatic influence on economic outcomes. The approach taken places emphasis on understanding the channels and mechanisms through which weather has an effect, and through which climate change *could* have an effect, on economic behaviour – rather than estimating the impact of future climate change – to better inform the design and implementation of policy. This thesis is composed of four papers that adopt this new paradigm, providing new insights into how weather affects economic outcomes today, how economic agents respond to and manage the economic consequences of changes in their natural environment, and providing explicit mechanisms through which the impacts of, and adaptation to, climate change *could* affect economic outcomes in the future.

Acknowledgments

The completion of this thesis has been supported by many wonderful people. First and foremost, I am eternally grateful for the supervision provided by Robin Burgess and John Van Reenen, without whom this accomplishment would not have been possible. I would also like to thank Eric Neumayer, Antoine Dechezleprêtre, and Alex Bowen for their advice, encouragement, and support at various stages of the Ph.D.

I have also benefited considerably from discussions with other faculty and colleagues in the Department of Economics, the Centre for Economic Performance, STICERD, and the Grantham Research Institute. In particular, I thank Yonas Alem, Gharad Bryan, Greg Fischer, Vernon Henderson, Ralf Martin, and Mirabelle Muûls for their guidance, support, and mentoring. In addition, I would like to thank Jon de Quidt, Thiemo Fetzer, Sam Marden, and Frank Pisch for their thoughts and suggestions, as well as for ensuring that I do not take myself too seriously.

Tremendous thanks are also due to the administrative staff at the LSE, with special mention to Gisela Lafico, Mark Wilbor, Linda Cleavely, Jane Dickson, Jo Miller, Ginny Pavey, Nic Warner and Joe Joannes, who have all made the process a lot smoother than it could have been.

Last, but by no means least, I thank my family and friends for putting up with me for the past 5 years. I am especially thankful to my wife Naomi, for her patience, confidence, and support throughout the process. I couldn't have done this without you.

Funding Acknowledgments

Financial support from the Grantham Foundation, the ESRC Centre for Climate Change Economics and Policy, the Bagri Fellowship, and the Multi Donor Trust Fund on Sustainable Urbanization of the World Bank and Department for International Development is gratefully acknowledged.

Contents

List of Tables

List of Figures

Introduction

This thesis seeks to contribute to our understanding of the relationship between economic and natural systems, with a focus on the economic consequences of local changes in the weather and natural climate. Economists and other social scientists have long been interested in this relationship. However, it is only recently that researchers have begun to engage with this relationship. This is due to the emergence of climate change as a major policy issue, advances in the statistical and econometric tools available to evaluate counterfactual outcomes and identify causal relationships, and advances in computer science and climatology that have improved our ability to parameterise and identify the components associated with climatological variations that are most relevant to the socio-economic outcomes of interest.

The research presented here exploits exogenous variation in weather realisations over time within a given spatial unit to plausibly identify the effects of weather on economic outcomes. The focus of the existing literature to date has been to assess the potential economic consequences of future climate change. However, to draw inferences about future climate change, a number of stringent assumptions have to be made, such as the complex mapping of weather onto climate and the endogenous responses of economic agents to longer-run climate change.

I propose a new paradigm to advance our understanding of climatic influence on economic outcomes. I argue that: (1) understanding the effects of short-run weather variation on economic activity is of first-order interest in the context of economic development, given the centrality of agriculture to the economic lives of the poor; (2) instead of estimating the impact of climate change, we should focus on understanding the channels and mechanisms through which climate change could have an effect – rather than estimating the impact of future climate change – so that we can better design and implement policy to influence future impacts. By focussing on the channels and mechanisms through which weather affects, and climate change could effect, economic behaviour, this approach avoids the need to make any assumptions about the mapping of weather onto climate or about any endogenous technological or demographic change that may arise as a consequence of climate change.

This thesis is composed of four papers that adopt this paradigm, seeking to identify specific economic channels and mechanisms through which weather affects economic outcomes (through labour markets, global supply chains, and future income uncertainty) to provide insights into how weather affects economic outcomes today, as well as providing explicit mechanisms through which climate change could affect economic outcomes in the future.

Chapter 1 explores whether the reallocation of workers across sectors can attenuate the economic consequences of weather-driven changes in agricultural productivity in India. I find that temperature increases are associated with a reduction in agricultural production and the demand for agricultural labour; however, this is offset by labour movements into the formal manufacturing sector. Having demonstrated that temperature increases drive changes in the sectoral composition of labour, I explore the consequences of these labour movements on firm and incumbent worker outcomes in the formal manufacturing sector. I find that temperature increases result in an increase in the number of casual workers, resulting in an expansion of manufacturing output. In addition, I estimate that the expansion of casual worker employment is associated with a reduction in the wages of casual workers, and an increase in the wages of permanent workers, suggesting that the tasks that casual and permanent workers engage in are complementary in production. In line with this interpretation, I also find that this expansion is associated with an increase in manufacturing productivity. Collectively, these results suggest that the reallocation of labour across sectors could significantly attenuate the economic consequences of weather-driven changes in agricultural productivity.

Chapter 2, completed with Yonas Alem, explores whether uncertainty about future income realisations has an effect on the individual well-being of smallholder farmers in rural Ethiopia. We measure and identify the effects of income uncertainty, separately from realised events, by exploiting exogenous variation in rainfall variability, which we show is a reasonable proxy for income uncertainty after controlling for both contemporaneous and historical weather events. We find that increases in rainfall variability are associated with a reduction in objective consumption and subjective well-being, suggesting that the welfare gains from managing both short-run weather events and long-run climatic change are likely to be substantially greater than estimates based solely on realised events.

Chapter 3 continues to explore the effects of income uncertainty, examining the effects of parental income uncertainty on child labour and human capital investments in rural Ethiopia. Exploiting rainfall variability as a proxy for parental income uncertainty I find that an increase in parental income uncertainty is associated with a reduction in the number of hours children spend working on the farm and an increase in the number of hours spent studying at home, suggesting that parents invest more in human capital as a response to increases in income uncertainty – a portfolio response. Consistent with these effects, I also estimate increases in the likelihood that a child attends school and an increase in the number of grades achieved. However, I show that these effects vary through the life cycle of the child. In the very early stages of a child's life an increase in parental income uncertainty is associated with a reduction in the likelihood that the child attends school – a precautionary response. However, this relationship weakens and reverses as the child grows older and the returns to education, and consequently diversification, increase. These results suggest that farmers are responsive to changes in future income uncertainty and are actively making decisions to mitigate the economic consequences of future income shocks through investing in the human capital of their children. However, the consequences of income uncertainty vary throughout the life cycle – an important consideration, given the time-sensitive and irreversible nature of human capital investments.

Chapter 4, completed with Ralf Martin, Mirabelle Muûls and Ulrich Wagner, sets out to explore how firms in France are affected directly by local weather effects, as well as indirectly through supply chain networks. Exploiting firm-level exposure to domestic and international weather events, we estimate both the local and linkage effects of weather on the economic performance of French manufacturing firms. We find that domestic exposure to higher temperature and greater rainfall is associated with reductions in production. On the demand side, we estimate that increases in rainfall downstream results in an expansion of production upstream, suggesting that firms are able to increase their market share in response to localised productivity shocks in downstream markets. On the supply side, we observe that, on average, firms downstream are unaffected by upstream increases in temperature or rainfall. However, we find that this effect is heterogeneous across firms, and that firms with a greater initial import share from developing countries experience a relative contraction in production in response to increases in temperature upstream. Interestingly, we find that this effect is attenuated for firms with a greater initial import-share from countries with greater access to air conditioning, indicating that the effect of temperature on production in these countries is due to thermal stress. These results suggest that localised productivity shocks can have significant economic effects across countries, and that if we fail to account for the interconnectedness of firms and sectors we may substantially underestimate the consequences of short-run weather and future climate change on economic activity.

These papers provide new insights into how weather affects economic outcomes today, how economic agents respond to and manage the economic consequences of changes in their natural environment, and provide explicit mechanisms through which the impacts of, and adaptation to, climate change could affect economic outcomes in the future.

While these papers do provide some insight into how future climate change could affect economic outcomes, they do not seek to estimate the impact of future climate change. There are several reasons for this. First, estimating the impact of future climate change involves significant out-of-sample prediction. While economists and are often expected to make policy recommendations based on limited data, the degree of extrapolation required to draw conclusions about the expected impact of future climate change is unprecedented. In the absence of serious government intervention to mitigate greenhouse gas emissions, global average temperatures are expected to increase by between 1.5°C and 4.8°C by 2081–2100 relative to 1986–2005 (Intergovernmental Panel on Climate Change, 2014).[1] Changes of this magnitude have never been experienced in human history, which means that there have been no warming events of this magnitude to estimate economic impacts. Second, any attempt to estimate the impact of future climate change must address the fact that the most significant changes in climate will occur in the distant future. There are significant scientific and economic uncertainties that are introduced when estimating the impacts of events in the distant future. The relevant question here is whether impact estimates based on recent experience can be extrapolated to predict economic outcomes decades from now. To do so requires researchers to make assumptions about endogenous demographic and technological change between now and then, and ignores the potential for unforeseen adaptation opportunities. In considering these factors, it is important to bear in mind the limited success that has been achieved in predicting events and outcomes in the near future, let alone 100 years from now. As such, even the most carefully constructed analysis of future climate change impacts should be considered no more than highly educated guesses.

The approach taken in this dissertation seeks to avoid these issues. This does not diminish the contribution that existing research has made, establishing the impacts of weather and climate change to be of first-order economic importance. Instead the research presented in this thesis seeks to build on this literature with a different focus, seeking to understand how climate change *could* affect social and economic outcomes, rather than estimating the impact of climate change. In so doing, this approach should provide greater insights into specific channels and mechanisms through which weather affects economic outcomes and behaviour today and through which climate change could affect economic outcomes and behaviour in the future, guiding the design and implementation of policy interventions, where necessary, to manage any present and future impacts.

[1] The period 1986–2005 is approximately 0.6°C warmer than the period 1850–1900.

Chapter 1

Weather, Labour Reallocation, and Industrial Production: Evidence from India

To what degree can the movement of workers across sectors mitigate the economic consequences of weather-driven changes in agricultural productivity? Combining worker-level, firm-level and district-level data with high-resolution meteorological data, I examine the effects of weather on economic activity in India. I estimate that increases in temperature are associated with a reduction in agricultural production, average wages and the employment share of workers engaged in agriculture. However, alongside this reduction in labour demand I find that there is an increase in the employment share of workers engaged in manufacturing, with no change in unemployment or migration, indicating that the reallocation of labour across sectors may play an important role in managing the economic consequences of weather-driven changes in productivity. Having demonstrated this, I examine the effects of these labour movements on firm and incumbent worker outcomes in the formal manufacturing sector. I find that workers are employed into casual manufacturing activities, with a corresponding increase in manufacturing output. In addition, this increase in the employment of casual workers also results in a decrease in the average wage of casual workers, and an increase in the average wage of permanent manufacturing workers and manufacturing productivity, suggesting that the tasks that casual workers and permanent workers engage in are complementary in production. Counterfactual estimates suggest that the reallocation of labour across sectors could significantly offset the aggregate economic losses associated with weather-driven changes in agricultural productivity.

1.1 Introduction

Agriculture plays a central role in the livelihoods of many people in developing countries. As a consequence, understanding the relationship between economic and natural systems provides a deeper understanding of the economic lives of the poor, and is of central importance to environment and development policy. While it is important to understand the effects of weather on agricultural markets, it is also of interest to understand the degree to which, if workers are able to move across sectors or space, the economic consequences of weather-driven agricultural productivity shocks might be attenuated, and the consequences of such movements on firm and incumbent worker outcomes in destination sectors or locations. Understanding the responsiveness of workers to weather-driven agricultural productivity shocks therefore yields important insights into the study of climatic influence on economic outcomes, the functioning of labour markets in developing countries, and the short-run behaviour of firms.

I combine worker-, firm- and district-level data with high-resolution meteorological data to understand the effects of weather on agricultural productivity, industrial production and local labour markets in India. First, and unsurprisingly, I identify that increases in temperature are associated with a reduction in agricultural production and the wages of agricultural workers, demonstrating the important role that weather plays in driving short-run agricultural productivity, and in the livelihoods of agricultural workers.[1] I provide evidence to support the premise that variation in the weather provides a change in the demand for labour in agriculture. I then exploit this variation as the key source of variation to examine the effects of agricultural labour demand changes on the local economy. Given the effect of temperature on agricultural production, it is important to understand what happens to workers in response to these changes. While I find that the weather is a strong driver of short-run agricultural productivity, I observe that the weather has no effect on agricultural prices, consistent with a "law of one price", suggesting that reductions in agricultural productivity should push workers out of agriculture and into other tradable sectors of the economy; however, this depends on the ability of other sectors to absorb workers in response to these short-run productivity shocks. I find that increases in temperature result in a reduction in the employment share of agriculture and a big shift of labour into the manufacturing sector. In addition, I estimate that there are no changes in unemployment, or in population through migration, in response to changes in the

[1] Deaton (1992); Paxson (1992); Rosenzweig and Binswanger (1993); Townsend (1994); Jayachandran (2006); Guiteras (2009); Taraz (2012); Kaur (2014); Mobarak and Rosenzweig (2014); Kala (2015).

weather, bounding local labour markets. These results suggest that the ability of non-agricultural sectors to absorb workers is one of the main factors that enable workers to manage weather-driven changes in agricultural productivity, and highlights the important role that market integration and diversification can play in attenuating the economic consequences of weather-driven changes in agricultural productivity.[2]

Having demonstrated that other sectors of the economy are major absorbers of labour in the face of weather-driven changes in agricultural productivity, it is of interest to understand what these workers do, and how they affect firm and incumbent worker outcomes when they move into the manufacturing sector. However, identifying these effects presents a number of empirical challenges. To interpret the effects of weather on manufacturing outcomes as being driven by labour reallocation, it is necessary that these outcomes not be affected by weather in any other way. This is a strong assumption, as there are potentially many channels through which the weather could affect other sectors, directly and/or through agricultural linkages.[3] Consequently, any estimate of the elasticity between weather and economic outcomes will provide the net effect of all the competing and complementary channels involved. Given this ambiguity, it is difficult to interpret empirical estimates of weather in a meaningful way. Where empirically relevant channels move in the same direction, we fail to arrive at a meaningful economic interpretation. Where multiple channels are competing, the effects may be missed entirely or selected interpretations underestimated.

To try and address this concern, I exploit variation in the propensity of firms to absorb labour in response to year-to-year changes in the weather, helping to identify the channel of interest – the labour reallocation effect. To do this, I construct a measure

[2]Matsuyama (1992); Foster and Rosenzweig (2004); Jayachandran (2006); Burgess and Donaldson (2010, 2012); Autor et al. (2013); Bustos et al. (2015); Costinot et al. (2015); Donaldson (2015a,b); Henderson et al. (2015); Hornbeck and Keskin (2015); Hornbeck and Moretti (2015); Mian and Sufi (2015).

[3]Changes in agricultural productivity could affect manufacturing outcomes in sectors that use agricultural products as inputs, propagating shocks through intermediaries (Acemoglu et al., 2012), and a reduction in agricultural income could reduce the consumption base for manufactured products with local demand (Soderbom and Rijkers, 2013; Henderson et al., 2012; Santangelo, 2015). Weather may also affect manufacturing production directly through its impact on factors of production. For example, an increase in temperature may reduce production through a reduction in the health or physical/cognitive ability of workers and managers, through an increase in absenteeism due to avoidance behaviour (Mackworth, 1946, 1947; Kenrick and McFarlane, 1986; Hsiang, 2010; Cachon et al., 2012a; Adhvaryu et al., 2015; Burgess et al., 2014a; Somonathan et al., 2015; Heal and Park, 2014; Graff Zivin and Neidell, 2014; Graff Zivin et al., 2015). Heavy rainfall may affect workers' ability to get to work (Bandiera et al., 2015b), disrupt supply chains. In addition, increased temperature, or a reduction in rainfall in areas dependent on hydroelectric power generation, is likely to put additional stress on an already fragile electricity infrastructure, reducing the supply of electrical power (Ryan, 2014; Alcott et al., 2015). Increases in temperature or reductions in rainfall may increase groundwater use, resulting in competition for water between agriculture and industry (Keskin, 2010). Finally, capital stocks and flows may be affected if weather affects capital depreciation, the relative productivity of inputs, or the level of investment in the economy if capital is locally constrained (Jina and Hsiang, 2015; Asher and Novosad, 2014).

of India's labour regulation environment that builds on Besley and Burgess (2004), who classify the rigidity of the labour market environment using state-level amendments to the Industrial Disputes Act of 1947 (hereafter IDA). In rigid labour market environments, firms face significant hiring and firing costs that, I argue, diminish the incentive to hire workers in response to transitory changes in the availability of labour (Oi, 1962; Nickell, 1978; Bentolila and Bertola, 1990; Hamermesh, 1993; Heckman, 2003; Besley and Burgess, 2004; Haltiwanger et al., 2008; Ahsan and Pagés, 2009; Adhvaryu et al., 2013; Amirapu and Gechter, 2014; Chaurey, 2015). By contrast, these costs are significantly lower in flexible labour market environments, where firms have more power over hiring decisions. However, this is not sufficient to identify the effects of labour reallocation. There may be many differences across space that could conflate the estimated effects of weather on manufacturing outcomes based on these differences. Consequently, I introduce firm-level exposure to the labour regulation environment, based on chapter 5b of the Industrial Disputes Act, which specifies the size that firms can reach before they are regulated under the IDA. For the identification strategy to have any viability, the other effects of weather must be constant across labour regulation environments.

First, I examine the effects of temperature on unregulated firms, directly testing the assumption that the other channels of weather are constant across labour regulation environments. I estimate that the net effects of temperature on unregulated firms are statistically indistinguishable across labour regulation environments.[4]

Second, I examine the effects of temperature on regulated firms. I estimate that, in rigid labour market environments, an increase in temperature is associated with a negative impact on firm productivity, the average wage of permanent workers, the number of casual workers hired, and the number of items the firm produces, consistent with – but not limited to – an emerging literature that suggests that increases in temperature have significant effects on labour productivity through a drag on physiological and cognitive ability.[5] However, in flexible labour market environments, I estimate that firms experience a relative increase in employment and output, with new entrants moving into casual manufacturing activities, offsetting the direct effects of temperature. These effects provide support for the premise that firms in flexible

[4]In further support of this identification strategy, I observe some (though limited) bunching in the firm-size distribution to the left of the regulatory threshold in the rigid labour market environments, but not in the case of the flexible labour market environments. Furthermore, when looking at the differential effects of temperature across labour regulation environments on GDP at the district level, based solely on spatial variation, I find no significant difference between flexible and rigid states in any other sectors, suggesting that, at a minimum, the net effect of other policy variation, heterogeneous weather effects and general equilibrium considerations wash out across labour regulation environments.

[5]Mackworth (1946, 1947); Kenrick and McFarlane (1986); Hsiang (2010); Cachon et al. (2012a); Adhvaryu et al. (2015); Somonathan et al. (2015); Heal and Park (2014); Graff Zivin and Neidell (2014); Graff Zivin et al. (2015).

labour market environments are more able to absorb workers in response to agricultural productivity shocks. Interestingly, I also estimate a relative increase in the average wage of permanent workers, manufacturing productivity (TFP and output per worker), and the number of items that the firm produces, suggesting that the tasks that casual workers and permanent workers engage in are complementary in production.

The absence of movement into permanent positions suggests that labour markets can be characterised, at least in the short run, as dualistic: workers earn different wages depending on the type of employment activities in which they engage (casual vs. permanent).[6] These results are consistent with an emerging literature that explores the impact of agricultural productivity shocks on local economic activity (Hornbeck, 2012; Hornbeck and Naidu, 2014; Bustos et al., 2015; Henderson et al., 2015; Hornbeck and Keskin, 2015; Marden, 2015). However, most of the research to date has focused on permanent changes in agricultural productivity due to permanent or long-run changes in technology or the environment. By focusing on short-run changes in the weather, it is plausible that other factors of production, such as capital or the allocation of land, are held constant, allowing me to identify the effects on manufacturing outcomes of labour reallocation, rather than the collective change in factors of production. In support of this premise, I find that increases in temperature have no effect on capital, management, or the number of plants.

These results indicate that the reallocation of labour across sectors could play an important role in attenuating the economic consequences of agricultural productivity shocks, provided that the direct effects of temperature on manufacturing could be mitigated.

Counterfactual estimates examining the impact of temperature on total GDP suggest that the reallocation of labour from agriculture into the regulated formal manufacturing sector could offset aggregate losses by up to 42%, provided that the direct effects of temperature on manufacturing productivity are mitigated. Collectively, these results highlight the role that liberalising goods and labour markets can play, as well as the importance of the local policy environment, in managing the economic consequences of weather-driven changes in agricultural productivity.

The remainder of the paper is structured as follows: section 1.2 examines the relationship between weather and agricultural production; section 1.3 investigates the degree to which workers are able to move across sectors and space in response to weather-induced labour demand shocks; section 1.4 explores the impact of labour reallocation on manufacturing outcomes; section 1.5 discusses the implications of

[6]Understanding whether the differences between casual and manufacturing workers in the manufacturing sector are driven by frictions or human capital differences is beyond the scope and capacity of the data and so remains an important question for future research.

these results, considering the degree to which labour reallocation across sectors could offset losses to agriculture; section 1.6 concludes.

1.2 The Effects of Weather on Agricultural Markets

As in many developing countries, agriculture plays an important role in India's economy. During the time period of this study – the beginning of the 21st century –, agriculture accounted for roughly 15–20% of GDP, 60–70% of land use, and 50–60% of employment – most of whom are landless labourers employed on daily contracts.

A key feature of India's agricultural landscape is its dependence on the timing and intensity of the monsoon (Rosenzweig and Binswanger, 1993).[7] Rainfall plays an important and salient role in the production of crops; the role of temperature, however, is a consideration often neglected in economic analysis. The monsoon's arrival in early summer is especially important for the kharif season, which corresponds with this period, but also for the rabi season, which begins at the end of the kharif season and continues through the cooler autumn and winter months before being harvested in the spring. Consequently, rabi yields are highly dependent on the degree to which rainfall can be stored in the soil. High temperatures prior to the monsoon affect the onset of the monsoon – a thermally driven phenomenon –, the degree to which rainfall drains from the soil, and soil temperature, which is important for seed germination and plant growth. High temperatures during the monsoon directly affect the kharif crop and increase the rate of evapotranspiration, which affects the availability of moisture in the soil, necessary for rabi crop production. Finally, high temperatures directly affect the rabi crop, even in the case in which irrigation is used.[8]

In this section I examine the effects of weather on two sets of agricultural outcomes. First, I examine the degree to which weather affects agricultural production in India, identifying the sign and magnitude of this relationship. Second, I examine the effects of weather on agricultural prices. This provides an insight into the expected response of labour following a change in agricultural productivity. A priori, it is ambiguous as to whether a reduction in agricultural productivity will result in an increase or decrease in the demand for labour. In a state of autarky, a reduction in agricultural production will result in an increase in prices as supply falls. Jayachandran (2006) shows that if workers have an inelastic labour supply and face a binding subsistence constraint for food (only relevant in the absence of trade), then a reduction

[7] Less than 30% of cultivated land is irrigated.

[8] While temperature is an important determinant of vapour pressure deficit, which irrigation can alleviate, around one third of the effects of temperature on yield losses arise due to an increase in the pace of crop development, which provides less time for the plant to develop and absorb nutrients and calories (Schlenker and Roberts, 2009). Fishman (2012) demonstrates these effects in the context of India by showing that higher temperatures still have a direct effect on rice yields – a crop known to be naturally resistant to higher temperatures – after controlling for irrigation.

in agricultural production will result in an increase in agricultural labour. Furthermore, an increase in prices could reduce the consumption base of the local economy, reducing demand for other commodities (Henderson et al., 2012; Soderbom and Rijkers, 2013; Santangelo, 2015). By contrast, if the local economy is open to trade, then consumption and production are separable. In a state of autarky, agricultural surplus is necessary for the movement of workers into non-agricultural production, as the local economy is responsible for feeding itself. Only when enough food is produced can the economy focus on producing other products. However, when an economy is open to trade, the local economy does not need to produce food itself. Food can be imported and paid for by the export revenues of other commodities. Consequently, instead of rural prosperity fuelling the movement of workers out of agriculture – the historical norm for many developed countries –, rural deprivation pushes workers out of agriculture into other sectors of the economy. In this case, prices in the local economy are exogenous, set on the global market. Consequently, a local change in production will not affect the price of tradable products, but will result in a change in local comparative advantage. Appendix A.1 presents a simple model based on Matsuyama (1992) demonstrating how the comparative statics vary based on market integration. By understanding the responsiveness of prices to changes in the weather, we can gain an insight into the degree to which Indian districts are integrated into other markets, either national or international, allowing us to postulate the direction in which labour might move following a change in agricultural productivity.

1.2.1 Data – Yields and Prices

Data on crop yields and farm-gate prices come from the ICRISAT Village Dynamics in South Asia Macro-Meso Database (henceforth VDSA), which is a compiled from a number of official government data sources. The data analysed cover 12 major crops across 302 districts in 19 states between 1960 and 2009.[9] For comparability with the other datasets I restrict my attention to the period 2001–2007. For each crop and district, the data provide the total area planted, total production in tonnes, and farm-gate prices. It is straightforward to calculate yields as total production divided by total area planted. I also calculate the value of production, defined as price multiplied by yield. Prices, by crop, are deflated to 2001 Rs. Panel A of Table 4.1 provides summary statistics for the VDSA data. It is interesting to note that even in the raw data there is very little between-district variation in prices.

[9]The 12 crops are Barley, Cotton, Finger Millet, Groundnut, Linseed, Maize, Pearl Millet, Rice, Rape and Mustard Seed, Sorghum, Sugarcane, and Wheat.

1.2.2 Data – Rainfall and Temperature

Rainfall and temperature data are collected from the ERA-Interim Reanalysis archive, which provides 6-hourly atmospheric variables for the period on a $0.25° \times 0.25°$ quadrilateral grid. Daily variables are calculated for each district centroid using inverse distance weighting from all grid points within 100km. The weight attributed to each grid point decreases quadratically with distance.[10] Although India has a large system of weather stations that provide daily readings dating back to the 19th century, the spatial and temporal coverage of ground stations that report temperature and rainfall readings has sharply deteriorated over time. Furthermore, there are many missing values in the publicly available series. If we were to base the construction of this data on a selection rule that requires data for 365 days of the year, the database would have very few observations. Reanalysis data provides a solution to these issues and to endogeneity concerns related to the placement of weather stations, variation in the quality of data collection, and variation in the quantity of data collected. By combining observational data, from ground stations and remote-sensing products (satellites), with global climate models, a consistent best estimate of atmospheric parameters can be produced over time and space (Auffhammer et al., 2013a). This results in an estimate of the climate system that is separated uniformly across a grid, that is more uniform in quality and realism that observations alone, and that is closer to the state of existence than any model could provide alone. This type of dataset is increasingly being used by economists, especially in developing countries, where the quality and quantity of weather data is limited (see Guiteras (2009); Schlenker and Lobell (2010); Burgess et al. (2014a); Kudamatsu et al. (2014a); Alem and Colmer (2015); Colmer (2015a).[11] Panel D of Table 4.1 provides summary statistics for the ERA-Interim Reanalysis Data.[12]

1.2.3 Empirical Specification – Yields and Prices

The unit of observation in this analysis is the crop $\times$ district level. In 2001, the average district population was 1.75 million and the average area was 5,462 km^2 (Census of

[10] The results are robust to alternative methods of construction, including: linear weights; cubic weights; the simple average of each point in the district; the average of each point in the district weighted by the area share of cultivated land; and the average of each point in the district weighted by population. Measures based on averages result in a smaller sample size, as some districts do not contain a data point and require the inverse distance weighting procedure.

[11] The results are broadly robust to additional rainfall and temperature datasets from both satellite (TRMM) and ground station (UDEL) sources.

[12] Further details on all data sources are available in appendix A.2.

India, 2001).[13] The main empirical specification for estimating the effect of weather on agricultural outcomes is based on the following model,

$$\log Y_{cdt} = f(w_{dt}) + \alpha_{cd} + \alpha_{ct} + \phi_s t + \varepsilon_{cdt}$$

where: Y_{cdt} represents the outcome of interest – yields, the value of production, or farm-gate prices; α_{cd} is a vector of crop $\times$ district fixed effects; and α_{ct} is a vector of crop $\times$ year fixed effects, absorbing all unobserved time-varying differences in the dependent variable that are common across districts. However, the assumption that shocks or time-varying factors are common across districts is unlikely to be valid, so I also include a set of flexible, state-specific time trends, $\phi_s t$.

The last term is the stochastic error term, ε_{cdt}. I follow the approach of Hsiang (2010) by assuming that the error term ε_{dt} is heteroskedastic and serially correlated within a district over time (Newey and West, 1987) and spatially correlated across contemporaneous districts (Conley, 1999). For each result I loop over all possible distances up to 2000km, selecting the parameter value that maximises the standard errors. I then repeat this exercise for serial correlation, consistently resulting in a kernel of 1 year.[14]

$f(w_{dt})$ is a function of rainfall and temperature. In the most basic specification, $f(w_{dt})$ is modelled as a function of daily average temperature and total rainfall:

$$f(w_{dt}) = \beta_1(Temperature_{dt}) + \beta_2(Rainfall_{dt})$$

Total rainfall is calculated for each state's monsoon period, beginning with the first month in which total monthly rainfall exceeds 100mm and ending with the first month that rainfall falls below 100mm.[15] As discussed, temperature is important for agricultural production both during and outside of the monsoon period. Consequently, I use crop calendars to define the relevant time period over which to construct the temperature variables. Temperature variables starting in March prior to the onset of the monsoon and ending when the crop is harvested.

[13]This is roughly twice the average area of a U.S. county (2,585 km^2) and nearly 18 times greater than the average population of a U.S. county (100,000). When compared to commuting zones and labour market areas in the U.S. – developed because county boundaries are not considered adequate confines for an area's local economy and labour market –, Indian districts are approximately 4 times the population size (401,932) and around half the area (11,396 km^2).

[14]Results are also robust when standard errors are clustered at the state level. Fisher et al. (2012) report that clustering at the state level in the U.S. provides equivalent results to directly accounting for spatial correlation using the Conley (1999) standard error adjustment. The average state size in India, when compared to the United States, is roughly similar when compared to states east of the 100th meridian, the historic boundary between (primarily) irrigated and (primarily) rainfed agriculture in the United States.

[15]Results are robust to alternative definitions of the monsoon and temperature variables, accounting for non-linearities.

1.2.4 Results – Yields and Prices

In Table 1.2 we observe that a 1°C increase in temperature is associated with a 17.4% reduction in yield and the value of production. We also see that a 100mm increase in rainfall is associated with a 0.82% increase in yield and a 0.71% increase in the value of production.

It is interesting to note that, in terms of its relative contribution, a one standard deviation change in temperature is shown to have twice the effect on yields that a one standard deviation change in rainfall has, highlighting the important role that temperature plays in Indian agriculture.[16] This suggests that the importance attributed to rainfall for agricultural production in India may have been overestimated by the omission of temperature in previous work, or that, over time, farmers have become more effective in managing the effects of rainfall shocks, given the salient nature of the monsoon. In addition, rainfall is storable and can be substituted with ground water resources (manually, or through the use of irrigation systems), whereas the effects of temperature are more difficult to address, requiring heat-resistant crop varieties.

As discussed above, it is also important to understand the degree to which weather affects agricultural prices. In column 3 we observe that, on average, neither temperature nor rainfall has a significant statistical or economic effect on agricultural prices.[17] This suggests that Indian districts are integrated with other markets, and can be considered as small, open economies. Consequently, a reduction in agricultural production should result in an outflow of labourers into other sectors of the economy due to a change in local comparative advantage. The next section formally tests this hypothesis.

Columns (4-6) show that the results are robust to examining the effects of weather on the main crop in each district, defined as the crop with the largest cultivated share in each district, averaged over the period 2001–2007 to account for mean reversion.

1.3 The Effects of Weather on Employment, Wages and Migration

Given the significance of weather as a driver of short-run agricultural productivity, it is of interest to understand how these effects feed into labour market outcomes, providing insights into the consequences of weather shocks on the economic lives of agricultural workers. In this section I examine the effects of weather on wages,

[16] These results are robust across weather data sets and over an extended period of analysis dating back to the 1960s.

[17] Allen and Atkin (2015) find a similar result looking at the effects of market access on agricultural prices in India between 1960 and 2010.

employment and unemployment within districts, as well as the effects of weather on migration, examining the degree to which weather shocks in other districts affect employment and unemployment within the destination district. This provides insights into the relative importance of labour movements across and within districts in response to changes in agricultural productivity, as well as bounding local labour markets in India.

1.3.1 Data – Wages and Employment

Data on wages, employment and migration come from the National Sample Survey Organisation (hereafter, the NSS employment survey). The NSS employment survey is a nationally representative household survey which collects information on employment and wages in rural and urban areas. For the purpose of this analysis I make use of NSS survey rounds 60, 61, 62 and 64, covering 2003–04, 2004–05, 2005–06, and 2007–08. The level of analysis using the NSS data is at the district $\times$ year level. The data cover 487 districts across rural and urban areas, corresponding to the sample of districts used in the manufacturing firm analysis. I calculate the average day wage in each district, focusing separately on agricultural labourers, non-agricultural workers in rural areas, and urban workers. I also calculate sectoral employment shares for each district. The analysis focusses on four sectors, broadly defined as agriculture, manufacturing, services, and construction. Employment shares are defined as the number of employees in each aggregated sector divided by the total number of employees in each district. In addition to employment shares I look at the unemployment rate, defined as the number of unemployed workers divided by the sum of all workers and the unemployed in each district. Panel B of Table 4.1 provides summary statistics for wages and Panel C provides summary statistics for employment. Agriculture accounts for an average of 54% of local employment, with manufacturing employing 19%, services 17%, and construction 7%. Examining the differences in wages across sectors, we observe that, on average, agricultural labourers receive significantly lower wages than non-agricultural wages. Whether these differences are driven by adjustment costs, human capital differences, compensating differentials associated with sector-specific amenities, or bargaining power, is unclear; however, examining the degree to which workers move across sectors provides some insight into the degree to which adjustment costs are a first-order concern.

1.3.2 Data – Migration

Round 64 of the NSS Employment Survey contains a special schedule on seasonal migration. This provides data on the origin district of seasonal migration; however, there is no detail on the destination of seasonal migrants. Instead, the NSS reports the destination of migrants in District ℓ_o in six relevant categories: rural or urban migration within the same District (m_{oo}); rural or urban migration between Districts in the same State ($\sum_{\ell_d \neq \ell_o \in S_o} m_{od}$); rural or urban migration between States ($\sum_{S_d \neq S_o} \sum_{\ell_d \neq \ell_o \in S_d} m_{od}$). Consequently, it is necessary to predict the district of destination for seasonal migrants who migrate to different districts. To do this, I draw inspiration from Imbert and Papp (2015) and use the 2001 Indian Population Census, extracting data on migrant workers by state of last residence. For each destination district ℓ_d, I observe: the number of migrant workers from the same district (M_{dd}); the number of migrant workers from other districts in the same state ($\sum_{\ell_o \neq \ell_d \in S_d} M_{do}$); the number of migrant workers from districts in other states ($\sum_{S_o \neq S_d} \sum_{\ell_o \neq \ell_d \in S_o} M_{do}$). I combine these data to estimate seasonal migration flows $\hat{m_{od}}$, using the following algorithm:

$$m_{od} = \begin{cases} m_{od} & \text{if } \ell_o = \ell_d \\ \frac{\sum_{\ell_o \neq \ell_d \in S_d} M_{do}}{\sum_{S_d} \sum_{\ell_o \neq \ell_d \in S_d} M_{do}} \sum_{\ell_d \neq \ell_o \in S_o} m_{od} & \text{if } \ell_o \neq \ell_d \text{ and } S_o = S_d \\ \frac{\sum_{S_o \neq S_d} \sum_{\ell_o \neq \ell_d \in S_o} M_{do}}{\sum_{S_d} \sum_{S_o \neq S_d} \sum_{\ell_o \neq \ell_d \in S_o} M_{do}} \sum_{S_d \neq S_o} \sum_{\ell_d \neq \ell_o \in S_d} m_{od} & \text{if } \ell_o \neq \ell_d \text{ and } S_o \neq S_d \end{cases}$$

I deviate from Imbert and Papp (2015) in two respects. First, by using migrant workers rather than the total population of permanent migrants. Second, by broadening my attention beyond urban destinations. Non-agricultural production is not restricted to urban areas, and so rural–urban migration is not the appropriate characterisation of migration flows in the context of this paper. Indeed, a number of papers provide evidence to suggest that non-agricultural production in India is decentralising, from urban to peri-urban and even rural areas, taking advantage of cheaper labour and vastly cheaper land prices (Ghani et al., 2012; Desmet et al., 2015; Colmer, 2015b). These adjustments provide stronger support for the identification assumption, on which this approach relies: that the proportion of NSS seasonal migrants who go from district ℓ_o to district ℓ_d, either in the same state or between states, is the same as the proportion of census migrant workers in district ℓ_d who come from another district ℓ_o, either in the same state or between states.

In Table 1.5 we observe that rural-origin migrants comprise the bulk of migration flows, accounting for nearly 90% of all seasonal migration. We observe that the manufacturing sector is the destination industry from rural areas for around 15% of within-district rural migrants, 5% of within-State rural migrants, and 21% of between-State rural migrants. However, most strikingly, we observe that there is very little seasonal migration in India – an observation that has been highlighted by a number of papers (Foster and Rosenzweig, 2008; Munshi and Rosenzweig, 2015; Morten, 2013).

The data presented here provide a number of insights, with implications for the effects of localised shocks in India. First, given the limited spatial mobility observed, workers are likely to be limited in their ability to mitigate the economic consequences of agricultural productivity shocks by moving across space. Second, this implies that sectoral shocks are likely to have a bigger effect on other sectors in the local economy, as employment adjustments are less diversified across space. Finally, this implies that localised productivity shocks elsewhere are unlikely to have a large effect on economic outcomes across space; however, the validity of this argument is decreasing as the spatial correlation of localised productivity shocks increases, and as the importance of a specific location for the supply of workers increases. I test this prediction by examining the effects of localised temperature shocks in origin districts on employment and wages in destination districts to understand the degree to which localised productivity shocks propagate through labour markets across space. As a consequence of this exercise it will be possible to identify the boundaries of economic activity in India, identifying local labour markets.

1.3.3 Empirical Specification – Employment, Wages, and Migration

In analysing the effect of weather on employment, wages, and migration, the unit of analysis in this section is at the district level. The main empirical specification for estimating the effect of weather on labour market outcomes is based on the following model,

$$Y_{dt} = f(w_{dt}) + \alpha_d + \alpha_t + \phi_s t + \varepsilon_{dt}$$

where: Y_{dt} represents the outcome of interest – sectoral employment shares and the log of average wages; α_d is a vector of district fixed effects, absorbing all unobserved district-specific time-invariant variation in the dependent variables; and α_t is a vector of crop $\times$ year fixed effects, absorbing all unobserved time-varying differences in the dependent variable that are common across districts. I also include a set of flexible, state-specific time trends, $\phi_s t$.

As in the analysis on agricultural outcomes, $f(w_{dt})$ is a function of rainfall and temperature. In the most basic specification, $f(w_{dt})$ is modelled as a function of daily average temperature measured over the agricultural year, and total rainfall measured over the state-specific monsoon period.

The last term is the stochastic error term, ε_{dt}. Standard errors are adjusted as in section 1.2.3.

The specification examining the degree to which productivity shocks in foreign districts affect labour market outcomes locally through migration differs slightly.

Using the bilateral migration flows discussed in section I construct a spatial weights matrix summarising the migratory relationship between each district.[18] As mentioned, migration flows between ℓ_o and ℓ_d produce and $o \times d$ matrix $\mathbf{M}_{o \times d}$,

$$\mathbf{M} = \begin{pmatrix} m_{11} & m_{12} & \cdots & m_{1D} \\ m_{21} & m_{22} & \cdots & m_{2D} \\ \vdots & \vdots & \ddots & \vdots \\ m_{D1} & m_{D2} & \cdots & m_{DD} \end{pmatrix}$$

Each weight m_{do} reflects the contribution of migration flows from district o to district d. In the case that all migration is spread equally between all districts, each entry in $M_{o \times d}$ will be equal to $1/d$. At the other extreme, the case in which all migration occurs within districts provides an identity matrix. Based on the data, migration patterns in India tend towards the identity matrix extreme, far from an equal distribution of migrants.

To identify the degree to which local labour demand shocks affect economic outcomes in destination sectors, I weight temperature and rainfall variation by the bilateral migration matrix, examining the migration-weighted effects of weather in district o on economic outcomes in district d through migration. The estimating equation is specified as follows,

$$Y_{dt} = \beta f(w_{dt}) + \gamma \sum_o \left[\frac{m_{od}}{M_d} f(w_{ot}) \right] + \alpha_d + \alpha_t + \phi_s t + \varepsilon_{dt}$$

where: Y_{dt} represents sectoral employment and unemployment shares of the labour force in destination district d; α_d is a vector of district fixed effects; α_t is a vector of year fixed effects; $\phi_s t$ a set of state-specific time trends.

$\sum_o \left[\frac{m_{od}}{M_d} f(w_{ot}) \right]$ captures the migration-weighted effects of weather in other districts.

By directly controlling for local weather effects, $f(w_{dt})$, to account for the correlation of weather across space, γ identifies the effects of weather variation in foreign districts on local labour market outcomes through migration.

1.3.4 Results – Wages, Employment, and Migration

I begin by examining the effects of local weather on average wages and employment shares within districts.

Table 1.3 presents the effects of temperature and rainfall on the average wage of workers in each sector within the local economy. The effect of weather on the average wage is ambiguous, as the overall effect depends on the change in composition of

[18] A similar matrix can be constructed using the straight-line distance between districts.

the workforce in each sector as well as the direct effects of temperature and rainfall changes. If, for a given level of demand, hot, dry weather reduces the supply of labour then the average wage will rise. If, for a given level of supply, hot dry weather reduces the demand for labour there is less work available and so the average wage will fall. We observe that a 1°C increase in the daily average temperature is associated with a 4% reduction in the average day wage for agricultural workers, consistent with a reduction in the demand for agricultural labour. This may be a function of a supply and demand effect if workers are less willing to work in the heat, counteracting the reduction in the average wage. However, it is clear that the demand effect dominates. While this acts as an insurance mechanism for farm owners, a reduction in the average wage combined with a reduction in the availability of work – on the intensive or the extensive margin – could have significant welfare effects on agricultural workers if they are limited in their ability to find other work. Interestingly, we see that rainfall does not have any significant statistical or economic effect on the average day wage, and is robust across other weather datasets. As discussed, rainfall is estimated to have less of an effect on agricultural production and so the impact on agriculture may not be significant enough to affect labour market outcomes. Associated with this consideration, it may simply be the case that, due to the relatively short time-series, there is not enough power to identify these effects.

In addition to the effects on agricultural wages, we observe that an increase in temperature is associated with a 7.5% reduction in the average day wage in manufacturing. This suggests either that temperature has a significant direct effect on the productivity of workers in manufacturing, or that agricultural workers are able to move across sectors in response to temperature increases. To understand the degree to which the latter channel may be relevant, I estimate the effects of weather on employment and unemployment shares, identifying the degree to which workers are able to move across sectors, and find jobs, in response to reductions in the demand for agricultural labour.

Figure 1.1 presents semi-parametric estimates of the relationship between temperature and the labour force share in agricultural and non-agricultural activities as well as unemployment. Semi-parametric approaches obtain local estimates of these relationships and display them visually to gain insights into the degree to which there may be non-linearities. Consistent with the inferences drawn from the effects of weather on agricultural prices, we observe that, as temperatures increase, the labour force share engaged in agriculture declines significantly. However, most interestingly, this is completely offset by an increase in the labour force share engaged in non-agricultural activities. These effects appear linear and symmetric, suggesting that workers are able to find jobs rather than become unemployed. This is further supported by the observation that the labour force unemployment share does not respond to increases or decreases in temperature. This suggests that workers face little impediment in the movement across sectors, especially given the contemporaneous nature of these shocks. Furthermore, it appears that the ability of sectors to absorb workers is one of the main ways that workers are able to manage weather-driven changes in agricultural productivity.

Table 1.4 presents the regression results for these effects, broken down by sectoral employment shares and unemployment. As demonstrated, we observe that an increase in the daily average temperature is associated with a significant reduction in the district share of agricultural employment (8.4%/ 1°C). Furthermore, we observe that this is offset by an increase in the share of manufacturing employment (5.54%/1°C) and a smaller increase in the share of services employment (3.89%/1°C).[19] Most importantly, we observe that there are no changes in unemployment, suggesting that workers are relatively unconstrained in their ability to move across sectors in response to temperature shocks. Again, rainfall is shown to have no significant effect on changes in the composition of employment in the local economy.[20] This is consistent with the premise that temperature has a greater effect on agricultural production than does rainfall – a premise that has found support in a number of other recent studies, emphasising the importance of temperature variation over rainfall as a driver of economic outcomes (Hsiang, 2010; Dell et al., 2012; Gray and Mueller, 2012; Burgess et al., 2014a; Mueller et al., 2014; Burke et al., 2015).

In addition to looking at the effects of weather on local economic activity, I also examine the effects of weather on migration. This allows me to examine whether short-run changes in the weather result in a reallocation of labour across space, distorting the definition of the local labour market and, consequently, the interpretation of the results, as well as being an outcome of interest in its own right.

Table 1.6 presents the results of this exercise. As before, the effects of local temperature shocks on local economic activity appear to be unchanged after controlling for the migration-weighted weather effects. In addition, I find that the migration-weighted weather effects have no impact on employment shares in destination markets, indicating that there is little migration across districts in response to temperature increases. Consequently, local labour markets in India can be bounded at the district level. The reason behind the limited migration remains unclear; however, the ability of other sectors to absorb workers in response to agricultural productivity shocks somewhat mitigates the need to move across space.

Collectively, the results presented in this section suggest that workers in India are relatively able to move across sectors in response to transitory labour demand shocks.

[19]Results available upon request show that the movement into manufacturing is driven by men, whereas the movement into services is driven by women.

[20]This result is robust across alternative weather data sets.

Under the assumption that the wage effects are driven solely by reductions in labour demand, the estimated labour demand elasticity is $\epsilon_{L_a w_a} = 2.05$, suggesting that labour demand in agriculture is very elastic. Looking at the cross-elasticity of labour demand, the results suggest that a 1% reduction in the agricultural wage is associated with a 1.35% increase in the manufacturing share of employment. These estimates indicate that labour is highly responsive to changes in the agricultural wage.

1.4 The Effects of Weather on Manufacturing

Having established that agricultural workers are able to move across sectors in response to weather-induced changes in agricultural productivity and that the manufacturing sector is a major absorber of these workers, it is of interest to understand what these workers do and how they affect the productivity of firms and the labour market outcomes of incumbent workers.

1.4.1 Data – Manufacturing Plants

Manufacturing plant microdata come from the Annual Survey of Industries (ASI) collected by the Ministry of Statistics and Program Implementation (MoSPI), Government of India. The ASI covers all registered industrial units that employ 10 or more workers and use electricity, or employ at least 20 workers and do not use electricity. The ASI frame is divided into two schedules: the census schedule, which is surveyed every year, and the sample schedule, which is randomly sampled every few years. The ASI has a much wider coverage than other datasets, such as the Census of Manufacturing Industries (CMI) and the Sample Survey of Manufacturing Industries (SSMI), and is comparable to manufacturing surveys in the United States and other industrialised countries. However, the ASI does not cover informal industry that falls outside the Factories Act, 1948. The formal sector accounts for approximately two-thirds of manufacturing output in India and is therefore not representative of all manufacturing activities. It is, however, representative of tradable manufacturing in India, since the informal sector trades very small volumes, if at all. See appendix A.2 for more details on the ASI data preparation. The data cover an average of 20,456 firms in 487 districts (defined using 2001 boundaries) that have positive agricultural output, in 22 states, observed between 2001 and 2007. This results in 143,197 firm-year observations.

The outcomes of interest are the log of total output, employment, and the average day wage (defined as the total wage bill/the total number of man days worked during the year). Employment outcomes are examined for both permanent (non-managerial) workers and contract workers. The distinction between contract workers and permanent workers is important for this analysis. Contract workers are on casual contracts and so a priori are the type of worker that one would expect to move between the agricultural and manufacturing sector.

Within the formal manufacturing sector, permanent workers in regulated firms earn, on average, 1.68 times more than contract workers. Consequently, there is greater common support between the wages of contract and agricultural workers within district – the average contract wage is 2.15 times the average agricultural wage. By contrast, the average permanent worker wage is 3.83 times greater than the average agricultural wage.

In addition, using worker-level data from the NSS (discussed in section 1.3), I run individual-level mincerian wage regressions to estimate the size of wage gaps after controlling for education, age, gender and district fixed effects. Table 1.7 shows that there is a significant wage gap between permanent manufacturing workers and agricultural workers, with permanent manufacturing workers earning 1.60 times more than agricultural workers.[21] Furthermore, we observe that the average wage gap between casual manufacturing workers and agricultural workers almost disappears, with casual manufacturing workers earning 13% more than agricultural workers. This suggests that labour markets can be characterised as dualistic: workers earn different wages depending on the activities in which they engage. However, labour markets do not appear to be dualistic across sectors (agriculture vs. non-agriculture) but rather in terms of the type of employment activities in which workers engage (casual vs. permanent).

In addition to the outcome variables described above, I construct two measures of productivity. The first is a simple measure: output per worker. While this is a crude measure of productivity, it provides a relatively useful measure of the average labour productivity of the firm. The second measure is an estimate of total factor productivity. Appendix A.2 provides an explicit model of TFP, in the context of a profit-maximising firm, that I use to construct my empirical estimates.

1.4.2 The Labour Regulation Environment in India

The combination of manufacturing microdata with meteorological data provides the basis of the main empirical analysis. However, this is not sufficient to identify how the movement of labour out of agriculture affects productivity and wages in the manufacturing sector. The key empirical challenge relates to the fact that, while weather is an important driver of short-run agricultural productivity, there are potentially many empirically relevant channels through which weather could affect manufacturing outcomes. Consequently, any estimate of the reduced form elasticity between weather

[21] This data includes employment in both the informal and formal sector.

and the outcomes of interest will provide the net effect of all empirically relevant channels.

To try to address this challenge, I set out to identify the labour reallocation channel, net of the remaining empirically relevant channels, by exploiting variation in the propensity of firms to absorb workers in response to transitory weather shocks. To do this, I exploit spatial variation and firm-level exposure to India's labour regulation environment. Regulated firms in rigid labour market environments should have little incentive to hire workers in response to year-to-year changes in the weather. By contrast, in more flexible labour market environments, regulated firms should be more able to absorb workers out of agriculture in response to temperature increases.

Industrial regulation in India has mainly been the result of central planning; however, the area of industrial relations is an exception to this, providing spatial variation in firms' incentives regarding the hiring and firing of workers following transitory changes in labour demand. The key piece of legislation used to measure state-level variation in sectoral mobility is the Industrial Disputes Act of 1947 (hereafter the IDA). The IDA regulates Indian Labour Law concerning trade unions, setting out conciliation, arbitration, and adjudication procedures to be followed in the case of an industrial dispute, and was designed to offer workers in the formal manufacturing sector protection against exploitation by employers. Up until the mid 1990s, the IDA was extensively amended at the state level, resulting in spatial variation in labour market rigidities. Besley and Burgess (2004) use these extensive state-level amendments (113 in total) to construct a measure of the labour regulation, environment studying its impact on manufacturing performance and urban poverty. By examining the amendments made in each state over time, states are coded as either neutral, pro-worker, or pro-employer. A pro-worker amendment is classified as one that decreases a firm's flexibility in the hiring and firing of workers; Pro-employer amendments are classified as increasing a firm's flexibility in hiring and firing. Importantly, neither the timing nor the direction of amendments is correlated with the weather.

The cumulation of these scores over time determines the state's labour regulation environment. Consequently, West Bengal, Maharashtra and Orissa are assigned as pro-worker states (rigid). Rajasthan, Tamil Nadu, Karnataka, Kerala and Andhra Pradesh are assigned as pro-employer states (flexible). The remaining states are assigned as neutral. This assignment captures spatial variation in the propensity of firms to take advantage of transitory labour supply changes arising from year-to-year changes in agricultural productivity.

However, state-level variation is not sufficient to identify the labour reallocation channel, as it may simply capture the heterogeneous effects of weather, general equilibrium effects, or other state-level variation, confounding the interpretation of the estimated coefficients. I therefore combine this spatial variation with firm-level exposure to the regulation based on chapter 5b of the IDA, which specifies the size that firms can become before the IDA has a binding effect. The firm-size threshold is 50 in West Bengal, 300 in Uttar Pradesh, and 100 elsewhere.[22]

A further consideration is whether the workers moving out of agriculture are likely to be affected by the IDA. A priori we would expect these workers to enter the regulated formal manufacturing sector as casual contract workers. This raises an important question about the degree to which the labour regulation environment impacts the employment of casual workers. Contract workers are not directly considered as workmen under the IDA and, consequently, are not de jure regulated within manufacturing firms. However, this does not mean that contract workers are not affected by the IDA (Bertrand et al., 2015; Chaurey, 2015). Contract workers are still de jure regulated by the IDA under the contractor that hires them. Consequently, the availability of these workers to firms in rigid labour market environments may be directly affected by the willingness of contractors to put these workers on the books in response to transitory changes in the weather. In addition, contract workers may be de facto affected by the IDA within firms. On the one hand, the exemption of contract workers from the IDA may provide an added incentive to hire contract workers in rigid labour markets, allowing employers to bypass some of the regulations in the IDA. If so, this would imply that the labour reallocation channel would be relatively larger in rigid labour market environments. Looking at the data, one observes, consistent with this argument, that the share of firms using contract workers – an extensive margin measure – is higher in rigid markets than in flexible markets (Figure 1.2). On the other hand, the use of contract workers has been vigorously, and in some cases violently, opposed by unions and permanent workers, suggesting that firms may face significant costs associated with hiring contract workers, especially in rigid labour market environments. Furthermore, the Contract Labour Regulation and Abolition Act of 1970 prohibits the use of contract labour if the work "is done ordinarily through regular workmen in that establishment." To the degree that this is enforced, this restricts the degree to which firms can bypass the IDA. Consequently, in rigid labour market environments, where it is expected that firms have to negotiate with unions over decisions that affect the labour force, the hiring of contract workers, in response to transitory changes in labour availability, may be restricted. In support of this premise we observe, on the intensive margin, that the share of workers employed as contract workers is higher in flexible labour market environments, suggesting that, conditional on hiring contract workers, firms in more flexible markets are able to hire more casual workers than firms in rigid labour markets (Figure 1.3).

[22] Results are robust to applying a uniform threshold across all states away from the regulated threshold, mitigating concerns that the results could be driven by the movement of firms around the size threshold.

Given that, on average, there is no difference in the total number of workers hired by firms across labour regulation environments, this implies that there is a higher proportion of contract workers in flexible than in rigid labour market environments (Table 1.8). In practice, whether there is a differential propensity to hire more casual workers in rigid or in flexible labour market environments is an empirical question. Most importantly, it does not affect the identification of the labour reallocation effect, which simply requires that there be a differential effect across labour regulation environments. More pertinent to the empirical focus of this paper, is whether the types of task in which movers engage in, when working in the manufacturing sector are substitutable or complementary to incumbent manufacturing workers. In this respect, the paper provides insights into the relationship between the average new entrant and average incumbent worker in the manufacturing sector.

Table 1.8 presents difference-in-means tests for firm-level characteristics across rigid and flexible labour market environments. We observe that, on average, there is no difference in the output, number of items produced, productivity, capital stock, or number of workers across labour market environments; however, there is a slightly higher share of of contract employment in flexible states, and the average wage of permanent workers is higher in rigid states. While I cannot rule out differences in the types of firm across labour regulation environments, this suggests that there are not significant first-order differences. A further concern could be that the elasticity of supply varies across the labour regulation environments, such that the availability of workers in response to increases in temperature differs across labour regulation environments. However, I observe no differences in the share of agricultural employment, agricultural GDP, total GDP, or transportation infrastructure, indicating again that this is not a first-order concern.

1.4.3 Empirical Specification – Manufacturing Outcomes

To identify the sign and magnitude of the labour reallocation channel, I interact the net effects of weather with a measure of the labour regulation environment, splitting the sample at the regulatory firm-size threshold. The estimation equation for both samples is written as follows,

$$\log Y_{ijdt} = \beta f(w_{dt}) + \gamma f(w_{dt}) \times \textsc{Flexibility} + \alpha_{jd} + \alpha_{jt} + \phi_s t + \varepsilon_{ijdt} \qquad (1.1)$$

The dependent variable, Y_{ijdt}, is the natural log of: total output (sales), the number of items the firm produces, employment (by worker type), the average day wage (by worker type), and the two measures of productivity described above – output per worker and measured TFP. The unit of analysis is firm i, in sector j, in district d, at time t.

District $\times$ industry (α_{jd}) fixed effects absorb all unobserved time-invariant variation within these dimensions; industry $\times$ year (α_{jt}) fixed effects control for sector-specific time-varying differences in the dependent variable that are common across districts; and a set of flexible state-specific time trends ($\phi_s t$) relaxes the assumption that shocks or time-varying factors that affect the outcome variables are common across districts.

As in the previous sections, $f(w_{dt})$ is a function of rainfall and temperature,

$$f(w_{dt}) = \beta_1(Temperature_{dt}) + \beta_2(Rainfall_{dt}) \tag{1.2}$$

where total rainfall is measured over the state-specific monsoon period and the daily average temperature is measured over the agricultural year.[23]

The problem associated with estimating the simple linear regression model absent the interaction term is that β captures the sum of all empirically relevant channels through which temperature affects the manufacturing outcomes.

The interaction term, $f(w_{dt}) \times$ FLEXIBILITY, captures the propensity of firms to absorb workers in response to negative agricultural productivity shocks. The main specification allows for a continuous measure of the labour regulation environment, based on Besley and Burgess (2004), bounded between 0 and 1. West Bengal is the baseline state, coded 0 as it is the most rigid labour regulation environment. Andhra Pradesh and Tamil Nadu are coded as 1, as they are the most flexible labour regulation environments.

The last term is the stochastic error term, ε_{dt}. Standard errors are adjusted as in section 1.2.3.

1.4.4 Results – Manufacturing

I begin by examining the effects of temperature on unregulated firms, providing a direct test for the identification assumption that the other empirically relevant temperature effects are constant across labour regulation environments.

Below the regulatory threshold, there should be no differential impact of temperature on firms across labour regulation environments. As a result, these estimates do not disentangle the labour reallocation effect, but rather test important identification assumption: that any additional channels through which weather could affect manufacturing outcomes are constant across labour regulation environments, i.e., if increases in temperature makes workers less productive, the effect in West Bengal is the same as in Andhra Pradesh. This also tests for the presence of any additional

[23] In appendix A.4 I test for alternative specifications, accounting for non-linearities in the temperature schedule.

spatial differences such as general equilibrium effects or other policy differences that are correlated with the spatial dimension of the labour regulation environment.

Table 1.9 presents results that provide direct evidence that this is the case. In Panel A we observe that an increase in temperature is associated with a net increase in manufacturing output (6.40%/1°C) and employment (3.72%/1°C), consistent with the results in section 1.3. However, most importantly, we observe in Panel B that there are no differential effects of temperature between labour regulation environment, providing support for the identification of the labour reallocation effect in regulated firms. I provide further support for the identification strategy in appendix A.3, where I show that there is some – though limited – evidence of bunching in the raw data just below the firm-size employment threshold for rigid labour market environments, but not for flexible labour market environments. In addition, I show that there is no differential effect of weather across labour market environments on other sectors of the economy, providing further support for the premise that there are no significant spatial differences between labour regulation environments that affect unregulated sectors of the economy. Furthermore, I also show that the results are robust to using a uniform size threshold across all states, away from the official thresholds, to mitigate any endogeneity concerns relating to selection around the size threshold.[24]

Table 1.10 presents results that compare the effects of temperature on regulated firms across labour regulation environments. In light of the results in Table 1.9, the interaction term provides plausible identification of the sign and magnitude of the labour reallocation effect on the manufacturing outcomes of regulated firms.

In Panel A we observe that the effect of temperature on regulated firms is zero, indicating either that temperature has no effect on the manufacturing outcomes of regulated firms, or that there are multiple competing effects. Panel B demonstrates that this zero is a net effect. In rigid labour market environments I estimate that effects of temperature on manufacturing outcomes are significant and negative, affecting labour productivity (-8.77%/1°C), measured TFPR (-8.78%/1°C), the average wage of permanent workers (-6.91%1°C), the employment of casual workers (-12.1%/1°C) and the number of items produced (-5.6%/11°C). This is consistent with, though not limited to, an expanding literature which suggests that high temperatures may have a direct negative effect on labour productivity (Mackworth, 1946, 1947; Hsiang, 2010; Graff Zivin and Neidell, 2014; Adhvaryu et al., 2015; Graff Zivin et al., 2015; Somonathan et al., 2015).

While it is beyond the scope of this paper, and indeed the capacity of the data, to identify the precise mechanisms through which this remaining net effect has an effect

[24]In practice these concerns would downward bias the effects of labour reallocation, as the inflow of workers to firms below the threshold could push them over the regulatory threshold reduce the average number of employees within regulated manufacturing firms.

on manufacturing outcomes, one thing is clear: if these effects can be mitigated, the realised impact of temperature on manufacturing output through labour reallocation will be significantly larger, offsetting the economic losses associated with temperature increases in agriculture.

Examining the relative effect of temperature in flexible labour market environment, compared to the effects of temperature in rigid labour market environments, the interaction term, Temperature $\times$ Flexibility shows that an increase in temperature is associated with a relative increase in total output (15%/1°C) and the employment of contract workers (17.1%/1°C). This is consistent with the premise that firms in more flexible labour market environments have a greater capacity to absorb workers in response to weather-driven changes in agricultural productivity. The hiring of movers into contract (casual) worker positions rather than permanent worker positions is consistent with anecdotal evidence and a priori reasoning. Related to the discussion of the de facto impact of the labour regulation environment on contract workers, we observe that there is a relative increase in the employment of contract workers in more flexible labour markets, suggesting that firms in more rigid labour markets still face restrictions in the hiring of contract workers, at least in response to short-run changes in the availability of workers.

From the workers' perspective, it is reasonable to suppose that agricultural workers on casual contracts would be more likely to find casual work in the manufacturing sector before moving into permanent work. In addition, the presence of centralised contractors that provide firms with casual labour significantly reduces search costs for these positions compared to permanent positions. This is consistent with the evidence provided by Bryan et al. (2014) in Bangladesh, Franklin (2015) in Ethiopia, and Hardy and McCasland (2015) in Ghana, who demonstrate that there are significant search costs associated with finding employment.

From the firms' perspective, the results are consistent with the premise that manufacturing firms hire workers on casual contracts as a screening process, rather than hiring movers into permanent contracts straight away. Employers face an adverse selection problem, as they can only discern a worker's true ability after a hiring decision has been made, especially in the absence of employment histories. By using contract workers, firms can learn more about a worker's productivity before deciding whether to hire them permanently. This is consistent with the evidence provided by Heath (2015) that suggests that firms in Bangladeshi garment factories hire workers through referrals to mitigate adverse selection and moral hazard concerns. In doing so, firms can punish the referral provider if the new entrant is unproductive. Hardy and McCasland (2015) also highlight the importance of worker screening in the hiring decisions of firms in Ghana.

While there appears to be little impediment to moving across sectors within casual tasks, the absence of employment into permanent manufacturing positions suggests that casual and permanent labour markets are segmented, at least in the short run. Local labour markets in developing countries can therefore be characterised as dualistic, not in terms of sectors (agriculture vs. non-agriculture), but rather in terms of the type of employment in which workers engage (casual vs. permanent). This raises an interesting question about the degree to which casual workers face adjustment costs in the movement into permanent positions. As noted, there is a significant wage gap between casual manufacturing workers and permanent workers. However, while this gap exists, it less clear how it should be interpreted. On the one hand, wage gaps may represent adjustment costs, implying that there are arbitrage opportunities to increase productivity if these costs could be reduced – a misallocation of talent (Banerjee and Duflo, 2007; Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009; Moretti, 2011; Bryan et al., 2014; Gollin et al., 2014; Hsieh et al., 2014; Bandiera et al., 2015a; Bryan and Morten, 2015; Munshi and Rosenzweig, 2015). On the other hand, average wage gaps may simply represent differences in human capital between casual and permanent workers, with low-skilled workers selecting into casual tasks and high-skilled workers selecting into permanent tasks (Roy, 1951; Heckman and Sedlacek, 1985; Heckman and Honore, 1990; Miguel and Hamory, 2009; Beegle et al., 2011; Lagakos and Waugh, 2013; Young, 2013, 2015). This interpretation would suggest that, while average wage gaps exist, marginal productivities are equalised across activities – an efficient allocation of talent. Both of these channels may be further confounded by differences in the bargaining power or amenities across tasks. As discussed, the evidence presented so far, alongside evidence from worker-level mincerian wage regressions, suggests that adjustment costs, to the degree that they exist, are very limited across sectors within casual activities.

Unfortunately, it is beyond the scope and capacity of the data to provide inferences about the relative contribution of these channels to the wage gap between casual and permanent manufacturing workers; however, in appendix A.5 I provide an upper bound on the gains from reallocation, under the assumption that the total wage gap is driven by adjustment costs.[25] Understanding the relative importance of the role that adjustment costs play in impeding the movement of workers out of casual employment and into permanent positions remains an important area for future research.

In addition to the employment and production results, which demonstrate a consistency between the worker-level and establishment-level datasets, I examine the effect of labour reallocation on productivity and the average wage of incumbent workers in the manufacturing sector. To begin, we observe that the average wage of casual

[25]It is important to note that the lower bound is zero.

workers falls in response to labour reallocation (-9.63%/1°C). Using the estimated effect of temperature on the number of casual workers employed and the effects of temperature on the average wage of casual workers, I provide an exogenous estimate of the own-quantity elasticity of labour supply,

$$\epsilon_{w_m^c L_m^c} \propto \frac{\partial \log w_m^c}{\partial Temperature} / \frac{\partial \log L_m^c}{\partial Temperature} = \frac{\partial \log w_m^c}{\partial \log L_m^c} = -0.56 \tag{1.3}$$

These results suggest that a 1% increase in the employment of casual workers, employed out of agriculture, is associated with a 0.56% reduction in the average wage of casual workers, indicating that on average, the tasks these workers engage in are substitutable in production.

More surprisingly, I estimate that the hiring of casual workers in response to an increase in temperature is associated with a relative increase in average labour productivity (14.8%/1°C) and measured TFPR (8.54%/1°C), as well as an increase in the average wage of permanent workers (9.09%/1°C) and the number of items that the firm produces (7.73%/11°C). This suggests that firms restructure production to take advantage of the increased supply of casual workers. Given that casual and permanent labour markets are segmented, it suggests that the tasks that the casual entrants and permanent workers engage in are complementary in production. In light of this, it is possible to provide an exogenous estimate of the elasticity of substitution, σ, between the new entrants into casual positions and the incumbent permanent workers. If $\sigma < 1$ the new entrant casual workers and incumbent permanent workers engage in tasks that are complementary in the production process. If $\sigma > 1$ then these workers engage in tasks that are substitutable in the production process.

$$\sigma \propto \frac{\partial \log w_m^p}{\partial Temperature} / \frac{\partial \log L_m^c}{\partial Temperature} = \frac{\partial \log w_m^p}{\partial \log L_m^c} = 0.53 \tag{1.4}$$

These results suggest that a 1% increase in the number of casual workers, employed out of agriculture, is associated with a 0.53% increase in the average wage of permanent manufacturing workers. To the degree that new entrants out of agriculture and incumbent casual workers are substitutable in tasks, this would indicate that, on average, contract and permanent workers in the regulated Indian manufacturing sector engage in complementary production tasks.

However, one concern may be that these effects are driven by accompanying changes in other factors of production, confounding the interpretation of the results. Yet, one of the attractive features of the empirical context and identification strategy is that the movement of workers across sectors is driven by short-run changes in the weather and so one may consider that other factors of production and the technology of the firm are held fixed. Table 1.11 directly tests this consideration. I begin by looking at the effects of temperature on capital and capital depreciation.

If capital were to increase alongside labour, then it would be difficult to attribute increases in productivity and permanent worker wages to the reallocation of labour alone. Consistent with the premise that the other factors of production are held fixed, we observe that there is no change in capital or capital depreciation in response to temperature changes, and that this effect does not vary across labour regulation environments. Second, I consider the effects of temperature on the number of managers and the wages of managers. While a crude measure of the organisational structure of the firm, this allows us to test whether productivity increases were driven by organisational change or whether the increase in permanent worker wages could be driven by the extraction of rents from the firm. If this were the case then we may also expect managers to share in these rents. We observe neither an increase in the number of managers nor changes in the average wage of managers, suggesting that neither changes in management nor rent extraction appear to provide first-order explanations for the results. Finally, I explore whether the firm expands the number of plants – a proxy for entry and exit considerations that are not directly observable in the data. Again, we observe that the firm does not open or close plants in response to changes in temperature, suggesting that there are unlikely to be significant changes in the number of firms or in the market structure in response to changes in temperature. These findings suggest that the results can be interpreted as being driven by the increase in casual workers, rather than changes in other factors of production or changes in the technology or management structure of the firm. In addition to the supporting evidence presented here, a number of additional robustness tests, specification extensions and results are presented in the appendices.

The above results highlight the problems associated with the identification and interpretation of reduced-form weather results, but demonstrate the insights that can be gleaned from attempting to isolate specific channels and mechanisms through which weather can affect economic outcomes. I show that increases in temperature are associated with reductions in firm productivity in rigid labour market environments, where firms are less able to absorb workers in response to weather-driven changes in agricultural productivity. However, in more flexible labour regulation environments I estimate that increase in temperature are associated with a relative increase in manufacturing production and the employment of casual workers, resulting in a restructuring of production. Yet, this reallocation does not result in a net increase in output for regulated firms, as the direct effects of temperature are associated with a counteracting reduction in productivity. Consequently, if the direct effects of temperature on manufacturing could be mitigated, the reallocation of labour across sectors in response to weather-driven changes in agricultural productivity could significantly offset the losses to agriculture.

1.5 Counterfactual Analysis

In this section I explore what my results imply for aggregate production in India, considering how much of the losses to agriculture could be offset provided that the direct effects of temperature on manufacturing could be mitigated. Due to the competing effects of temperature on manufacturing, the net effect is zero, and so labour reallocation only offsets losses in manufacturing.[26]

Setting the direct effects of temperature to zero, I first focus on estimating how much of the losses to agriculture are offset by the labour reallocation effect, restricting my attention to regulated formal manufacturing firms. Second, I allow the estimated effects to be extrapolated to the rest of the formal manufacturing sector, treating unregulated formal manufacturing firms as having the same flexibility as Andhra Pradesh and Tamil Nadu.[27]

I begin by estimating the baseline effects of temperature on GDP, using data on sectoral GDP for each district, provided by Indicus analytics, focusing on agriculture, manufacturing, construction and services. Table 1.12 presents the results of this exercise, showing that a 1°C increase in temperature is associated with a reduction in agricultural GDP (-11.04%/1°C), a reduction in total manufacturing GDP (-4.75%/1°C), and no change in services or construction GDP. On average, a 1°C increase in temperature is associated with a 2.47% reduction in total GDP. Taking as given the estimated effects of temperature on manufacturing output for the regulated and unregulated formal manufacturing sector (presented in panel A of tables 1.9 and 1.10), the residual effect of temperature on the informal sector, necessary to induce a 4.75% reduction in total manufacturing GDP, is -22.25% (table 1.13).

First, I consider how much these losses could be offset provided that the direct effects of temperature on manufacturing could be mitigated. I split total manufacturing GDP into three components: the informal manufacturing sector (34% of GDP), the regulated formal manufacturing sector (22% of GDP) and the unregulated formal manufacturing sector (44% of GDP). I assume that the direct effects of temperature are offset for the formal manufacturing sector. This is because the informal manufacturing sector is largely non-tradable, so the effects may be driven by changes in local demand (Soderbom and Rijkers, 2013; Henderson et al., 2012, 2015; Santangelo,

[26] Below the regulatory threshold the net effect is positive indicating that there is a greater inflow of workers to these firms. As the other channels are constant across labour regulation environments (Table 1.9) the effect of labour reallocation on unregulated firms is expected to be the same as the estimated effect above the regulatory threshold, plus the net effect on unregulated firms.

[27] One concern relating to the validity of this exercise is that firms may only be hiring workers in response to the reductions in productivity associated with the direct effects of temperature. Consequently, in the absence of direct temperature effects, firms may have limited capacity to expand production in the short run. This could also explain the limited effects of rainfall on labour reallocation provided that the direct effects of rainfall on manufacturing outcomes are more limited.

2015). The effect of temperature on the unregulated formal manufacturing sector is assumed to be 6.4%/1°C (based on Panel A of Table 1.9). The effect of temperature on the regulated formal manufacturing sector is assumed to be the labour regulation environment-weighted effect from the main results taken from Panel B of table 1.10; i.e., 15%/1°C in Andhra Pradesh and Tamil Nadu and 0%/1°C in West Bengal. The effect of temperature on the informal manufacturing sector is assumed to be -21%/1°C based on the calculation described above. In this counterfactual environment, a 1°C increase in temperature is associated with a 2.20% reduction in total GDP, an 11% reduction in losses. The second counterfactual extrapolates the labour reallocation effect to the unregulated formal manufacturing sector based on the assumption that the other channels through which temperature affects manufacturing are constant across labour regulation environments. This assumption is supported by panel B of table 1.9. In this case the labour reallocation effect driven by temperature is assumed to increase by 15%/1°C, resulting in an increase in output of 21.4%/1°C. In this counterfactual environment a one degree increase in temperature is associated with a -1.43% reduction in total GDP – a 42.2% reduction in losses. Table 1.13 presents the results of these exercises. Figure 1.4 presents the distribution of each result.

These results suggest that there could be significant gains from mitigating the direct effects of temperature on manufacturing, and that the movement of workers across sectors could significantly offset the aggregate effects of temperature on local economic activity.

1.6 Conclusion

One of the salient features of economic life in developing countries is the centrality of agriculture to employment. Consequently, understanding the relationship between economic and natural systems can provide important insights into the lives of the poor, many of whom are dependent on agriculture as a source of income. While the literature has largely focused on outcomes within agricultural markets, understanding the degree to which workers are able to move across sectors in response to changes in labour demand helps to understand climatic influence on economic outcomes, the functioning of labour markets in developing countries, and the short-run behaviour of firms.

Consistent with a large literature examining the effects of weather on agricultural production, I estimate that temperature is a strong driver of short-run agricultural productivity. However, I also estimate that there are no effects on agricultural prices, consistent with a "law of one price", indicating that Indian districts are integrated with other markets. A priori, this suggests that reductions in agricultural productivity should result in an outflow of labourers into other sectors due to a local change in comparative advantage.

Consistent with this premise, I present evidence to suggest that agricultural workers in India are relatively able to move across sectors within local labour markets when temperature increases, moving chiefly into the manufacturing sector. This movement completely offsets the reduction in agricultural employment, with no increases in unemployment, or population through migration, indicating that the ability of other sectors to absorb workers is a key channel through which workers can manage agricultural productivity shocks. These results highlight the role that market integration and diversification can play in attenuating the aggregate consequences of idiosyncratic productivity shocks.

Having established these facts, I explore how the movement of labour across sectors affects economic outcomes in the formal manufacturing sector, which is representative of tradable industry in India. The principal challenge associated with identifying the effects of labour reallocation is that there are many channels through which temperature could affect manufacturing outcomes. Consequently, the effects of weather on manufacturing outcomes provides a net effect of all the empirically relevant channels, without a clear economic interpretation.

To discern the impact of labour reallocation, I interact the net effects of temperature with spatial and firm-level exposure to India's labour regulation environment, providing variation in the propensity of firms to absorb labour in response to short-run changes in labour availability.

For unregulated firms I estimate that the net effect of temperature on production and employment is positive and significant consistent with the premise that workers are able to move across sectors in response to weather-driven changes in agricultural productivity. In addition, I estimate that there are no differences in the effects of temperature across labour regulation environments indicating that any additional channels, through which temperature affects manufacturing, are constant across labour regulation environments. By contrast, for regulated firms in rigid labour market environments, an increase in temperature is associated with a reduction in productivity and the average wage of permanent workers, consistent with a literature that suggests that temperature is an important determinant of labour productivity. However, I demonstrate that there is a relative increase in the employment of casual workers, production and productivity in labour regulation environments where regulated firms have a greater propensity to absorb workers in response to weather-driven changes in agricultural productivity. Furthermore, there are no changes in capital, management, or the number of plants, suggesting that the employment of casual workers allows productive workers to engage in more productive tasks, resulting in a restructuring of production. These results support the premise that the ability of firms to absorb workers is a key channel through which workers are able to manage agricultural productivity shocks, indicating that the local policy environment can play an important role by affecting the ability of firms to absorb labour.

In considering the aggregate consequences of these effects, I consider the degree to which labour reallocation could offset losses to agriculture, provided that the direct effects of temperature on manufacturing could be mitigated. Using data on district level GDP, by sector, I estimate that an increase in temperature is associated with a significant loss in total GDP (-2.47%/1°C), driven largely by losses to agriculture. Restricting the reallocation of labour to the regulated formal manufacturing sector, I find that these losses could be offset by 11%. However, extrapolating the estimated effects of reallocation to the rest of the formal manufacturing sector, these losses could be offset by up to 42.2%. This exercise suggests that there could be significant gains to mitigating the direct effects of temperature on manufacturing, not only for the manufacturing sector but for aggregate production.

My findings have three main implications. First, regarding the labour market decisions of the poor, my results suggest that workers in agriculture are highly responsive to changes in the agricultural wage, resulting in movements across sectors within casual employment activities. The significant movement of workers across sectors within casual employment activities alongside the absence of movement into permanent manufacturing positions suggests that labour markets can be characterised as dualistic: whereby workers earn different wages depending on their type (Lewis, 1954). However, labour markets do not appear to be dualistic across sectors (agriculture vs. non-agriculture) but rather in terms of the type of employment activities in which workers engage (casual vs. permanent). Consequently, when engaged in casual employment, the delineation of activities by sector has little relevance, with workers engaging in activities across sectors in rural or urban areas of the local labour market. However, as workers move up the skill ladder into permanent jobs, the delineation of employment by sector begins to have relevance. Many important research opportunities remain that might improve our understanding about whether workers face constraints that impede their movement out of casual and into permanent employment, and whether these constraints are amenable to policy.

Second, regarding the behaviour of firms, my results suggest that firms in India can act as major absorbers of labour, even in the short-run, highlighting the importance of diversification in the management of idiosyncratic productivity shocks. In addition, I demonstrate that even sectors considered to be considerably less climate-sensitive than agriculture can be significantly affected by changes in the weather. Understanding the relationship between environmental conditions and firm behaviour remains a fruitful area of research; especially the management and innovation of firms in the face of short-run and long-run environmental change.

Finally, regarding climatic influence on economic outcomes, my results show that workers can adapt to temperature increases by moving across sectors, and that the ability of firms to absorb these movements is one of the main ways that workers manage the effects of weather-driven changes in agricultural productivity. Consequently, we may overestimate the damages associated with future climate change if we do not take into account the adaptation responses of economic agents.

However, my results also demonstrate the difficulties associated with using weather data to draw inferences about economic behaviour. The reduced-form elasticity between weather and economic outcomes is a function of all competing and complementary channels and so fails to have a clear economic interpretation. By focusing on understanding precisely the channels and mechanisms through which weather affects economic activity rather than estimating the magnitude of weather effects, we can draw deeper insights about the design and implementation of policies to mitigate the economic consequences of environmental change. In relation to understanding the potential impacts of climate change this approach avoids the need to consider the complex relationship between weather (short run) and climate (long run), as well as avoiding stringent assumptions about any endogenous demographic and economic changes that may arise in response to climate change. Instead I focus on understanding how climate change *could* affect economic outcomes so that we can better design and implement policy, where necessary, to mitigate the economic consequences of climate change, today and in the future.

Figures and Tables

Table 1.1: Summary Statistics - Agriculture in India (2001–2007)

	Mean	Std. Dev. (within)	Std. Dev. (between)
Panel A: Agricultural Data			
Yield	1.488	0.390	1.442
Value (Rs.)	17,457.43	8,533.22	18,519.98
Production ('000 Tonnes)	79.456	40.280	206.430
Area ('000 Hectares)	41.94	12.234	84.946
Price (Rs./Tonne)	12,742.24	4,121.103	7,862.196
Number of Crops	9.928	0.764	1.875
Average Crop Share	0.104	0.0234	0.182
Average Share of Main Crop	0.551	0.0406	0.186
Panel B: Wage Data			
Average Day Wage: Agriculture	62.006	25.016	49.052
Average Day Wage: Manufacturing	94.315	48.415	41.775
Average Day Wage: Services	188.366	48.637	43.442
Average Day Wage: Construction	82.367	30.894	31.225
Panel C: Employment Data			
District Employment Share: Agriculture	0.547	0.115	0.189
District Employment Share: Manufacturing	0.193	0.069	0.104
District Employment Share: Services	0.170	0.056	0.084
District Employment Share: Construction	0.070	0.038	0.043
Unemployment Rate	0.103	0.041	0.060
Panel D: Meteorological Data			
Daily Average Temperature (°C)	24.847	0.271	4.185
Degree Days ($t_L = 17, t_H = \infty$)	3,103.204	90.757	809.660
Degree Days ($t_L = 0, t_H = 17$)	5,995.568	22.590	704.792
Monsoon Rainfall (mm)	927.297	206.509	482.657

Table 1.2: The Effects of Weather on Agricultural Outcomes

	Agricultural Outcomes					
	Log Yield (All Crops)	Log Value (All Crops)	Log Price (All Crops)	Log Yield (Main Crop)	Log Value (Main Crop)	Log Price (Main Crop)
Daily Average Temperature (°C)	-0.174***	-0.174***	0.000552	-0.270***	-0.267***	0.00319
	(0.0508)	(0.0442)	(0.0145)	(0.0968)	(0.0831)	(0.0314)
Monsoon Rainfall (100mm)	0.00819**	0.00711**	-0.00108	0.0208***	0.0145**	-0.00632*
	(0.00400)	(0.00359)	(0.00189)	(0.00720)	(0.00665)	(0.00369)
Crop × District FE	Yes	Yes	Yes	Yes	Yes	Yes
Crop × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
State-Year Trends	Yes	Yes	Yes	Yes	Yes	Yes
Observations	9,723	9,723	9,723	1,523	1,523	1,523

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Standard errors are adjusted to reflect spatial dependence (up to 1,000km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

Table 1.3: The Effects of Weather on Average Wages

	Agriculture	Manufacturing	Services	Construction
Daily Average	-0.0408*	-0.0751**	0.0415	-0.0127
Temperature (°C)	(0.0209)	(0.0337)	(0.0272)	(0.0211)
Monsoon Rainfall (100mm)	-0.00292	-0.0116*	0.000721	0.00552
	(0.0286)	(0.0641)	(0.0414)	(0.0408)
Observations	1755	1748	1824	1701

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Standard errors are adjusted to reflect spatial dependence (up to 360km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1–7 years.

Table 1.4: The Effects of Weather on the District Share of Employment - By Sector

	District Employment Shares				
	Agriculture	Manufacturing	Construction	Services	Unemployment
Daily Average	-0.0839***	0.0554***	-0.00969**	0.0389***	-0.00427
Temperature (°C)	(0.0134)	(0.00946)	(0.00441)	(0.00575)	(0.00458)
Monsoon Rainfall	-0.0235	0.0226	-0.0196***	0.0179*	0.00264
(100 mm)	(0.0228)	(0.0144)	(0.00717)	(0.00918)	(0.00770)
Average Share	0.546	0.200	0.071	0.163	0.112
Observations	1831	1831	1831	1831	1831

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Standard errors are adjusted to reflect spatial dependence (up to 630km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

Table 1.5: The Seasonal Migrant Share of the Destination Population

	(1) Mean	(2) Std. Dev	(3) Observations
Panel A: $\frac{\textit{All Migrants}}{\textit{Destination Workers}}$			
Within District	0.00753	0.0101	585
Different District, Same State	0.000270	0.00374	16,796
Different District, Different State	0.00364	0.0305	324,844
Panel B: $\frac{\textit{Rural Origin Migrants}}{\textit{Destination Workers}}$			
Within District	0.0072	0.0102	571
Different District, Same State	0.00029	0.0046	16,530
Different District, Different State	0.0034	0.0302	312,937

Notes: An individual is a seasonal migrant if they spent more than one month, but less than 6 months away from the household. Panel A presents the share of the 2001 Census population that are migrants from all locations, i.e., rural-urban, rural-rural, urban-rural, urban-urban. Panel B presents the share of the 2001 Census population that are migrants from rural areas, i.e., rural-urban, rural-rural. The migration data is constructed from the 2001 Population Census and the National Sample Survey round 64 (2007–2008). All shares are winsorized at 1.

Figure 1.1: Semi-Parametric Estimates of the Relationship between Temperature and the Labour Force Share of Economic Activity

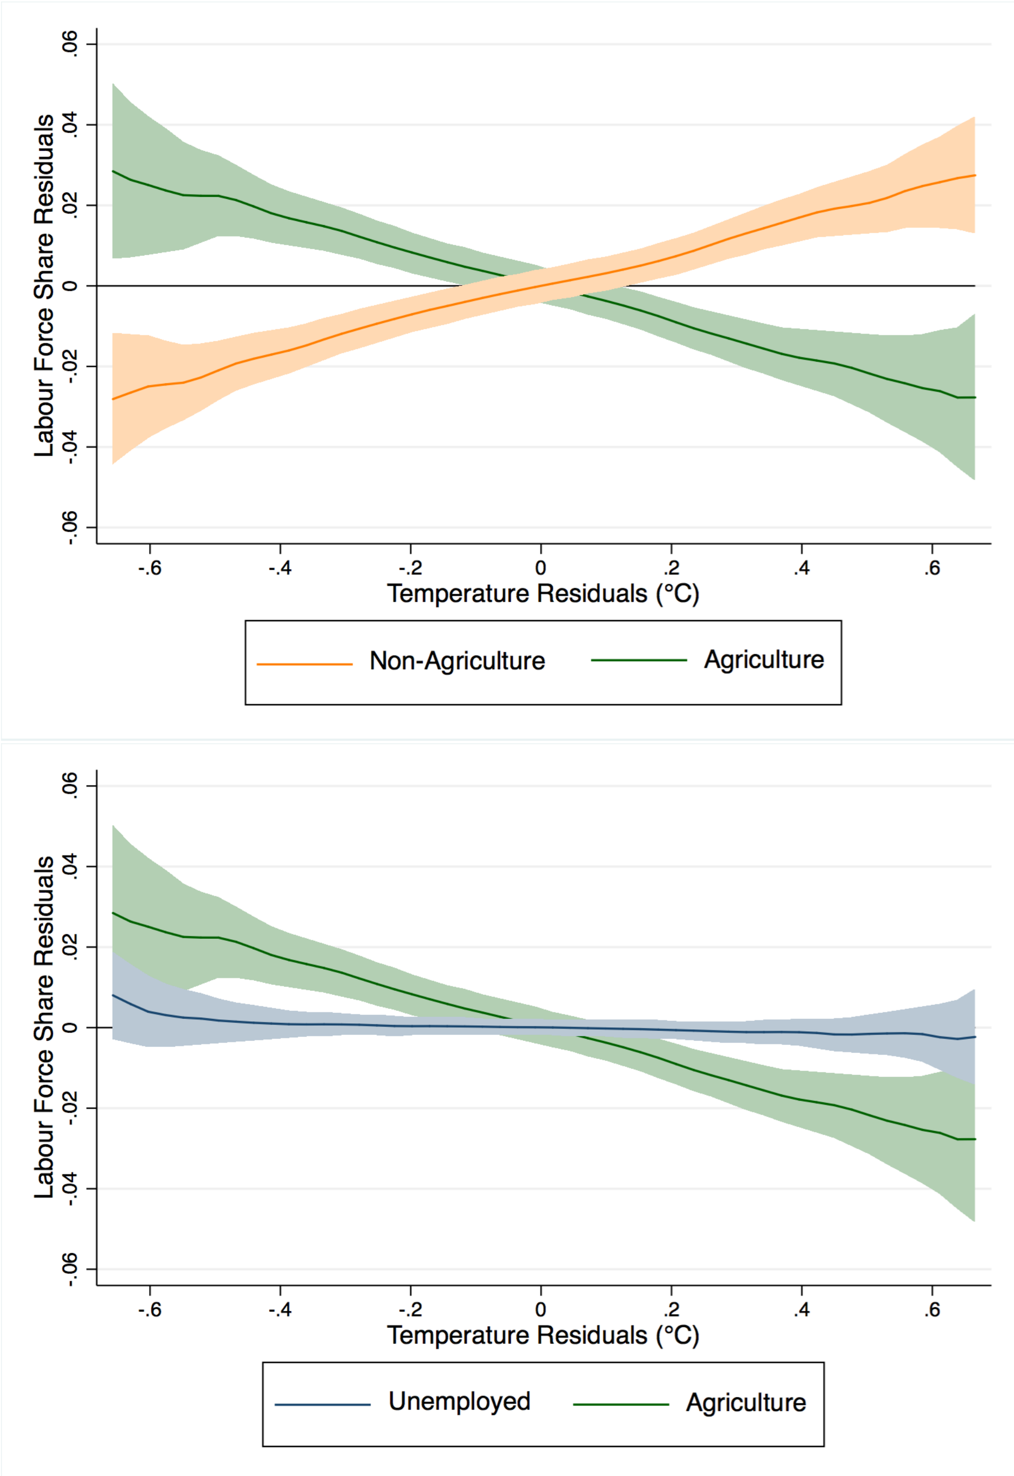

Notes: Each variable is regressed on district and year fixed effects as well as monsoon rainfall. The figures above are the result of loess regressions of the residuals from this exercise.

Table 1.6: The Effects of Weather in Foreign Districts on the Share of Employment in Destination Districts - By Sector

	Destination District Employment Shares				
	Agriculture	Manufacturing	Construction	Services	Unemployment
Panel A: All Origins					
Local Daily Average Temperature (°C)	-0.0843*** (0.0134)	0.0561*** (0.00945)	-0.0101** (0.00462)	0.0392*** (0.00591)	-0.00693 (0.00495)
Local Monsoon Rainfall (100 mm)	-0.0314 (0.0236)	0.0279* (0.0150)	-0.0179** (0.00788)	0.0192* (0.0100)	0.00536 (0.00802)
Foreign Daily Average Temperature (°C)	-0.0228 (0.0675)	0.0151 (0.0409)	-0.00156 (0.0173)	-0.00301 (0.0311)	0.0685*** (0.0212)
Foreign Monsoon Rainfall (100 mm)	0.0806 (0.0873)	-0.0505 (0.0523)	-0.0287 (0.0271)	-0.0156 (0.0430)	0.00698 (0.0344)
Panel B: Within State					
Local Daily Average Temperature (°C)	-0.0773*** (0.0133)	0.0527*** (0.00910)	-0.0117*** (0.00447)	0.0386*** (0.00598)	-0.00777 (0.00507)
Local Monsoon Rainfall (100 mm)	-0.0395* (0.0238)	0.0311** (0.0148)	-0.0176* (0.00909)	0.0178* (0.0106)	0.00531 (0.00918)
Foreign Daily Average Temperature (°C)	-0.0417* (0.0225)	0.0223 (0.0142)	0.00872 (0.00728)	0.00364 (0.0104)	0.0210*** (0.00701)
Foreign Monsoon Rainfall (100 mm)	0.00484 (0.0291)	-0.00157 (0.0183)	0.000809 (0.00995)	0.00338 (0.0134)	0.00921 (0.0120)
Average Share	0.546	0.200	0.071	0.163	0.112
Observations	1,827	1,827	1,827	1,827	1,827

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Standard errors are adjusted to reflect spatial dependence (up to 630km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

Table 1.7: Average Wage Gap (Agriculture vs. Manufacturing)

	India Wide	Within District	Within District Skill Adjusted
Average Wage Gap (Casual Manufacturing Workers)	1.336	1.226	1.137
Average Wage Gap (permanent Manufacturing Workers)	2.401	2.021	1.604
Average Wage in Agriculture (Rs.)	56.204	58.498	55.036
District Fixed Effects	No	Yes	Yes
Individual Controls	No	No	Yes
Observations	44,713	44,713	44,713

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Individual level controls include age, education, and gender. Estimates are based on individual-level mincerian wage regressions on the working-age population (14-65) controlling for a sector dummy (β) specifying whether the individual is engaged in agricultural, casual manufacturing labour, or permanent manufacturing employment. The wage gap is calculated as $\exp(\beta)$. Individual controls include level of education, age, and gender.

Figure 1.2: The share of Firms that employ Contract Workers – by Labour Regulation Environment

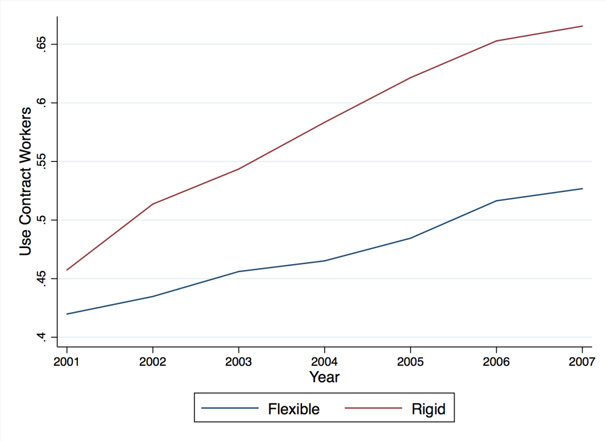

Figure 1.3: The share of Workers that are employed as Contract Workers – by Labour Regulation Environment

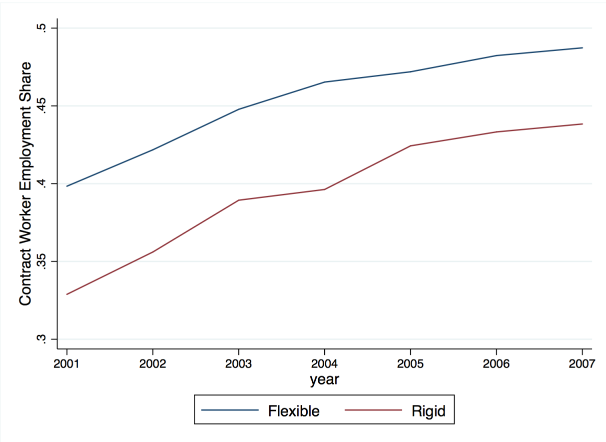

Table 1.8: Descriptive Statistics - Difference in Means Tests

	Rigid States	Flexible States	Difference
ASI Data			
Total Output	1,766.109	1,764.365	1.744
(Million Rs.)	(315.551)	(315.554)	(483.139)
Items Produced	2.557	2.174	-0.382
	(0.235)	(0.057)	(0.242)
Output per Worker	3.386	2.560	0.826
(Million Rs.)	(0.561)	(0.295)	(0.634)
TFPR	6.091	6.033	0.057
(Log)	(0.038)	(0.040)	(0.009)
Total Number of Workers	441.982	445.962	-3.979
(Non-mangers)	(31.199)	(51.643)	(60.335)
Employment	178.796	296.675	-117.879
(Contract Workers)	(11.793)	(74.758)	(75.682)
Average Day Wage	150.607	143.610	6.996
(Contract Workers)	(6.220)	(3.921)	(7.353)
Employment	348.667	316.588	32.078
(Regular Workers)	(45.051)	(19.200)	(48.971)
Average Day Wage	290.552	216.674	73.897*
(Regular Workers)	(39.758)	(15.072)	(2.274)
Contract Employment	0.502	0.548	-0.045*
Share	(0.023)	(0.015)	(0.017)
Capital	1,308.902	1,293.227	-15.675
(Million Rs.)	(274.278)	(315.569)	(418.106)
Share of Agricultural Employment	0.557	0.545	0.012
	(0.028)	(0.021)	(0.035)
Share of Agricultural GDP	0.164	0.162	0.002
	(0.016)	(0.011)	(0.020)
Total GDP	111.735	48.443	63.292
(Billion Rs.)	(37.838)	(6.840)	(38.451)
Average Travel Times to Other Districts	33.743	34.148	-0.405
(hours)	(1.809)	(1.768)	(2.530)

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Rigid States = 1, Neutral and Flexible States = 0. The sample is restricted to regulated firms. Standard errors are clustered at the State Level. Bilateral travel times data was provided by Allen and Atkin (2015).

Table 1.9: Main Results: Manufacturing Firm Outcomes – Unregulated Firms

	Production and Employment				Productivity and Wages			
	Total Output	Items Produced	Employment Contract	Employment Permanent	Output Per Worker	TFPR	Day Wage Contract	Day Wage Permanent
Panel A: Net Effect								
Daily Average Temperature (°C)	0.0640**	-0.00931	-0.0332	0.0372**	0.0226	0.0109	0.00471	0.0149*
	(0.0259)	(0.00893)	(0.0275)	(0.0149)	(0.0234)	(0.0122)	(0.00986)	(0.00837)
Panel B: Differential Effect								
Daily Average Temperature (°C)	0.0771	-0.0173	0.0915	0.0871**	-0.0514	-0.000815	-0.0225	-0.0106
	(0.0577)	(0.024)	(0.0818)	(0.0370)	(0.0589)	(0.0303)	(0.0311)	(0.0275)
Temperature × Flexible	-0.0198	0.012	-0.191	-0.0754	0.111	0.0178	0.0417	0.0385
	(0.0792)	(0.033)	(0.120)	(0.0534)	(0.0804)	(0.0405)	(0.0449)	(0.0378)
Rainfall Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed Effects	Sector × District, Sector × Year, and State-Year Time Trends							
Observations	94,085	94,085	30,831	85,943	94,085	80,942	30,831	85,943

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Panel A presents the results from regression the outcome variable on the level of temperature and monsoon rainfall. Panel B presents the differential effects of temperature across labour regulation environments. Standard errors are adjusted to reflect spatial dependence (up to 1,800km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

Table 1.10: Main Results: Manufacturing Firm Outcomes – Regulated Firms

	Production and Employment				Productivity and Wages			
	Total Output	Items Produced	Employment Contract	Employment Permanent	Output Per Worker	TFPR	Day Wage Contract	Day Wage Permanent
Panel A: Net Effect								
Daily Average Temperature (°C)	0.0100	-0.0101	-0.0199	0.0211	0.00101	-0.0368***	-0.0168	-0.0147
	(0.0284)	(0.0101)	(0.0274)	(0.0188)	(0.0253)	(0.0140)	(0.0129)	(0.00980)
Panel B: Differential Effects								
Daily Average Temperature (°C)	-0.0798	-0.056***	-0.121*	-0.000179	-0.0877*	-0.0878***	0.0401	-0.0691***
	(0.0492)	(0.0205)	(0.0644)	(0.0336)	(0.0456)	(0.0274)	(0.0316)	(0.0200)
Temperature × Flexibility	0.150**	0.0773**	0.171*	0.0356	0.148**	0.0854**	-0.0963**	0.0909***
	(0.0763)	(0.0311)	(0.0933)	(0.0456)	(0.0731)	(0.0368)	(0.0462)	(0.0315)
Rainfall Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed Effects	Sector × District, Sector × Year, and State-Year Time Trends							
Observations	49,112	49,112	24,557	47,363	49,112	44,305	24,557	47,363

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Panel A presents the results from regression the outcome variable on the level of temperature and monsoon rainfall. Panel B presents the differential effects of temperature across labour regulation environments. Standard errors are adjusted to reflect spatial dependence (up to 1,800km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

Table 1.11: Additional Manufacturing Firm Outcomes – Regulated Firms

	Capital	Capital Depreciation	Employment Managers	Day Wage Managers	Number of Plants
Daily Average Temperature (°C)	0.107 (0.0693)	0.0535 (0.0511)	-0.00699 (0.0411)	-0.0312 (0.0295)	0.000919 (0.0143)
Temperature × Flexible	-0.0449 (0.105)	-0.0180 (0.0756)	0.0664 (0.0640)	0.0107 (0.0454)	-0.00197 (0.0220)
Rainfall Controls	Yes	Yes	Yes	Yes	Yes
Fixed Effects	Sector × District, Sector × Year, and State-Year Time Trends				
Observations	48,898	38,900	48,498	48,498	49,112

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. All dependent variables are in log terms. Standard errors are adjusted to reflect spatial dependence (up to 1,800km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

Table 1.12: The Effects of Weather on GDP

	Total GDP	Agricultural GDP	Services GDP	Manufacturing GDP	Construction GDP
Daily Average Temperature °C	-0.0247* (0.0128)	-0.110** (0.0452)	-0.00211 (0.00944)	-0.0475* (0.0246)	0.0280 (0.0226)
Rainfall Controls	Yes	Yes	Yes	Yes	Yes
Fixed Effects	District, Year, and State-Year Time Trends				
Observations	6,792	6,763	6,792	6,264	6,780

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. All dependent variables are in logs. Standard errors are adjusted to reflect spatial dependence (up to 400km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

Table 1.13: Counterfactuals Estimates

	Baseline	Counterfactual 1	Counterfactual 2
Informal (34%)	-22.25%	-22.25%	-22.25%
Unregulated Formal (44%)	6.40%	6.40%	21.4%
Regulated Formal (22%)	0%	15% × Flexibility_s	15% × Flexibility_s
Total Manufacturing Effect (Flexibility = 1)	-4.75%	-1.44%	5.15%
Total Effect (Aggregate)	-2.47%	-2.22%	-1.43%
Losses Offset (%)	–	11%	42.2%

Figure 1.4: The Counterfactual Distribution of GDP Changes in Response to a 1°C Increase in Temperature

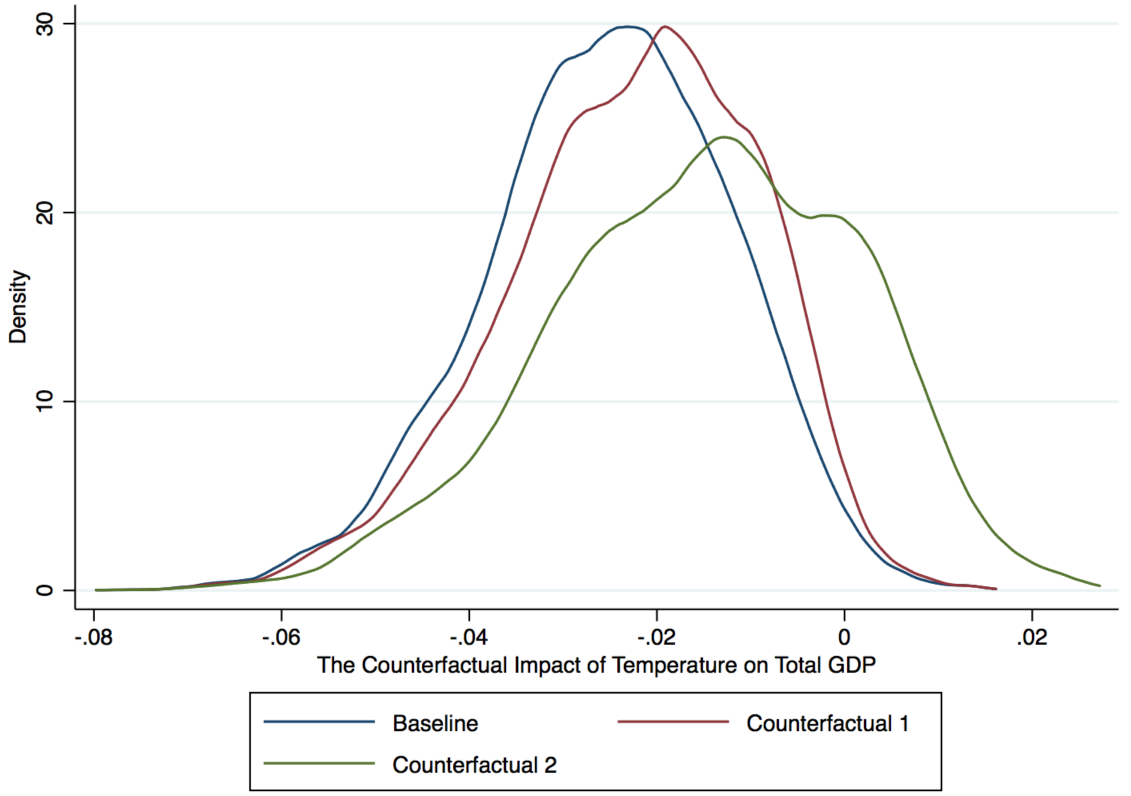

Chapter 2

Consumption Smoothing and the Welfare Cost of Uncertainty

When agents are unable to smooth consumption and have distorted beliefs about the likelihood of future income realisations, uncertainty about future states of the world has a direct effect on individual welfare. However, separating the effects of uncertainty from the effects of realised events and identifying the welfare effects of uncertainty both present a number of empirical challenges. Combining individual-level panel data from rural Ethiopia with high-resolution meteorological data, we estimate the empirical relevance of uncertainty to objective consumption and subjective well-being. We estimate that an increase in income uncertainty – proxied by rainfall variability, after controlling for both contemporaneous and historical weather events – is associated with a reduction in objective consumption and subjective well-being. Decomposing the effects of uncertainty on subjective well-being we estimate that 80% of the effect is driven by the direct effects of uncertainty on well-being, with the remaining 20% being driven through consumption. These results suggest that the welfare gains from managing both short-run weather events, and long-run climate change are likely to be substantially greater than estimates based solely on realised shocks.

2.1 Introduction

Economists have long recognised that an individual's sense of well-being depends not only on their average income or expenditures, but on the risk they face as well. While a number of studies have attempted to isolate the effects of uncertainty on macroeconomic outcomes in both developed and developing countries (Koren and Tenreyro, 2007; Baker and Bloom, 2013; Bloom, 2009, 2014), efforts to identify the effects of uncertainty at the microeconomic level have so far been limited. In many developing countries, where insurance and credit market failures are commonplace, the consequences of uncertainty on individual welfare are likely to be exacerbated, providing a context within which it is possible to measure and identify the effects of uncertainty on individual behaviour and welfare.

A significant body of research in development economics has focussed on estimating the response of household consumption to income fluctuations (Townsend, 1994; Udry, 1994; Morduch, 1995; Fafchamps and Lund, 2003; Suri, 2011; Morten, 2013; Bryan et al., 2014; Kinnan, 2014). A related literature has sought to understand the effects of weather on agricultural productivity and economic behaviour in agrarian societies (Paxson, 1992; Jayachandran, 2006; Guiteras, 2009; Schlenker and Lobell, 2010; Burgess et al., 2014b; Kudamatsu et al., 2014b; Colmer, 2016). These literatures have demonstrated that, in the presence of insurance and credit market failures, households are exposed to consumption risk and must rely on imperfect risk sharing mechanisms, and that, given the nature of partial insurance, welfare gains exist from further consumption smoothing. However, these gains may be underestimated when focussing solely on the ex post consequences of income shocks.

A separate literature has consistently documented that individuals perform poorly in assessing probabilities and overestimate the likelihood of success as a result of distorted beliefs (Weinstein, 1980; Alpert and Raiffa, 1982; Buehler et al., 1994; Rabin and Schrag, 1999; Brunnermeier and Parker, 2005; Brunnermeier et al., 2013). In the presence of partial insurance, theory suggests that uncertainty relating to future income will have an additional direct impact on welfare beyond the ex post realisation of income shocks.

This paper aims to understand the empirical relevance of future income uncertainty to household welfare in rural Ethiopia – one of the least developed countries in Africa, characterised by its high vulnerability to inclement weather. If households are able to effectively smooth consumption, then uncertainty about future income flows should only have an indirect effect on individual well-being, through the decisions that households make to smooth consumption, and there should be no direct effect of uncertainty on individual well-being. However, if households are exposed to consumption risk, then uncertainty about future income may have a direct effect on individual well-being.

We present a simple model in the spirit of Brunnermeier and Parker (2005), which shows how expected future utility, or anticipated utility, can have a direct, contemporaneous effect on utility if farmers face imperfect insurance and have imperfect information about the probability that a future income shock is realised. This forecasting error, arising from imperfect information, creates a wedge between an individual's subjective probability and the objective probability of an income shock being realised, such that individuals underestimate the likelihood of a bad outcome. In this model, forward-looking farmers who care about expected future utility will make investments to maximise future utility, which may indirectly affect contemporaneous utility; however, these same farmers will also have higher contemporaneous utility if they are optimistic about the future (anticipatory utility), introducing a trade-off between risk management investments and the benefits of optimism. An increase in uncertainty about future income makes farmers less optimistic about the future. Consequently, in the presence of imperfect insurance, the model predicts that farmers living in areas with greater income uncertainty will have lower well-being than comparable farmers living in areas with lower income uncertainty. An attractive feature of this framework is that it tends towards a model of rational expectations as an individual's subjective probability tends towards the objective probability. In this instance, expectations about the future no longer enter directly into current utility and future income uncertainty only affects utility indirectly through the actions that farmers take to manage this uncertainty. This highlights the potential welfare gains that increased access to information can provide (Rosenzweig and Udry, 2013, 2014).

However, while the theoretical predictions of uncertainty are clear, measuring and identifying the effects of future income uncertainty on individual well-being poses a significant empirical challenge. We seek to make progress in addressing this challenge. Using panel data on smallholder farmers from rural Ethiopia combined with high-resolution meteorological data, we exploit plausibly exogenous variation in rainfall variability (the second moment of the rainfall distribution) after controlling for contemporaneous and historical rainfall shocks (the first moment) to measure and identify the effects of income uncertainty on individual well-being.

We begin by providing supporting evidence for the premise that rainfall variability is a reasonable proxy for future income uncertainty. First, we demonstrate that an increase in rainfall variability, measured over the previous five years, is associated with a reduction in realised rainfall in the season following the survey. In addition, we provide support for the premise that farmers are aware of this signal. Using data on risk perceptions, we show that an increase in rainfall variability is associated with an increase in the belief that the rains will fail. Furthermore, we show that an increase in the belief that the rains will fail is associated with lower rainfall realisations in the following season; however, once we control for rainfall variability farmers beliefs no longer enter significantly into the relationship, indicating that beliefs about future income are largely driven by rainfall variability. Finally, we demonstrate that rainfall variability has no direct effect on agricultural production. Each of these considerations supports the premise that rainfall variability is a reasonable proxy for future income uncertainty, helping us to disentangle the effects of future income uncertainty from the effects of realised events.

In the main empirical exercise we examine the effects of rainfall variability on consumption and subjective well-being. We observe that an increase in inter-annual rainfall variability (a proxy for income uncertainty after controlling for contemporaneous and historical income shocks) has a negative effect on objective realised consumption as well as life satisfaction – a more evaluative measure of subjective well-being. We estimate that the effect of rainfall variability on life satisfaction has a direct effect above and beyond the effects through realised consumption. This direct effect accounts for 80% of the total effect on life satisfaction. The remaining 20% is captured by the effect of uncertainty, mediated through consumption – the indirect effect. Our results indicate that the welfare gains associated with managing both short-run weather events and long-run climate change are likely to be substantially greater than estimates based solely on realised shocks.

The remainder of the paper is structured as follows: section 2.2 provides a brief review of the literature; section 2.3 presents the theoretical framework that provides the structure for our empirical analysis; section 2.4 presents the data and economic context; section 2.5 provides supporting evidence for the premise that rainfall variability is a suitable proxy for future income uncertainty; section 2.6 presents the identification strategy and main empirical specification; section 2.7 discusses our results; the final section summarises the implications of these results and concludes.

2.2 Literature Review

Uncertainty is a nebulous concept. Knight (1921) created the modern definition of *uncertainty*. He began by defining the related concept of *risk*, which, he argued, covers a known probability distribution over a set of events. By contrast, *uncertainty* captures people's inability to forecast the likelihood of events happening when an individual's prior is infinitely diffuse. Bayesian uncertainty, a related concept, captures how diffuse an individual's prior is. For farmers in rain-dependent agrarian communities, an increase in rainfall variability makes forming expectations about rainfall realisations more difficult, affecting decisions about which crops to plant and which inputs, and how much of each input, to use in the production process. The concepts of risk and uncertainty are strongly related and, in many cases, the term risk may be applied in the context of uncertainty when outcomes involve a loss. Empirically, the measurement of uncertainty is challenging because it is not directly observed.

Interest in the economic consequences of uncertainty has seen a resurgence in recent years (Bloom, 2014). This has been driven in part by policy attention following the role that uncertainty played in shaping the Great Recession, alongside an increase in the availability of measures of uncertainty through more readily available proxies and increased computing power.

Given the difficulties associated with measuring uncertainty, it should be clear that there is no perfect measure but there is a broad range of proxies – such as the volatility of the stock market or GDP – because when a data series becomes more volatile it is harder to forecast (Ramey and Ramey, 1995; Koren and Tenreyro, 2007; Bloom, 2009; Carriere-Swallow and Céspedes, 2013; Bloom, 2014). Given these measures, much of the literature has focussed on macroeconomic outcomes in developed countries. However, risk and uncertainty is pervasive in developing countries and affects decision-making and welfare at the individual level as well as the macroeconomic level. As such, we introduce a new proxy for uncertainty – rainfall volatility – that is suited to understanding the consequences of uncertainty on individual welfare in agrarian societies.

A central challenge in this literature is identifying the effects of uncertainty. Specifically, the challenge relates to disentangling the impact of uncertainty from the impact of realised events. A very large literature has attempted to estimate the response of household consumption to income fluctuations (Townsend, 1994; Udry, 1994; Morduch, 1995; Fafchamps and Lund, 2003; Suri, 2011; Morten, 2013; Bryan et al., 2014; Kinnan, 2014). However, very little energy has been spent on identifying the effects of income uncertainty. We argue that by controlling for contemporaneous and historical rainfall shocks, any residual variation in rainfall variability acts as a suitable proxy for the effects of income uncertainty on smallholder farmers. Similar to the reasoning behind the use of stock market and GDP volatility, increased volatility in rainfall patterns makes it harder to forecast, increasing uncertainty about future income realisations. By stripping out variation associated with realised income effects, namely the level of rainfall, the volatility parameter should plausibly distinguish the effects of uncertainty from realised events. In doing so we also contribute to an expanding literature which seeks to understand climatic influence on economic outcomes (See Dell et al. (2014) for a recent review of this literature). This literature has examined how variation in the weather affects economic behaviour – of particular relevance to economic activity in developing countries given the importance that agricultural production and employment plays for the economic lives of the poor. In addition, to estimating the effects of weather, this literature seeks to use these estimates to inform our understanding about the potential damages associated with future climate change. However, by focussing only on realised shocks the costs of climate change may be underestimated. To the degree that climate change may result in an increase in rainfall variability, understanding how changes in rainfall variability affect individual welfare helps us to inform the effects of future income uncertainty today, as well as an additional channel through which climate change may affect economic activity in the future.

An additional challenge that arises when moving from the macroeconomic level to the microeconomic level is how to calculate the effects of uncertainty on individual welfare. The past decade has seen rapid growth in research on, and policy interest in, subjective well-being. In addition to "objective" measures of welfare, such as income and consumption, subjective measures of welfare are increasingly being used to elicit measures of experienced utility (Kahneman et al., 1997; Frey and Stutzer, 2002; Layard, 2005; Kahneman and Krueger, 2006; Dolan and Kahneman, 2008; Benjamin et al., 2012; Aghion et al., 2015; De Neve et al., 2015) to value non-market goods (Welsch, 2002, 2006; Rehdanz and Maddison, 2011; Carroll et al., 2009; Frey et al., 2010; Levinson, 2012; Feddersen et al., 2015; Baylis, 2016) and to evaluate government policy (Gruber and Mullainathan, 2005; Diener et al., 2009; Dolan et al., 2011; Levinson, 2013). Well-being is a broad measure of welfare that encompasses all aspects of the human experience. Researchers in this expanding field of economics use subjective measures of well-being to analyse and evaluate the impact of economic and non-economic factors on people's experienced utility.

Whether uncertainty about the future has a direct effect on well-being is ambiguous. The degree to which it does relates to the concept of anticipatory utility. Anticipatory utility has been a widely debated subject in academic and policy circles dating back to the time of Hume (1711–1776), Bentham (1789), Marshall (1891) and Jevons (1905). In "Principles of Economics", Marshall writes,

> *"...when calculating the rate at which a future benefit is discounted, we must be careful to make allowance for the pleasures of expectation."* (Marshall, 1981, p.178).

The other side of the coin is that future losses are also incorporated into utility. More recently, work in behavioural economics has explored the importance of anticipatory utility on decision-making (Lowenstein, 1987; Geanakoplos et al., 1989; Caplin and Leahy, 2001; Yariv, 2001; Brunnermeier and Parker, 2005; Brunnermeier et al., 2013). The next section introduces a model, based on Brunnermeier and Parker (2005), that formalises this concept providing some structure to the empirical analysis conducted in the proceeding sections.

2.3 Theoretical Motivation

In this section we present a model, based on the optimal expectations framework by Brunnermeier and Parker (2005), in which beliefs about future states of the world can enter directly into the current utility function; that is, agents care about both current utility and expected future utility. While all forward-looking agents who care about expected future utility will make investments to maximise future utility, if an agent's subjective probability about a future utility shock differs from the true probability, then their beliefs about the future will affect utility today.[1] For example, agents will have higher current utility if they are optimistic about the future; i.e., their subjective probability is lower than the true probability. In the context of this paper, farmers living in areas with lower climate variability may have lower subjective probabilities regarding the likelihood of a negative income shock being realised in the next period and so may have higher current utility. The framework presented provides a theoretical mapping between utility and life satisfaction, and motivates our empirical strategy.

2.3.1 Utility Maximisation Given Beliefs

Consider a world in which uncertainty about future income can be described as a binary state $s_t \in \{0,1\}$, where $s_t = 1$ indicates that the farmer is going to experience a negative income shock and $s_t = 0$ indicates that he will not. Let $p(s_t|\underline{s}_{t-1})$ denote the true probability that state $s_t \in \{0,1\}$ is realised following state history $\underline{s}_{t-1} = (s_1, s_2, \ldots, s_{t-1}) \in \{0,1\}$. We depart from the standard neoclassical model in so far as agents are endowed with subjective probabilities that may not coincide with the true state. These subjective probabilities are relevant for the decision making of the agent. Conditional and unconditional subjective probabilities are denoted $\hat{p}(s_t|\underline{s}_{t-1})$ and $\hat{p}(s_t)$ respectively.

At time t, the farmer receives some level of income which is consumed, c_t. For tractability, we assume there are no savings, so income is equal to consumption in each period,

$$\hat{\mathbb{E}}[U(c_1, c_2, \ldots, c_T|\underline{s}_t] \tag{2.1}$$

where $U(\cdot)$ is strictly increasing and strictly quasi-concave, and $\hat{\mathbb{E}}$ is the subjective expectations operator associated with $\hat{p}$, which depends on the information available to farmer i at time t.

[1]If an agent's subjective probability about a future utility shock is equal to the true probability, then the investments of forward-looking agents who care about expected future utility may still affect contemporaneous utility indirectly, through investment decisions. This is a consideration we explore in the empirical analysis.

The farmer maximises utility of consumption subject to his budget constraint:

$$\begin{aligned} c_{t+1} &= f(c_t, s_{t+1}), \quad (2.2) \\ g(c_{T+1}) &\geq 0 \text{ given } c_0 \quad (2.3) \end{aligned}$$

where $f(\cdot)$ provides the evolution of income, which is continuous and differentiable in c, $g(\cdot)$ gives the endpoint condition, and c_0 is the initial level of consumption. The optimal consumption is denoted $c^*(\underline{s}_t, \hat{p})$.

When the subjective probability of an income shock does not coincide with the true probability, the utility of the farmer, $\hat{\mathbb{E}}[U(\cdot)|\underline{s}_t]$, depends on expected future utility or anticipated utility, such that the subjective conditional belief has a direct impact on utility. To clarify this further, consider the standard model with time-separable utility flows and exponential discounting. In this case, utility at time t,

$$\hat{\mathbb{E}}[U(\cdot)|\underline{s}_t] = \beta^{t-1}\left(\sum_{\tau=1}^{t-1}\beta^{\tau}u(c_{t-\tau}) + u(c_t) + \hat{\mathbb{E}}\left[\sum_{\tau}^{T-t}\beta^{\tau}u(c_{t+\tau}|\underline{s}_t)\right]\right) \quad (2.4)$$

is the sum of memory utility from past consumption, utility from current consumption, and anticipatory utility from future consumption. Empirically, we identify these factors by controlling for past weather realisations (memory utility), contemporaneous weather (current consumption), and climate variability (anticipatory utility).

2.3.2 Optimal Beliefs and Life Satisfaction

The subjective beliefs of farmers are a complete set of conditional probabilities following any history of events, $\hat{p}(s_t|\underline{s}_{t-1})$; that is, the subjective probability that a shock will occur in the future depends on the history of shocks in the past. In this way, locations with a more variable climate may be more likely to experience a shock in the future. Subjective probabilities must satisfy four properties.

Assumption 1. Subjective probabilities are restricted in the following ways:

i $\sum_{s_t \in S}\hat{p}(s_t|\underline{s}_{t-1}) = 1$

ii $\hat{p}(s_t|\underline{s}_{t-1}) \geq 0$

iii $\hat{p}(s'_t) = \hat{p}(s'_t|s'_{t-1})\hat{p}(s'_{t-1}|s'_{t-2})\ldots\hat{p}(s'_1)$

iv $\hat{p}(s'_t|s'_{t-1}) = 0$ if $\hat{p}(s'_t|s'_{t-1}) = 0$

Assumption 1(i) states simply that subjective probability must add up to one; assumptions 1(i) - (iii) state that the law of iterated expectations holds for subjective probabilities; and assumption 1(iv) states that in order to believe something is possible, it must be possible.

The optimal beliefs for the farmer are the subjective probabilities that maximise the farmer's lifetime well-being and are defined as the expected time-average of the farmer's utility.

Definition 1. Optimal expectations (OE) are a set of subjective probabilities $\hat{p}^{OE}(s_t|\underline{s}_{t-1})$ that maximise lifetime well-being

$$\mathcal{W} = \mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T}\hat{\mathbb{E}}[U(c_1^*,\ldots,c_T^*|\underline{s}_t)]\right] \tag{2.5}$$

If farmers have rational expectations, (i.e, $\hat{p}(s_t|\underline{s}_{t-1}) = p(s_t|\underline{s}_{t-1})$) then the well-being and utility derived from the actions that farmers take will coincide. In this case, utility at time t only depends on present consumption (i.e., memory utility) and anticipatory utility does not enter into the utility function. This could be the case, for example, if an exact weather forecast or actuarially fair insurance is both available and effective. However, if subjective probabilities differ from the true probability that a shock will occur, then there will be a wedge between well-being and the farmer's utility, in this case memory utility, and anticipatory utility will enter into the utility function as in equation 4 and 5.

2.4 Data

The analysis conducted in this paper uses household survey data from rural Ethiopia. For the rural analysis, two rounds of a panel data set – the Ethiopian Rural Household Survey (ERHS) – that covers households from 18 villages in rural Ethiopia is used. The ERHS was conducted by Addis Ababa University in collaboration with the Centre for the Study of African Economies (CSAE) at the University of Oxford and the International Food Policy Research Institute (IFPRI) in seven rounds between 1989 and 2009. The sampling was constructed carefully to represent the major agro-ecological zones of Ethiopia. Households from six villages affected by drought in central and southern Ethiopia were surveyed for the first time in 1989. In 1994 the sample was expanded to cover 15 villages across the major regions of Ethiopia (Tigray, Amhara, Oromia, and Southern Nations Nationalities and People's Region), representing 1,477 households. Further rounds were completed in 1995, 1997, 1999, 2004, and 2009. The additional villages incorporated in the sample were chosen to account for the diversity in farming systems throughout the country. Stratified random sampling was used within each village, based on the gender of household heads.

This paper makes use of the final two rounds (2004 and 2009) as only these years contain questions on subjective well-being. One of the surprising features of the data set is the limited attrition compared to other household surveys in developing countries. Attrition of the panel has been low at 1-2 percent of households per round since the survey first began, indicating substantial persistence in the social structure of villages in rural Ethiopia (Dercon and Hoddinott (2009)).

In addition to the household survey data, rainfall and temperature data has been constructed from 6-hourly precipitation reanalysis data at the village level from the ERA-Interim data archive supplied by the European Centre for Medium-Term Weather Forecasting (ECMWF).[2] Previous studies have relied on the use of meteorological data provided by the Ethiopian meteorological service and the number of missing observations is a concern. This has been exacerbated by the serious decline in the past few decades in the number of weather stations around the world that are reporting. Lorenz and Kuntsmann (2012a) show that, since 1990, the number of reporting weather stations in Africa has fallen from around 3,500 to around 500. With 54 countries in the continent, this results in an average of fewer than 10 weather stations per country. Looking at publicly available data, the number of stations in Ethiopia included by the National Oceanic and Atmospheric Administration's (NOAA) National Climatic Data Centre (NCDC) is 18; however, if we were to apply a selection rule that required observations for 365 days, this would yield a database with zero observations. For the two years for which we have economic data (2004 and 2009), weather station data is available for 50 days in Addis Ababa in 2004 and is available for all 18 stations for an average of 200 days (minimum of 67 days, maximum of 276 days) in 2009. This is likely to result in a huge increase in measurement error when this data is used to interpolate across the 63 zones and 529 woredas (districts) reported in 2008. If this measurement error is classical, i.e., uncorrelated with the actual level of rainfall measured, then our estimates of the effect of these variables will be biased towards zero. However, given the sparsity of stations across Ethiopia (an average of 0.03 stations per woreda), the placement of stations is likely to be correlated with agricultural output; that is, weather stations are placed in more agriculturally productive areas, where the need for weather information is higher. As a result, we might expect that estimates using weather stations are systematically biased upwards. For these reasons, the use of remote-sensing data on a uniform grid has great value in areas with low station density.

The ERA-Interim reanalysis data archive provides 6-hourly measurements for a very rich set of atmospheric parameters, from 1st January, 1979 until the present day, on a global grid of quadrilateral cells defined by parallels and meridians at a resolution of 0.25 x 0.25 degrees (equivalent to 28km x 28km at the equator).[3] Reanalysis

[2]See Dee et al. (2011) for a detailed discussion of the ERA-Interim data.

[3]To convert degrees to km, multiply 28 by the cosine of the latitude, e.g, at 40 degrees latitude 0.25 x 0.25 degree cells are 28 x cos(40) = 21.4 km x 21.4 km.

data is constructed through a process whereby climate scientists use available observations as inputs into climate models to produce a physically consistent record of atmospheric parameters over time (Auffhammer et al., 2013a). This results in an estimate of the climate system that is separated uniformly across a grid, making it more uniform in quality and realism than observations alone, and one that is closer to the state of existence than any model would provide alone. This provides a consistent measure of atmospheric parameters over time and space. This type of data is increasingly being used by economists, since it fills in the data gap apparent in developing countries, where the collection of consistent weather data is lower down the priority list in governmental budgets (see Dell et al. (2014) for a review of its recent applications in the literature).

By combining the household data with the ERA-interim data, we create a panel that allows for microeconomic analysis of weather and climate in rural Ethiopia.

The outcome variables of interest from the economic data are objective real per capita consumption in adult equivalent units, c_{it}, and subjective life satisfaction, $\mathcal{W}_{it} = \hat{\mathbb{E}}[U(\cdot)|\underline{s}_t]$, asked of both the head and spouse of the household.

Real per capita consumption is constructed in the following way. First, all food consumption in the past 7 days is valued and scaled up to a month. In addition, expenditures on items purchased by the household in a typical month are added. On top of this, the value of own production is imputed by multiplying the quantity produced by the median price paid by other households in the same district. Finally, consumption expenditures are spatially deflated to ensure comparability over time and space. This is very important given the significant inflation observed between 2004 and 2009 due to rapid increases in world grain prices and internal monetary policy (Durevall et al., 2013), with average inflation peaking at 55.2% and food price inflation at 92% (Central Statistics Agency, 2009).

Figure **??** plots the distribution of log real consumption per capita. To estimate the degree of consumption dispersion, we calculate the unconditional log difference between the 90th and 10th percentile household. From this calculation we estimate that per capita consumption in the 90th percentile household is approximately 7 times greater than in the 10th percentile household, indicating substantial consumption inequality in rural Ethiopia.[4]

[4] $\exp(1.935) = 6.92$.

Figure 2.1: The Distribution of Consumption in Rural Ethiopia

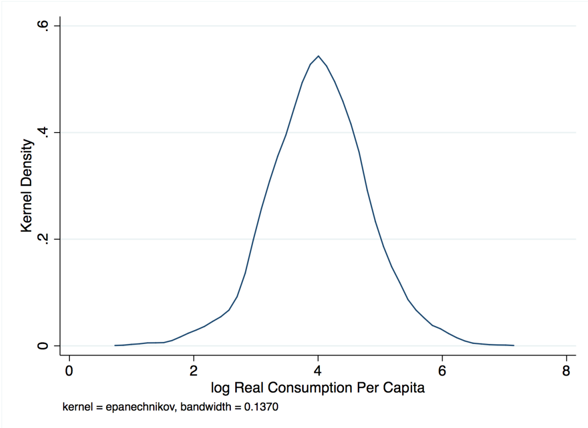

Our measure of subjective well-being in rural Ethiopia is constructed using responses related to the level of agreement with the following statement as the dependent variable: "I am satisfied with my life." A score of one is described as "Very Dissatisfied" and a score of seven is described as "Very Satisfied". These questions are similar to the standard questions used in cross-country surveys such as the World Values Survey and the Eurobaromoter Survey. We also demonstrate the robustness of our results to alternative measures of subjective well-being.

Figure ?? plots the distribution of responses for this question. The distribution of life satisfaction in rural Ethiopia is bimodal, in contrast to the distributions observed in developed countries. In addition, the average level of life satisfaction is substantially lower than the average levels reported in developed countries, where responses are shown to be skewed to the right with a long left tail. The average response is the middle group: "neither satisfied nor dissatisfied".

Figure 2.2: The Distribution of Life Satisfaction in Rural Ethiopia

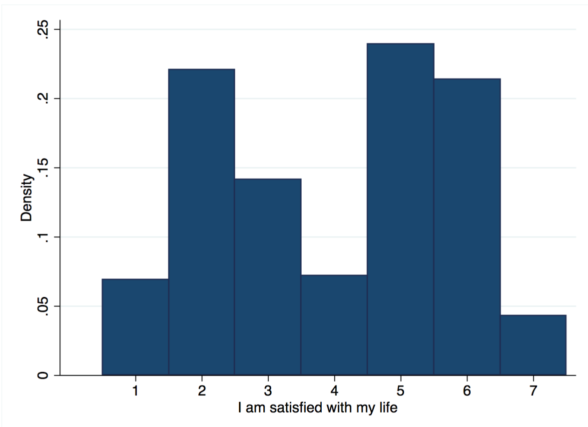

Our explanatory variables are motivated by the theoretical model in section 2.3. We calculate measures of memory utility, contemporaneous utility, and a proxy for uncertainty, which aims to isolate the effects of anticipatory utility.

Rainfall and temperature measures for each village are estimated by taking all data points within 100km of the village or city centroid and then interpolating through a process of inverse distance weighting. The weight attributed to each grid point decreases quadratically with distance.

The main variable of interest is a proxy for future income uncertainty – rainfall variability. Starting from a measure of total annual rainfall for each village, we calculate the coefficient of variation for rainfall (CV), measured as the standard deviation divided by the mean for the previous five years, the time period between each survey round.[5] In addition, we construct a second measure of rainfall variability defined as the standard deviation of rainfall, measured over the previous five years.

Table **??** presents the descriptive statistics of the variables of interest for the period analysed.

2.5 Rainfall Variability and Future Income Uncertainty

The use of rainfall variability as a proxy for uncertainty is driven by the importance of agriculture for subsistence consumption and livelihoods in rural parts of Sub-Saharan Africa, where access to irrigation is sparse. The consideration of uncertainty as a determinant of welfare is distinct from the literature, which examines the effects of weather shocks on welfare. If farmers form expectations about the climatic conditions of their area, we might expect that they plant crops that are suited to that area. Any deviation from the conditions on which this optimal cropping decision is based, such as more or less rainfall, may not be welfare-improving. The formation of these expectations is key for production. For this reason, we use rainfall variability, which, we argue, affects the farmers' ability to forecast the likelihood of future rainfall realisations, increasing uncertainty about future income.

To support this claim, we demonstrate, using the full panel of weather data between 1979 and 2012, that an increase in rainfall variability measured over the previous five years is associated with a reduction in realised rainfall in the season following the household survey. Table 2.2 shows that a one unit increase in the coefficient of variation is associated with a 3–4mm reduction in rainfall during the main agricultural season in Ethiopia (Meher). A one standard deviation increase in rainfall variability during this period (9.197 units) would be associated with a 27–36mm reduction in the main rainy season (Meher), around 15% of the average change in Meher rainfall each year. This suggests that increases in rainfall variability should be associated with an increase in future income uncertainty.

A second consideration is whether farmers are aware of this signal. This is obviously very difficult to test; however, using data from the 2009 round on risk percep-

[5] The results are robust to measurement over alternative time periods.

tions we are able to provide some support for the claim that farmers may use rainfall variability to shape their expectations about the likelihood of future income shocks.

Table 2.3 presents descriptive statistics for the variables of interest relating to farmers' risk perceptions. we use data from three questions. The first question asks *"How often do the Meher rains fail for your land?"* Farmers report that on average the rains fail every 2.867 years. This implies that the rains fail roughly twice between each survey round, highlighting the significant volatility of the climate in rural Ethiopia, providing support for the premise that rainfall variability is an important driver of future income uncertainty in the short- to medium-run. The second question asks *"Are there any signs (e.g. in the temperature, the behaviour of animals, the rain received in the last few years, the number of the year) that you might receive timely or sufficient rain?"*. 55% of farmers report some perceived sign of timely or sufficient rains. The third variable is based on an exercise requiring the farmer to place beans on two counters to ascertain the likelihood that the rains will fail. The farmers are asked *"Given the number of times the rains fail, and any signs you have observed, indicate by placing beans on the relevant square how likely it is that you think the Meher rains will fail this year. The more sure you are that they will fail the more beans should be placed on the low rainfall square. If you think there is an equal chance that they will fail place the beans equally between the two squares."* The farmers receive 20 beans in total. In table 2.3 0% corresponds to 0 beans, and 100% corresponds to 20 beans, being allocated to the square indicating that the rains will fail. On average, farmers perceive that there is a 59.3% chance that the rains in the coming season will fail.

Using these data, we first look at the correlation between rainfall variability and the likelihood of there being positive signs and the belief that the rains will fail. In Table 2.4 we show that a one standard deviation increase in rainfall variability (14.614 points) is associated with a 5.84% reduction in the likelihood of farmers perceiving that there are any positive signs. In addition, we find that a one unit increase in rainfall variability is associated with a 4.23% increase in farmers beliefs' that the rains will fail.

In addition, we show that an increase in the belief that the Meher rains will fail is associated with a reduction in the following season's Meher rainfall. However, it is interesting to note that once rainfall variability is controlled for – negatively associated with the following season's Meher rainfall –, farmers' beliefs no longer have any predictive power in determining future Meher rainfall. This suggests that rainfall variability is a major determinant of farmers' beliefs about the likelihood that the rains will fail, supporting the premise that rainfall variability is a good proxy for future income uncertainty.

A final consideration is whether rainfall variability has a direct effect on agricultural production, resulting in a realised income shock. If rainfall variability is to be an effective proxy for future income uncertainty, then it should have no direct effect on contemporaneous income. Using data on each household's agricultural production, we calculate agricultural yields, defined as the cultivated area-weighted production divided by total cultivated area.[6]

Using this data, we estimate the effects of rainfall variability on agricultural production using the specification,

$$\begin{aligned} \log(\text{Yield}_{hvt}) &= \beta_1 \text{Rainfall Variability}_{v\underline{t}} + \beta_2 f(\text{w}_{vt}) + \beta_3 \overline{f(\text{w}_{v\underline{t-5}})} \\ &+ \alpha_h + \alpha_m + \alpha_t + \epsilon_{hvt} \end{aligned} \tag{2.6}$$

where Rainfall Variability$_{v\underline{t}}$ is my proxy for future income uncertainty – the coefficient of variation or the standard deviation of rainfall measured over the previous 5 years –, $f(\text{w}_{vt})$ is a function of contemporaneous weather variables, and $\overline{f(\text{w}_{v\underline{t-5}})}$ is function of historical weather variables measured over the previous 5 years. In addition, we include a vector of household fixed effects, year fixed effects and month of survey fixed effects.

In column (1) of Table 2.5 Iestimate the effects of rainfall variability absent contemporaneous and historical rainfall controls. we find that an increase in rainfall variability is associated with a contraction in agricultural yields. However, once contemporaneous and historical rainfall controls are included (columns 2 and 3), rainfall variability has no effect on agricultural yields, supporting the premise that rainfall variability provides a plausible proxy for future income uncertainty, rather than acting as a realised income shock. As discussed, one of the key difficulties associated with the measurement and identification of uncertainty is separating the effects of uncertainty from realised events. These results provide support for my measure of future income uncertainty, allowing us to proceed with the main empirical exercise.

2.6 Empirical Strategy

Having provided supporting evidence for our measure of future income uncertainty, we now proceed to estimate the effects of rainfall variability on objective consumption and subjective well-being.

We begin by examining the effects of rainfall variability on objective consumption to explore the degree to which households respond to changes in future income uncertainty. As demonstrated rainfall variability has no direct effect on agricultural production and so does not affect income directly; however, consumption expenditures may still respond to household decision-making in the face of future income

[6]The crops used are white teff, black teff, barley, wheat, maize, and sorghum.

uncertainty. The effect of uncertainty on contemporaneous consumption is theoretically ambiguous: consumption expenditures may increase if farmers increase their spending on inputs that mitigate the economic consequences of future rainfall shocks (to the degree that such investments are available); consumption may decrease if farmers exhibit decreasing absolute risk aversion and engage in precautionary saving (to the degree that saving is possible); or uncertainty about future income may have no effect on present consumption if farmers are unable to smooth consumption over time, through precautionary saving or defensive investments. To examine the relevance of these effects we being by examining the effects of rainfall variability on real consumption per capita using the following specification,

$$\log C_{it} = \beta_1 \text{Rainfall Variability}_{v\underline{t}} + \beta_2 f(\mathrm{w}_{vt}) + \beta_3 \overline{f(\mathrm{w}_{vt-5})} + \alpha_h + \alpha_m + \alpha_t + \epsilon_{ivt} \tag{2.7}$$

where Rainfall Variability$_{v\underline{t}}$ is the variable of interest – our proxy for future income uncertainty –, $f(w_{vt})$ is a function of contemporaneous rainfall and temperature variables, and $\overline{f(\mathrm{w}_{vt-5})}$ is function of historical rainfall and temperature variables defined over the previous 5 years (the time period over which our measure of rainfall variability is measured). In addition, we control for household (α_h), year (α_t), and month of survey (α_m) fixed effects.

In addition, to estimating the effects of future income uncertainty on objective consumption, we also examine the effects of future income uncertainty on subjective well-being,

$$\mathcal{W}_{it} = \beta_1 \text{Rainfall Variability}_{v\underline{t}} + \beta_2 f(\mathrm{w}_{vt}) + \beta_3 \overline{f(\mathrm{w}_{vt-5})} + \alpha_i + \alpha_m + \alpha_t + \epsilon_{ivt} \tag{2.8}$$

$\mathcal{W}_{ivt}$ is our measure of subjective well-being – life satisfaction.[7]

Individual fixed effects, α_i, allow us to address any issues associated with time-invariant unobserved individual heterogeneity, which has been shown to be an important determinant of subjective well-being (Argyle, 1999; Diener and Lucas, 1999; Ferrer-i Carbonell and Frijters, 2004).[8] In addition to individual fixed effects, we con-

[7]Results are robust to using an ordered probit model with random effects to account for an ordinal measurement of life satisfaction rather than the cardinal measurement implied by the linear regression model. The use of linear regression models implies that the spacing between different outcomes, e.g., "Very Satisfied" and "Dissatisfied", or "Satisfied" and "Very Satisfied", are uniform. The use of an ordered probit model assumes that the respondent's well-being, $\mathcal{W}_{ivt}$, is an unobserved latent outcome conventionally proxied by a self-reported life satisfaction response, $\mathcal{W}^*_{ivt}$, on an ordinal scale. However, because it is not possible to formulate a fixed effects ordered probit model since the fixed effects are not conditioned out of the likelihood, we must use random effects.

[8]Results (available upon request) are also robust using household or village fixed effects. The results are consistent in sign and magnitude across models.

trol for year fixed effects to control for aggregate shocks, economic development, and macroeconomic policies. We also include survey month fixed effects to control for seasonal variation in the timing of the survey.

The last term in equations 2.7 and 2.8 is the stochastic error term, ϵ_{ivt}. We follow the approach of Hsiang (2010) by assuming that the error term may be heteroskedastic and spatially correlated across contemporaneous districts (Conley, 1999). We loop over all possible distances between 10km and 1000km selecting the parameter value that provides the most conservative standard errors.

The focus of our empirical exercise is to identify the effects of uncertainty on individual welfare. In section 2.3 we provided a theoretical mapping for our empirical exercise in the form of equation 2.4. In terms of our empirical analysis we assume that there exists a mapping between rainfall and consumption in agrarian societies such that,

$$U(rain_{it}) = \underbrace{\sum_{\tau=1}^{t-1} \beta^{\tau} u(rain_{it-\tau}) +}_{\text{historical shocks}} \underbrace{\beta^{t} u(rain_{it})}_{\text{contemporaneous shocks}} + \underbrace{\hat{\mathbb{E}}\left[\sum_{\tau=t+1} \beta^{\tau} u(rain_{it+\tau|\underline{s}_t})\right]}_{\mathbb{E}(\text{future shocks})} \tag{2.9}$$

We argue that once historical and contemporaneous effects have been controlled for any residual variation in rainfall likely captures the expectation of future effects, i.e. the effect of uncertainty through anticipation utility.

In any analysis of uncertainty, measurement and identification is highly challenging. The evidence provided in section 2.5 suggests that our proxy for future income uncertainty is plausible. However, there are additional statistical considerations that may inhibit our ability to disentangle the effects of uncertainty from the realised events. As the first moment and second moment of the rainfall distribution are correlated ($\rho = 0.28$) it is important to control for first-moment effects to isolate the effects of uncertainty, to the degree that they are empirically relevant, from income effects. We do this by controlling for historical and contemporaneous rainfall.

A second concern may be that non-linearities in the relationship between rainfall and income imply that accounting for the first moment isn't sufficient to remove all the residual variation associated with income from the error term. As a consequence, our measure of uncertainty may be driven by realised income shock effects. We account for this by allowing contemporaneous and historical weather events to enter quadratically into our estimating equation.

A further concern is the high degree of correlation between atmospheric parameters. As temperature may also be an important factor in explaining variation in income – and is highly correlated with rainfall – we also account for contemporaneous and historical temperature effects to further remove as much residual variation in income as possible. We can never be certain that our measure of uncertainty is free from any residual variation associated with income; however, we argue that accounting for the above considerations should allay any first-order concerns. These considerations are supported by the evidence presented in section 2.5.

2.7 Results

2.7.1 Uncertainty and Objective Consumption

First we examine the effect of future income uncertainty on real consumption per capita – an objective measure of household welfare. The results of this exercise are presented in Table 2.6. Columns 1-3 estimate the effects of rainfall variability on real consumption per capita for the years 2004 and 2009. Columns 4-6 apply to the extended panel covering the years, 1995, 1997, 1999, 2004 and 2009.

The effect of uncertainty on contemporaneous consumption is theoretically ambiguous: consumption expenditures may increase if farmers increase their spending on inputs that mitigate the economic consequences of future rainfall shocks (to the degree that such investments are available); consumption may decrease if farmers exhibit decreasing absolute risk aversion and engage in precautionary saving (to the degree that saving is possible); or uncertainty about future income may have no effect on present consumption if farmers are unable to smooth consumption over time, through precautionary saving or defensive investments.

Across all specifications and time frames we observe that an increase in rainfall variability is associated with a contraction in real consumption per capita. Column (1) presents the results of regressing the log of real consumption per capita against our measures of income uncertainty – the coefficient of variation (panel A) and standard deviation for rainfall (panel B) – finding a second negative relationship consistent with a precautionary savings interpretation. However, this specification does not control for any contemporaneous or historical weather effects and so the interpretation is confounded by realised income effects. After controlling for these effects the estimated coefficient drops from -0.0189 to -0.0149. Column (3) incorporates nonlinearities into the contemporaneous and historical rainfall and temperature variables by including quadratic terms into the linear specification. In doing so the estimated coefficient increases from -0.0149 to -0.0271.

To understand the magnitude of these effects, a one standard deviation increase in rainfall variability (6.560 units) would be associated with a 9.7-17.7% reduction in real consumption per capita. Such a change is equivalent to 14-25% of the within-household standard deviation change in consumption during the period of study.

In addition to estimating the effects parametrically, it is also interesting to examine the degree to which there are non-linearities in the relationship between rainfall variability and consumption. We do this by estimating this relationship semi-parametrically and present the results visually as in Hsiang et al. (2013). The idea behind this approach is to obtain local estimates of the relationship being studied and to display them in a way that visually weights the degree of regression uncertainty underlying the relationship. This procedure has two steps. First, all variation associated with contemporaneous and historical weather effects, as well as the fixed effects is absorbed from the data, ensuring that the scales are identical. Then a LOESS regression of the residuals of rainfall variability and real consumption per capita is estimated repeatedly using a bootstrapping procedure. The residuals for rainfall variability on the horizontal-axis are subdivided into 500 grid points. Each bootstrapped regression is evaluated at the grid-point along the horizontal-axis. This results in a set of fitted values for each grid point. In the second step the fitted values are plotted. For each horizontal grid point, a kernel density is estimated. The colouring of the graph relates to two considerations. First the overall colour intensity at each grid point along the horizontal axis related to the overall mass of data that is available in that part of the distribution. This colouring is then stretched out vertically in relation to the density of the fitted values. 95% confidence intervals are plotted as dashed lines. The results of this exercise are plotted in Figure 2.5.

2.7.2 Uncertainty and Subjective Well-Being

Uncertainty and Life Satisfaction

In addition to examining the effects of uncertainty on objective consumption we are interested in understanding the broader effects of uncertainty on individual well-being, beyond the effects on household consumption, by examining the effects of rainfall variability on life satisfaction in rural Ethiopia. Our theory predicts that when there is partial insurance and the subjective probability of an income shock does not coincide with the true probability, the utility of the farmer depends on anticipatory utility, such that the subjective conditional belief has a direct effect on utility. Given the importance of rainfall as a driver of income in agrarian societies, an increase in rainfall variability increases uncertainty about future income flows. By controlling for contemporaneous and historical rainfall events, we argue that any residual variation in rainfall variability captures the direct effects of income uncertainty on individual well-being. Table 2.7 presents our main results, examining the effects of uncertainty on life satisfaction.

Column (1) of Table 2.7 presents the results of regressing life satisfaction against our measures of uncertainty – the coefficient of variation (panel A) and standard deviation for rainfall (panel B). The relationship between uncertainty and life satisfaction is negative and significant. However, this specification does not control for any contemporaneous or historical weather effects and so if these factors are correlated with our measure of uncertainty, this will confound the interpretation of this effect. Columns 2-5 demonstrate the relevance of this omitted variable bias. When we control for contemporaneous and historical shocks (following the main empirical specification in equation 6) the coefficient on rainfall variability declines from -0.0332 to -0.0261 (Column 2).

Column 3 further tests the robustness and interpretation of our results by controlling directly for the logarithm of real consumption per capita. While this is a bad control (Angrist and Pischke, 2009), it allows us to test the degree to which the direct effects of uncertainty on individual well-being are mediated by the effects of uncertainty through consumption, unbundling the channels through which rainfall variability affects individual well-being. We estimate that controlling for consumption mediates the estimated effect of uncertainty on life satisfaction by approximately 20%, reducing the sign of the coefficient from -0.0261 to -0.0217. Consequently, we argue that the effects of uncertainty on individual well-being are largely explained by direct effects, rather than effects driven through consumption. This is consistent with a wide literature exploring the psychic costs of income shocks and poverty (van den Bos et al., 2009; Hare et al., 2009; Delgado and Porcellie, 2009; Doherty and Clayton, 2011), suggesting that income uncertainty can have a significant direct effect on individual welfare, above and beyond the effects of realised changes in income. Consequently, the welfare gains associated with managing both short-run weather events as well as long-run climatic change are likely to be substantially greater than estimates based solely on realised shocks.

Columns 4 and 5 accounts for potential non-linearities in the contemporaneous or historical rainfall or temperature controls that could be correlated with our measure of rainfall variability. We do so by including a quadratic term for the contemporaneous and historical rainfall and temperature controls. It is of interest to note that by doing so the estimated effect of uncertainty increases from -0.0261 to -0.0742.

As in the analysis of rainfall variability on real consumption per capita, we also the degree to which there are non-linearities in the relationship between rainfall variability and life satisfaction, estimating the relationship semi-parametrically. The results of this exercise are plotted in Figure 2.6.

To understand the magnitude of these effects, a one standard deviation increase in rainfall variability (6.560 units) would be associated with a 0.171-0.486 point reduction in life satisfaction. This is equivalent to 18-50% of the within-individual standard deviation change in life satisfaction during the period of study. Mediating the effects of uncertainty through consumption (columns 3 and 5) reduces this magnitude to be equivalent to 15-47% of the within-individual standard deviation change in life satisfaction during the period of study, highlighting again that the effects of rainfall variability – a proxy for future income uncertainty once contemporaneous and historical income shocks have been controlled for – are largely driven by direct effects on individual well-being, rather than through the actions or consumption of the farmers themselves.

2.7.3 Supporting Evidence

In addition to the main results discussed above this section presents additional evidence to support the premise of this paper: that rainfall variability acts as a plausible proxy for future income uncertainty. We begin by examining the effects of rainfall variability on present affect (happiness), where we argue that future income uncertainty should have less of an effect. In addition, we examine the effects of rainfall variability on alternative, evaluative measures of subjective well-being to provide support for our main effects.

Uncertainty and Happiness

Within the subjective well-being literature, it is generally considered that questions based on the life satisfaction scale are more evaluative measures, whereas questions related to happiness are a better measure of present affect (Benjamin et al., 2013; Levinson, 2013).[9] While both measures of subjective well-being are highly correlated ($\rho = 0.426$) we might expect that rainfall variability should have a smaller effect on happiness (contemporaneous well-being) than life satisfaction (evaluative well-being) if it is capturing the effects of future income uncertainty.

Table 2.8 presents the results from this analysis. In column (1) we estimate that rainfall variability has a significant effect on happiness; however, once we control for contemporaneous and historical weather effects this estimate becomes insignificant, both statistically and in terms of the size of the coefficient. When we account for potential non-linearities in the contemporaneous and historical rainfall or temperature controls, through the inclusion of quadratic terms, the estimated coefficients become statistically significant; however, the magnitude of the coefficients are still very small. Semi-parametric estimates of this relationship are plotted in Figure 2.7.

These results are consistent with our predictions that future income uncertainty should have less impact on contemporaneous affect, compared to more evaluative measures of well-being.

[9] The happiness question is, “Taken all together, how would you say things are for you these days? Would you say you are: Not too happy; Pretty happy; Very happy?”

Alternative Evaluative Measures of Subjective Well-Being

To provide further support for this argument we find that rainfall variability has a similar effect on an alternative evaluative measure of life satisfaction, the Cantril Ladder Scale. The Cantril Scale asks respondents to imagine a ladder with steps numbered from zero at the bottom to 10 at the top. The top of the ladder represents the best possible life, and the bottom of the ladder represents the worst possible life. The respondent is asked *"Which step of the ladder would you say you personally feel you stand at this time?"* As discussed different measures of subjective well-being provide different perspectives on the process by which respondents reflect on, or experience, their lives. The Cantril ladder lies closer to the end of the continuum representing more evaluative judgements of life, similar to the measure of life satisfaction. The results estimating the effects of rainfall variability of the respondents Cantril Ladder score are reported in Table 2.9 and Figure 2.8. These estimates are more similar to the estimated effects on life satisfaction, than the insignificant estimate of rainfall variability on happiness (present affect).

2.8 Conclusion

The ability to manage consumption risk is a significant determinant of individual and household welfare in developing countries, where households live in an uncertain environment with limited access to formal financial markets. While the realised effects of income shocks are well understood, this paper has explored the empirical relevance of future income uncertainty to the welfare of smallholder farmers in rural Ethiopia.

We first presented a simple model based on Brunnermeier and Parker (2005), which demonstrates how expected future utility, or anticipated utility, can have a direct contemporaneous effect on utility, if farmers face imperfect insurance and have imperfect information about the probability that a future income shock is realised. The model predicts that farmers living in areas with greater income uncertainty will have lower well-being than comparable farmers living in areas with less income uncertainty.

However, a central challenge is measuring and identifying the effects of uncertainty. Specifically, the challenge relates to disentangling the impact of uncertainty from the impact of realised events. Using panel data on smallholder farmers in rural Ethiopia combined with high-resolution atmospheric data, we exploited plausibly exogenous variation in inter-annual rainfall variability – a proxy for income uncertainty after controlling for contemporaneous and historical rainfall shocks – to examine the effects of income uncertainty on objective consumption and subjective well-being.

Empirically, we began by providing support for the premise that rainfall variability is a reasonable proxy for future income uncertainty. First we demonstrated that an increase in rainfall variability, measured over the previous five years, is associated with a reduction in realised rainfall in the season following the survey. In addition, we provided support for the premise that farmers are aware of this signal. Using data on risk perceptions, we showed that an increase in rainfall variability is associated with an increase in the belief that the rains will fail. Furthermore, we showed that an increase in the belief that the rains will fail is associated with lower rainfall realisations in the following season; however, once we control for rainfall variability, farmers' beliefs no longer enter significantly into the relationship, indicating that beliefs about future rainfall realisations are largely driven by rainfall variability. Finally, we demonstrated that rainfall variability has no direct effect on agricultural production. Collectively, this evidence supports the premise that rainfall variability is a reasonable proxy for future income uncertainty, helping us to disentangle the effects of future income uncertainty from the effects of realised events.

While there are no effects of rainfall variability on agricultural production, we estimate that a one standard deviation increase in rainfall variability would be associated with a 9.7–17.7% reduction in real consumption per capita, indicating the presence of a precautionary saving channel. In addition, we estimate that a one standard deviation increase in rainfall variability would be associated with a 0.171–0.486 point reduction in life satisfaction. We decompose this effect into the direct effects of income uncertainty and those mediated by the effects of income uncertainty through changes in consumption, finding that only 20% of the estimate effect is mediated by consumption. Finally, we demonstrate the effects of rainfall variability on alternative measures of subjective well-being. We show that rainfall variability has no impact on happiness, a measure of present affect, but that it does have an effect on the Cantril Ladder Scale, a more evaluative measure of subjective well-being similar to life satisfaction.

Our results suggest three things: first, that the returns to consumption smoothing and the welfare gains associated with managing both short-run weather events and long-run climate change are likely to be substantially greater than estimates based solely on consumption fluctuations and realised shocks; second, that in understanding the consequences of environmental change, it is important to understand how expectations about future states of the world affect economic behaviour, as well as the consequences of realised change. Finally, that the inclusion of subjective welfare measures alongside objective measures will better allow researchers and policy makers to understand the economic lives of the poor and evaluate broader welfare effects associated with policy interventions, mitigating omitted variable bias in cost–benefit analyses.

Table 2.1: Descriptive Statistics

	MEAN	STD. DEV. (Within)	STD. DEV. (Between)	OBS.
Panel A: *Outcome Measures*				
LIFE SATISFACTION (score/max)	0.567	0.137	0.223	3,869
LOG REAL CONSUMPTION PER CAPITA	3.970	0.392	0.695	3,869
Panel B: *Weather Data*				
Proxies for Future Income Uncertainty				
RAINFALL VARIABILITY (σ/μ)	21.577	3.171	7.077	3,869
RAINFALL VARIABILITY (σ, mm)	299.849	46.592	82.833	3,869
Contemporaneous Weather				
TOTAL RAINFALL (mm)	1,424.25	197.930	456.294	3,869
AVERAGE TEMPERATURE (°C)				
Historical Weather (5-year Averages)				
AVERAGE TOTAL RAINFALL (mm)	1,435.39	66.762	315.435	3,869
AVERAGE TEMPERATURE (°C)				

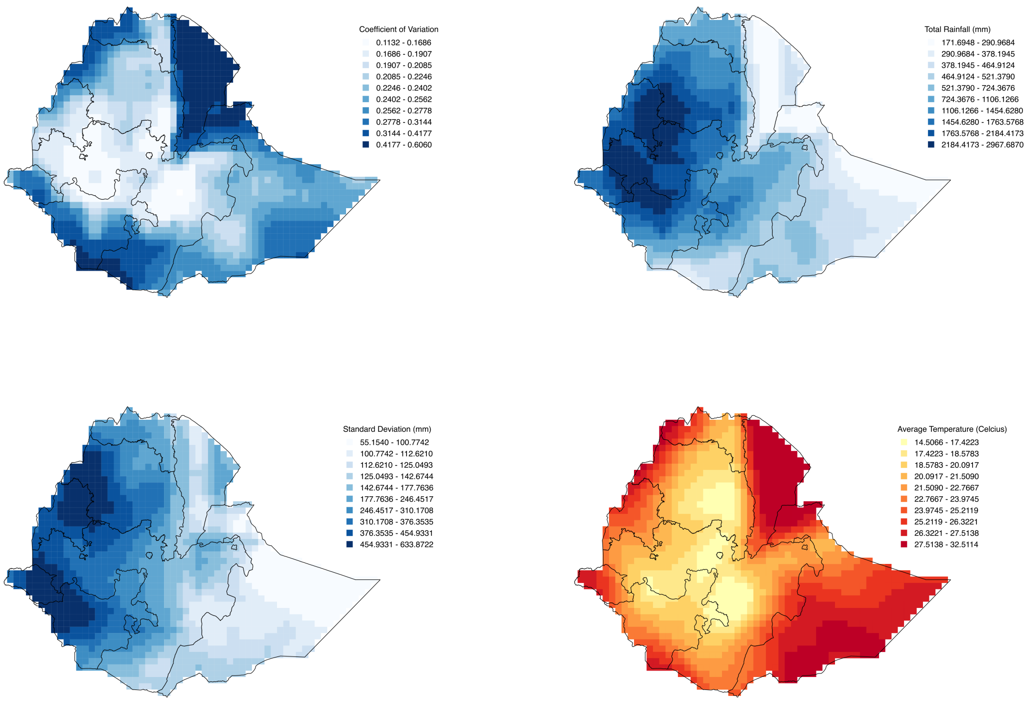

Figure 2.3: Spatial Variation in Rainfall and Temperature (1979–2012). Top Left = Coefficient of Variation; Top Right = Total Rainfall (mm); Bottom Left = Std Dev. Rainfall (mm); Bottom Right = Average Temperature (°C)

.

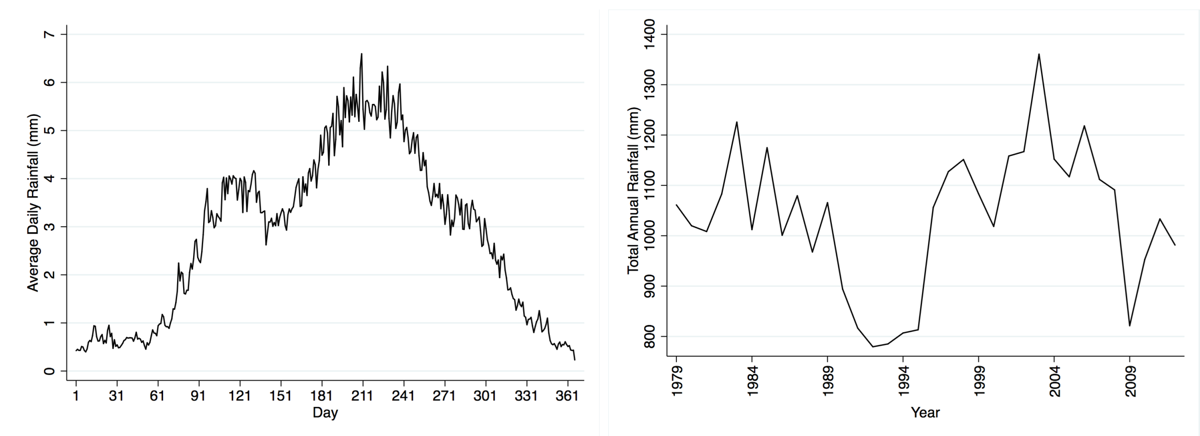

Figure 2.4: Temporal Variation in Rainfall (1979–2013). Top = Within-year Distribution (1979-2012 average). Bottom = Between-year Distribution

.

Table 2.3: Descriptive Statistics: Household Perceptions of Rainfall Risk

Question	Mean	Std. Dev.	Obs.
How often do the rains fail? Every . . . years.	2.867	1.650	1,212
Any Signs that rain may be timely or sufficient?	0.552	0.497	1,282
How likely is it that the rains will fail this year?	59.3%	16.3%	1,310

Note 1: The data is only available for 2009.

Note 2: Question: How often do the Meher rains fail for your land? Once every . . . years.

Note 3: Question: Are there any signs (e.g. in the temperature, the behaviour of animals, the rain received in the last years, the number of the year) that you might receive timely or sufficient rain? Yes = 1, No = 0.

Note 4: Question: Given the number of times the rains fail, and any signs you have observed, indicate by placing beans on the relevant square how likely it is that you think the Meher rains will fail this year. The more sure you are that they will fail the more beans should be placed on the square. If you think there is an equal chance that they will fail place the beans equally between the two squares. (20 beans in total) 0 beans = 0% likely, 20 beans = 100% likely.

Table 2.2: The Effects of Past Rainfall Variability on Future Rainfall

	(1) Meher Rainfall (mm)	(2) Meher Rainfall (mm)	(3) Meher Rainfall (mm)
Panel A: Coefficient of Variation:			
Rainfall Variability (σ/μ)	-2.875**	-4.024***	-4.171***
	(1.091)	(0.888)	(1.132)
Panel B: Standard Deviation:			
Rainfall Variability (σ, 100 mm)	-8.323	-29.014***	-29.657**
	(10.586)	(9.099)	(11.202)
Fixed Effects		Village and Year	
Weather Controls	No	Yes	Yes
Quadratic Weather Controls	No	No	Yes
Observations	480	480	480

Note 1: Specification: $\text{Rainfall}_{vt} = \beta_1 \text{CV}_{v\underline{t}} + \beta_2 \text{Rainfall}_{vt-1} + \beta_3 \overline{Rainfall}_{vt-5} + \alpha_v + \alpha_t + \varepsilon_{vt}$.

Note 2: Meher Rainfall is defined as the total cumulative rainfall received between June and November for the following season. See table **??** for definitions of the explanatory variables.

Table 2.5: The Effects of Rainfall Variability on Agricultural Yields

	(1) log Yields	(2) log Yields	(3) log Yields
Panel A: Coefficient of Variation:			
Rainfall Variability (σ/μ)	-0.0347***	-0.0118	-0.00173
	(0.00897)	(0.0114)	(0.0126)
Panel B: Standard Deviation:			
Rainfall Variability (σ, 100 mm)	-0.257***	-0.114	-0.073
	(0.0724)	(0.0785)	(0.107)
Fixed Effects	Household and Year		
Weather Controls	No	Yes	Yes
Quadratic Weather Controls	No	No	Yes
Observations	2,334	2,334	2,334

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. The unit of analysis is at the household level. Our proxies for uncertainty are the coefficient of variation for rainfall (Panel A) and the standard deviation of rainfall (Panel B), measured over the previous 5 years, the time period between each survey round. Historical measures of atmospheric parameters correspond to this period. Contemporaneous and historical rainfall is measured in hundreds of mm. Contemporaneous and historical temperature is measured in °C. Standard errors are adjusted to reflect spatial dependence, as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cut-off of 150 km. The distance is selected following a decision rule in which we choose the distance that provides the most conservative standard errors, looped over all distances between 10 and 1000km.

Table 2.4: Household Perceptions of Rainfall Risk – Correlations

	(1) Positive Signs	(2) Positive Signs	(3) Belief that rains will fail	(4) Belief that rains will fail
Rainfall Variability (σ/μ)	0.0003	-0.00404**	0.00414***	0.00291**
	(0.00159)	(0.0017)	(0.00104)	(0.00131)
Constant	0.546***	0.689***	0.526***	0.523***
	(0.0485)	(0.123)	(0.0231)	(0.0551)
Weather Controls	No	Yes	No	Yes
Observations	1,282	1,282	1,282	1,282

	(1) Meher Rainfall (mm)	(2) Meher Rainfall (mm)	(3) Meher Rainfall (mm)	(4) Meher Rainfall (mm)
Positive Signs	9.107		11.240	9.719
	(14.454)		(14.352)	(13.492)
Belief that rains will fail		-131.566**	-131.227**	68.918
		(41.865)	(41.956)	(44.328)
Rainfall Variability (σ/μ)				-6.050***
				(0.319)
Constant	682.688***	770.427***	759.458***	738.988***
	(10.855)	(25.566)	(27.165)	(25.603)
Observations	1,282	1,310	1,282	1,282

Note 1: The data is only available for 2009.
Note 2: Specification 1: $Y_{iv} = \beta_0 + \beta_1 \text{CV}_{vt} + \beta_2 \text{Rainfall}_{vt-1} + \beta_3 \overline{Rainfall}_{vt-5} + \varepsilon_{iv}$. Standard errors are clustered at the village level. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01.
Note 3: Specification 2: Meher $\text{Rainfall}_v = \beta_0 + \beta_1$ Positive $\text{Signs}_{iv} + \beta_2 \text{Belief}_{iv} + \varepsilon_{iv}$. Standard errors are clustered at the household level. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01.

Note 4: Meher Rainfall is defined as the total cumulative rainfall received between June and November for the following season. See table **??** for definitions of the explanatory variables.

Table 2.7: Uncertainty and Life Satisfaction

	Life Satisfaction				
	(1)	(2)	(3)	(4)	(5)
Panel A: Coefficient of Variation:					
Rainfall Variability (σ/μ)	-0.0332***	-0.0261***	-0.0217***	-0.0742***	-0.0685***
	(0.00792)	(0.00554)	(0.00557)	(0.00500)	(0.00535)
Panel B: Standard Deviation:					
Rainfall Variability (σ, 100 mm)	-0.00245***	-0.00158***	-0.00128**	-0.00523***	-0.00480***
	(0.000700)	(0.000531)	(0.000537)	(0.000395)	(0.000398)
Fixed Effects	Individual, Year, Month				
Weather Controls	No	Yes	Yes	Yes	Yes
Quadratic Weather Controls	No	No	No	Yes	Yes
Consumption Control	No	No	Yes	No	Yes
Observations	4,064	4,064	4,064	4,064	4,064

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. The unit of analysis is at the individual level. Our proxies for uncertainty are the coefficient of variation for rainfall (Panel A) and the standard deviation of rainfall (Panel B), measured over the previous 5 years, the time period between each survey round. Historical measures of atmospheric parameters correspond to this period. Contemporaneous and historical rainfall is measured in hundreds of mm. Contemporaneous and historical temperature is measured in °C. Standard errors are adjusted to reflect spatial dependence, as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cut-off of 150 km. The distance is selected following a decision rule in which we choose the distance that provides the most conservative standard errors, looped over all distances between 10 and 1000km.

Figure 2.6: Semi-parametric Estimates of the Relationship between Rainfall Variability and Life Satisfaction.

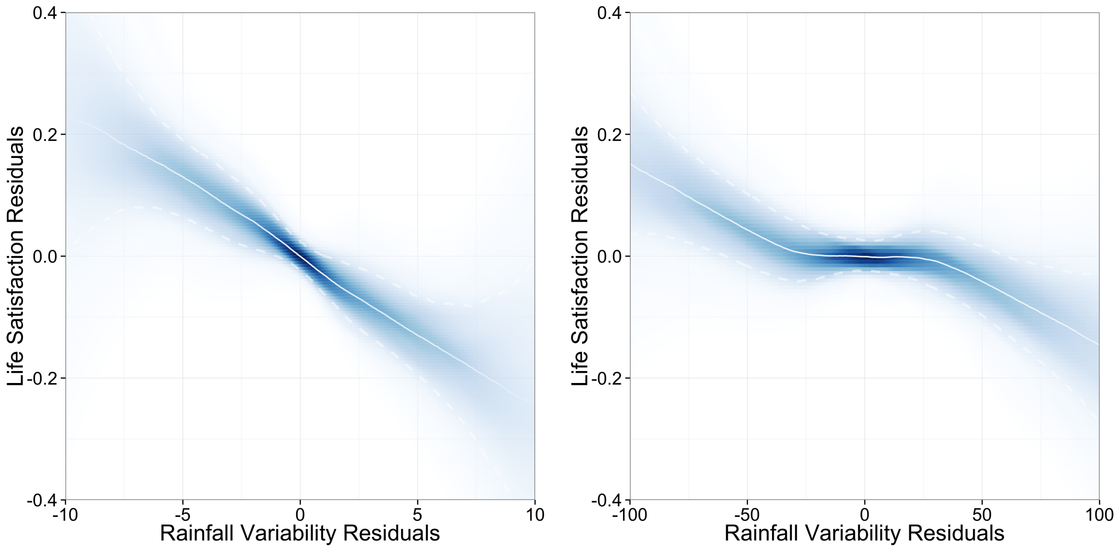

Notes: The measures of rainfall variability are the coefficient of variation (left) and the standard deviation of rainfall, measured over the previous 5 years, the time period between each survey round. Each variable is regressed on individual, year and month fixed effects as well as contemporaneous and historical rainfall and temperature controls. The figures above are the result of loess regressions of the residuals from this exercise.

Table 2.6: Uncertainty and Consumption

	Log Real Consumption Per Capita					
	(1) (2004 - 2009)	(2) (2004 - 2009)	(3) (2004 - 2009)	(4) (1995 - 2009)	(5) (1995 - 2009)	(6) (1995 - 2009)
Panel A: Coefficient of Variation:						
Rainfall Variability (σ/μ)	-0.0189***	-0.0149***	-0.0271***	-0.00999***	-0.0121***	-0.0118***
	(0.00291)	(0.00456)	(0.00129)	(0.00382)	(0.00378)	(0.00450)
Panel B: Standard Deviation:						
Rainfall Variability (σ, 100 mm)	-0.143***	-0.106***	-0.207***	-0.0676***	-0.0893***	-0.0873**
	(0.0252)	(0.0345)	(0.0200)	(0.0247)	(0.0272)	(0.0386)
Fixed Effects	Village, Year, Month					
Weather Controls	No	Yes	Yes	No	Yes	Yes
Quadratic Weather Controls	No	No	Yes	No	Yes	Yes
Observations	2,721	2,721	2,721	2,721	2,721	2,721

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. The unit of analysis is at the household level. Our proxies for uncertainty are the coefficient of variation for rainfall (Panel A) and the standard deviation of rainfall (Panel B), measured over the previous 5 years, the time period between each survey round. Historical measures of atmospheric parameters correspond to this period. Contemporaneous and historical rainfall is measured in hundreds of mm. Contemporaneous and historical temperature is measured in °C. Standard errors are adjusted to reflect spatial dependence, as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cut-off of 150 km. The distance is selected following a decision rule in which we choose the distance that provides the most conservative standard errors, looped over all distances between 10 and 1000km.

Figure 2.5: Semi-parametric Estimates of the Relationship between Rainfall Variability and Log Real Consumption Per Capita.

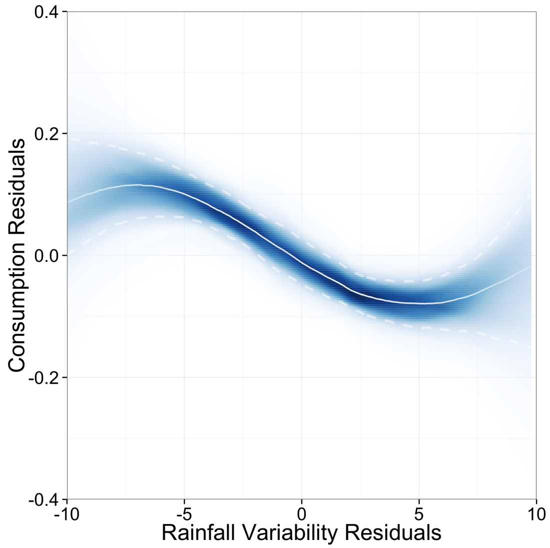

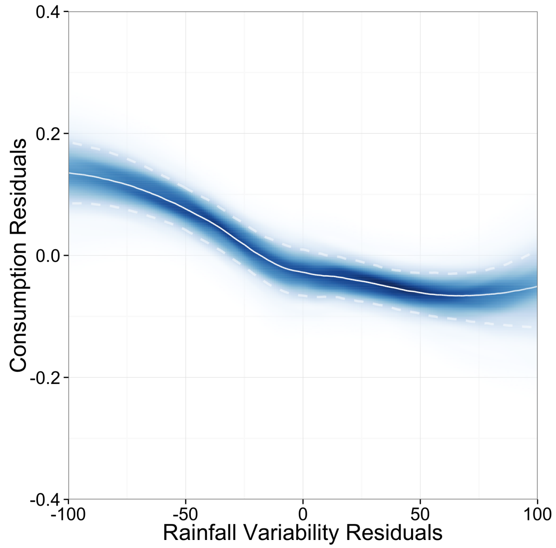

Notes: The measures of rainfall variability are the coefficient of variation (left) and the standard deviation of rainfall, measured over the previous 5 years, the time period between each survey round. Each variable is regressed on village, year and month fixed effects as well as contemporaneous and historical rainfall and temperature controls. The figures above are the result of loess regressions of the residuals from this exercise.

Table 2.8: Uncertainty and Happiness

	Happiness				
	(1)	(2)	(3)	(4)	(5)
Panel A: Coefficient of Variation:					
Rainfall Variability (σ/μ)	-0.00400***	-0.00438	-0.00331	-0.0128**	-0.0114*
	(0.00108)	(0.00332)	(0.00343)	(0.00622)	(0.00609)
Panel B: Standard Deviation:					
Rainfall Variability (σ, 100 mm)	-0.0319***	-0.0349	-0.0278	-0.123***	-0.112***
	(0.00963)	(0.0230)	(0.0242)	(0.0443)	(0.0435)
Fixed Effects	Individual, Year, Month				
Weather Controls	No	Yes	Yes	Yes	Yes
Quadratic Weather Controls	No	No	No	Yes	Yes
Consumption Control	No	No	Yes	No	Yes
Observations	4,064	4,064	4,064	4,064	4,064

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. The unit of analysis is at the individual level. Our proxies for uncertainty are the coefficient of variation for rainfall (Panel A) and the standard deviation of rainfall (Panel B), measured over the previous 5 years, the time period between each survey round. Historical measures of atmospheric parameters correspond to this period. Contemporaneous and historical rainfall is measured in hundreds of mm. Contemporaneous and historical temperature is measured in °C. Standard errors are adjusted to reflect spatial dependence, as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cut-off of 150 km. The distance is selected following a decision rule in which we choose the distance that provides the most conservative standard errors, looped over all distances between 10 and 1000km.

Figure 2.7: Semi-parametric Estimates of the Relationship between Rainfall Variability and Happiness.

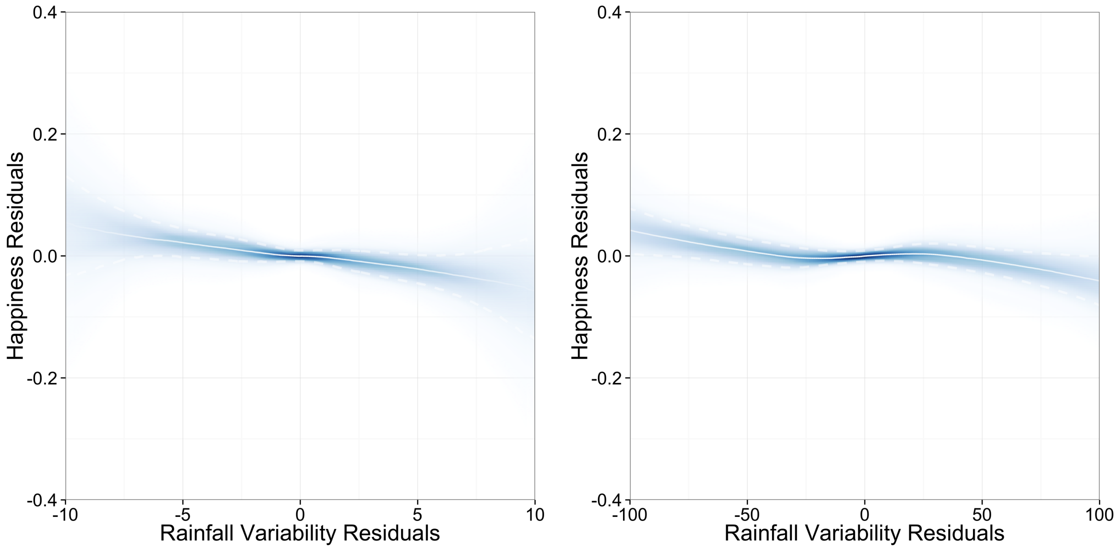

Notes: The measures of rainfall variability are the coefficient of variation (left) and the standard deviation of rainfall, measured over the previous 5 years, the time period between each survey round. Each variable is regressed on individual, year and month fixed effects as well as contemporaneous and historical rainfall and temperature controls. The figures above are the result of loess regressions of the residuals from this exercise.

Table 2.9: Uncertainty and the Cantril Ladder Scale

	Cantril Ladder Scale				
	(1)	(2)	(3)	(4)	(5)
Panel A: Coefficient of Variation:					
Rainfall Variability (σ/μ)	-0.0370***	-0.0163**	-0.0129*	-0.0344**	-0.0284*
	(0.0142)	(0.00738)	(0.00716)	(0.0138)	(0.0148)
Panel B: Standard Deviation:					
Rainfall Variability (σ, 100 mm)	-0.300***	-0.141**	-0.118*	-0.287***	-0.243**
	(0.114)	(0.0673)	(0.0651)	(0.107)	(0.115)
Fixed Effects		Individual, Year, Month			
Weather Controls	No	Yes	Yes	Yes	Yes
Quadratic Weather Controls	No	No	No	Yes	Yes
Consumption Control	No	No	Yes	No	Yes
Observations	4,060	4,060	4,060	4,060	4,060

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. The unit of analysis is at the individual level. Our proxies for uncertainty are the coefficient of variation for rainfall (Panel A) and the standard deviation of rainfall (Panel B), measured over the previous 5 years, the time period between each survey round. Historical measures of atmospheric parameters correspond to this period. Contemporaneous and historical rainfall is measured in hundreds of mm. Contemporaneous and historical temperature is measured in °C. Standard errors are adjusted to reflect spatial dependence, as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cut-off of 150 km. The distance is selected following a decision rule in which we choose the distance that provides the most conservative standard errors, looped over all distances between 10 and 1000km.

Figure 2.8: Semi-parametric Estimates of the Relationship between Rainfall Variability and the Cantril Ladder Scale.

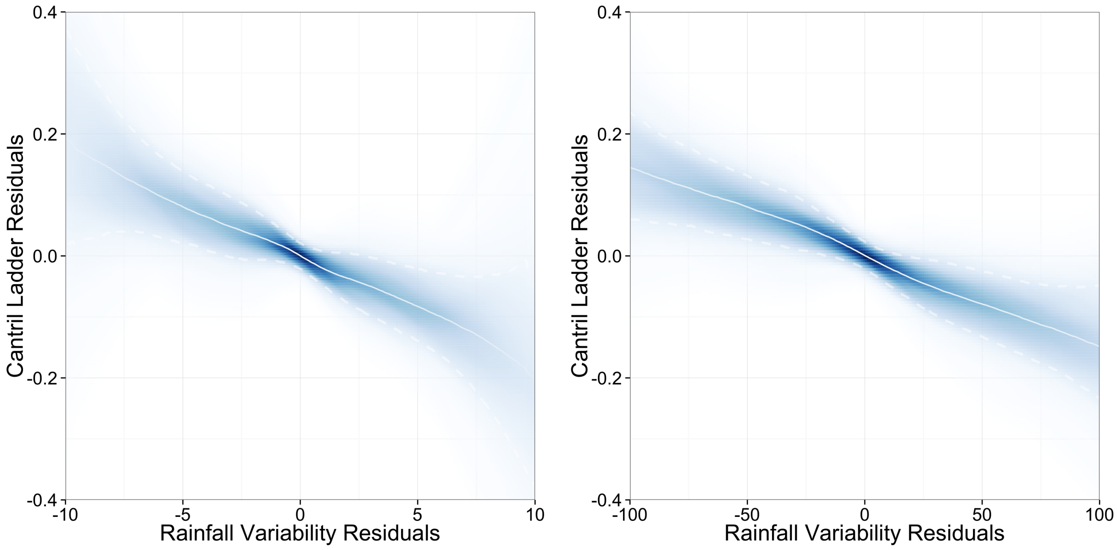

Notes: The measures of rainfall variability are the coefficient of variation (left) and the standard deviation of rainfall, measured over the previous 5 years, the time period between each survey round. Each variable is regressed on individual, year and month fixed effects as well as contemporaneous and historical rainfall and temperature controls. The figures above are the result of loess regressions of the residuals from this exercise.

Chapter 3

Parental Income Uncertainty, Child Labour, and Human Capital Accumulation

How does parental income uncertainty affect child labour and human capital investments in village economies? Theoretically, the relationship between income uncertainty and human capital is ambiguous: on the one hand, a precautionary response could reduce investments in human capital; on the other hand, a portfolio motive could increase investments in human capital as households attempt to diversify their income streams. Using child-level panel data from rural Ethiopia, I estimate the effects of parental income uncertainty – proxied by rainfall variability, after controlling for contemporaneous and historical weather events – on child labour and educational outcomes. I find that an increase in uncertainty at the time of the survey is associated with a reduction in the number of hours children spend working on the farm and an increase in the number of hours spent studying at home, suggesting that parents invest more in human capital as a response to increases in income uncertainty. Consistent with such a response, I find that an increase in parental income uncertainty is also associated with an increase in the likelihood that a child attends school and an increase in the number of grades achieved. However, consistent with the precautionary motive, I estimate that an increase in parental income uncertainty during the early stages of a child's life cycle – when the portfolio response is weakest – is associated with a reduction in the likelihood that the child attends school. This relationship weakens and reverses as the child grows older and the returns to education, and consequently diversification, increase.

3.1 Introduction

It has long been established that human capital accumulation is an important driver of economic growth (Mankiw et al., 1992); however, less clear is how economic conditions determine human capital accumulation. While a large literature has sought to understand the relationship between realised economic conditions and human capital accumulation in developing countries (Jacoby and Skoufias, 1997; Jensen, 2000; Schady, 2004; Thomas et al., 2004; Beegle et al., 2006; Krueger, 2007; Maccini and Yang, 2009; Shah and Steinberg, 2015), there has been little examination of how economic uncertainty affects the investment decisions of households. This paper seeks to understand the impact of parental income uncertainty on human capital accumulation in developing countries, exploring its affects on the child labour and education decisions of smallholder farmers in rural Ethiopia – a context where insurance and credit market failures impede the ability of households to smooth consumption, but where access to education is relatively unconstrained, with the largest opportunity cost relating to foregone labour.

Theoretically, the relationship between parental income uncertainty and human capital accumulation is ambiguous. On the one hand, households may reduce investments in human capital, allocating more time to child labour on the farm in an effort to mitigate the economic consequences of productivity shocks in the event that they are realised – a "precautionary motive". On the other hand, an increase in income uncertainty could be associated with an increase in human capital investments as households adjust the time allocation of children away from risky activities on the farm towards less risky investments – a "portfolio motive". In section 3.3 I introduce a simple model of human capital investment under uncertainty to formalise these concepts and guide the interpretation of the empirical analysis.

The main empirical challenge associated with understanding the effects of uncertainty on economic behaviour is in separating the effects of uncertainty from the realisation of economic events. One of the key features of the economic environment in rural Ethiopia is the role that rain-fed agriculture plays as a source of income. Ethiopia is one of the least developed countries in Africa and is characterised by its high vulnerability to weather shocks. With limited irrigation opportunities, farmers are dependent on the timing and intensity of rainfall to produce crops for subsistence production. Consequently, in the absence of effective consumption smoothing technologies, uncertainty about future states of the world can have a significant effect on decision-making and individual welfare, above and beyond the realisation of income shocks (Alem and Colmer, 2015). Using panel data on the children of smallholder farmers in rural Ethiopia combined with high-resolution meteorological data, I measure and identify the effects of income uncertainty on child labour and educational investments by exploiting plausibly exogenous variation in rainfall variability (the second moment of the rainfall distribution), which is a plausible measure of parental income uncertainty after controlling for contemporaneous and historical rainfall shocks, (the first-moment).

I begin by providing supporting evidence for the premise that rainfall variability is a reasonable proxy for parental income uncertainty. First, I demonstrate that an increase in rainfall variability, measured over the previous five years (the time between survey rounds) is associated with a reduction in realised rainfall in the rainy season following the survey. I also provide support for the premise that farmers are aware of this signal. Using data on risk perceptions, I show that an increase in rainfall variability is associated with an increase in the belief that the rains will fail. Furthermore, I show that an increase in the belief that the rains will fail is correlated with lower rainfall realisations in the following season, indicating that farmers' beliefs about rainfall realisations correlate pro-cyclically with actual rainfall realisations; however, once I control for rainfall variability, farmers' beliefs no longer enter significantly into this relationship, suggesting that beliefs about future income are largely driven by rainfall variability. Finally, I present evidence to suggest that rainfall variability has no direct effect on agricultural production. All of these considerations provide support for the premise that rainfall variability is a reasonable proxy for parental income uncertainty, helping to isolate the effects of future income uncertainty from realised events.

Having shown that rainfall variability is a plausible proxy for parental income uncertainty, after controlling for contemporaneous and historical weather events, I examine the effects of this mechanism on the time allocation of children, as well as their participation in educational activities.

In line with the "portfolio motive" I find that, on average, an increase in parental income uncertainty is associated with a reallocation of time away from labour on the farm towards time spent on educational activities, particularly home study. Consistent with this reallocation, I also find that an increase in parental income uncertainty is associated with an increase in the likelihood that children attend school and the number of grades attained. These results suggest that farmers invest in the human capital, and adjust the time allocation, of their children away from risky activities on the farm towards less risky investments.

However, parental income uncertainty may have different effects on human capital investments at different points of the child's life cycle. For example, the precautionary motive may have a stronger effect in the earlier stages of a child's life, as parents discount the future benefits of human capital investments. To understand whether this is the case, I estimate the effects of historical rainfall variability (defined as the coefficient of variation for rainfall in the 5 years before a child reached their n^{th} birthday) on the likelihood that they are attending school at the time of the survey. In doing so, I estimate the effects of parental income uncertainty at each age of the child from birth until their age at the time of the survey. I find that in the early stages of the child's life cycle – when the portfolio motive is weakest – an increase in parental income uncertainty is associated with a reduction in the likelihood that children attend school at the time of the survey. As investment decisions in human capital are time-sensitive and irreversible, short-run changes in the local economic environment may lead households to sacrifice valuable investments, with a long-run and persistent impact (Udry, 1994; Jacoby and Skoufias, 1997; Duflo, 2000; Maccini and Yang, 2009; Banerjee and Mullainathan, 2010). However, as the child grows older this relationship weakens and reverses as the returns to education, and consequently income diversification, increase.

Together these results indicate that the impact of short-run weather events, as well as long-run climatic change, are likely to be greater than estimates based solely on realised shocks, and that, on average, households appear to be engaging in defensive investments to mitigate exposure to future income shocks. However, the impact of changes in parental income uncertainty on human capital investments depends on where in the life cycle the child is at the time. I demonstrate evidence that in the early stages of the life cycle, increases in parental income uncertainty result in a reduction in human capital investments.

The remainder of the paper is organised as follows: section 3.2 provides a brief background and literature review; section 3.3 presents a simple theoretical model which formalises the theoretical channels, helping to guide the empirical findings; section 3.4 summarises the data; section 3.5 provides supporting evidence for the premise that rainfall variability is a suitable proxy for parental income uncertainty; section 3.6 describes the empirical specification and outlines the identification strategy; section 3.7 discusses the results; the final section summarises the implications of these results and concludes.

3.2 Literature Review

In traditional rural economies children may be withdrawn from school as a response to income shocks, with long-run and irreversible impacts on human capital and, consequently, lifetime earnings. However, economic uncertainty may also affect human capital investments, above and beyond the realisation of events. Given that investments in human capital are irreversible, "real options" may be generated, whereby uncertainty about the future may result in a "wait and see" approach to investment (Dixit and Pindyck, 1994). An analogous channel arises when considering uncertainty from a consumption perspective. When there is greater economic uncertainty, consumers may delay investments in durable consumption. Instead households have a greater incentive to accumulate precautionary savings to smooth consumption against future risk (Kimball, 1991; Paxson, 1992; Carroll, 1997; Carroll and Kimball, 2001). Consequently, in the presence of uninsured risk, prudent households are likely to save more than in the absence of uncertainty.

One channel through which this effect may be observed, first discussed by Cain (1982), and more recently by Fitzsimons (2007), is the decision to not enrol children in school. Education is an irreversible investment with delayed, and potentially increasing, marginal returns. Fafchamps and Pender (1997) further argue that the precautionary motive for holding liquid assets impedes productive investment, such as investments in education, even if households are able to self-finance them. As a result, it is argued that the effect of the precautionary motive on irreversible and illiquid investments, such as education, is augmented. There are two mechanisms that can be explored here. The first mechanism arises from the decision not to enrol children in school to increase saving through reduced educational expenditures with an assumed increase in the labour supply of children, or increased idleness. The second mechanism results from a risk-management and productivity motive by which households may invest more in the land, taking more care over the land-preparation and cultivation stages in efforts to reduce the likelihood of crop failure in the event of an adverse shock. This mechanism works along the intensive margin of labour supply. The question that remains is whether an increase in labour supply along the intensive margin is sufficient to also affect decisions on the extensive margin. Within this literature it is assumed that parents optimally invest in the number and quality of children, determined by investments in human capital, to maximise household welfare. Under this assumption, Fitzsimons (2007) argues that children have an instantaneous earnings potential in addition to the benefit of reduced educational expenditures. Consequently, we might expect that higher levels of economic uncertainty may increase the number of hours worked by children and reduce investments in education.

If increased uncertainty results in reduced human capital investments then there may be long-run welfare costs through worsened later-life outcomes and opportunities (Strauss and Thomas, 1998; Maccini and Yang, 2009; Banerjee and Mullainathan, 2010; Antilla-Hughes and Hsiang, 2013). Even delays in educational attainment may have large effects if children are not able to reach a level of education that has real returns. Finally, these costs may be further exacerbated if households reduce investments in children based on gender (Sen, 1990; Duflo, 2005; Antilla-Hughes and Hsiang, 2013).

However, it is unclear whether this precautionary motive will dominate the decision-making of households. An opposing motive may result in an increase in human capital investments if such investments would allow households to diversify away from risky farming activities – a portfolio motive. Uncertainty about future income has been shown to be a driver of income source diversification in rural developing economies (Dercon and Krishnan, 1996; Dercon, 2002; Fafchamps, 2003). Given that educational investments affect the economic opportunities to which children have access in the future, the choice to diversify between time spent on farming and education, increasing investments in formal human capital, may be a response to increases in uncertainty. In effect, this can be seen as a diversification of human capital. In traditional agricultural economies such as rural Ethiopia, the empirical context of this paper, child labour can provide an opportunity for children to accumulate human capital informally through learning-by-doing (Rosenzweig and Wolpin, 1985; Grootaert and Kanbur, 1995; Rosenzweig, 1995; Fafchamps and Quisumbing, 2007; Lilleør, 2015). Consequently, the portfolio motive can be seen as a diversification of human capital.

Given these competing effects, whether uncertainty increases or decreases investments in human capital, is an empirical question: if the precautionary effect dominates, then an increase in uncertainty will reduce investments in formal education; if the portfolio effect dominates, then an increase in uncertainty will increase investments in formal education. However, the mapping of time allocation between time spent on education and time spent on child labour activities is an important consideration.

While it is commonly believed that an increase in time spent on child labour activities will translate into a reduction in educational investments, this may not be the case. In many studies – this study included – we observe that children are capable of both working and attending school.[1] Ravallion and Wodon (2000) argue that poor families can protect the educational investments of working children because there are other things that children do besides school and work.

"One cannot assume that the time these children spend working must come at the expense of formal time at school, although there may be displacement of informal (after-school) tutorials or homework."[2]

[1] 75% of the sample both attend school and work, either on the farm or in the home. 49% attend school and work on the farm. 59% attend school and work in the home.

[2] Jayachandran (2013) demonstrates that the displacement of informal schooling may have significant welfare effects of its own. If schools offer for-profit tutoring to their own students, this gives teachers a perverse incentive to teach less during school to increase demand for tutoring. Consequently, those who do not participate in out of school human capital investments could be adversely affected. It is not possible to test this implication due to the absence of test scores.

As a result, it is unclear whether reductions in educational investments will necessarily translate into increased child labour, and vice versa. Consequently, it is important to directly assess the time-use of children to better understand the reallocation of time across activities in response to changes in economic conditions, rather than focussing solely on whether children engage in education and child labour. By using data on the time allocation of children, this paper is able to understand both intensive-margin and extensive-margin responses to economic uncertainty.

3.3 Theoretical Framework

To motivate the empirical work and consider the theoretical implications of uncertainty on human capital accumulation, I introduce a two-period model of human capital investment, where parents invest in the human capital of their children in the first period and in turn receive a payoff from the child's accumulated human capital in the second period. The number of children and parents are normalised to one, and the parent maximises the total utility of the household.

Let c_t be consumption in period t, and $u(c_t)$ be the flow of utility from consumption during this period, where $u'(c_t) > 0$ and $u''(c_t) < 0$, $\forall\ t$. The household solves the following expected utility maximisation problem,

$$U(c_1, c_2) = u(c_1) + \beta \mathbb{E}[u(c_2)]$$

The parent works and earns an exogenous income $w_t L$. By contrast, the child's time is allocated between school and child labour. This decision is made by the parent, jointly with the decision over consumption. I abstract from borrowing and saving decisions such that consumption is equal to income in each period. Consequently, during the first period, the household income is equal to the earnings of the parent $w_1 L$, plus the returns to child labour, $(1 - s_1)w_1 h_1)$, where, $s_1 \in [0, 1]$ denotes the share of time that children spend in school accumulating human capital. I assume that child labour income in period 1 is not a function of human capital (Rosenzweig, 1980; Jacoby and Skoufias, 1997; Fitzsimons, 2007).

In the second period, the household receives a payoff from the child having accumulated human capital. Following Heckman and Cunha (2007), I assume that the income received in period 2 is a function of the human capital in the previous period, plus any investments made in the previous period. Normalising the initial human capital of the child, h_1, to one, the payoff received to the household is conditional on schooling investments $h_2 = f(s_1)$. I assume that $\frac{\partial f}{\partial s_1} \geq 0$, and $\frac{\partial^2 f}{\partial s_1^2} \leq 0$, implying that schooling results in weakly more human capital and that there are diminishing marginal returns to schooling. During the second period the household income is equal to the earnings of the parent, w_2L, which are uncertain in period 1, and the returns to child labour, w_2h_2).

Consequently, the parent solves the following problem,

$$\max_{c_1, s_1 \in [0,1]} u(c_1) + \beta \mathbb{E}[u(c_2)] \tag{3.1}$$

subject to,

$$\begin{aligned} c_1 &\leq w_1(L + (1 - s_1)) \\ c_2 &\leq w_2(L + h_2) \end{aligned}$$

Households maximise current and discounted future utility, subject to the constraint that total consumption, c, cannot be higher than the labour market income in both periods. Since utility is increasing in consumption, and there is no borrowing or saving, it will always be the case that $c_1 = w_1(L + (1 - s_1))$ and $c_2 = w_2(L + h_2)$. Consequently, I can substitute this into the maximisation problem to get,

$$\max_{s_1 \in [0,1]} u(w_1(L + (1 - s_1)) + \beta \mathbb{E}[u(w_2(L + h_2)]$$

At an interior optimum, households equalise the marginal utility of consumption by foregoing investments in schooling with the marginal benefit of human capital in later periods:

$$w_1 u'(c_1) = \beta \mathbb{E}[u'(c_2)] w_2 f'(s_1) \tag{3.2}$$

That is, households trade off the marginal benefit of additional utility from consumption for the net long-term benefit of investments in human capital. I am interested in understanding the effect that parental income uncertainty has on the optimal level of schooling; i.e., as income uncertainty increases, do parents invest more or less in schooling and, consequently, do overall levels of human capital increase or decrease?

3.3.1 The effect of parental income uncertainty on human capital

To isolate the effects of income uncertainty on human capital investment, I explore the effects of an increase in income dispersion, defined as a combination of additive and multiplicative shifts in the distribution of parental income (Sandmo, 1970; Fitzsimons, 2007; Fafchamps, 2009).

I rewrite period two parental income as, $\gamma(w_2L) + \theta$, the expected value of which is, $\mathbb{E}[\gamma(w_2L) + \theta]$. Here γ is the multiplicative shift parameter and θ is the additive one. Under the assumption of non-negative income, a multiplicative shift around zero will increase the mean. Consequently, for there to be a mean-preserving increase in risk, this must be counteracted by a negative shift in the additive parameter, holding the expected value constant. This requires that the differential of the expected value of future income is 0; i.e., $\mathbb{E}[(w_2 L)\mathrm{d}\gamma + \mathrm{d}\theta] = 0$. This implies that $\mathrm{d}\theta/\mathrm{d}\gamma = -\mathbb{E}[w_2 L] = -\xi$.

With this I examine the effect of an increase in income dispersion, γ, on the optimal choice of schooling, s_1^*, and the resulting level of human capital, h_2^*,

$$\left.\frac{\partial s_1^*}{\partial \gamma}\right|_{\frac{\mathrm{d}\theta}{\mathrm{d}\gamma}=-\xi} \propto \frac{\beta u''(c_1)\mathbb{E}[u''(c_2)(w_2 L - \xi)] \cdot [w_1 - w_2 f(s_1^*)]}{|H|}$$

where $|H|$ is the determinant of the Hessian matrix, known to be positive based on the second-order conditions. Following Sandmo (1970), it is straightforward to show that decreasing absolute risk aversion is a sufficient condition for the first component, $u''(c_1)\mathbb{E}[u''(c_2)(w_2 L - \xi)$, to be negative, such that increased uncertainty about future income decreases investments in human capital. However, the sign of the second term $[w_1 - w_2 f(s_1^*)]$ is ambiguous. If the effects of schooling on the wages of children are high enough, then this may offset the adverse effects of uncertainty on human capital investments. Consequently, whether increased uncertainty will increase or decrease investments in human capital is theoretically ambiguous.

3.4 Data

The main analysis uses two rounds of the Ethiopian Rural Household Survey collected by the University of Addis Ababa, the Centre for the Study of African Economics (CSAE) at the University of Oxford, and the International Food Policy Research Institute (IFPRI), covering 15 village communities in rural Ethiopia. This paper makes use of the latest two rounds of this panel, from 2004 and 2009. These years are included as they contain consistent measures of child labour across time.[3] The villages in the survey represent the diversity of farming systems throughout Ethiopia and capture climate differences across the country. Stratified random sampling is used within each village, based on whether households have male or female heads.

When analysing the time allocation of children, the dependent variables are defined as the total hours spent working in economic activities and domestic chores per week. In terms of the activities that children engage in, economic activities generally consist of farming activities, including land preparation, tending crops, processing

[3]The total survey consists of 7 rounds between 1989 and 2009. In 1989, households from six villages in central and southern Ethiopia were interviewed. In 1994, however, the sample was expanded to cover 15 villages across the country, representing 1477 households. Further rounds were completed in 1995, 1997, 1999, 2004 and 2009.

crops, and looking after livestock. I also look at domestic chores, for two reasons: first, child labour is not restricted to economic activities, and so it is interesting to consider the substitution of time among activities in response to an increase in uncertainty; second, in rural areas it may be difficult to distinguish between time spent on household chore activities and time spent on chores relating to farming activities.

In Table 4.1 we observe that the average number of hours spent per week across all child labour activities is 28 hours per week. This is a non-trivial amount of time and could result in a trade-off with time spent on educational and leisure activities.[4] This time is split evenly between labour on the farm and domestic chores.

In addition to the time spent on child labour activities I examine the time spent on human capital accumulation. I observe that children spend around 8 hours per week studying at home. Unfortunately, it is not possible to observe time spent studying in school. I therefore construct a residual measure that incorporates all other activities including leisure, sleep, and time spent studying in school to provide a full representation of time-use. While there are limitations to the conclusions that can be drawn from this measure, one may argue that all these components play an important role in accumulating human capital and so provide a useful and more broad measure for the intensive margin of human capital accumulation in the absence of more direct measures of the time spent in school. Unsurprisingly, children spend the majority of their time on these other activities, accounting for 132 hours per week. Being close to the equator, where there is roughly 12 hours of daylight each day throughout the year, sleep and nighttime activities may account for 84 hours (63%) of this residual with the remaining 48 hours (57%) being allocated among education and leisure.

In addition to the time allocated to these activities, I also construct binary variable measures of whether children engage in child labour activities – the extensive margin –, and more importantly for human capital accumulation, whether children attend school, and the number of grades that they have achieved. We observe that 92% of children engage in child labour activities. In addition, around 58% report that they attend school and have achieved an average of 3 grades.

To analyse the effects of parental income uncertainty on child labour and human capital investments, I combine the economic data with meteorological data constructed at the village level from the ERA-Interim data archive supplied by the European Centre for Medium-Term Weather Forecasting (ECMWF).[5] Where previous studies have relied on the use of meteorological data provided by the Ethiopian meteorological service, the number of missing observations, or observations which are recorded as zero on days when there are no records, is of concern. The ERA-Interim

[4] Ravallion and Wodon (2000) argue on a priori grounds that it would not be difficult for parents to assure that a child working 20 hours per week could still attend school.

[5] See Dee et al. (2011) for a detailed discussion of the ERA-Interim data.

reanalysis data archive provides daily measurements of many atmospheric parameters, of which precipitation and daily average temperature are exploited for the purposes of this paper. The data is available from 1 January 1979 until the present day, on a global grid of quadrilateral cells defined by parallels and meridians at a resolution of 0.25 x 0.25 degrees (equivalent to 28km x 28km at the equator).[6] Rainfall and temperature variables for each village are calculated through a process of inverse distance weighting, taking all data points within 100km of the village. The weight attributed to each grid point decreases quadratically with distance.

Reanalysis data is constructed through a process whereby model information and observations are combined to produce a consistent global best estimate of atmospheric parameters over a long period of time by optimally fitting a dynamic model to each period simultaneously (Auffhammer et al., 2013b). Models propagate information from areas with more observational data for areas in which observational data are scarce. This results in an estimate of the climate system that is separated uniformly across a grid, that is more uniform in its quality and realism than observations alone, and that is closer to the state of existence than any model would provide alone. This provides a consistent measure of atmospheric parameters over time and space. This type of data is increasingly being used by economists, especially in developing countries, where the quality and quantity of weather data is more limited (see Dell et al. (2014) for a review of its recent applications in the literature).

Figures 3.1 and 3.2 provide details on the spatial and temporal distribution of atmospheric parameters in Ethiopia.[7]

[6] To convert degrees to km, multiply 28 by the cosine of the latitude, e.g. at 40 degrees latitude $0.25° \times 0.25°$ cells are $28 \times \cos(40) = 21.4$ km $\times$ 21.4 km.

[7] It is important to note that all climate data, whether reanalysis or observational data, are subject to caveats and concerns. Reanalysis data is unlikely to match observational data perfectly. It is limited to some degree by resolution, even where observational data is present. Furthermore, reanalysis data are partly computed using climate models that are imperfect and contain systematic biases. This brings up further concern to issues of accuracy. However, in areas with limited observational data such as Ethiopia, reanalysis data is known to provide estimates that are better than observational data alone could provide. In addition, there are also statistical reasons as to why reanalysis data may be preferable. Previous studies have relied on the use of meteorological data provided by the Ethiopian meteorological service and the number of missing observations is a serious concern. This is exacerbated by the serious decline in the past few decades in the number of weather stations around the world that are reporting. Lorenz and Kuntsmann (2012b) show that since 1990 the number of reporting weather stations in Africa has fallen from around 3,500 to around 500. With 54 countries in the continent, this results in an average of fewer than 10 weather stations per country. Looking at publicly available data, the number of stations in Ethiopia included by the National Oceanic and Atmospheric Administration's (NOAA) National Climatic Data Centre (NCDC) is 18; however, if we were to apply a selection rule that required observations for 365 days this would yield a database with zero observations. For the two years for which I have economic data (2004 and 2009), weather station data is available for 50 days from one station (Addis Ababa) in 2004 and is available for all 18 stations for an average of 200 days (minimum of 67 days, maximum of 276 days) in 2009. This is likely to result in a huge increase in measurement error when this data is used to interpolate across the 63 zones and 529 Woredas (districts) reported in 2008. If this measurement error is classical, i.e., uncorrelated with the actual level of rainfall measured, then the estimates of the effect of these variables will be biased towards zero. However, given the sparse density of stations across Ethiopia (an average of 0.03 stations

3.5 Rainfall Variability and Parental Income Uncertainty

The use of rainfall variability as a proxy for uncertainty is driven by the importance of agriculture for subsistence consumption and livelihoods in rural parts of Sub-Saharan Africa, where access to irrigation is sparse. The consideration of uncertainty as a determinant of welfare is distinct from the literature, which examines the effects of weather shocks on welfare. If farmers form expectations about the climatic conditions of their area, we might expect that they plant crops that are suited to that area. Any deviation from the conditions on which this optimal cropping decision is based, such as more or less rainfall, may not be welfare-improving. The formation of these expectations is key for production. For this reason, I use rainfall variability, which, I argue, affects the farmers' ability to forecast the likelihood of future rainfall realisations, increasing uncertainty about future income.

To support this claim, I demonstrate, using the full panel of weather data between 1979 and 2012, that an increase in rainfall variability measured over the previous five years is associated with a reduction in realised rainfall in the season following the household survey. Table 3.2 shows that a one unit increase in the coefficient of variation is associated with a 3–4mm reduction in rainfall during the main agricultural season in Ethiopia (Meher). A one standard deviation increase in rainfall variability during this period (9.197 units) would be associated with a 27–36mm reduction in the main rainy season (Meher), around 15% of the average change in Meher rainfall each year. This suggests that increases in rainfall variability should be associated with an increase in future income uncertainty.

A second consideration is whether farmers are aware of this signal. This is obviously very difficult to test; however, using data from the 2009 round on risk perceptions I am able to provide some support for the claim that farmers may use rainfall variability to shape their expectations about the likelihood of future income shocks.

Table 3.3 presents descriptive statistics for the variables of interest relating to farmers' risk perceptions. I use data from three questions. The first question asks *"How often do the Meher rains fail for your land?"* Farmers report that on average the rains fail every 2.867 years. This implies that the rains fail roughly twice between each survey round, highlighting the significant volatility of the climate in rural Ethiopia, providing support for the premise that rainfall variability is an important driver of future income uncertainty in the short- to medium-run. The second question asks *"Are there any signs (e.g. in the temperature, the behaviour of animals, the rain received in the last few years, the number of the year) that you might receive timely or sufficient rain?"*. 55% of farmers report some perceived sign of timely or sufficient rains. The third

per Woreda), the placement of stations is likely to be correlated with agricultural output, i.e., weather stations are placed in more agriculturally productive areas, where the need for weather information is greater.

variable is based on an exercise requiring the farmer to place beans on two counters to ascertain the likelihood that the rains will fail. The farmers are asked *"Given the number of times the rains fail, and any signs you have observed, indicate by placing beans on the relevant square how likely it is that you think the Meher rains will fail this year. The more sure you are that they will fail the more beans should be placed on the low rainfall square. If you think there is an equal chance that they will fail place the beans equally between the two squares."* The farmers receive 20 beans in total. In table 3.3 0% corresponds to 0 beans, and 100% corresponds to 20 beans, being allocated to the square indicating that the rains will fail. On average, farmers perceive that there is a 59.3% chance that the rains in the coming season will fail.

Using these data, I first look at the correlation between rainfall variability and the likelihood of there being positive signs and the belief that the rains will fail. In Table 3.4 I show that a one standard deviation increase in rainfall variability (14.614 points) is associated with a 5.84% reduction in the likelihood of farmers perceiving that there are any positive signs. In addition, I find that a one unit increase in rainfall variability is associated with a 4.23% increase in farmers beliefs' that the rains will fail.

In addition, I show that an increase in the belief that the Meher rains will fail is associated with a reduction in the following season's Meher rainfall. However, it is interesting to note that once rainfall variability is controlled for – negatively associated with the following season's Meher rainfall –, farmers' beliefs no longer have any predictive power in determining future Meher rainfall. This suggests that rainfall variability is a major determinant of farmers' beliefs about the likelihood that the rains will fail, supporting the premise that rainfall variability is a good proxy for future income uncertainty.

A final consideration is whether rainfall variability has a direct effect on agricultural production, resulting in a realised income shock. If rainfall variability is to be an effective proxy for future income uncertainty, then it should have no direct effect on contemporaneous income. Using data on each household's agricultural production, I calculate agricultural yields, defined as the cultivated area-weighted production divided by total cultivated area.[8]

Using this data, I estimate the effects of rainfall variability on agricultural production using the specification,

$$\begin{aligned} \log(\textsc{Yield}_{hvt}) &= \beta_1 \textsc{Rainfall Variability}_{v\underline{t}} + \beta_2 f(\mathrm{w}_{vt}) + \beta_3 \overline{f(\mathrm{w}_{v\underline{t-5}})} \\ &+ \alpha_h + \alpha_m + \alpha_t + \epsilon_{hvt} \end{aligned} \tag{3.3}$$

[8]The crops used are white teff, black teff, barley, wheat, maize, and sorghum.

where $\text{Rainfall Variability}_{v\underline{t}}$ is my proxy for future income uncertainty – the coefficient of variation or the standard deviation of rainfall measured over the previous 5 years –, $f(\mathrm{w}_{vt})$ is a function of contemporaneous weather variables, and $\overline{f(\mathrm{w}_{v\underline{t-5}})}$ is function of historical weather variables measured over the previous 5 years. In addition, I include a vector of household fixed effects, year fixed effects and month of survey fixed effects.

In column (1) of Table 3.5 Iestimate the effects of rainfall variability absent contemporaneous and historical rainfall controls. I find that an increase in rainfall variability is associated with a contraction in agricultural yields. However, once contemporaneous and historical rainfall controls are included (columns 2 and 3), rainfall variability has no effect on agricultural yields, supporting the premise that rainfall variability provides a plausible proxy for future income uncertainty, rather than acting as a realised income shock. As discussed, one of the key difficulties associated with the measurement and identification of uncertainty is separating the effects of uncertainty from realised events. These results provide support for my measure of future income uncertainty, allowing us to proceed with the main empirical exercise.

3.6 Empirical Strategy

To analyse the effects of parental income uncertainty on human capital accumulation and child labour, I construct use the proxy for income uncertainty discussed in the previous section – rainfall variability. Starting with a measure of total annual rainfall for each village, I calculate the coefficient of variation for rainfall (CV), measured as the standard deviation divided by the mean for the previous five years, the time between survey rounds. The results are robust to alternative time periods and to using the standard deviation of rainfall rather than the coefficient of variation. However, one of the major advantages of the CV is that it is scale invariant, providing a comparable measure of variation across villages that may face very different levels of rainfall. Given the empirical context, where rain-fed agriculture plays a central role in subsistence consumption and the livelihoods of people, the use of rainfall variability provides a plausible measure of income uncertainty. However, the focus on income uncertainty is distinct from the focus of the literature to date, which focuses on realised income shocks. Yet, while these effects are conceptually distinct, empirically disentangling the effects of income uncertainty from realised income shocks is more complicated.

As the first moment and second moment of the rainfall distribution are correlated, it is important to control for first-moment effects to isolate the effects of uncertainty, to the degree that they are empirically relevant, from income effects. I therefore control for historical and contemporaneous weather effects, arguing that any residual variation in the rainfall distribution once these considerations have been controlled for is likely to capture uncertainty about future income realisations.

3.6.1 Uncertainty at the Time of the Survey

Using this identification strategy, I examine the effects of parental income uncertainty – proxied by rainfall variability – on the child labour and educational decisions of smallholder farmers,

$$Y_{ivt} = \beta_1 \text{Rainfall Variability}_{v\underline{t}} + \beta_2 f(w_{vt}) + \beta_3 \overline{f(\mathrm{w}_{v\underline{t-5}})} + \alpha_i + \alpha_a + \alpha_m + \alpha_t + \epsilon_{ihvt} \tag{3.4}$$

Y_{ivt} is the outcome variable of interest, defined as the log number of hours in each activity, a binary variable, indicating participation in each activity, and the number of grades attained. The key explanatory variable of interest is Rainfall Variability$_{v\underline{t}}$ – a proxy for parental income uncertainty after controlling for contemporaneous, $f(w_{dt})$, and historical, $\overline{f(\mathrm{w}_{vt-5})}$, weather events. In addition, I control for individual fixed effects (α_i), to allow comparison within child over time, year fixed effects (α_t) to control for aggregate shocks and uncertainty, month of survey fixed effects (α_m) to control for seasonal variation in the timing of the survey, and cohort fixed effects (α_a) to allow comparison within age-cohort.[9]

The last term in equation 3.4 is the stochastic error term, ϵ_{ihvt}. I follow the approach of Hsiang (2010) by assuming that the error term may be heteroskedastic and serially correlated within a district over time (Newey and West, 1987) and spatially correlated across contemporaneous villages (Conley, 1999). For each outcome of interest, I loop over all possible distances up to 1,000km, selecting the parameter value

[9] In addition to estimating linear models I also estimate a fixed-effects Poisson Quasi-Maximum Likelihood Estimator model (QMLE) to evaluate the time allocation data.[10]

The model is estimated using the following specification:

$$\mathbb{E}(Y_{ivt}) = \mu_i(\exp(\beta \text{Rainfall Variability}_{v\underline{t}} + \beta_2 f(w_{vt}) + \beta_3 \overline{f(\mathrm{w}_{v\underline{t-5}})} + \alpha_a + \alpha_t + \alpha_m) \tag{3.5}$$

where subscripts index individuals (i), households (h), village (v) and year (t). There are a number of benefits to using the Poisson QMLE instead of the standard Poisson MLE. For example, the use of the QMLE does not require that the data follow a Poisson distribution. All that is required is that the conditional mean of the variable of interest be correctly specified. A further benefit in the context of Poisson models is the mitigation of concerns surrounding under- and over-dispersion. This is because, unlike the MLE, the Poisson QMLE does not assume equi-dispersion. All that is required for optimality of the Poisson QMLE is that the conditional variance is proportional to the conditional mean. Furthermore, the Poisson QMLE will still be consistent in the case where the conditional variance is not proportional to the conditional mean. This means that I can work using a fixed-effects framework without needing to use models such as the negative-binomial or zero-inflation Poisson MLE to deal with consistency issues.

that provides the most conservative standard errors. I then repeat this process for serial correlation.[11]

In addition to estimating the effects parametrically, it is also interesting to examine the degree to which there are non-linearities in the relationship between rainfall variability and consumption. Ido this by estimating this relationship semi-parametrically. The idea behind this approach is to obtain local estimates of the relationship being studied and to display them in a way that visually weights the degree of regression uncertainty underlying the relationship. This procedure has two steps. First, all variation associated with contemporaneous and historical weather effects, as well as the fixed effects, is absorbed from the data, ensuring that the scales are identical. Then a LOESS regression of the residuals of rainfall variability and real consumption per capita is estimated repeatedly using a bootstrapping procedure. The residuals for rainfall variability on the horizontal axis are subdivided into 500 grid points. Each bootstrapped regression is evaluated at the grid-point along the horizontal axis. This results in a set of fitted values for each grid point. In the second step the fitted values are plotted. For each horizontal grid point, a kernel density is estimated. The colouring of the graph relates to two considerations. First, the overall colour intensity at each grid point along the horizontal axis is related to the overall mass of data that is available in that part of the distribution. This colouring is then stretched out vertically in relation to the density of the fitted values.

3.6.2 The Effects of Parental Income Uncertainty through the Life Cycle

In addition to examining the effects of parental income uncertainty at the time of the surveys, I also examine how parental income uncertainty may affect investments in human capital at different stages of the life cycle, estimating the effects of rainfall variability at each age in the child's life, from birth to the current age of the child:

$$Y_{ivt} = \beta_1 \textsc{Rainfall Variability}_{iv\underline{t_a}} + \beta_2 f(w_{ivt_a}) + \beta_3 \overline{f(\mathrm{w}_{iv\underline{t_{a-5}}})} + \alpha_v + \alpha_{t_a} + \epsilon_{ivt_a} \quad (3.6)$$

where Y_{ivt} is a dummy variable indicating whether the child is attending school at the time of the survey, and $\textsc{Rainfall Variability}_{ivt_a}$ is a proxy for parental income uncertainty measured over the 5 years prior to the year, t, in which child i was age a. For example, if the child was aged 5 in the year 2000, I would be examining the effect of rainfall variability between 1995 and 1999 on the likelihood of attending school at the time of the survey in either 2004 or 2009. $f(w_{ivt_a})$ and $\overline{f(\mathrm{w}_{iv\underline{t_{a-5}}})}$ are functions of

[11] Results are robust to clustering at the village level.

contemporaneous and historical weather variables during the same time period. α_v is a village fixed effect controlling for the persistent effects of rainfall variability on the locations in which these children live. For example, the effects of rainfall variability on the long-run income of households should be common to all individuals in the same area and so should be absorbed by the village fixed effects. α_{t_a} is a year fixed effect, corresponding to the year in which child i was aged a. In estimating this model, I am able to explore whether parental income uncertainty at different times during the child's life has a differential effect on human capital investments at the time of the survey. This approach is similar to the one used by Maccini and Yang (2009), where they explore the persistent effects of in-utero rainfall shocks on later life outcomes.

Again, I explore the degree to which there are non-linearities in the relationship between rainfall variability and consumption by estimating this relationship semi-parametrically.

3.7 Results

In this section I first present the main results, examining the effects of parental income uncertainty – proxied by climate variability – on child labour and education outcomes, before examining whether parental income uncertainty at different times in the child's life cycle has a differential effect on human capital accumulation.

3.7.1 The Effects of Rainfall Variability on Children's Time Use.

First I examine the effect of parental income uncertainty on the time allocation of children in rural Ethiopia. As discussed, the effect of parental income uncertainty on child labour and human capital accumulation is theoretically ambiguous: on the one hand, households may reduce investments in human capital, allocating more time to child labour on the farm in an effort to mitigate the economic consequences of productivity shocks in the event that they are realised – the "precautionary motive". On the other hand, households may increase investments in human capital, adjusting the time allocation of children away from risky activities on the farm towards less risky investments to diversify the household income portfolio.

The results of this exercise are presented in Table 3.6. The results of the semi-parametric exercise are plotted in Figure 3.3. We observe that an increase in rainfall variability – my proxy for future income uncertainty after controlling for contemporaneous and historical weather events – is associated with a reduction in child labour on the farm and an increase in the amount of time spent on home study, consistent with the "portfolio motive". These results are robust across both measures of rainfall variability, as well as to both the OLS and Poisson QMLE model, reported in Table 3.7.

To understand the magnitude of these effects, a one standard deviation increase in rainfall variability (6.560 units) would be associated with a 14.896% reduction in child labour on the farm – approximately 2 hours less farm work per week when compared to the average number of hours that children engage in (14.1 hours). Similarly, a one standard deviation increase in rainfall variability (6.560 units) would be associated with a 13.251% increase in the amount of time children spend on home study – 1–2 hours more home study per week, when compared to the average number of hours that children engage in (8.1 hours). These results suggest that households respond to an increase in parental income uncertainty by adjusting the time allocation of children away from risky farm activities towards time spent accumulating human capital, a strategy used to diversify the income portfolio of the household.

3.7.2 The Effects of Rainfall Variability on Children's Schooling

In addition to exploring how parental income uncertainty affects children's time use, it is also important to understand whether it has a direct effect on schooling. To explore this, I estimate the effects of rainfall variability on school attendance and grades achieved, as well as extensive margin participation in child labour activities.

The results of this exercise are presented in Table 3.8. The results of the semi-parametric exercise are plotted in Figure 3.4. We observe that an increase in rainfall variability – my proxy for future income uncertainty after controlling for contemporaneous and historical weather events – is associated with an increase in the likelihood that children attend school, as well as an increase in the number of grades achieved. These results are consistent with the "portfolio motive", providing further support for the premise that households invest more in accumulating human capital to diversify the income portfolio of the household. We also observe that an increase in rainfall variability is associated with a reduction in the likelihood that children engage in child labour on the farm and an increase in the likelihood that children engage in home study. These results are robust across both measures of rainfall variability, as well as to conditional Logit model specifications, reported in Table 3.9.

To understand the magnitude of these effects, a one standard deviation increase in rainfall variability (6.560 units) would be associated with a 5 percentage point increase in the likelihood that children attend school, and a 0.138 point increase in grades. These results suggest that households respond to an increase in parental income uncertainty by increasing investments in human capital, consistent with the "portfolio motive". This premise is further supported by the result that a one standard deviation increase in rainfall variability is associated with a 4 percentage point decrease in the likelihood that children engage in child labour, and a 6 percentage point increase in the likelihood that children engage in home study.

3.7.3 Heterogeneous Effects: Rainfall Variability through the Life Cycle

Having estimated the average effects of a change in parental income uncertainty, it is also interesting to understand whether changes in parental income uncertainty have a different effect on human capital investments, depending on the stage in life that the child has reached. To explore the potential for heterogeneity in the response to changes in parental income uncertainty, I estimate the effects of historical rainfall variability, measured at each age in the child's life, following an approach similar to Maccini and Yang (2009). This gives an estimate of the effects of parental income uncertainty at birth up until the age of the child at the time of the survey.

The results of this exercise are presented in Table 3.10. The results of the semi-parametric exercise are plotted in Figure 3.5. We observe that, in the early stages of the child's life, an increase in rainfall variability is associated with a decrease in the likelihood that the child attends school at the time of the survey, consistent with the "precautionary motive" dominating. However, around key decision-making periods in the educational calendar, it is interesting to see that an increase in rainfall variability is associated with an increase in the likelihood that children attend school. This relationship is observed at age 5 and 6, around the time when educational decisions are first being made, at age 15 when decisions about secondary education are being made, and at age 18 when decisions about national exams are being made. These results may suggest that the salient nature of human capital investments at these times attach greater weight on the "portfolio motive".

These results suggest that the average effect estimated in Tables 3.6 and 3.8 mask the heterogeneous motivations discussed in section 3.3. While, on average, an increase in parental uncertainty is associated with increased investments in human capital, this response appears to depend on the child's stage in the life cycle, with increases in parental income uncertainty being associated with a reduction in human capital investments during the early stages of the life cycle.

3.8 Conclusion

While a significant body of work has sought to understand the effects of realised income shocks on economic behaviour in developing economies, there is very little evidence relating to the economic consequences of future income uncertainty. The consequences of these changes may be particularly relevant in the context of irreversible investments, such as human capital, where short-run changes in the economic environment may have persistent effects on later-life outcomes and opportunities. As human capital accumulation is an important driver of economic growth, understanding the effects of future income uncertainty on human capital investments provides insights into the underlying drivers of economic development, as well as the economic consequences of future income uncertainty.

However, measuring and identifying the effects of uncertainty, separately from realised events, presents a number of empirical challenges. To help mitigate these concerns, I evaluate the empirical relevance of rainfall variability – a proxy for future income uncertainty, after controlling for contemporaneous and historical weather events – on the human capital investments of smallholder farmers in rural Ethiopia – one of the least developed countries in Africa that is characterised by a high vulnerability to inclement weather.

Theoretically, I show that the relationship between parental income uncertainty and human capital investment is ambiguous. On the one hand, households may reduce investments in the human capital of their children, allocating more time to child labour on the farm in an effort to mitigate the economic consequences of future income shocks in the event that they are realised – a precautionary motive. On the other hand, an increase in parental income uncertainty may result in an increase in human capital investments as parents seek to expand the economic opportunities of their children, and in doing so diversify the income portfolio of the household – a portfolio motive.

Empirically, I begin by providing supporting evidence for the premise that rainfall variability is a plausible proxy for parental income uncertainty. I begin by demonstrating that an increase in rainfall variability, measured over the previous five years (the time between survey rounds), is associated with a reduction in realised rainfall in the season following the survey. Importantly, I also provide evidence to suggest that farmers are aware of this signal. Using data on risk perceptions, I show that an increase in rainfall variability is associated with an increase in the belief that the rains will fail. Furthermore, an increase in this belief is associated with lower rainfall realisations in the following season. However, once I control for rainfall variability, farmers' beliefs about future income become statistically insignificant, suggesting that these beliefs are largely driven by rainfall variability. Finally, I show that rainfall variability has no direct effect on agricultural production. All of these results provide support for the idea that rainfall variability is a reasonable proxy for future income uncertainty, helping to disentangle the effects of parental income uncertainty from realised weather events.

Having demonstrated that rainfall variability is a reasonable proxy for parental income uncertainty – after controlling for contemporaneous and historical weather events –, I evaluate the degree to which parental income uncertainty affects human capital accumulation. I begin by examining the effects of rainfall variability on the time use of children. By examining the amount of time that children spend on child labour and human capital activities, I am able to explore intensive margin adjustments to human capital accumulation. Consistent with the "portfolio motive", I estimate that an increase in rainfall variability is associated with a reduction in time spent on farming activities and an increase in time spent on home study. To test the robustness of this interpretation, I also examine the effects of rainfall variability on extensive margin participation adjustments in schooling attendance, child labour, and the number of grades attained. Consistent with the intensive margin adjustments, I find that an increase in rainfall variability is associated with an increase in the likelihood that a child is in school, a reduction in the likelihood that a child engages in child labour, and an increase in the the number of grades attained. All of these results point to the existence of a portfolio motive in response to increases in parental income uncertainty.

However, this average effect masks heterogeneity in the response of parents. I uncover this through an examination of the effects of rainfall variability at different stages in the child's life cycle. In the early stages of a child's life, an increase in rainfall variability is associated with a reduction in the likelihood that the child attends school, consistent with the precautionary motive. As the child grows older this relationship weakens and reverses. This is consistent with the portfolio motive, capturing the idea that as a child grows older, the returns to education, and consequently diversification, increase.

These results highlight two things: first, that farmers are responsive to changes in future income uncertainty, and are actively making decisions to mitigate the economic consequences of future income shocks through investing in the human capital of their children; second, that in understanding the consequences of environmental change, it is important to understand how expectations about future states of the world affect economic behaviour, as well as the consequences of realised change.

Table 3.1: Descriptive Statistics

	MEAN	STD. DEV. (within)	STD. DEV. (between)	OBS.
Panel A: *Time Allocation (Hours)*				
CHILD LABOUR (TOTAL)	28.001	9.504	17.269	6,698
CHILD LABOUR (FARMING)	14.101	7.305	14.797	6,698
CHILD LABOUR (DOMESTIC CHORES)	13.856	6.395	12.791	6,698
HOME STUDY	8.106	5.013	9.208	6,698
LEISURE	131.936	10.579	18.966	6,698
Panel B: *Participation (%)*				
CHILD LABOUR (TOTAL)	92.012	15.164	23.806	6,698
CHILD LABOUR (FARMING)	61.182	24.424	44.472	6,698
CHILD LABOUR (DOMESTIC CHORES)	71.185	23.393	40.743	6,698
HOME STUDY	57.420	26.897	44.161	6,698
Panel C: *Education*				
ATTENDING SCHOOL (%)	58.465	26.547	44.173	6,698
GRADES ATTAINED	3.486	1.202	2.589	6,698
Panel D: *Weather Data*				
Proxies for Parental Income Uncertainty:				
RAINFALL VARIABILITY (σ/μ)	18.036	5.402	10.258	6,698
RAINFALL VARIABILITY (σ, mm)	248.933	75.460	124.382	6,698
Contemporaneous Weather:				
TOTAL RAINFALL (mm)	1,446.788	183.891	453.637	6,698
AVERAGE TEMPERATURE (°C)	19.029	0.258	1.679	6,698
Historical Weather (5-year Averages):				
TOTAL RAINFALL (mm)	1,430.608	65.800	340.656	6,698
AVERAGE TEMPERATURE (°C)	19.050	0.140	1.663	6,698

NOTES: The proxies for parental income uncertainty are defined as the coefficient of variation (σ/μ) and standard deviation ((σ) of rainfall over the previous 5 years, the time between each survey round. The time allocated to each activity is defined as the number of hours spent per week on each activity. Leisure is defined as the residual number of hours after accounting for the time allocated to child labour and home study. This residual accounts for the amount of time children spent eating, sleeping, playing and in school. Grades Attained is the number of grades that have been completed at the time of the survey ranging from 0 grades to 11 grades in total. Idleness is defined as a dummy variable equal to 1 if a child neither engages in child labour or attends school.

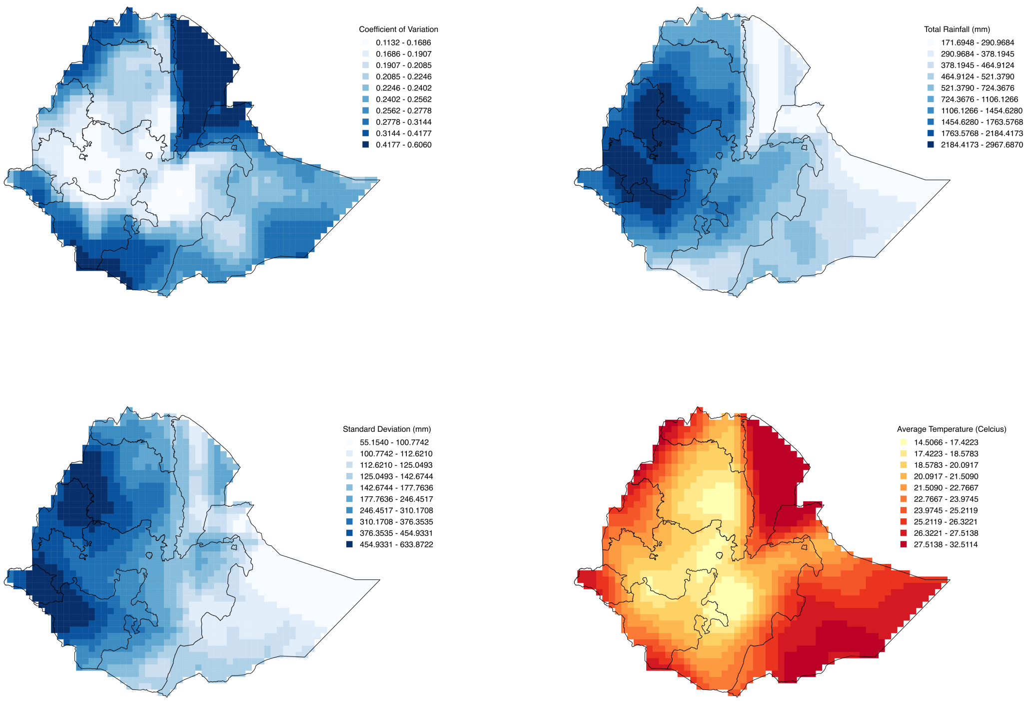

Figure 3.1: Spatial Variation in Rainfall and Temperature (1979–2012). Top Left = Coefficient of Variation; Top Right = Total Rainfall (mm); Bottom Left = Std Dev. Rainfall (mm); Bottom Right = Average Temperature (°C)

.

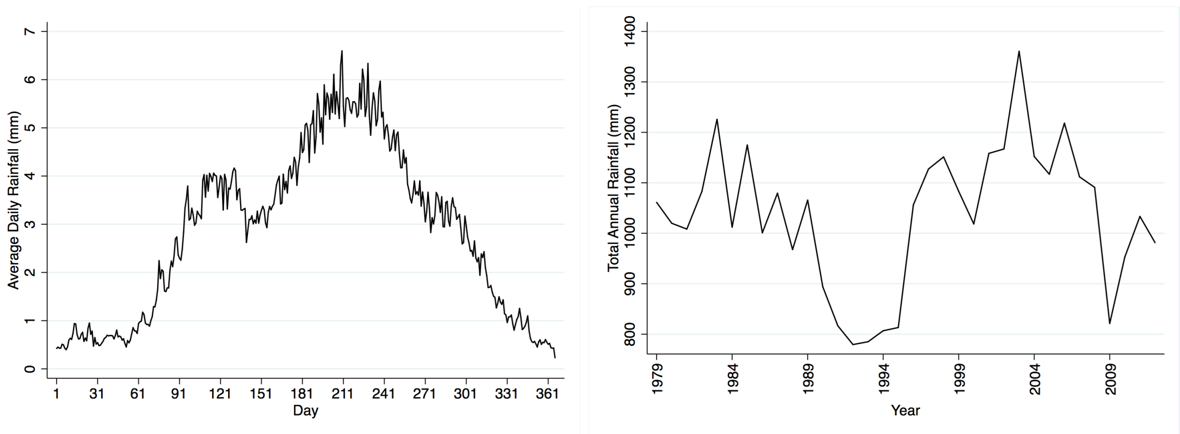

Figure 3.2: Temporal Variation in Rainfall (1979–2013). Top = Within-year Distribution (1979-2012 average). Bottom = Between-year Distribution

.

Table 3.3: Descriptive Statistics: Household Perceptions of Rainfall Risk

Question	Mean	Std. Dev.	Obs.
How often do the rains fail? Every . . . years.	2.867	1.650	1,212
Any Signs that rain may be timely or sufficient?	0.552	0.497	1,282
How likely is it that the rains will fail this year?	59.3%	16.3%	1,310

Note 1: The data is only available for 2009.

Note 2: Question: How often do the Meher rains fail for your land? Once every . . . years.

Note 3: Question: Are there any signs (e.g. in the temperature, the behaviour of animals, the rain received in the last years, the number of the year) that you might receive timely or sufficient rain? Yes = 1, No = 0.

Note 4: Question: Given the number of times the rains fail, and any signs you have observed, indicate by placing beans on the relevant square how likely it is that you think the Meher rains will fail this year. The more sure you are that they will fail the more beans should be placed on the square. If you think there is an equal chance that they will fail place the beans equally between the two squares. (20 beans in total) 0 beans = 0% likely, 20 beans = 100% likely.

Table 3.2: The Effects of Past Rainfall Variability on Future Rainfall

	(1) Meher Rainfall (mm)	(2) Meher Rainfall (mm)	(3) Meher Rainfall (mm)
Panel A: Coefficient of Variation:			
Rainfall Variability (σ/μ)	-2.875**	-4.024***	-4.171***
	(1.091)	(0.888)	(1.132)
Panel B: Standard Deviation:			
Rainfall Variability (σ, 100 mm)	-8.323	-29.014***	-29.657**
	(10.586)	(9.099)	(11.202)
Fixed Effects		Village and Year	
Weather Controls	No	Yes	Yes
Quadratic Weather Controls	No	No	Yes
Observations	480	480	480

Note 1: Specification: $\text{Rainfall}_{vt} = \beta_1 \, \text{CV}_{v\underline{t}} + \beta_2 \, \text{Rainfall}_{vt-1} + \beta_3 \, \overline{Rainfall}_{vt-5} + \alpha_v + \alpha_t + \varepsilon_{vt}$.

Note 2: Meher Rainfall is defined as the total cumulative rainfall received between June and November for the following season. See table **??** for definitions of the explanatory variables.

Table 3.5: The Effects of Rainfall Variability on Agricultural Yields

	(1) log Yields	(2) log Yields	(3) log Yields
Panel A: Coefficient of Variation:			
Rainfall Variability (σ/μ)	-0.0347***	-0.0118	-0.00173
	(0.00897)	(0.0114)	(0.0126)
Panel B: Standard Deviation:			
Rainfall Variability (σ, 100 mm)	-0.257***	-0.114	-0.073
	(0.0724)	(0.0785)	(0.107)
Fixed Effects	Household and Year		
Weather Controls	No	Yes	Yes
Quadratic Weather Controls	No	No	Yes
Observations	2,334	2,334	2,334

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. The unit of analysis is at the household level. Our proxies for uncertainty are the coefficient of variation for rainfall (Panel A) and the standard deviation of rainfall (Panel B), measured over the previous 5 years, the time period between each survey round. Historical measures of atmospheric parameters correspond to this period. Contemporaneous and historical rainfall is measured in hundreds of mm. Contemporaneous and historical temperature is measured in °C. Standard errors are adjusted to reflect spatial dependence, as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cut-off of 150 km. The distance is selected following a decision rule in which I choose the distance that provides the most conservative standard errors, looped over all distances between 10 and 1000km.

Table 3.4: Household Perceptions of Rainfall Risk – Correlations

	(1) Positive Signs	(2) Positive Signs	(3) Belief that rains will fail	(4) Belief that rains will fail
Rainfall Variability (σ/μ)	0.0003	-0.00404**	0.00414***	0.00291**
	(0.00159)	(0.0017)	(0.00104)	(0.00131)
Constant	0.546***	0.689***	0.526***	0.523***
	(0.0485)	(0.123)	(0.0231)	(0.0551)
Weather Controls	No	Yes	No	Yes
Observations	1,282	1,282	1,282	1,282
	(1) Meher Rainfall (mm)	(2) Meher Rainfall (mm)	(3) Meher Rainfall (mm)	(4) Meher Rainfall (mm)
Positive Signs	9.107		11.240	9.719
	(14.454)		(14.352)	(13.492)
Belief that rains will fail		-131.566**	-131.227**	68.918
		(41.865)	(41.956)	(44.328)
Rainfall Variability (σ/μ)				-6.050***
				(0.319)
Constant	682.688***	770.427***	759.458***	738.988***
	(10.855)	(25.566)	(27.165)	(25.603)
Observations	1,282	1,310	1,282	1,282

Note 1: The data is only available for 2009.
Note 2: Specification 1: $Y_{iv} = \beta_0 + \beta_1 \text{CV}_{vt} + \beta_2 \text{Rainfall}_{vt-1} + \beta_3 \overline{Rainfall}_{vt-5} + \varepsilon_{iv}$. Standard errors are clustered at the village level. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01.
Note 3: Specification 2: Meher $\text{Rainfall}_v = \beta_0 + \beta_1$ Positive $\text{Signs}_{iv} + \beta_2 \text{Belief}_{iv} + \varepsilon_{iv}$. Standard errors are clustered at the household level. Significance levels are indicated as * 0.10 ** 0.05 *** 0.01.
Note 4: Meher Rainfall is defined as the total cumulative rainfall received between June and November for the following season. See table **??** for definitions of the explanatory variables.

Figure 3.3: Semi-parametric Estimates of the Relationship between Rainfall Variability and Children's Time Use

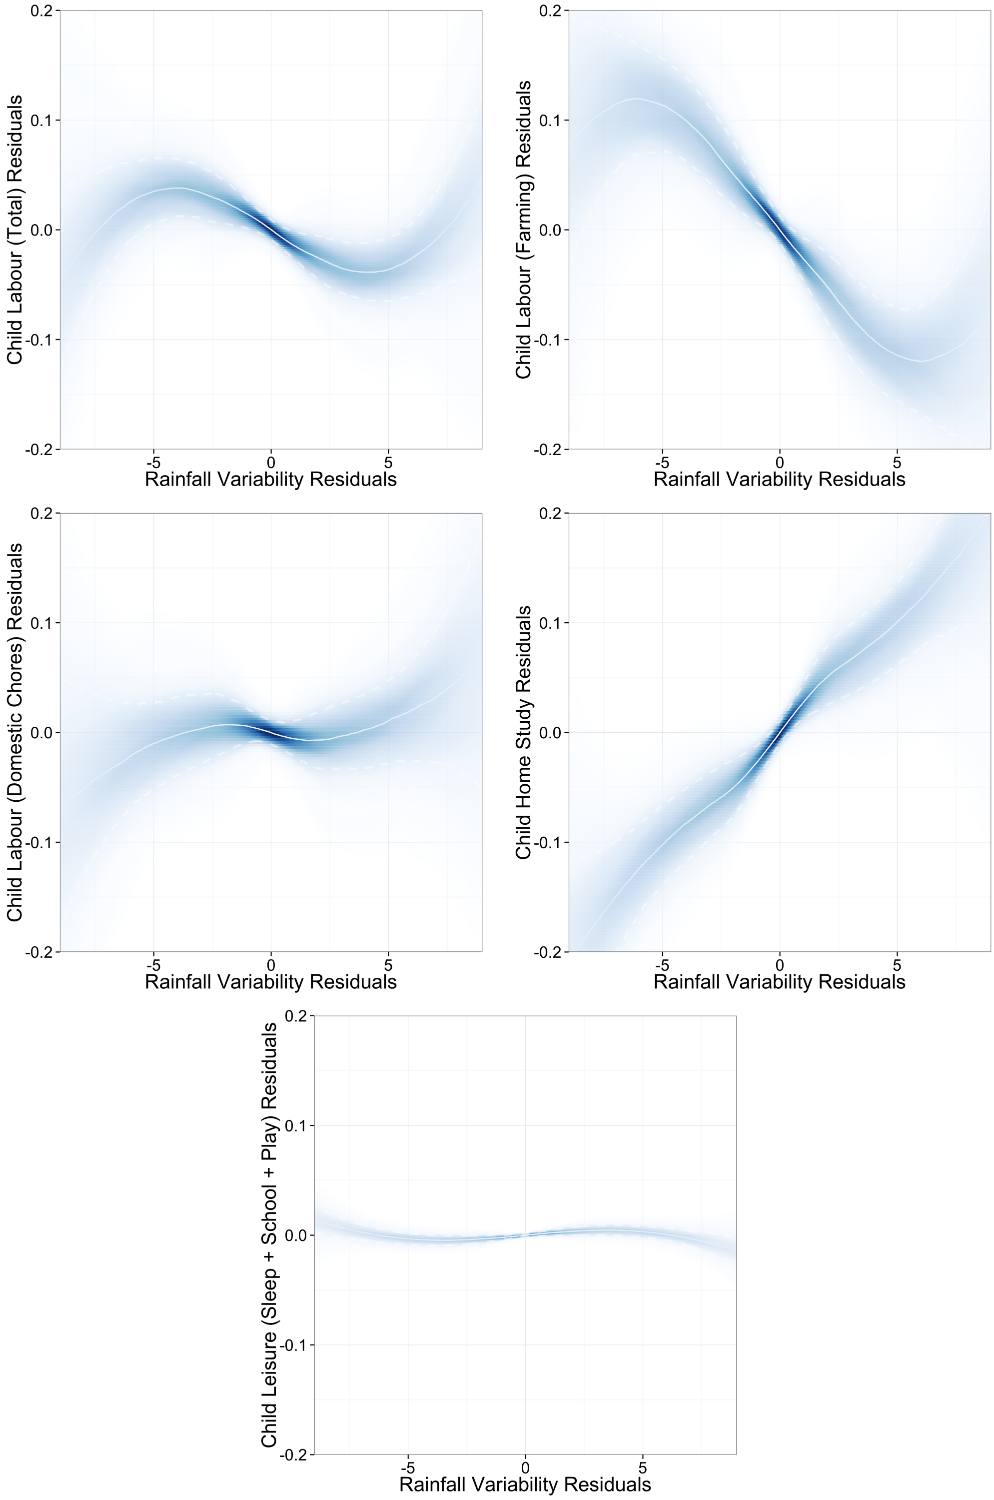

Notes: The measure of rainfall variability plotted here is the coefficient of variation, measured over the previous 5 years, the time period between each survey round. Each variable is regressed on individual, age, year and month fixed effects as well as contemporaneous and historical rainfall and temperature controls. The figures above are the result of loess regressions of the residuals from this exercise.

Table 3.6: The Effects of Parental Income Uncertainty on Children's Activities (Time Allocation) – OLS

	OLS				
	(1) log(Total Child Labour)	(2) log(Farm Labour)	(3) log(Domestic Chores)	(4) log(Home Study)	(5) log(Leisure)
Panel A: Coefficient of Variation:					
Rainfall Variability (σ/μ)	-0.00647***	-0.0224***	0.00196	0.0202***	0.000433
	(0.00164)	(0.00472)	(0.00473)	(0.00458)	(0.000569)
Panel B: Standard Deviation:					
Rainfall Variability (σ, 100 mm)	-0.0366**	-0.171***	0.0373	0.115***	0.00453
	(0.0171)	(0.0332)	(0.0238)	(0.0328)	(0.00390)
Fixed Effects	Individual, Year, Month				
Weather Controls	Yes	Yes	Yes	Yes	Yes
Observations	6,698	6,531	6,531	6,531	6,531

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. The unit of analysis is at the individual child level. My proxies for uncertainty are the coefficient of variation for rainfall (Panel A) and the standard deviation of rainfall (Panel B), measured over the previous 5 years, the time period between each survey round. Historical measures of atmospheric parameters correspond to this period. Contemporaneous and historical rainfall is measured in hundreds of mm. Contemporaneous and historical temperature is measured in °C. Standard errors are adjusted to reflect spatial dependence, as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cut-off of 150 km. The distance is selected following a decision rule in which I choose the distance that provides the most conservative standard errors, looped over all distances between 10 and 1000km. Results are also robust to applying family-wise errors rates (FWER) using the stepwise procedure of Romano and Wolf (2005).

Figure 3.4: Semi-parametric Estimates of the Relationship between Rainfall Variability and Participation in Educational and Child Labour Activities

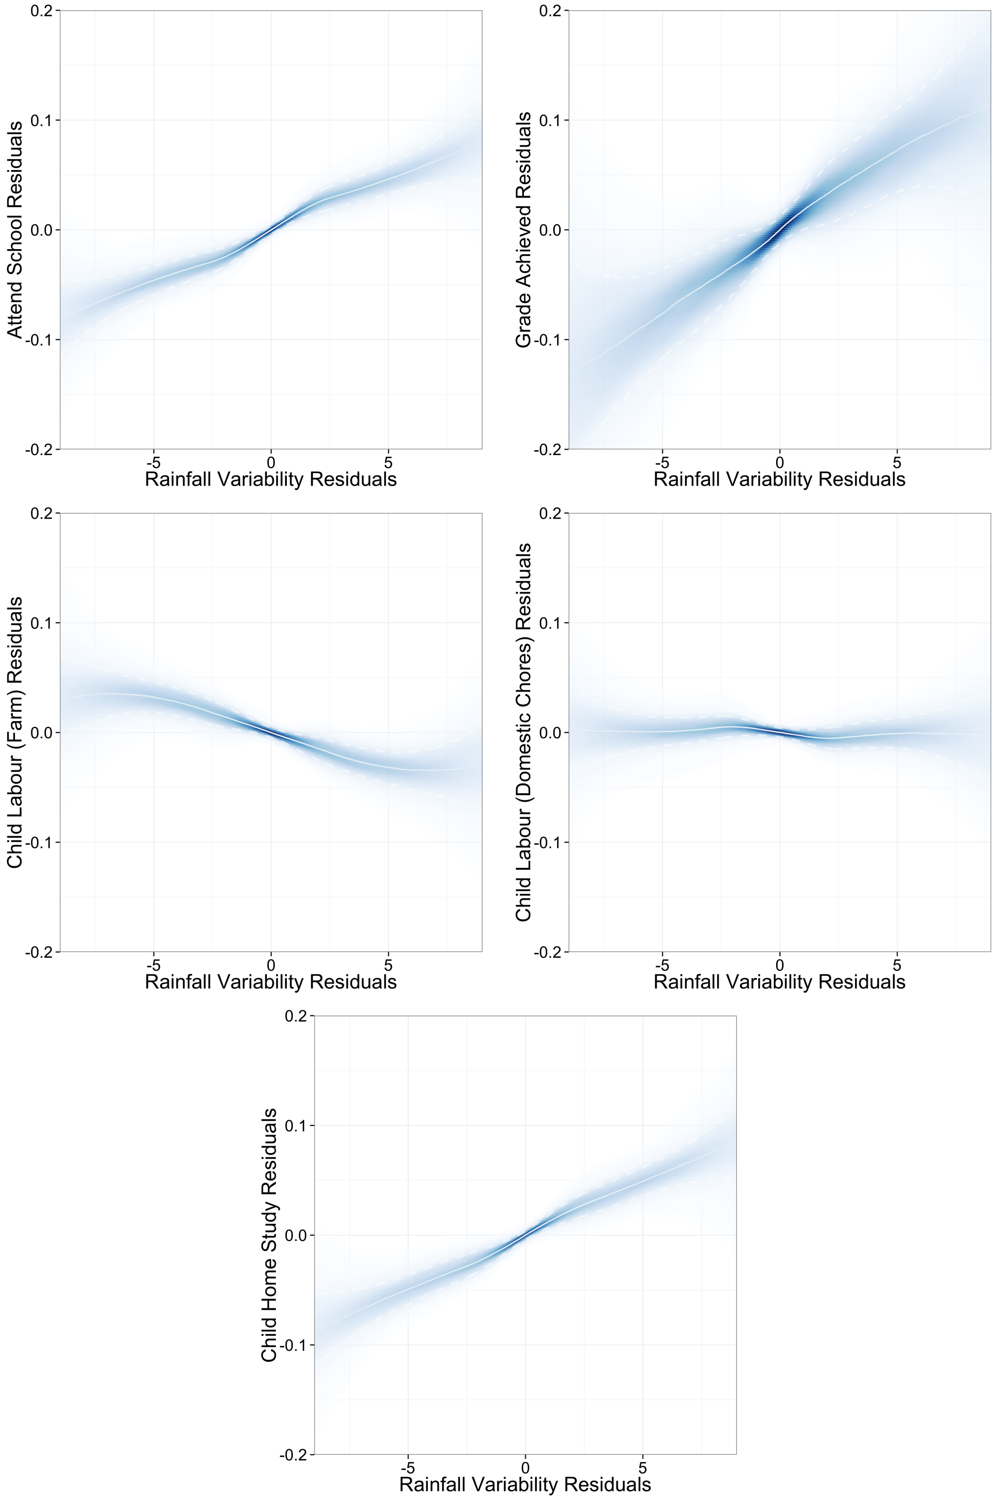

Notes: The measure of rainfall variability plotted here is the coefficient of variation, measured over the previous 5 years, the time period between each survey round. Each variable is regressed on individual, age, year and month fixed effects as well as contemporaneous and historical rainfall and temperature controls. The figures above are the result of loess regressions of the residuals from this exercise.

Table 3.7: The Effects of Parental Income Uncertainty on Children's Activities (Time Allocation) – Poisson QMLE

	Poisson QMLE				
	(1) log(Total Child Labour)	(2) log(Farm Labour)	(3) log(Domestic Chores)	(4) log(Home Study)	(5) log(Leisure)
Panel A: Coefficient of Variation:					
Rainfall Variability (σ/μ)	-0.00691***	-0.0177***	0.00619***	0.00981***	0.000173
	(0.000600)	(0.00219)	(0.000648)	(0.00177)	(0.000112)
Panel B: Standard Deviation:					
Rainfall Variability (σ, 100 mm)	-0.0462***	-0.148***	0.0630***	0.0421***	0.00248***
	(0.00478)	(0.0139)	(0.00217)	(0.00967)	(0.000749)
Fixed Effects	Individual, Year, Month				
Weather Controls	Yes	Yes	Yes	Yes	Yes
Observations Observations	4,134	3,210	3,556	3,432	3,962

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. The unit of analysis is at the individual child level. My proxies for uncertainty are the coefficient of variation for rainfall (Panel A) and the standard deviation of rainfall (Panel B), measured over the previous 5 years, the time period between each survey round. Historical measures of atmospheric parameters correspond to this period. Contemporaneous and historical rainfall is measured in hundreds of mm. Contemporaneous and historical temperature is measured in °C. Standard errors are adjusted to reflect spatial dependence, as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cut-off of 150 km. The distance is selected following a decision rule in which I choose the distance that provides the most conservative standard errors, looped over all distances between 10 and 1000km.

Table 3.8: The Effects of Parental Income Uncertainty on Participation in Educational and Child Labour Activities – OLS

	OLS				
	(1) Attend	(2) Grade Attained	(3) Farm Labour	(4) Domestic Chores	(5) Home Study
Panel A: Coefficient of Variation:					
Rainfall Variability (σ/μ)	0.00867***	0.0211***	-0.00631***	-0.000116	0.00957***
	(0.000845)	(0.000273)	(0.000929)	(0.00182)	(0.00101)
Panel B: Standard Deviation:					
Rainfall Variability (σ, 100 mm)	0.0600***	0.144***	-0.0467***	0.00774	0.0667***
	(0.00781)	(0.0137)	(0.00787)	(0.00951)	(0.00920)
Fixed Effects	Individual, Year, Month				
Weather Controls	Yes	Yes	Yes	Yes	Yes
Observations	6,698	5,155	6,698	6,698	6,698

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. My proxies for uncertainty are the coefficient of variation for rainfall (Panel A) and the standard deviation of rainfall (Panel B), measured over the previous 5 years, the time period between each survey round. Historical measures of atmospheric parameters correspond to this period. Contemporaneous and historical rainfall is measured in hundreds of mm. Contemporaneous and historical temperature is measured in °C. Standard errors are adjusted to reflect spatial dependence, as modelled in Conley (1999). Spatial autocorrelation is assumed to linearly decrease in distance up to a cut-off of 150 km. The distance is selected following a decision rule in which I choose the distance that provides the most conservative standard errors, looped over all distances between 10 and 1000km. Results are also robust to applying family-wise errors rates (FWER) using the stepwise procedure of Romano and Wolf (2005).

Table 3.9: The Effects of Parental Income Uncertainty on Participation in Educational and Child Labour Activities – Logit

	Conditional Logit Fixed Effects – Odds Ratios			
	(1) Attend	(2) Farm Labour	(3) Domestic Chores	(4) Home Study
Panel A: Coefficient of Variation:				
Rainfall Variability (σ/μ)	1.034***	0.959***	0.998	1.0543***
	(0.0134)	(0.0148)	(0.0162)	(0.0132)
Panel B: Standard Deviation:				
Rainfall Variability (σ, 100 mm)	1.233**	0.726***	1.0490	1.442***
	(0.122)	(0.0881)	(0.117)	(0.140)
Fixed Effects	Individual, Year, Month			
Weather Controls	Yes	Yes	Yes	Yes
Observations	1,888	1,598	1,466	1,938

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. My proxies for uncertainty are the coefficient of variation for rainfall (Panel A) and the standard deviation of rainfall (Panel B), measured over the previous 5 years, the time period between each survey round. Historical measures of atmospheric parameters correspond to this period. Contemporaneous and historical rainfall is measured in hundreds of mm. Contemporaneous and historical temperature is measured in °C. Standard errors are clustered at the village level.

Table 3.10: The Effects of Parental Income Uncertainty at different times during the Child's Life Cycle

Age	Attend	Observations	Age	Attend	Observations	Age	Attend	Observations	Age	Attend	Observations
Age 0	-0.00267*** (0.000992)	4,679	Age 5	0.00205* (0.00117)	4,679	Age 10	0.00155 (0.00124)	3,676	Age 15	0.00346** (0.00146)	2,182
Age 1	-0.00208* (0.00116)	4,679	Age 6	0.00226* (0.00118)	4,573	Age 11	0.00141 (0.00114)	3,363	Age 16	-0.000679 (0.00150)	1,788
Age 2	-0.00114 (0.00129)	4,679	Age 7	0.000946 (0.00124)	4,392	Age 12	0.00252** (0.00121)	3,174	Age 17	0.000357 (0.00183)	1,416
Age 3	-0.00138 (0.00127)	4,679	Age 8	0.00181 (0.00140)	4,158	Age 13	0.000705 (0.00108)	3,332	Age 18	0.00286* (0.00173)	1,148
Age 4	-0.0000229 (0.00118)	4,679	Age 9	0.00107 (0.00127)	3,919	Age 14	0.00106 (0.00134)	2,511	Age 19	-0.00356 (0.00386)	695

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Each line corresponds to an independent regression where the measure of rainfall variability is defined as the coefficient of variation for rainfall for the 5 years preceding the year at which the child was the specified age. Historical rainfall and temperature measures, used as control variables, correspond to this period. Standard errors are adjusted to reflect spatial dependence, as modelled in Conley (1999). The distance is selected following a decision rule in which I choose the distance that provides the most conservative standard errors, looped over all distances between 10 and 1000km. Results are also robust to applying family-wise errors rates (FWER) using the stepwise procedure of Romano and Wolf (2005).

Figure 3.5: Semi-parametric Estimates of the Relationship between Rainfall Variability and School Attendance at different stages of the Life Cycle

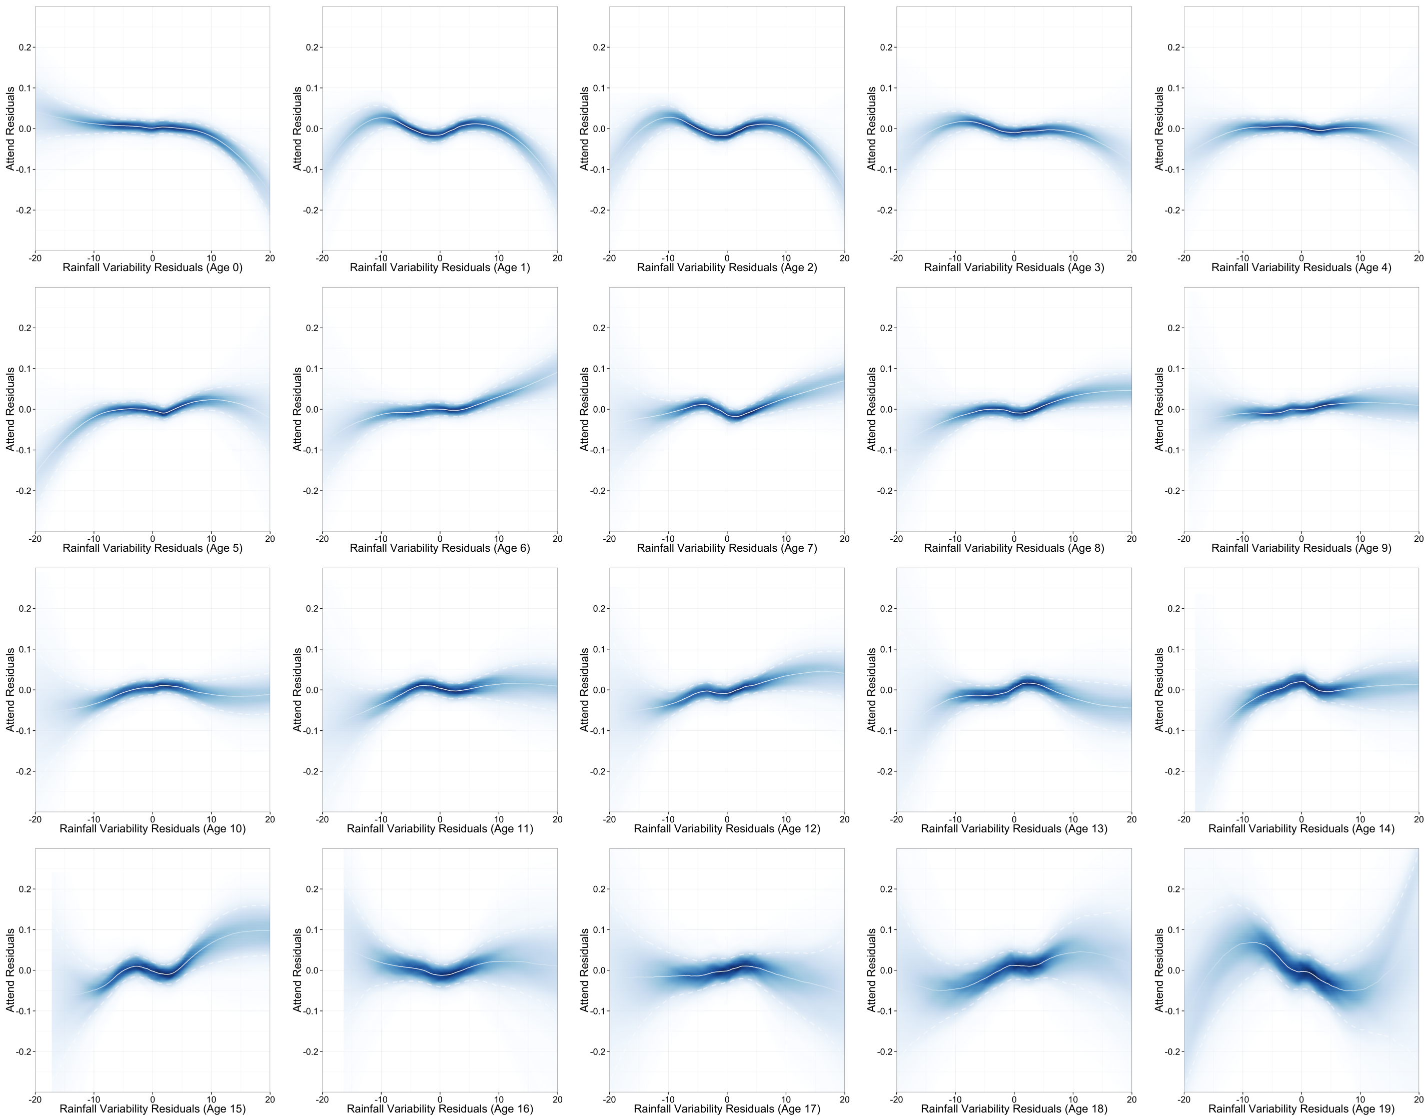

Notes: The measure of rainfall variability plotted here is the coefficient of variation, measured over the previous 5 years prior to the year in which the child was aged a. Each variable is regressed on village (α_v) and year of age (α_{t_a}) fixed effects as well as historical rainfall and temperature controls corresponding to the same time period. The figures above are the result of loess regressions of the residuals from this exercise.

Chapter 4

The Transmission of Localized Productivity Shocks in a Globalized World

This paper seeks to quantify the direct effects of local productivity shocks on firms, as well as the indirect effects through supply chain networks. We construct a unique dataset linking data on the economic performance and global trade transactions of French manufacturing firms with global weather data. Exploiting firm-level variation in exposure to both domestic and foreign weather shocks, we estimate both the local and linkage effects of weather on the economic performance of manufacturing firms in France. We observe that domestic exposure to higher temperatures and greater rainfall is associated with a reduction in the production of manufacturing firms. On the demand side, we estimate that increases in rainfall downstream results in an expansion in the production of upstream firms, suggesting that firms are able to increase their market share in response to localised productivity shocks in downstream markets. On the supply side, we observe that, on average, weather variation upstream has little effect on downstream production. However, we show that this effect is heterogeneous across firms, finding that firms with a greater initial import share from developing countries experience a relative contraction in production in response to increases in temperature upstream. This effect is attenuated for firms with a greater initial import-share from countries with greater access to air conditioning, indicating that the effect of temperature on production in these countries is due to thermal stress. These results suggest that localised productivity shocks can have significant economic effects across countries, and that if we fail to account for the interconnectedness of firms and sectors we may substantially underestimate the consequences of short-run weather and future climate change on economic activity.

4.1 Introduction

Economic activity has become increasingly integrated and increasingly global since the middle of the 20th century. A significant feature of this increased integration has been an increase in the fragmentation of production through vertical specialisation and integration, allowing firms to exploit reductions in the cost of production and take advantage of the productivity increases associated with specialisation – a global division of labour. However, this fragmentation also increases exposure to risk, with the potential to affect the economic performance and reputation of firms.[1] Indeed, a study conducted by the insurance company Zurich and the Chartered Institute of Purchasing and Supply (CIPS) found that 75% of business (based on a survey of 500 firms across 71 countries) experienced at least one supply chain disruption during 2012.

However, it is not only the economic environment that has changed over this period. It is now well established that we are committed to some form of climate change. Current projections suggest that an increase in global mean temperature of at least 2°C is very likely (IPCC, 2007; Intergovernmental Panel on Climate Change, 2012, 2014). This could lead to sea level rise and changed weather patterns, including more violent and more frequent storms, and an increase in the incidence of heat waves (Barriopedro et al., 2011).

What is less clear is the economic consequences of these effects, and how they are distributed across countries, sectors, and individuals. This paper does not attempt to attribute our estimates of extreme weather on manufacturing productivity to climate change, nor speak to the future impact of climate change per se. Rather, we attempt to understand the channels and mechanisms through which extreme weather events have historically had an effect on economic activity, to aid in the design of effective interventions to minimise the economic consequences of future events. In this respect, we seek to contribute to a growing literature that explores the empirical relevance of weather on economic and social outcomes (Dell et al., 2014). However, our understanding of how weather affects sectors of the economy, other than agriculture, is very limited (Colmer, 2016).

This paper explores the effects of local productivity shocks on manufacturing productivity, when production is globally fragmented. We introduce a theory of production that distinguishes three channels through which these productivity shocks affect firm performance: first, through local productivity shocks at domestic production lo-

[1] In recent years, firms have increasingly suffered from shocks such as factory disasters in Bangladesh, the horse meat scandal in the UK, extreme weather events, or the discovery of labour trafficking and bonded labour.

cations; second, through upstream effects (i.e., productivity shocks at the production locations of intermediate inputs); finally, through downstream effects (i.e., productivity shocks at the location of the firm's customers). Then, combining high-resolution atmospheric data with a unique dataset of firm-level trade transactions, we exploit exogenous variation in exposure to weather events both domestically and internationally to estimate the direct and linkage effects of local productivity shocks on the productivity of manufacturing firms in France.

We find evidence that domestic variation in temperature and rainfall affect manufacturing production – the direct effects of local productivity shocks. This helps to demonstrate that weather variation can significantly impact manufacturing production, even in a developed country context, and is consistent with an expanding literature exploring the effects of weather, and other environmental factors, on worker and firm productivity (Cachon et al., 2012b; Graff Zivin and Neidell, 2014; Adhvaryu et al., 2015; Somonathan et al., 2015; Colmer, 2016).

In terms of international exposure, we estimate linkage effects through exports (demand-driven) and imports (supply-driven). On the demand side, we observe that increases in rainfall downstream results in expansions in the production of domestic firms, suggesting that when firms in downstream countries experience localised productivity shocks, upstream firms are able to increase their market share by serving downstream demand. On the supply side, we observe that, on average, weather variation upstream has very little effect on downstream production, indicating that, while variation in the weather may have direct effects on production, damages do not appear to be propagated through the production network. However, we show that this effect is heterogeneous across firms, finding that firms with a greater initial import-share from developing countries experience a contraction in production in response to increases in temperature upstream. To understand the meaning of this "development effect", we consider three channels through which developing countries may be more affected by increases in temperature: (1) physical exposure due to their location in the tropics; (2) economic sensitivity and lower adaptive capacity due to weak infrastructure; (3) weaker institutions or governance that impede the ability of firms downstream to manage supply chain disruptions.

First, we find that firms exposed to a greater share of trading partners from countries in hot climate – defined as countries in the top tercile of the global climate distribution – are not differentially affected by increases in temperature upstream. Furthermore, this channel neither mediates nor moderates the development effect. This suggests that the development effect is not driven by physical exposure.

Second, we find that firms exposed to trading partners with a lower "ease of doing business" rank – based on the World Bank's Doing Business reports, which provide comparable cross-country data on the quantity and quality of business regulatory environments – experience a contraction in manufacturing production in response to temperature increases upstream. This suggests that governance and institutional factors are important in managing supply-chain disruptions. This control partially mediates the development effect, but does not moderate it.

Finally, we show that the firms exposed to trading partners with a higher import-weighted air conditioning availability per capita – defined as the trade value of air conditioning equipment per capita in US dollars – experience an increase in manufacturing production in response to temperature increases upstream. While this control does not mediate the development effect, we show that the interaction of air conditioning availability with the share of imports from developing countries does attenuate the development effect completely, suggesting that the temperature effects estimated domestically, within France, and upstream in developing countries, may be the result of thermal stress, and that adaptive capacity plays an important role in mediating the effects of upstream weather effects.

Our results indicate that localised productivity shocks can have significant economic effects across countries, and that by failing to account for the interconnectedness of firms and sectors we may substantially underestimate the consequences of short-run weather and future climate change on economic activity.

In addition, while adaptation to climate change has largely been considered a private good, in contrast to climate change mitigation – a public good –, it is interesting to consider that climate change adaptation may have a public good component and that, consequently, there may be underinvestment in adaptation. As such, there may be an incentive for firms and governments to invest in adaptation beyond their geographic boundaries to attenuate their exposure to climate change through production linkages.

The next section will present existing literature on this topic. Section 4.3 presents the theoretical framework that provides the foundation for our empirical analysis, Section 4.4 presents our empirical approach, Section 4.5 presents the data and its construction from the raw data and Section 4.6 provides the results from the econometric analysis. Section 4.7 concludes.

4.2 Literature Review

The findings of this paper contribute to three areas of literature. First, we contribute to a scarce literature on the fragmentation of trade and the impact of supply chain disruptions. With the rise of global supply chains and just-in-time production strategies, the manufacturing sector is more vulnerable to productivity shocks in other countries (Baldwin, 2013). Costinot et al. (2016) show that the impact of local shocks depend not only on their average level, but also on the dispersion of these shocks across the world. As a result, the impact of productivity shocks on aggregate are ambiguous and will depend on how they affect comparative advantage both within and between countries. Trade has the potential to dampen or exacerbate the impact of productivity shocks, depending on whether shocks increase heterogeneity across countries (increasing trade), or result in more homogeneity across countries (reducing trade). This paper argues that, in the context of manufacturing, the degree of interdependence between nations will also be important – it is not only changes in comparative advantage that matter, but also the level of fragmentation in the production of goods through vertical specialisation and integration.

Second, we add to a large emerging literature that analyses climatic influence on economic outcomes. The impact of major weather events on European economies has been the focus of a number of recent studies and governmental reports. The 2003 heat wave in France had some moderate economic but tragic public health consequences.[2] The 2007 central England summer floods cost the economy over £3 billion (ASC, 2010) and the harsh winter of 2009 cost £1 billion and prompted a government review into the resilience of England's transport system (DfT, 2010). The weather may also have positive impacts on the economy.Subak et al. (2000) argue that the hot summer of 1995 increased expenditure in tourism by around £239 million. These case studies suggest that the relationship between economic and natural systems can have important economic consequences.

There is also a wealth of econometric studies that have gone beyond estimating the impacts of one-off events in individual countries. These studies apply panel data methods to examine the effects of weather and extreme events, over time and within a given spatial location, on economic performance (See Dell et al. (2014) for a recent review of this literature). Within this literature, there has been recent interest in examining the impact of weather on sectors of the economy that are less "climate-sensitive". They propose a number of channels through which inclement weather could affect production.

First, inclement weather could cause delays to inbound delivery of parts from suppliers, due to accidents, congestion or delayed shipments (Brodsky and Hakkert, 1988; Golob and Recker, 2003). Conditional on the importance of intermediary goods, the network structure of supply chains (Baldwin and Venables, 2013), and the degree to which the firm is able to store inventories, such effects may be exacerbated or mitigated. In the same way, we might expect this channel to impact employees. If workers are unable to, or choose not to, travel to work, we would expect a reduction in productivity (Bandiera et al., 2015b).

Weather could also impact worker productivity directly, through physiological or cognitive channels. Indoor temperatures may reduce labour productivity, even if

[2]http://www.senat.fr/rap/r03-195/r03-195.html

firms have air conditioning units, as the weather outside could influence the emotional or psychological state of employees, which in turn may impact their productivity (Simonsohn, 2010; Cachon et al., 2012b; Metcalfe et al., 2012; Lee et al., 2012; Baylis, 2016). These effects may be exacerbated through the hierarchy of the firm if management decisions are affected by reductions in cognitive ability (Bandiera et al., 2015b; Adhvaryu et al., 2014).

Other factors of production may also be affected. For example, capital stocks and flows may be affected if weather affects capital depreciation, the relative productivity of inputs (if productivity shocks are not hicks neutral), or the level of investment in the economy if capital is locally constrained. In areas with fragile electricity infrastructures, an increase in temperature or reduction in rainfall in areas dependent on hydroelectric power generation may put additional stress on, what may already be, a fragile electricity infrastructure, reducing the supply of electrical power (Ryan, 2014; Alcott et al., 2015). This may affect labour productivity indirectly due to a reduction in air conditioning, capital directly, as capital needs electricity for production, or may increase the cost of production where firms rely on generators.

The main limitation of this literature, whether we are examining the effects of weather variation or natural disasters, in the short-run or long-run, is that most estimates fail to help us to understand the empirically relevant mechanisms through which these factors affect economic outcomes. Given the multiple channels through which these events could affect economic outcomes, the estimated coefficients provide merely a net effect, with little economic interpretation. In the event that multiple channels have the same sign, we may overestimate the importance of channels through confirmation bias, interpreting the effects as propagating through channels that conform to our prior beliefs, reducing the effectiveness of interventions if other unexpected channels are more important. If the effects from different channels move in opposite directions, the net effect could substantially underestimate the economic impact, reducing the perceived benefits of intervention. Without a better understanding of the channels and mechanisms through which these effects flow, it is very difficult to design effective policies to mitigate the damages from future events and exploit economic opportunities where present.

The final literature relating to the objectives of this paper explores the origins of aggregate fluctuations. The observation that economy-wide shocks, such as inflation or conflict, have little role in explaining most fluctuations has led economists to explore the role that idiosyncratic shocks play in explaining aggregate fluctuations (Lucas, 1977; Jovanovic, 1987; Cochrane, 1994; Carvalho, 2010; Gabaix, 2011; Acemoglu et al., 2012; Carvalho and Gabaix, 2013; Carvalho, 2014). The idea that idiosyncratic shocks could explain aggregate fluctuations has long been dismissed due to a "diversification argument": in an economy consisting of n sectors hit by independent shocks, aggregate fluctuations would have a magnitude proportional to $\frac{1}{\sqrt{n}}$, a negligible effect at high levels of disaggregation. It is observed that aggregate output concentrates around its mean at a very rapid rate. Consequently, microeconomic shocks should average out and only have negligible aggregate effects.

However, this argument has been shown to be invalid when firms and sectors are interconnected in an asymmetric way, propagating local productivity shocks throughout the rest of the economy, resulting in significant aggregate fluctuations (Acemoglu et al., 2012). This idea is founded on a rapidly expanding literature that explores the role that networks play in economic activity (Jackson, 2008; Atalay, 2014; Carvalho and Gabaix, 2013). Gabaix (2011) also shows that the diversification argument may not apply when the firm size distribution is sufficiently heavy-tailed. Drawing on the results from this literature, the next section presents a theoretical framework to provide a basis for the estimating equations presented in the empirical section.

4.3 Theoretical Framework

4.3.1 A Simple Model

We consider an economy where final demand is driven by consumers that have love-of-variety preferences for a set of F final goods. Final good consumers are distributed across G countries. Preferences of consumers in country g are

$$U_g = \sum_{f \in \mathbb{F}} \frac{1}{1 - \frac{1}{\eta_f}} \Lambda_{fg}^{1-\frac{1}{\eta_f}} Q_{fg}^{1-\frac{1}{\eta}}$$

where Λ_{fg} is a good-specific demand shock within country g. Consequently, the demand for good f by consumers in country g becomes

$$Q_{fg} = \kappa_g^{-\eta_f} \Lambda_{fg}^{\eta_f - 1} P_f^{-\eta_f}$$

where κ_g is a Lagrange multiplier. Consumption goods are produced using a set $\mathbb{I}$ of I intermediate goods using Cobb-Douglas production technology:

$$Q_f = \prod_{i \in \mathbb{I}} Q_{fi}^{\alpha_{fi}} \tag{4.1}$$

We assume that these production functions are constant returns so that $\sum_{i \in \mathbb{I}} \alpha_{fi} = 1$ for all f. Initially, we assume that there is just one intermediate per final good so that $Q_f = Q_1$. This intermediate good 1 requires a set of other intermediates $\mathbb{I}$ as well as labour and is produced by a Cobb-Douglas technology:

$$Q_1 = A_1 L_1^{\alpha_{1L}} \prod_{j \in \mathbb{I}} Q_{1j}^{\alpha_{1j}}$$

where A_1 is a Hicks-neutral productivity shock and $\alpha_{11} = 0$. Also assume again that this production function is constant returns: $\alpha_{1L} + \sum_{j \in \mathbb{I}} \alpha_{1j} = 1$.

The unit cost of producing good 1 is consequently

$$c_1 = A_1^{-1} \left(\frac{P_{L1}}{\alpha_{L1}} \right)^{\alpha_{L1}} \Pi_j \left(\frac{P_j}{\alpha_{1j}} \right)^{\alpha_{1j}}$$

where P_{L1} is the wage firm 1 is facing and P_j are the prices of the intermediate goods $j \neq 1$.

Intermediate inputs other than good 1 use only labour:

$$Q_j = A_j L_j \text{ for } j \neq 1$$

Hence unit costs become,

$$c_j = \frac{P_{Lj}}{A_j} \text{ for } j \neq 1$$

Suppose that all intermediate and labour markets are competitive. Hence $P_j = \frac{P_{Lj}}{A_j}$ for $j \neq 1$ and,

$$P_1 = A_1^{-1} \left(\frac{P_{L1}}{\alpha_{L1}} \right)^{\alpha_{L1}} \Pi_j \left(\frac{P_{Lj}}{A_j \alpha_{1j}} \right)^{\alpha_{1j}}$$

Finally, we require that in the final good, and all intermediate goods, markets clear. Note that total global demand for final good 1 becomes,

$$Q_1 = \frac{P_f^{1-\eta_f} \sum_{g \in \mathbb{G}} \kappa_g^{-\eta_f} \Lambda_{fg}^{\eta_f - 1}}{P_1} \frac{1}{\mu_f} = \frac{\left(\mu_f P_1\right)^{1-\eta_f} \sum_{g \in \mathbb{G}} \kappa_g^{-\eta_f} \Lambda_{fg}^{\eta_f - 1}}{P_1} \frac{1}{\mu_f}$$

where μ_f is the markup over marginal costs implied by the demand elasticity: $\mu_f = \frac{1}{1 - \frac{1}{\eta_f}}$. Alternatively, we can write,

$$P_1 Q_1 = R_1 = \left(\mu_f P_1\right)^{1-\eta_f} \sum_{g \in \mathbb{G}} \kappa_g^{-\eta_f} \Lambda_{fg}^{\eta_f - 1} \frac{1}{\mu_f}$$

Using the mean value theorem, we can log linearise this expression as,

$$\Delta r_1 \approx \left(\eta_f - 1\right) \left[\sum_{g \in G} \frac{\overline{R_{1g}}}{R_1} \Delta \lambda_{fg} \right] + \left(\eta_f - 1\right) \left[\Delta a_1 + \frac{\overline{P_{L1} L_1}}{R_1} \Delta p_{L1} - \sum_{i \neq 1} \frac{\overline{P_i Q_{1i}}}{R_1} \left(\Delta p_{Li} - \Delta a_i \right) \right] \tag{4.2}$$

where Δr_1 is the (log) change between two periods in revenue and similarly for $\Delta\lambda_{fg}$, Δp_{Li} and Δa_i. $\frac{R_{1g}}{R_1}$ is the share of revenue from country g (or export share in revenue). $\frac{P_{L1}L_1}{R_i}$ is the expenditure share on labour and $\frac{P_i Q_{1i}}{R_1}$ the expenditure share for intermediate factor i by firm 1. A bar indicates the average share over the two comparison periods.

Now suppose that weather events can affect this economy in two ways: via demand or supply shocks; i.e.,

$$\lambda_{fg} = \gamma_\lambda W_{fg} + \varepsilon_{\lambda g}$$

$$a_i = \gamma_a W_i + \varepsilon_{ai}$$

where W represents a set of weather variables. We can use this to rewrite 4.2 as

$$\Delta r_1 \approx \beta_{DOWN}\Delta W_{DOWN} + \beta_{LOC}\Delta W_1 + \beta_{UP}\Delta W_{UP} + \varepsilon_1 \tag{4.3}$$

where:

$$\beta_{DOWN} = \gamma_{\lambda DOWN}\left(\eta_f - 1\right)$$

$$\beta_{UP} = \gamma_{aUP}\left(\eta_f - 1\right)$$

$$\beta_{LOC} = \gamma_{aLOC}\left(\eta_f - 1\right)$$

$$\Delta W_{DOWN} = \sum_{g\in G}\overline{\frac{R_{1g}}{R_1}}\Delta W_{fg}$$

$$\Delta W_{UP} = \sum_{i\neq 1}\overline{\frac{P_i Q_{1i}}{R_1}}\Delta W_i$$

$$\varepsilon_1 = \left(\eta_f - 1\right)\left[\sum_{g\in G}\overline{\frac{R_{1g}}{R_1}}\Delta\varepsilon_{\lambda g}\right] + \left(\eta_f - 1\right)\left[\Delta\varepsilon_{a1} + \overline{\frac{P_{L1}L_1}{R_1}}\Delta p_{L1} - \sum_{i\neq 1}\overline{\frac{P_i Q_{1i}}{R_1}}\left(\Delta p_{Li} - \Delta\varepsilon_{ai}\right)\right]$$

Hence, using a sales share–weighted average of downstream weather indicators and an intermediate input factor–weighted average of upstream weather indicators allows us to determine the marginal impact of different weather events. A number of points are worth making about equation 4.3. First, note that the regression coefficients capture, the effect of weather variation through productivity and demand shocks, as well as the elasticity of demand. If demand is highly inelastic ($\eta_f \to 1$) so that firms can pass trough the negative (or positive) effects of weather to consumers, we would expect no marginal impact on revenues. Second, correct identification of these weather effects will depend on the characteristics of ε_1. It is unlikely that ε_1 is independent of the weather. For instance, wages could be affected by weather shocks as well. Third, the structure of the economy that led to equation 4.3 is highly restrictive. For instance, while we allowed good 1 to require a range of other intermediates, we imposed that all these other intermediates only use labour as an input. In reality, input–output linkages are probably more complex. Good 1 could itself be a required input in the production of some of its own inputs. Equally, some goods might be used both as inputs in the production of other goods and for final products. This can potentially introduce serious measurement error in the construction of our weather indicators where we are currently able to trace only nodes of first degree; e.g., we see that a French firm imports from Germany and therefore consider that the firm is exposed to weather shocks in Germany. However, if the German firm it is importing from imports a large fraction of their intermediates from Asia this would not be captured. Equally, we would understate the exposure of the firm to local French shocks if it turned out that the German firm is importing a considerable amount of intermediates from France. In future work we will explore these issues by taking on a sectoral approach, relying on country-specific input–output data as well as bilateral trade data. To handle this data we need a more general model, which we develop below.

4.3.2 A more general model

We build on the basic model introduced above. In particular, we impose no further restrictions on the composite consumption goods described in equation 4.1.

Moreover, we make no asymmetric assumption between the different intermediate goods; i.e., all intermediate goods are produced according to a general production function of the form,

$$Q_i = A_i L^{\alpha_{iL}} \prod_{j \in \mathbb{I}} Q_{ij}^{\alpha_{ij}}$$

for all $i \in \mathbb{I}$ where we assume constant returns to scale: $\alpha_{iL} + \sum_{j \in \mathbb{I}} \alpha_{ij} = 1$ for all i. We denote by $\mathbf{M}_M$ the $i \times i$ matrix collecting all factor intensities α_{ij}.

Again we assume that all intermediate inputs are competitively supplied. The unit cost of supplying final products will be,

$$c_f = \Pi_i \left(\frac{P_i}{\alpha_{fi}} \right)^{\alpha_{fi}}$$

where P_i is the price of the intermediate i. Similarly, the unit cost of intermediate products will be,

$$c_i = A_i^{-1} \left(\frac{P_{Li}}{\alpha_{Li}} \right)^{\alpha_{Li}} \Pi_j \left(\frac{P_j}{\alpha_{ij}} \right)^{\alpha_{ij}} \tag{4.4}$$

We can solve for the equilibrium in this economy in two steps. First, because of perfect competition in all intermediate markets, it is the case that,

$$P_i = c_i \text{ for all } i \in \mathbb{I}$$

As a consequence, equation 4.4 implies a system of equations in log terms:

$$\mathbf{p} = \mathbf{A} + \mathbf{M}_M \mathbf{p} \tag{4.5}$$

where $\mathbf{A}$ is a vector consisting of elements $-a_i - \sum_j \alpha_{ij} \ln \alpha_{ij} + \alpha_{Li} (p_{Li} - \ln \alpha_{Li})$. Second, we can then solve equation 4.5 and express all intermediate prices as (log) linear combinations of the productivity shocks a_i:

$$\mathbf{p} = \mathbf{S_P A}$$

where $\mathbf{S}_P = (\mathbf{I}_J - \mathbf{M}_M)^{-1}$.

We can further solve for the revenue of all producers by requiring that both intermediate and final goods markets are in equilibrium. For this, note that demand for intermediate input i from final product f is,

$$Q_{fi} = \frac{P_f^{1-\eta_f} \sum_{g \in \mathbb{G}} \kappa_g^{-\eta_f} \Lambda_{fg}^{\eta_f - 1}}{\mu_f P_i} \alpha_{fi} = \frac{c_f^{1-\eta_f} \sum_{g \in \mathbb{G}} \kappa_g^{-\eta_f} \Lambda_{fg}^{\eta_f - 1}}{P_i} \mu_f^{-\eta_f} \alpha_{fi} \text{ for } f \in \mathbb{F} \text{ and } i \in \mathbb{I}$$

whereas factor demand from intermediates becomes,

$$Q_{ij} = \frac{P_i Q_i}{P_j} \alpha_{ij} \text{ for } i, j \in \mathbb{I}$$

Hence, the equilibrium conditions imply the following equations:

$$P_i Q_i = \sum_{f \in \mathbb{F}} c_f^{1-\eta_f} \sum_{g \in \mathbb{G}} \kappa_g^{-\eta_f} \Lambda_{fg}^{\eta_f - 1} \mu_f^{1-\eta_f} \frac{\alpha_{fi}}{\mu_f} + \sum_{j \in \mathbb{I}} P_j Q_j \alpha_{ji} \text{ for all } i \in \mathbb{I}$$

which we can write as,

$$\mathbf{R} = \mathbf{A}_R + \mathbf{M}_\mathbf{M}^{'} \mathbf{R}$$

where $\mathbf{A_R}$ is a vector of the expenditure on intermediate good i derived directly from final good f consumption.

So that,

$$\mathbf{R} = \mathbf{S}_R \mathbf{A}_R$$

where $\mathbf{S}_R = \left(\mathbf{I}_I - \mathbf{M}_M^{\prime}\right)^{-1}$; i.e., the revenue of a particular intermediate is a linear combination of final good revenues.

Consequently,

$$\begin{aligned} R_i &= \sum_{j \in \mathbb{I}} s_{Rij} \sum_{f \in \mathbb{F}} c_f^{1-\eta_f} \sum_{g \in \mathbb{G}} \kappa_g^{-\eta_f} \Lambda_{fg}^{\eta_f - 1} \mu_f^{1-\eta_f} \frac{\alpha_{fj}}{\mu_f} \\ &= \sum_{j \in \mathbb{I}} s_{Rij} \sum_{f \in \mathbb{F}} \sum_{g \in \mathbb{G}} c_f^{1-\eta_f} \kappa_g^{-\eta_f} \Lambda_{fg}^{\eta_f - 1} \mu_f^{1-\eta_f} \frac{\alpha_{fj}}{\mu_f} \end{aligned}$$

Note that,

$$\frac{\partial R_i}{\partial \lambda_{fg}} = \left(\eta_f - 1\right) R_{fg} \frac{\alpha_{fj}}{\mu_f} \sum_{j \in \mathbb{I}} s_{Rij}$$

where $R_{fg} = c_f^{1-\eta_f} \kappa_g^{-\eta_f} \Lambda_{fg}^{\eta_f - 1} \mu_f^{1-\eta_f}$ is the expenditure by consumers from country g on final good f.

$$\begin{aligned} \frac{\partial R_i}{\partial a_k} &= \sum_{j \in \mathbb{I}} s_{Rij} \sum_{f \in \mathbb{F}} \sum_{g \in \mathbb{G}} \left(1 - \eta_f\right) c_f^{1-\eta_f} \kappa_g^{-\eta_f} \Lambda_{fg}^{\eta_f - 1} \mu_f^{1-\eta_f} \frac{\alpha_{fj}}{\mu_f} \frac{\partial \ln c_f}{\partial a_k} \\ &= \sum_{j \in \mathbb{I}} s_{Rij} \sum_{f \in \mathbb{F}} \left(1 - \eta_f\right) \frac{\alpha_{fj}}{\mu_f} \frac{\partial \ln c_f}{\partial a_k} \mu_f^{1-\eta_f} \sum_{g \in \mathbb{G}} c_f^{1-\eta_f} \kappa_g^{-\eta_f} \Lambda_{fg}^{\eta_f - 1} \\ &= \sum_{j \in \mathbb{I}} \sum_{f \in \mathbb{F}} \left(1 - \eta_f\right) s_{Rij} \frac{\alpha_{fj}}{\mu_f} \frac{\partial \ln c_f}{\partial a_k} R_f \end{aligned}$$

where R_f is the total (global) revenue/expenditure on good f.

$$\frac{\partial \ln c_f}{\partial a_k} = - \sum_{n \in \mathbb{I}} \alpha_{fn} s_{Pnk}$$

Hence we can write the (log) change in revenue between two periods, equivalent to equation 4.2 as,

$$
\begin{aligned}
\Delta r_{it} \approx & \qquad \overline{\sum_{f\in\mathbb{F}}\sum_{g\in\mathbb{G}}\left(\eta_f-1\right)\frac{R_{fg}}{R_i}\frac{\alpha_{fj}}{\mu_f}}\sum_{j\in\mathbb{I}}s_{Rij}\Delta\lambda_{fgit} \\
& -\sum_{k\in\mathbb{I}}\overline{\sum_{j\in\mathbb{I}}\sum_{f\in\mathbb{F}}\left(1-\eta_f\right)s_{Rij}\frac{\alpha_{fj}}{\mu_f}\sum_{n\in\mathbb{I}}\alpha_{fn}s_{Pnk}\frac{R_f}{R_i}}\Delta a_{kit} \\
& -\sum_{k\in\mathbb{I}}\overline{\sum_{j\in\mathbb{I}}\sum_{f\in\mathbb{F}}\left(1-\eta_f\right)s_{Rij}\frac{\alpha_{fj}}{\mu_f}\sum_{n\in\mathbb{I}}\alpha_{fn}s_{Pnk}\frac{R_f}{R_i}\alpha_{kL}}\Delta p_{Lkt}
\end{aligned}
$$

i.e. the change in revenue is a linear combination of the wage, demand, and productivity shocks for all products.

4.4 Econometric Approach

Most existing studies on the impact of weather on economic activity rely on country- or, at best, industry-level data. Here we rely on a large longitudinal sample of firm-level data for France. Not only does this significantly increase the sample size, and consequently power of our statistical estimation, but most importantly it allows us to construct a counterfactual based on firm-level exposure to trade and domestic weather patterns. In addition, the use of firm-level data presents us with various avenues for examining the mechanisms and channels through which weather events impact on economic activity.

In this study we distinguish 3 broad channels:

1. Upstream disturbances: losses due to weather disturbances for suppliers of business i
2. Production disturbances: losses due to weather disturbances at the production locations of business i
3. Downstream disturbances: losses due to weather disturbances at the locations of customers of business i.

The structure of our regression model is as follows:

$$
y_{ijt} = W'_{it}\beta_W + \alpha_i + \alpha_{jt} + \epsilon_{ijt} \tag{4.6}
$$

where i indexes firms, j sector, t time, and W_{it} is a vector of variables capturing weather events, α_{jt} and α_i are sector $\times$ year and firm fixed effects. The firm-level outcomes y_{ijt} considered below are the logarithm of value added, employment, and labour productivity, defined as value added per worker.

4.4.1 Construction of Weather Indices

The central element of our strategy is the construction of weather variables that vary at the firm level. Corresponding to our three disturbance channels, we construct three types of weather indices. First, to capture upstream disturbances, we construct import-weighted averages of global weather events. For a weather outcome in country c at time t, W_{ct}, we construct the upstream weather index as

$$W_{i,t}^{IMPORTS} = \sum_c \phi_{ic\tau}^{IM} \mathbf{W}_{ct}$$

where $\phi_{ic\tau}^{IM}$ represents the share of firm i's imports from country c at time τ defined as:

$$\phi_{ic\tau}^{IM} = \frac{M_{ic\tau}}{M_{i\tau}}$$

with firm i importing the amount $M_{ic\tau}$ from country c at time τ. To ensure the exogeneity of the treatment variable, we fix the ratio between trade and the firm's revenue at the earliest year, τ, that data is available for each firm. This is important because it fixes not only the endogenous choice of suppliers, but also endogenous movements in prices.

Second, to capture local production disturbances, we construct weather variables for the location within France of firm i. This allows us to include a vector of weather variables, $\mathbf{W}_{ijt}^{D}$, measured at firm level. The downside of this measure is that we cannot take into account the various locations of multi-plant firms.

Finally, to capture downstream disturbances, we construct an export-weighted index of global weather variables similar to the index constructed for imports described above:

$$W_{i,t}^{EXPORTS} = \sum_c \phi_{ic\tau}^{EX} \mathbf{W}_{ct}$$

where $\phi_{ic\tau}^{EX}$ represents the share of firm i's exports to country c at time τ defined as:

$$\phi_{ic\tau}^{EX} = \frac{X_{ic\tau}}{X_{i\tau}}$$

with firm i exporting the amount $X_{ic\tau}$ from country c at time τ.

While our import- and export-based measures cannot capture domestic upstream and downstream disturbances, this does not necessarily imply that we fail to pick up such effects. To the extent that upstream and downstream domestic weather events are correlated with weather events at the production locations of firms – W_{it}^{D} – these would be picked up by the production location effects – W_{it}^{D}. This should not lead to systematic errors of type I, rather it might lead to errors of type II; i.e., we have no reason to expect that our approach would identify weather events spuriously. On the other hand, we might not be fully capturing some effects that are nevertheless present.

4.4.2 Estimation

To establish the effects of weather on the outcome variables, we need to be reasonably confident that our error terms are independent of the weather variables. We employ the fixed-effects estimator to recover consistent estimates of the parameters of equation (4.6). This controls for fixed unobserved heterogeneity across firms. Moreover, we include sector-by-year dummies which account for a possible correlation between weather and the business cycle or sector-specific shocks. It might also be the case that more productive businesses locate systematically in areas with certain weather; e.g., cooler weather in the north. We account for that by allowing for fixed effects and only exploiting variation within firms. Hence, our identification rests on the assumption that any factors driving location-specific deviations in productivity are not correlated with the weather variables.

4.5 Data

In order to analyse the impact of weather on manufacturing firms, we combine three datasets: financial data from the Enquete Annuelle des Entreprises (EAE), firm-level trade transactions data from the French Institute of Statistics (INSEE), and weather data constructed from the ERA-Interim Reanalysis data.

4.5.1 L'Enquete Annuelle des Entreprise (EAE)

The Enquete Annuelle des Entreprise provides annual business survey data, available for the period 1993 to 2007. Firm- and plant-level data is collected for firms with more than 20 employees and all the plants of those firms. At the firm level, the dataset provides balance-sheet data including turnover, employment, capital, and aggregate wages, as well as information about firm location and industry classification.

4.5.2 Trade data

Firms in France submit declarations of an almost comprehensive record of annual exports and imports by destination/origin country at the eight-digit product level to

French Customs and data is available from 1996 to 2010.[3] Reporting thresholds exist for compulsory declarations inside and outside the European Union. Outside the EU, exports or imports are only reported if their annual total is above 1000 euros or 1000 kg. Within the EU, these thresholds vary through time and by flow, but many firms report their transactions despite being under the thresholds. In order to harmonise these different thresholds, we consider as non-traders firms whose total exports or imports within or outside the EU are less than 150,000 euros.

4.5.3 Weather data

In addition to the firm-level data and trade transaction data, weather data has been constructed from the ERA-Interim data archive provided by the European Centre for Medium-Term Weather Forecasting (ECMWF).[4] This dataset is the latest product provided by the ECMWF, offering a consistent set of high quality weather variables over time and space. For the purpose of our analysis, we use daily precipitation and average temperature, measured globally on a 0.75° (latitude) $\times$ 0.75° (longitude) grid (equivalent to 83 km $\times$ 83 km at the equator). This data provides a complete record of daily average temperatures and precipitation from 1979–2012 and is constructed using a process called Reanalysis, by which observational and satellite data is combined with sophisticated climate models to construct a consistent global best estimate of atmospheric parameters over time and space. This process results in an estimate of the climate system across a grid that is more uniform in its quality and realism than observational data (weather stations or satellites) alone, and is closer to the state of existence than any model would provide by itself. This source of data is hugely beneficial for our study, given that we are interested in both domestic and overseas weather variation. In this context the use, of observational data would result in inconsistent measurement across countries, due to differences in reporting practices and resources committed to the accurate collection of weather data. This would be of particular concern when examining the effect of trade shocks in developing countries, where the quality of observational weather data is considerably worse than in developed countries. While reanalysis data is partly computed using climate models that are imperfect and contain systematic biases, it allows us to compare our estimates of the domestic and international effects as they are nested in the same data.

The variables used to estimate the domestic effect of weather variation are matched to the postcode of each firm through a process of inverse distance weighting. For the domestic variables we use a version of the ERA-Interim that has been validated at

[3]Similar data from other countries has been used to analyse firm trading behaviour and gains from trade. The export dimension of this particular dataset for France is also used by Eaton et al. (2011) and Mayer et al. (2014) as well as others.

[4]See Dee et al. (2011) for a detailed discussion of the ERA-Interim data.

a 0.25° (latitude) × 0.25° (longitude) grid, providing greater resolution within country. A weighted annual average of the daily average temperature and precipitation variables is taken for all grid points within 100 km of the geographic centre of each postcode. The weighting used is the inverse of the distance squared from the postcode centroid.

4.5.4 Descriptive statistics

By combining the three datasets described above, we construct a unique firm-level trade-transaction dataset, which allows us to explore how firms are impacted by and respond to both domestic and foreign weather shocks.

We match our sample from the EAE for which employment, value added, and postcode information are available to the domestic French weather data from the ERA-Interim Reanalysis. For the trading firms among these, the trade transaction data from the French INSEE is merged in, yielding a final sample of 205,856 firm-years.

The resulting firm-level dataset's main variables are described in Table 4.1. Our final dataset contains an unbalanced panel of 205,856 firm-years that show a wide heterogeneity of size in terms of employment and value added. We observe that the value of imports is approximately one-fifth of total value added and that the value of exports is approximately one-third of total value added. In terms of international engagement, we observe that, on average, a firm imports around 20% of intermediates from developing countries, and exports 23% of products to developing countries.

In terms of local exposure to the weather, we observe that firms within France face an average temperature of 11.24°C and 828 mm of rainfall each year. The within-firm standard deviation for each variable is 0.452°C and 120 mm respectively. In terms of international exposure to the weather, we observe that firms face higher temperatures through their export relationships compared to their import relationships. The import-weighted average temperature exposure is 8.257°C and the export-weighted average temperature exposure is 9.579°C. The within-firm standard deviation for each variable is 0.378°C and 0.358°C respectively. In terms of rainfall, we observe that firms are exposed to more rainfall through their import partners than through their export partners. The import-weighted total rainfall exposure is 1,183 mm, while the export-weighted total rainfall exposure is 1,172 mm. The within-firm standard deviation for each variable is 103 mm and 116 mm respectively.

4.6 Results

In this section we examine how firms are affected by variation in the weather, both locally and through trade. We then explore whether there is any variation in the estimated trade effects based on the characteristics of trading partners. In particular, we examine whether the level of development moderates the estimated effect.

Table 4.2 indicates that local increases in temperature and rainfall are associated with contractions in value added and employment. In terms of the magnitude of these effects, a one standard deviation increase in temperature (0.452°C) is associated with a 1.74% contraction in value added and a 1.83% contraction in employment. The standardised effects of rainfall are smaller. A one standard deviation increase in rainfall (120 mm) is associated with a 0.53% contraction in value added and a 0.35% contraction in employment.

While we are able to estimate the relationship between weather and the economic production of firms, there is nothing in these estimates that gives us a deeper understanding of the channels and mechanisms that drive these effects. Consequently, we are unable to say anything about how firms may act to better manage these effects. However, having shown that weather has a significant effect on manufacturing production in a developed country context, we are able to elucidate one channel through which firms may be affected by weather, namely through the disruption of production through trade. To explore this channel, we estimate the effects of weather in trading partner countries.

First we examine the effects of weather in import partner countries. Table 4.2 shows that, on average, increases in temperature or rainfall in import-partner countries have no effect on production downstream, suggesting that firms may be able to manage the effects, indicating that there is little transmission of weather effects through supply chains. This may capture the ability of other countries to manage the effects of temperature, such that temperature does not matter for production in trading partner countries. Alternatively, this may capture the ability of firms in France to manage their supply chains, such that the consequences of temperature increases in import partner countries are not passed on to firms in France.

Second, we examine the effects of weather in export-partner countries. These impacts seek to capture any demand-side responses. Table 4.2 shows that, on average, an increase in temperature in export partner countries has no effect on production upstream. However, we estimate that increases in rainfall in export partner countries result in an increase in value added and the number of employees upstream. A one standard deviation increase in downstream rainfall (116 mm) is associated with a 0.26% increase in value added and a 0.156% increase in employment. These results are consistent with the interpretation that if weather effects reduce production locally, then this may increase demand for production from areas outside of the local market, resulting in an increase in exports from upstream markets.

To explore whether the estimated effects are constant, we begin by exploiting variation in the characteristics of trading partner countries to understand the existence of any heterogeneous effects.

4.6.1 Heterogeneity in the Development of Trading Partner Countries

We begin by exploring whether the level of development is a moderating factor in propagating the effects of weather events downstream, or affecting the demand for products upstream. A recent literature has documented significant effects of temperature on manufacturing production in developing countries – much larger that the estimated effects shown to affect manufacturing production in France (Dell et al., 2012; Adhvaryu et al., 2015; Somonathan et al., 2015; Colmer, 2016). Consequently, one may hypothesise that developing countries are less able to manage these effects due to weaker infrastructure and lower adaptive capacity, arising from institutional and governance factors. We initially explore the potential for this heterogeneous response by constructing for each firm the share of trade from developing countries in the baseline year:

$$\textsc{Development Share}^{G}_{i\tau} = \sum_{c} \phi^{G}_{ic\tau} \textsc{Developing}_{c} \text{ where } G \in \{Import, Export\}$$

where $\phi^{G}_{ic\tau}$ represents the share of firm i's imports/exports to country c at time τ,

$$\phi^{G}_{ic\tau} = \frac{G_{ic\tau}}{G_{i\tau}}$$

and $\textsc{Developing}_{c}$ is a dummy variable, equal to one, if the country is defined as a developing country based on the World Bank classification.

In table 4.3 we observe that, as the share of imports from developing countries increases, there is a relative contraction in value added and the number of employees. Evaluated at the average development share (0.20657), we estimate that firms with this import portfolio exposed to a one standard deviation increase in upstream temperatures (0.378) would experience a relative contraction of 0.125% in value added, and a relative contraction of 0.05% in employment. Interestingly, we observe that firms that have an import portfolio absent any developing countries experience an expansion of value added and employment when their trading partners experience warmer weather.

However, it is unclear as to how these effects should be interpreted, or rather what mechanisms underly this heterogeneity. Developing countries may be more likely to pass through the consequences of weather-driven supply chain disruptions downstream because they are: (1) more physically exposed due to their location in the tropics; (2) more economically sensitive and have lower adaptive capacity due to weak infrastructure; (3) impeded by institutional or governance factors that affect the ability of firms downstream to manage supply chain disruptions.

We seek to unbundle these channels by exploring variation in the estimated effect through physical exposure to higher temperatures, the regulatory business environment, and the availability of air conditioning. We do this through three additional exercises. First, we explore whether firms exposed to a greater share of trading partners from countries in hot climates are differentially affected by increases in upstream temperatures, and whether this mediates or moderates the development effect – an examination of the physical exposure channel. Second, we examine whether firms exposed to trading partners with a lower "ease of doing business" rank – based on the World Bank's Doing Business reports, which provide comparable cross-country data on the quantity and quality of business regulatory environments – are differentially affected by increases in temperature upstream, and whether this mediates or moderates the development effect – an exploration of the institutions and governance channel. Finally, we explore whether firms exposed to trading partners with a higher availability of air conditioning machines per capita – defined as the trade value of air conditioning equipment per capita in US dollars – are differentially affected by increases in upstream temperatures, and whether this mediates or moderates the development effect – an examination of adaptive capacity channel.

Understanding the Development Effect: Heterogeneity in the Climate of Trading Partner Countries

To understand what drives or moderates this development effect, we first examine whether firms exposed to a greater share of trading partners from countries with hot climate countries are differentially affected by increases in upstream temperatures – an exploration of the physical exposure channel.

To do this, we construct a measure using the same approach used to construct the development share variable, interacting the trade share between firm i and country c with a binary variable indicating whether country c is in the top tercile of the global climate distribution:

$$\textsc{Hot}^{G}_{i\tau} = \sum_{c} \phi^{G}_{ic\tau} \textsc{Hot}_{c} \text{ where } G \in \{Import, Export\}$$

where $\phi^G_{ic\tau}$ represents the share of firm i's imports/exports to country c at time τ,

$$\phi^G_{ic\tau} = \frac{G_{ic\tau}}{G_{i\tau}}.$$

$\textsc{Hot}_c$ is a dummy variable, equal to one, if the country is in the top tercile of the global temperature distribution, based on its 30-year average temperature.

In column (1), (4), and (7) of Table 4.4 we find that firms with a higher trade share from countries in hot climates are not differentially affected by increases in temperature upstream. Evaluated at the mean share of trade from hot climates (0.01327), a one standard deviation increase in temperature (0.378°C) would be associated with a 0.002% reduction in value added. In columns (2), (5), and (8) we examine whether controlling for the trade share from hot climates mediates the development effect, finding that there is little change in the development share coefficient. Finally, in columns (3), (6), and (9) we explore whether the development share coefficient is moderated if a firm also has a higher share of trade from hot climates; i.e., are firms more or less affected by the development share effect if they also have a higher share of trade from hot climates? We find that the share of trade from hot climates does not significantly moderate the development effect.

Understanding the Development Effect: Heterogeneity in the Business Environment of Trading Partner Countries

Next we examine whether firms exposed to trading partners with more difficult business regulation environments are differentially affected by increases in temperatures upstream, and whether this mediates or moderates the development effect – an exploration of the institutions and governance channel. To do this, we interact the trade share between firm i and country c with country c's rank in the World Bank's Doing Business report – which provides comparable cross-country data on the quantity and quality of business regulatory environments. This variable provides an indicator of the ease of doing business with country c,

$$\textsc{Doing Business Rank}^G_{i\tau} = \sum_c \phi^G_{ic\tau} \mathrm{DB}_c \text{ where } G \in \{Import, Export\}$$

where $\phi^G_{ic\tau}$ represents the share of firm i's imports/exports to country c at time τ,

$$\phi^G_{ic\tau} = \frac{G_{ic\tau}}{G_{i\tau}}$$

and $\mathrm{DB}_c = \frac{\textsc{Doing Business Rank}_c}{\max(\textsc{Doing Business Rank})}$ a continuous variable, normalised to one if the country rank is equal to the worst ranked country. A higher value is associated with a lower ease of doing business.

In columns (1), (4) and (7) of Table 4.5 we find that firms with a higher trade share from countries with a more difficult business environment are differentially affected by increases in upstream temperatures, experiencing a relative contraction in value added and employment. Evaluated at the mean rank (0.130), a one standard deviation increase in temperature (0.378°C) would be associated with a relative contraction in value added of 0.33%, and a relative contraction in employment of 0.14%. In columns (2), (5) and (8) we examine whether controlling for the trade share–weighted ease of doing business rank mediates the development effect, finding that the size and significance of both the development share and ease of doing business coefficients are mediated slightly. Finally, in columns (3), (6) and (9) we explore whether the development share coefficient is moderated if a firm also has a higher share of trade from countries with a more difficult business environment; i.e., are firms more or less affected by the development share effect if they also have a higher share of trade from countries with a more difficult business environment? Again, we find that the share of trade from countries with more difficult business environments does not significantly moderate the development effect.

Understanding the Development Effect: Heterogeneity in the Adaptive Capacity of Trading Partner Countries

Finally, we explore whether firms exposed to trading partners with a higher availability of air conditioning machines per capita are differentially affected by increases in upstream temperatures, and whether this mediates or moderates the development effect – an examination of the adaptive capacity channel. We do this by interacting the trade share between firm i and country c with a proxy for country c's availability of air conditioning machines – defined as the trade value of air conditioning equipment per capita in US dollars.[5] We do this due to the absence of comparable cross-country data on air conditioning availability. Our measure is a variant of the one used in Heal and Park (2015), who exploit the value of air conditioning imports per capita for each country. We argue that using imports as a proxy for air conditioning availability may miss countries that have a comparative advantage in producing air conditioning machines, and consequently do not import air conditioning. By also including the exports of air conditioning machines, we capture a more balanced measure of air conditioning availability based on domestic production and imports.

$$\textsc{Air Conditioning Availability}^{G}_{i\tau} = \sum_{c} \phi^{G}_{ic\tau} \text{AC } \textsc{per capita}_{c} \text{ where } G \in \{Import, Export\}$$

[5]Thanks go to Lucas Davis for this suggestion.

where $\phi^G_{ic\tau}$ represents the share of firm i's imports/exports to country c at time τ,

$$\phi^G_{ic\tau} = \frac{G_{ic\tau}}{G_{i\tau}}$$

and AC per capita$_c$ is the total trade value (imports + exports) of air conditioning machines per capita, measured in US dollars.

Table 4.6 reports the results of this exercise. In columns (1), (4) and (7) we observe that firms with a higher trade share from countries with greater access to air conditioning are differentially affected by increases in upstream temperatures, experiencing a relative expansion in value added and employment. Evaluated at the mean value of AC per capita (\$16.94), a one standard deviation increase in temperature (0.378°C) is associated with a relative expansion in value added and employment of 0.42% and 0.30% respectively. In columns (2), (5) and (8) we examine whether controlling for the trade share–weighted ease of doing business rank mediates the development effect, finding that the both size and significance of the development share are not mediated. However, in columns (3), (6) and (9) we explore whether the development share coefficient is moderated if a firm also has a higher share of trade from countries with more air conditioning availability; i.e., are firms more or less affected by the development share effect if they also have a higher share of trade from countries with more air conditioning availability? Interestingly, we find that the coefficient capturing the differential effect of temperature in countries with a higher trade share of air conditioning becomes insignificant, and that the development share effect is mediated by the triple interaction of upstream temperature, air conditioning availability and the share of trade from developing countries. Evaluating this effect at the mean development share (0.2657) and mean value of AC per capita (\$16.94), we estimate that a one standard deviation increase in temperature (0.378°C) is associated with a 0.016% increase in value added. However, in the absence of air conditioning, a one standard deviation increase in temperature (0.378°C) is associated with a 0.29% contraction in value added. In addition, we find that the availability of air conditioning per capita has no differential effect on firms, independently of the effect through developing countries, highlighting the role that thermoregulation can play in mitigating the effects of temperature on production, as well as the role that (the absence of) adaptive capacity may play in explaining the development effect.

4.7 Conclusion

This paper set out to explore the effects of local productivity shocks on manufacturing productivity when production is globally fragmented. We introduce a theory of production that distinguishes three channels through which productivity shocks could affect firm performance: first, through local productivity shocks at domestic production locations; second, through upstream effects (i.e., productivity shocks at the production locations of intermediate inputs); and finally, through downstream effects (i.e. productivity shocks at the location of the firms customers).

Empirically, we explore these mechanisms by combining high-resolution atmospheric data with a unique dataset of firm-level trade transactions. We exploit exogenous variation in exposure to weather events both domestically and internationally to estimate the direct and linkage effects of local productivity shocks on the productivity of manufacturing firms in France.

Within France, we find evidence that local increases in temperature and rainfall negatively affect manufacturing production – the direct effects of local productivity shocks. These findings are consistent with an expanding literature exploring that explores the effects of weather and other environmental factors on worker and firm productivity (Cachon et al., 2012b; Graff Zivin and Neidell, 2014; Adhvaryu et al., 2015; Somonathan et al., 2015; Colmer, 2016).

In terms of international exposure, we estimate linkage effects through exports (demand-driven) and through imports (supply-driven). On the demand side, we show that increases in rainfall downstream result in expansions in the production of upstream firms, suggesting that firms are able to increase their market share in response to localised productivity shocks in downstream markets.

On the supply side, we observe that, on average, weather variation upstream has very little effect on domestic production, indicating that damages do not appear to be propagated through the production network. However, we show that this effect is heterogeneous across firms, finding that firms with a greater initial import share from developing countries experience a relative contraction in production in response to increases in temperature upstream. To understand the meaning of this "development effect" we explore whether developing countries are more affected to temperature increases through three channels: (1) physical exposure due to their location in the tropics; (2) economic sensitivity and lower adaptive capacity, due to weak infrastructure; (3) weaker institutions or governance that impede the ability of firms downstream to manage supply chain disruptions.

Exploring these channels, we find that firms with a greater exposure to trading partners with more air conditioning per capita – measured using the trade value of air conditioning machines per capita – attenuates the development effect. This suggests that the temperature effects estimated upstream in developing countries – and perhaps domestically within France – are the result of thermal stress, and that (the lack of) adaptive capacity is an important consideration in explaining why upstream temperature increases in developing countries may affect manufacturing outcomes downstream.

Our results have two main implications. First, they suggest that localised productivity shocks can have significant economic effects through production linkages, and that by failing to account for the interconnectedness of firms and sectors we may substantially underestimate the consequences of short-run weather and future climate change on economic activity. Second, while adaptation to climate change has largely been considered a private good, in contrast to climate change mitigation (a public good), our results suggest that climate change adaptation may have a public good dimension, and that consequently there may be underinvestment in adaptation. As such, there may be an incentive for firms and governments to invest in adaptation beyond their geographic boundaries to attenuate their exposure to climate change through production linkages.

Table 4.1: Descriptive Statistics

	Mean	Std. Dev. (within)	Std. Dev. (between)	Obs.
Panel A: *Firm Data*				
Value Added (000,000's)	28.294	161.756	259.917	205,856
Employment	132	158	544	205,856
Labour Productivity	176.466	695.144	1,282.415	205,856
Panel B: *Trade Data*				
Total Value of Imports (000,000's)	6.771	50.799	99.264	205,856
Total Value of Exports (000,000's)	8.788	24.486	148.480	205,856
Share of Imports (Developing)	20.657%	–	33.468%	205,856
Share of Exports (Developing)	23.538%	–	35.995%	205,856
Share of Imports (Hot)	1.327%	–	9.493%	205,856
Share of Exports (Hot)	5.270%	–	19.013%	205,856
$\frac{\text{Ease of Doing Business Rank}}{\max(\text{Doing Business Rank})}$ (Imports)	0.130	–	0.106	205,856
$\frac{\text{Ease of Doing Business Rank}}{\max(\text{Doing Business Rank})}$ (Exports)	0.169	–	0.179	205,856
AC Availability Per Capita, $US (Imports)	16.946	–	8.500	205,856
AC Availability Per Capita, $US (Exports)	14.075	–	9.607	205,856
Panel C: *Domestic Weather Data*				
Daily Average Temperature (°C)	11.239	0.452	1.239	205,856
Annual Rainfall (mm)	828.490	120.452	136.529	205,856
Panel D: *Trade-Weighted Weather Data*				
Upstream Daily Average Temperature (°C)	8.257	0.378	4.445	205,856
Upstream Annual Rainfall (mm)	1,183.384	103.282	317.187	205,856
Downstream Daily Average Temperature (°C)	9.579	0.358	5.756	205,856
Downstream Annual Rainfall (mm)	1,172.448	116.597	389.509	205,856

Figure 4.1: Spatial Weather Variation in France

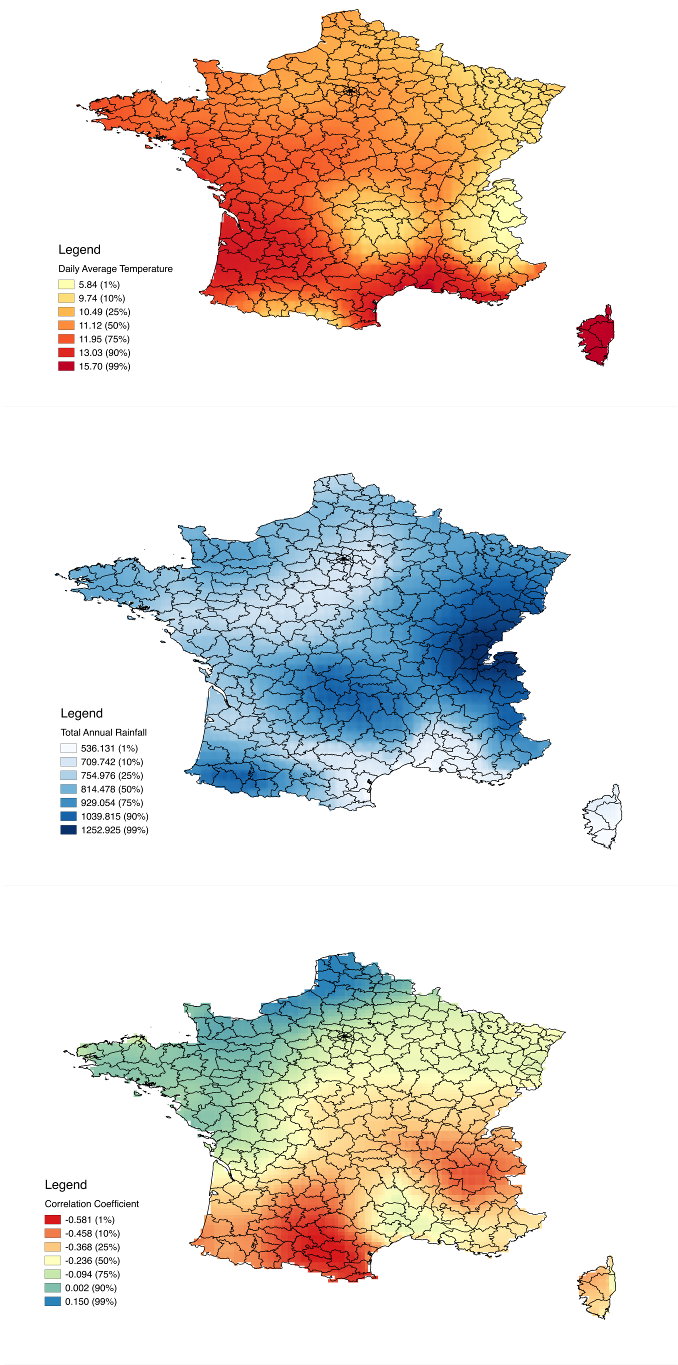

Figure 4.2: Spatial Weather Variation around the World

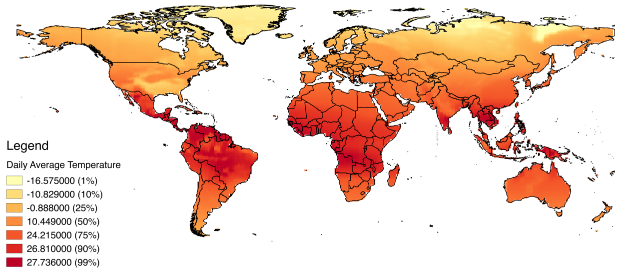

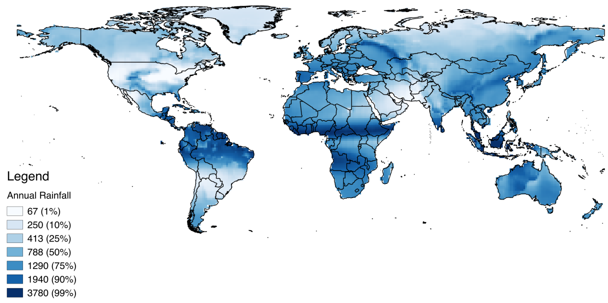

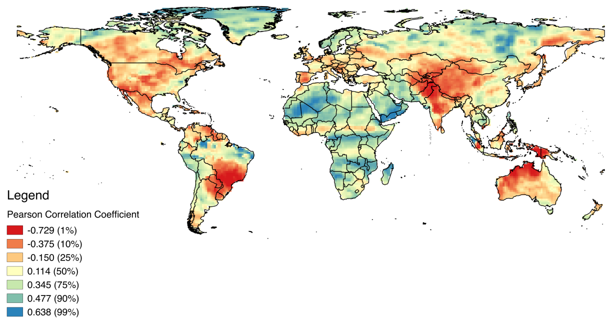

Table 4.2: The Effects of Weather on French Manufacturing Firms

	(1) Value Added	(2) No. of Employees	(3) Value Added Per Worker
Domestic Effects:			
LOCAL TEMPERATURE (°C)	-0.0386***	-0.0405***	0.00195
	(0.00562)	(0.00420)	(0.00445)
LOCAL RAINFALL (100 mm)	-0.00442***	-0.00296***	-0.00146
	(0.00120)	(0.000721)	(0.00106)
Import Partner Effects:			
UPSTREAM TEMPERATURE (°C)	0.00506	0.00319	0.00186
	(0.00358)	(0.00219)	(0.00303)
UPSTREAM RAINFALL (100 mm)	-0.00190	0.000214	-0.00212**
	(0.00117)	(0.000627)	(0.00108)
Export Partner Effects:			
DOWNSTREAM TEMPERATURE (°C)	-0.00266	-0.00329	0.000628
	(0.00338)	(0.00219)	(0.00277)
DOWNSTREAM RAINFALL (100 mm)	0.00230**	0.00135***	0.000946
	(0.000895)	(0.000500)	(0.000814)
FIXED EFFECTS		FIRM, SECTOR × YEAR	
Observations	205,856	205,856	205,856

NOTES: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. The unit of analysis is at the firm level. Local temperature is measured as the average daily temperature over the calendar year. Local rainfall is measured as the total rainfall accumulated over the calendar year. Downstream temperature is defined as the average daily temperature over the calendar year measured in each export partner country, weighted by the export-share with that country in the baseline year. Downstream rainfall is defined as the total rainfall accumulated over the calendar year measured in each export partner country, weighted by the export-share with that country in the baseline year. Upstream temperature is defined as the average daily temperature over the calendar year measured in each import partner country, weighted by the import-share with that country in the baseline year. Upstream rainfall is defined as the total rainfall accumulated over the calendar year measured in each import partner country, weighted by the import-share with that country in the baseline year. Standard errors are clustered at the commune level.

Table 4.3: The Effects of Weather in Developing Countries on French Manufacturing Firms

	(1) Value Added	(2) No. of Employees	(3) Value Added Per Worker
Domestic Effects:			
Local Temperature (°C)	-0.0381***	-0.0403***	0.00221
	(0.00564)	(0.00422)	(0.00447)
Local Rainfall (100 mm)	-0.00443***	-0.00301***	-0.00142
	(0.00120)	(0.000722)	(0.00106)
Import Partner Effects:			
Upstream Temperature (°C)	0.00866**	0.00480**	0.00387
	(0.00400)	(0.00244)	(0.00340)
Upstream Temperature × Development Share	-0.0161**	-0.00768*	-0.00843
	(0.00725)	(0.00456)	(0.00610)
Upstream Rainfall (100 mm)	-0.00243*	0.000684	-0.00311**
	(0.00141)	(0.000753)	(0.00129)
Upstream Rainfall × Development Share	0.00126	-0.00167	0.00294
	(0.00199)	(0.00109)	(0.00183)
Export Partner Effects:			
Downstream Temperature (°C)	-0.00203	-0.000658	-0.00138
	(0.00401)	(0.00256)	(0.00326)
Downstream Temperature × Development Share	-0.00175	-0.00813*	0.00638
	(0.00689)	(0.00440)	(0.00563)
Downstream Rainfall (100 mm)	0.00330***	0.00263***	0.000671
	(0.0000117)	(0.00000694)	(0.0000104)
Downstream Rainfall × Development Share	-0.00209	-0.00257***	0.000484
	(0.00166)	(0.000949)	(0.00143)
Fixed Effects		Firm, Sector × Year	
Observations	205,856	205,856	205,856

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. The unit of analysis is at the firm level. Local temperature is measured as the average daily temperature over the calendar year. Local rainfall is measured as the total rainfall accumulated over the calendar year. Downstream temperature is defined as the average daily temperature over the calendar year measured in each export partner country, weighted by the export-share with that country in the baseline year. Downstream rainfall is defined as the total rainfall accumulated over the calendar year measured in each export partner country, weighted by the export-share with that country in the baseline year. Upstream temperature is defined as the average daily temperature over the calendar year measured in each import partner country, weighted by the import-share with that country in the baseline year. Upstream rainfall is defined as the total rainfall accumulated over the calendar year measured in each import partner country, weighted by the import-share with that country in the baseline year. Development share is defined as the share of imports/exports from developing countries in the baseline year. This is aggregated over all trading partners for each firm. Standard errors are clustered at the commune level.

Table 4.4: The Heterogeneous Effects of Weather in Trading Partner Countries on French Manufacturing Firms

	(1) Value Added	(2) Value Added	(3) Value Added	(4) Employment	(5) Employment	(6) Employment	(7) Labour Productivity	(8) Labour Productivity	(9) Labour Productivity
Domestic Effects:									
LOCAL TEMPERATURE (°C)	-0.0387***	-0.0382***	-0.0382***	-0.0406***	-0.0404***	-0.0404***	0.00190	0.00218	0.00219
	(0.00562)	(0.00564)	(0.00564)	(0.00420)	(0.00422)	(0.00422)	(0.00445)	(0.00447)	(0.00447)
Import Partner Effects:									
UPSTREAM TEMPERATURE (°C)	0.00549	0.00880**	0.00895**	0.00310	0.00472*	0.00473*	0.00239	0.00408	0.00423
	(0.00359)	(0.00400)	(0.00400)	(0.00220)	(0.00244)	(0.00244)	(0.00304)	(0.00340)	(0.00340)
UPSTREAM TEMPERATURE × HOT	-0.0404	-0.0354	-0.0723	0.0178	0.0210	0.0201	-0.0582*	-0.0563*	-0.0924
	(0.0370)	(0.0373)	(0.0783)	(0.0222)	(0.0224)	(0.0434)	(0.0323)	(0.0326)	(0.0725)
UPSTREAM TEMPERATURE × DS		-0.0147**	-0.0155**		-0.00782*	-0.00785*		-0.00688	-0.00761
		(0.00729)	(0.00731)		(0.00460)	(0.00462)		(0.00615)	(0.00615)
UPSTREAM TEMPERATURE × HOT × DST			0.0578			0.000714			0.0571
			(0.0924)			(0.0530)			(0.0807)
Export Partner Effects:									
DOWNSTREAM TEMPERATURE	-0.00357	-0.00257	-0.00232	-0.00360	-0.000970	-0.000903	0.0000287	-0.00160	-0.00142
	(0.00346)	(0.00403)	(0.00405)	(0.00219)	(0.00255)	(0.00257)	(0.00284)	(0.00328)	(0.00329)
DOWNSTREAM TEMPERATURE × HOT	0.0444*	0.0477**	0.0377	0.0205	0.0264*	0.0386	0.0239	0.0213	-0.000929
	(0.0231)	(0.0233)	(0.0358)	(0.0155)	(0.0157)	(0.0260)	(0.0190)	(0.0189)	(0.0273)
DOWNSTREAM TEMPERATURE × DS		-0.00375	-0.00469		-0.00913**	-0.00957**		0.00538	0.00489
		(0.00691)	(0.00699)		(0.00446)	(0.00455)		(0.00556)	(0.00564)
DOWNSTREAM TEMPERATURE × HOT × DS			0.0141			-0.0216			0.0356
			(0.0504)			(0.0358)			(0.0398)
FIXED EFFECTS	FIRM, SECTOR × YEAR								
RAINFALL CONTROLS	YES								
Observations	205,856	205,856	205,856	205,856	205,856	205,856	205,856	205,856	205,856

NOTES: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. The unit of analysis is at the firm level. Local temperature is measured as the average daily temperature over the calendar year. Downstream temperature is defined as the average daily temperature over the calendar year measured in each export partner country, weighted by the export-share with that country in the baseline year. Upstream temperature is defined as the average daily temperature over the calendar year measured in each import partner country, weighted by the import-share with that country in the baseline year. HOT is defined as the share of imports/exports, in the baseline year, from countries with a climate that is in the highest tercile of the global distribution. DS is defined as the share of imports/exports from developing countries in the baseline year. This is aggregated over all trading partners for each firm. Standard errors are clustered at the commune level.

Table 4.5: The Heterogeneous Effects of Business Environment in Trading Partner Countries on French Manufacturing Firms

	(1) Value Added	(2) Value Added	(3) Value Added	(4) Employment	(5) Employment	(6) Employment	(7) Labour Productivity	(8) Labour Productivity	(9) Labour Productivity
Domestic Effects:									
Local Temperature (°C)	-0.0387***	-0.0383***	-0.0381***	-0.0406***	-0.0404***	-0.0404***	0.00185	0.00209	0.00227
	(0.00564)	(0.00566)	(0.00565)	(0.00421)	(0.00422)	(0.00423)	(0.00447)	(0.00448)	(0.00448)
Import Partner Effects:									
Upstream Temperature (°C)	0.0134***	0.0152***	0.0172***	0.00656**	0.00762**	0.00939***	0.00685	0.00762*	0.00781*
	(0.00497)	(0.00511)	(0.00551)	(0.00309)	(0.00317)	(0.00346)	(0.00422)	(0.00435)	(0.00468)
Upstream Temperature × DB Rank	-0.0677**	-0.0568**	-0.0753**	-0.0294*	-0.0271	-0.0451**	-0.0383	-0.0297	-0.0302
	(0.0277)	(0.0282)	(0.0340)	(0.0165)	(0.0169)	(0.0208)	(0.0238)	(0.0242)	(0.0290)
Upstream Temperature × DS		-0.0133*	-0.0244**		-0.00647	-0.0160**		-0.00679	-0.00840
		(0.00736)	(0.0118)		(0.00463)	(0.00758)		(0.00619)	(0.00989)
Upstream Temperature × DB Rank × DS			0.0826			0.0730*			0.00957
			(0.0639)			(0.0413)			(0.0531)
Export Partner Effects:									
Downstream Temperature	-0.00770	-0.00692	-0.00384	-0.00574*	-0.00353	-0.00357	-0.00196	-0.00338	-0.000272
	(0.00478)	(0.00503)	(0.00527)	(0.00302)	(0.00319)	(0.00346)	(0.00398)	(0.00418)	(0.00431)
Downstream Temperature × DB Rank	0.0361*	0.0371*	0.0183	0.0184	0.0213*	0.0252*	0.0178	0.0159	-0.00687
	(0.0188)	(0.0191)	(0.0214)	(0.0124)	(0.0125)	(0.0145)	(0.0154)	(0.0156)	(0.0178)
Downstream Temperature × DS		-0.00385	-0.0147		-0.00932**	-0.00874		0.00547	-0.00592
		(0.00701)	(0.0104)		(0.00447)	(0.00694)		(0.00571)	(0.00825)
Downstream Temperature × DB Rank × DS			0.0556			-0.0114			0.0670*
			(0.0430)			(0.0301)			(0.0352)
Fixed Effects	Firm, Sector × Year								
Rainfall Controls	Yes								
Observations	205,856	205,856	205,856	205,856	205,856	205,856	205,856	205,856	205,856

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. The unit of analysis is at the firm level. Local temperature is measured as the average daily temperature over the calendar year. Downstream temperature is defined as the average daily temperature over the calendar year measured in each export partner country, weighted by the export-share with that country in the baseline year. Upstream temperature is defined as the average daily temperature over the calendar year measured in each import partner country, weighted by the import-share with that country in the baseline year. **DB Rank** is defined as the share of imports/exports, in the baseline year, interacted with the Doing Business Report Rank for that country (normalised such that a value of 1 is equal to the worst ranked country). This is aggregated over all trading partners for each firm. Standard errors are clustered at the commune level.

Table 4.6: The Heterogeneous Effects of Adaptive Capacity in Trading Partner Countries on French Manufacturing Firms

	(1) Value Added	(2) Value Added	(3) Value Added	(4) Employment	(5) Employment	(6) Employment	(7) Labour Productivity	(8) Labour Productivity	(9) Labour Productivity
Domestic Effects:									
LOCAL TEMPERATURE (°C)	-0.0395***	-0.0390***	-0.0390***	-0.0413***	-0.0410***	-0.0411***	0.00174	0.00198	0.00207
	(0.00565)	(0.00566)	(0.00566)	(0.00422)	(0.00423)	(0.00423)	(0.00448)	(0.00449)	(0.00449)
Import Partner Effects:									
UPSTREAM TEMPERATURE (°C)	-0.00428	-0.00129	0.00196	-0.00480	-0.00311	-0.00380	0.000521	0.00181	0.00576
	(0.00601)	(0.00626)	(0.00658)	(0.00376)	(0.00386)	(0.00411)	(0.00506)	(0.00530)	(0.00552)
UPSTREAM TEMPERATURE × AC PER CAPITA	0.000563*	0.000628**	0.000443	0.000475**	0.000487***	0.000525**	0.0000886	0.000141	-0.0000828
	(0.000294)	(0.000298)	(0.000324)	(0.000186)	(0.000189)	(0.000209)	(0.000249)	(0.000252)	(0.000272)
UPSTREAM TEMPERATURE × DS		-0.0180**	-0.0467***		-0.00891*	-0.00118		-0.00913	-0.0456***
		(0.00731)	(0.0168)		(0.00459)	(0.00996)		(0.00617)	(0.0155)
UPSTREAM TEMPERATURE × AC PER CAPITA × DS			0.00173*			-0.000416			0.00215**
			(0.000925)			(0.000532)			(0.000856)
Export Partner Effects:									
DOWNSTREAM TEMPERATURE	-0.0114**	-0.0105**	-0.0147***	-0.0121***	-0.00933***	-0.0121***	0.000697	-0.00116	-0.00260
	(0.00477)	(0.00518)	(0.00548)	(0.00325)	(0.00344)	(0.00361)	(0.00389)	(0.00421)	(0.00444)
DOWNSTREAM TEMPERATURE × AC PER CAPITA	0.000668***	0.000680***	0.000961***	0.000673***	0.000688***	0.000884***	-0.00000547	-0.00000795	0.0000771
	(0.000250)	(0.000252)	(0.000281)	(0.000181)	(0.000183)	(0.000204)	(0.000202)	(0.000202)	(0.000222)
DOWNSTREAM TEMPERATURE × DS		-0.00358	0.0267*		-0.00964**	0.00989		0.00606	0.0168
		(0.00695)	(0.0146)		(0.00451)	(0.00910)		(0.00564)	(0.0121)
DOWNSTREAM TEMPERATURE × AC PER CAPITA × DS			-0.00206**			-0.00138***			-0.000675
			(0.000842)			(0.000520)			(0.000699)
FIXED EFFECTS	FIRM, SECTOR × YEAR								
RAINFALL CONTROLS	YES								
Observations	205,856	205,856	205,856	205,856	205,856	205,856	205,856	205,856	205,856

NOTES: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. The unit of analysis is at the firm level. Local temperature is measured as the average daily temperature over the calendar year. Downstream temperature is defined as the average daily temperature over the calendar year measured in each export partner country, weighted by the export-share with that country in the baseline year. Upstream temperature is defined as the average daily temperature over the calendar year measured in each import partner country, weighted by the import-share with that country in the baseline year. **AC per capita** is defined as the share of imports/exports, in the baseline year, interacted with the Trade Value of Air Conditioning Machines per capita for that country. This is aggregated over all trading partners for each firm. Standard errors are clustered at the commune level.

Bibliography

ACEMOGLU, D., V. CARVALHO, A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2012): "The Network Origins of Aggregate Fluctuations," *Econometrica*, 80, 1977–2016.

ADHVARYU, A., A. CHARI, AND S. SHARMA (2013): "Firing Costs and Flexibility: Evidence from Firms' Labor Adjustments to Shocks in India," *Review of Economics and Statistics*, 95, 725–740.

ADHVARYU, A., N. KALA, AND A. NYSHADHAM (2014): "Management and Shocks to Worker Productivity: Evidence from Air Pollution Exposure in an Indian Garment Factory," *Mimeo*.

——— (2015): "The Light and the Heat: Productivity Co-benefits of Energy-saving Technology," *Mimeo*.

AGHION, P., U. AKCIGIT, A. DEATON, AND A. ROULET (2015): "Creative Destruction and Subjective Wellbeing," *NBER Working Paper 21069*.

AHSAN, A. AND C. PAGÉS (2009): "Are All Labor Regulations Equal? Evidence from Indian Manufacturing," *Journal of Comparative Economics*, 37, 62–75.

ALCOTT, H., A. COLLARD WEXLER, AND S. O'CONNEL (2015): "How Do Electricity Shortages Affect Industry? Evidence from India," *American Economic Review*.

ALEM, Y. AND J. COLMER (2015): "Consumption Smoothing and the Welfare Cost of Uncertainty," *Mimeo*.

ALLEN, D. AND D. ATKIN (2015): "Volatility, Insurance, and the Gains from Trade," *Mimeo*.

ALPERT, M. AND H. RAIFFA (1982): "A Progress Report on the Training of Probability Assessors," in *Judgement Under Uncertainty: heuristics and biases*, ed. by D. Kahneman, P. Slovic, and A. Tversky, Cambridge University Press, 294–305.

AMIRAPU, A. AND M. GECHTER (2014): "Indian Labour Regulations and the Cost of Corruption: Evidence from the Firm Size Distribution," *Mimeo*.

ANGRIST, J. AND J. PISCHKE (2009): *Mostly Harmless Econometrics: an empiricist's companion*, Princeton University Press, Princeton, N.J., USA.

ANTILLA-HUGHES, J. AND S. HSIANG (2013): "Destruction, Disinvestment, and Death: Economic and Human Losses following Environmental Disaster," *Mimeo*.

ARGYLE, M. (1999): "Causes and Correlates of Happiness." in *Well-Being: The Foundations of Hedonic Psychology: scientific perspectives on employment and suffering*, ed. by D. Kahneman, E. Diener, and N. Schwarz, New York: Russel Sage Foundation.

ASC (2010): "Adaptation Sub-Committee: How Well Prepared is the UK for Climate Change?" Tech. rep., Committee on Climate Change.

ASHER, S. AND P. NOVOSAD (2014): "The Impact of Agricultural Output on Local Economic Activity in India," *Mimeo*.

ATALAY, E. (2014): "How Important Are Sectoral Shocks?" *Mimeo*.

AUFFHAMMER, M., S. HSIANG, W. SCHLENKER, AND A. SOBEL (2013a): "Using Weather Data and Climate Model Output in Economic Analyses of Climate Change," *Review of Environmental Economics and Policy*.

——— (2013b): "Using Weather Data and Climate Model Output in Economic Analyses of Climate Change," *Review of Environmental Economics nad Policy*.

AUTOR, D., D. DORN, AND G. HANSON (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*.

BAKER, S. AND N. BLOOM (2013): "Does Uncertainty Reduce Growth? using disasters as natural experiments," *NBER Working Paper 19475*.

BALDWIN, R. (2013): "Trade and Industrialization after Globalization's Second Unbundling: How Building and Joining a Supply Chain Are Different and Why It Matters," in *Globalization in an Age of Crisis: Multilateral Economic Cooperation in the Twenty-First Century*, ed. by R. Feenstra and A. Taylor, University of Chicago Press.

BALDWIN, R. AND T. VENABLES (2013): "Spiders and Snakes: Offshoring and Agglomeration in the Global Economy," *Journal of International Economics*, 90, 245–254.

BANDIERA, O., R. BURGESS, N. DAS, S. GULESCI, I. RASUL, AND M. SULAIMAN (2015a): "The Misallocation of Labor in Village Economies," *Mimeo*.

BANDIERA, O., A. PRAT, AND R. SADUN (2015b): "Managing the Family Firm: Evidence from CEOs at Work," *Mimeo*.

BANERJEE, A. AND E. DUFLO (2007): "The Economic Lives of the Poor," *Journal of Economic Perspectives*, 21, 141–168.

BANERJEE, A. AND S. MULLAINATHAN (2010): "The Shape of Temptation: Implications for the Economic Lives of the Poor," *NBER Working Paper No. 15973*.

BARRIOPEDRO, D., E. FISCHER, J. LUTERBACHER, R. TRIGO, AND R. GARCIA-HERRERA (2011): "The Hot Summer of 2010: Redrawing the Temperature Record Map of Europ," *Science*, 332.

BAYLIS, P. (2016): "Temperature and Temperament: Evidence from a Billion Tweets," *Mimeo*.

Beegle, K., J. De Weerdt, and S. Dercon (2011): "Migration and Economic Mobility in Tanzania: Evidence from a Tracking Survey," *Review of Economics and Statistics*, 93, 1010–1033.

Beegle, K., R. H. Dehejia, and R. Gatti (2006): "Child labour and agricultural shocks," *Journal of Development Economics*, 81, 80–96.

Benjamin, D., O. Heffetz, M. Kimball, and A. Rees-Jones (2012): "What Do You Think Would Make You Happier? What Do You Think You Would Choose?" *American Economic Review*, 102, 2083–2110.

Benjamin, D., M. Ori Heffetz, M. Kimball, and N. Szembrot (2013): "Aggregating Local Preferences to Guide Marginal Policy Adjustments," *American Economic Review: Paper and Proceedings*.

Bentham, J. (1789): *Principles of Morals and Legislation*, Oxford: Clarendon Press.

Bentolila, S. and G. Bertola (1990): "Firing Costs and Labour Demand: How Bad is Eurosclerosis," *The Review of Economic Studies*, 57, 381–402.

Bertrand, M., C. Hsieh, and N. Tsivanidis (2015): "Contract Labor and Firm Growth in India," *Mimeo*.

Besley, T. and R. Burgess (2004): "Can Labor Regulation Hinder Economic Performance? Evidence from India," *Quarterly Journal of Economics*, 19, 91–134.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685.

——— (2014): "Fluctuations in Uncertainty," *Journal of Economic Perspectives*, 28, 153–176.

Brodsky, H. and A. S. Hakkert (1988): "Risk of a Road Accident in Rainy Weather," *Accident Analysis & Prevention*, 20, 161–176.

Brunnermeier, M., F. Papakonstantinou, and J. Parker (2013): "Optimal Time-Inconsistent Beliefs: misplanning, procratistination, and commitment," *Mimeo*.

Brunnermeier, M. and J. Parker (2005): "Optimal Expectations," *American Economic Review*, 95, 1092–1118.

Bryan, G., S. Chowdhury, and A. Mobarak (2014): "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh," *Econometrica*, 82, 1671–1758.

Bryan, G. and M. Morten (2015): "Economic Development and the Spatial Allocation of Labor: Evidence from Indonesia," *Mimeo*.

Buehler, R., D. Griffin, and M. Ross (1994): "Exploring the Planning Fallacy: why people underestimate their task completion times," *Journal of Personality and Social Psychology*, 67, 366–381.

Burgess, R., O. Deschenes, D. Donaldson, and M. Greenstone (2014a): "The Unequal Effects of Weather and Climate Change: Evidence from Mortality in India," *Mimeo*.

Burgess, R., O. Deschênes, D. Donaldson, and M. Greenstone (2014b): "The Unequal Effects of Weather and Climate Change: evidence from mortality in India," *Mimeo*.

Burgess, R. and D. Donaldson (2010): "Can Openness Mitigate the Effects of Weather Shocks? Evidence from India's Famine Era," *American Economic Review: Paper and Proceedings*, 100, 449–453.

——— (2012): "Railroads and the Demise of Famine in Colonial India," *Mimeo*.

Burgess, R., M. Hansen, B. Olken, P. Potapov, and S. Sieber (2011): "The Political Economy of Deforestation in the Tropics," *NBER Working Paper No. 17417*.

Burke, M., S. Hsiang, and E. Miguel (2015): "Global Non-linear Effect of Temperature on Economic Production," *Nature*.

Bustos, P., B. Caprettini, and J. Ponticelli (2015): "Agricultural Productivity and Structural Transformation: Evidence from Brazil," *Mimeo*.

Cachon, G., S. Gallino, and M. Olivares (2012a): "Severe Weather and Automobile Assembly Productivity," *Mimeo*.

Cachon, G. P., S. Gallino, and M. Olivares (2012b): "Severe Weather and Automobile Assembly Productivity," *Mimeo*.

Cain, M. (1982): "Perspectives on Family and Fertility in Developing Countries," *Population Studies*, 36, 159–175.

Caplin, A. and J. Leahy (2001): "Psychological Expected Utility Theory and Anticipatory Feelings," *Quarterly Journal of Economics*, 116, 55–79.

Carriere-Swallow, Y. and L. Céspedes (2013): "The Impact of Uncertainty Shocks in Emerging Economies," *Journal of International Economics*, 90, 316–325.

Carroll, C. (1997): "Buffer-stock saving and the life cycle/permanent income hypothesis," *Quarterly Journal of Economics*, 107, 1–55.

Carroll, C. and M. Kimball (2001): "Liquidity Constraints and Precautionary Saving," *Working Paper 2001 Johns Hopkins University*.

Carroll, N., P. Frijters, and M. Shields (2009): "Quantifying the Costs of Drought: New Evidence from Life Satisfaction Data," *Journal of Population Economics*, 22, 445–461.

Carvalho, V. (2010): "Aggregate Fluctuations and the Network Structure of Intersectoral Trade," *Mimeo*.

——— (2014): "From Micro to Macro via Production Networks," *Journal of Economic Perspectives*, 28, 23–48.

Carvalho, V. and X. Gabaix (2013): "The Great Diversification and its Undoing," *American Economic Review*, 103, 1697–1727.

Central Statistics Agency, T. F. D. R. o. E. (2009): "Country and Regional Level Consumer Price Indices," *Data*.

Chaurey, R. (2015): "Labor Regulations and Contract Labor Use: Evidence from Indian Firms," *Journal of Development Economics*, 114, 224–232.

Cochrane, J. (1994): "Shocks," *Carnegie-Rochester Conference Series on Public Policy*.

Colmer, J. (2015a): "Uncertainty, Child Labour, and Human Capital Accumulation," *Mimeo*.

——— (2015b): "Urbanization, Growth, and Development: Evidence from India," *Mimeo*.

——— (2016): "Weather, Labour Reallocation, and Industrial Production: Evidence from India," *Mimeo*.

Conley, T. (1999): "GMM Estimation with Cross Sectional Dependence," *Journal of Econometrics*, 92, 1–45.

Costinot, A., D. Donaldson, and C. Smith (2015): "Evolving Comparative Advantage and the Impact of Climate Change in Agricultural Markets: Evidence from a 1.7 Million Fields around the World." *Journal of Political Economy*, forthcoming.

——— (2016): "Evolving Comparative Advantage and the Impact of Climate Change in Agricultural Markets: Evidence from 1.7 Million Fields around the World," *Journal of Political Economy*.

De Neve, J., G. Ward, F. De Keulenaer, B. Van Landeghem, G. Kavetsos, and N. M.I. (2015): "The Asymmetric Experience of Positive and Negative Economic Growth: global evidence using subjective well-being data," *CEP Discussion Paper No. 1304*.

Deaton, A. (1992): "Saving and Income Smoothing in Cote d'Ivoire," *Journal of African Economies*, 1, 1–24.

Dee, D. P., S. M. Uppala, A. J. Simmons, P. Berrisford, P. Poli, S. Kobayashi, U. Andrae, M. A. Balmaseda, G. Balsamo, P. Bauer, P. Bechtold, A. C. Beljaars, L. M. van de Berg, J. Bidlot, N. Bormann, C. Delsol, R. Dragani, M. Fuentes, A. J. Geer, L. Haimberger, S. B. Healy, H. Hersbach, E. V. Holm, L. Isaksen, P. Kallberg, M. Kohler, M. Matricardi, A. P. McNally, B. Monge-Sanz, J. Morcrette, B. Park, C. Peubey, P. de Rosnay, C. Tavolato, J. Thepaut, and F. Vitart (2011): "The ERA-Interim reanalysis: configuration and performance of the data assimilation system." *Quarterly Journal of the Royal Meteorological Society*, 137, 553–597.

Delgado, M. and A. Porcellie (2009): "Acute Stress Modulates Risk Taking in Financial Decision Making," *Psychological Science*, 20, 278–283.

Dell, M., B. Jones, and B. Olken (2012): "Temperature Shocks and Economic Growth: Evidence from the Last Half Century," *American Economic Journal: Macroeconomics*, 4, 66–95.

——— (2014): "What Do We Learn from the Weather? The New Climate-Economy Literature," *Journal of Economic Literature*, 52, 740–798.

Dercon, S. (2002): "Income Risk, Coping Strategies, and Safety Nets," *The World Bank Research Observer*, 17, 141–166.

Dercon, S. and J. Hoddinott (2009): "The Ethiopian Rural Household Surveys 1989 - 2004: introduction," *IFPRI*.

Dercon, S. and P. Krishnan (1996): "Income Portfolios in Rural Ethiopia and Tanzania: Choices and Constraints," *Journal of Development Studies*, 32, 850–875.

Desmet, K., E. Ghani, S. O'Connel, and E. Rossi-Hansberg (2015): "The Spatial Development of India," *Journal of Regional Science*.

DfT (2010): "The Resilience of England's Transport Systems in Winter," Tech. rep., Department for Transport – an Independent Review.

Diener, E. and R. Lucas (1999): "Personality and Subjective Well-being," in *Well-Being: The Foundations of Hedonic Psychology: scientific perspectives on employment and suffering.*, ed. by D. Kahneman, E. Diener, and N. Schwarz.

Diener, E., R. Lucas, U. Schimmack, and J. Helliwell (2009): *Well-Being for Public Policy*, New York: Oxford University Press.

Dixit, A. and R. Pindyck (1994): *Investment under Uncertainty*, Princeton University Press.

Doherty, T. and S. Clayton (2011): "The Psychological Impacts of Global Climate Change," *American Psychologist*, 66, 265–276.

Dolan, P. and D. Kahneman (2008): "Interpretations of Utility and Their Implications for the Valuation of Health," *Economic Journal*, 118, 215–234.

Dolan, P., R. Layard, and R. Metcalfe (2011): "Measuring Subjective Well-Being for Public Policy: recommendations on measures," *CEP Special Paper No. 23*.

Donaldson, D. (2015a): "The Gains from Market Integration," *Annual Review of Economics*, forthcoming.

——— (2015b): "Railroads of the Raj: Estimating the Impact of Transportation Infrastructure," *American Economic Review*, forthcoming.

Duflo, E. (2000): "Child Health and Household Resources in South Africa: Evidence from the Old Age Pension Program," *American Economic Review*, 90, 393–398.

——— (2005): "Gender equality in development." *Mimeo*.

Durevall, D., J. Loening, and Y. Ayalew Birru (2013): "Inflation Dynamics and Food Prices in Ethiopia," *Journal of Development Economics*, 104, 89–106.

Eaton, J., S. Kortum, and F. Kramarz (2011): "An Anatomy of International Trade: Evidence from French Firms," *Econometrica*, 79, 1453–1498.

Fafchamps, M. (2003): *Rural Poverty, Risk and Development*, Edward Elgar Publishing.

Fafchamps, M. and S. Lund (2003): "Risk-sharing Networks in Rural Phillipines," *Journal of Development Economics*, 71, 261–287.

Fafchamps, M. and J. Pender (1997): "Precautionary Saving, Credit Constraints and Irreversible Investments: Theory and Evidence from Semi-arid India," *Journal of Business and Economic Statistics*, 15, 180–194.

Fafchamps, M. and A. Quisumbing (2007): "Household Formation and Marriage Markets in Rural Areas," in *Handbook of Development Economics*, ed. by T. Schultz and J. Strauss, Elsevier, vol. 4, 3187–3247.

Feddersen, J., R. Metcalfe, and M. Wooden (2015): "Subjective Wellbeing: why weather matters," *Journal of Royal Statistical Society*.

Ferrer-i Carbonell, A. and P. Frijters (2004): "How Important is Methodology for the Estimates of the Determinants of Happiness?" *Economic Journal*, 114, 641–659.

Fisher, A., M. Hanemann, M. Roberts, and W. Schlenker (2012): "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Comment," *American Economic Review*, 102, 3749–60.

Fishman, R. (2012): "Climate Change, Rainfall Variability and Adaptation through Irrigation: Evidence from Indian Agriculture," *Mimeo*.

Fitzsimons, E. (2007): "The Effects of Risk on Education in Indonesia," *Economic Development and Cultural Change*, 56, 1–25.

Foster, A. and M. Rosenzweig (2004): "Agricultural Productivity Growth, Rural Economic Diversity, and Economic Reforms: India 1970 - 2000," *Economic Development and Cultural Change*, 52, 509–542.

——— (2008): "Economic Development and the Decline of Agricultural Economic Development and the Decline of Agricultural Employment," *Handbook of Development Economics*, 4.

Franklin, S. (2015): "Location, Search Costs, and Youth Unemployment: A Randomized Trial of Transport Subsidies in Ethiopia," *Mimeo*.

Frey, B., S. Luechinger, and A. Stutzer (2010): "The Life Satisfaction Approach to Environmental Valuation," *Annual Review of Resource Economics*, 2, 139–160.

Frey, B. and A. Stutzer (2002): *Happiness and Economics*, Princeton University Press, Princeton, N.J., USA.

Gabaix, X. (2011): "The Granular Origins of Aggregate Fluctuations," *Econometrica*, 79, 733–772.

Geanakoplos, J., D. Pearce, and E. Stacchetti (1989): "Psychological Games and Sequential Rationality," *Games and Economic Behaviour*, 1, 60–79.

Ghani, E., A. Goswami, and W. Kerr (2012): "Is India's Manufacturing Sector Moving Away from Cities?" *NBER Working paper 17992*.

Gollin, D., D. Lagakos, and M. Waugh (2014): "The Agricultural Productivity Gap," *Quarterly Journal of Economics*, 129, 939–993.

Golob, T. F. and W. W. Recker (2003): "Relationships Among Urban Freeway Accidents, Traffic Flow, Weather, and Lighting Conditions." *Journal of Transportation Engineering*, 129, 342–353.

Graff Zivin, J., S. Hsiang, and M. Neidell (2015): "Temperature and Human Capital in the Short- and Long-Run," *NBER Working Paper 21157*.

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

Gray, C. and V. Mueller (2012): "Natural disasters and population mobility in Bangladesh," *Proceedings of the National Academy of Sciences*.

Grootaert, C. and R. Kanbur (1995): "Child Labor: An Economic Perspective," *International Labor Review*, 134, 187–203.

Gruber, J. and S. Mullainathan (2005): "Do Cigarette Taxes Make Smokers Happier?" *The B.E. Journal of Economic Analysis and Policy*, 5.

Guiteras, R. (2009): "The Impact of Climate Change on Indian Agriculture," *Mimeo*.

Haltiwanger, J., S. Scarpetta, and H. Schweiger (2008): "Assessing Job Flows Across Countries: The Role of Industry, Firm Size, and Regulations," *NBER Working Paper No. 13920*.

Hamermesh, D. (1993): *Labour Demand*, Princeton University Press.

Hardy, M. and J. McCasland (2015): "Are Small Firms Labor Constrained? Experimental Evidence from Ghana," *Mimeo*.

Hare, T., C. Camerer, and A. Rangel (2009): "Self-Control in Decision-Making Involves Modulation of the vmPFC Valuation System," *Science*, 324, 646–648.

Hausman, J. and Z. Griliches (1984): "Econometric Models for Count Data with an Application to the Patents *R&D* Relationship," *Econometrica*, 52, 909–938.

Heal, G. and J. Park (2014): "Feeling the Heat: Temperature, Physiology, and the Wealth of Nations," *Mimeo*.

——— (2015): "Goldilocks Economies? Temperature Stress and the Direct Impacts of Climate Change," *NBER Working Paper No. 21119*.

Heath, R. (2015): "Why Do Firms Hire Using Referrals? Evidence from Bangladeshi Garment Factories," *Mimeo*.

Heckman, J. and B. Honore (1990): "The Empirical Content of the Roy Model," *Econometrica*, 58, 1121–1149.

Heckman, J. and G. Sedlacek (1985): "Heterogeneity, Aggregation, and Market Wage Functions: An Empirical Model of Self-Selection in the Labor Market," *Journal of Political Economy*, 93, 1077–1125.

Heckman, J. a. (2003): "Law and Employment: Lessons from Latin American and the Caribbean," *NBER Working Paper No. 10129*.

Henderson, J., A. Storeygard, and U. Deichman (2015): "Has Climate Change Driven Urbanization in Africa?" *Mimeo*.

Henderson, J., A. Storeygard, and D. Weil (2012): "Measuring Economic Growth from Outer Space," *American Economic Review*, 102, 994–1028.

Hornbeck, R. (2012): "The Enduring Impact of the American Dust Bowl: Short- and Long-run Adjustments to Environmental Catastrophe," *American Economic Review*.

Hornbeck, R. and P. Keskin (2015): "Does Agriculture Generate Local Economic Spillovers? Short-Run and Long-Run Evidence from the Ogallala Aquifer," *American Economic Journal: Economic Policy*, 7, 192–213.

Hornbeck, R. and E. Moretti (2015): "Who Benefits from Productivity Growth? The Local and Aggregate Impacts of Local TFP Shocks on Wages, Rents, and Inequality," *Mimeo*.

Hornbeck, R. and S. Naidu (2014): "When the Levee Breaks: Black Migration and Economic Development in the American South," *American Economic Review*, 104, 963–990.

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

Hsieh, C., E. Hurst, C. Jones, and P. Klenow (2014): "The Allocation of Talent and US Economic Growth," *Mimeo*.

Hsieh, C. and P. Klenow (2009): "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124, 1403–1448.

Imbert, C. and J. Papp (2015): "Short-term Migration, Rural Workfare Programs and Urban Labor Markets: Evidence from India," *Mimeo*.

Intergovernmental Panel on Climate Change (2012): "Managing the Risks of Extreme Events and Disasters to Advance Climate Change Adaptation. A Special Report of Working Groups I and II of the IPCC," Tech. rep.

——— (2014): "Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change," Tech. rep.

IPCC (2007): "Climate Change 2007: The Physical Science Basis. Contribution of Working Group 1 to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change," Tech. rep.

Jackson, M. (2008): *Social and Economic Networks*, Princeton University Press.

Jacoby, H. and E. Skoufias (1997): "Risk, financial markets, and human capital in a developing country," *Review of Economic Studies*, 64, 311–335.

Jayachandran, S. (2006): "Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries," *Journal of Political Economy*, 114, 538–575.

——— (2013): "Incentives to Teach Badly: After-School Tutoring in Developing Countries," *Mimeo*.

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

Jevons, W. (1905): *The Principles of Economics*, Macmillan and Co.

Jina, A. and S. Hsiang (2015): "Geography, Depreciation, and Growth," *American Economic Review: Paper and Proceedings*.

Jovanovic, B. (1987): "Micro Shocks and Aggregate Risk," *Quarterly Journal of Economics*, 102, 395–409.

Kahneman, D. and A. Krueger (2006): "Developments in the Measurement of Subjective Well-Being," *Journal of Economic Perspectives*, 20, 3–24.

Kahneman, D., P. Wakker, and R. Sarin (1997): "Back to Bentham? exploration of experienced utility," *The Quarterly Journal of Economics*, 112, 375–406.

Kala, N. (2015): "Ambiguity Aversion and Learning in a Changing World: The Potential Effects of Climate Change From Indian Agriculture," *Mimeo*.

Kaur, S. (2014): "Nominal Wage Rigidity in Village Labor Markets," *Mimeo*.

Kenrick, D. and S. McFarlane (1986): "Ambient Temperature and Horn Honking: A Field Study of the Heat/Aggression Relationship," *Environment and Behaviour*, 18, 179–191.

Keskin, P. (2010): "Thirsty Factories, Hungry Farmers: Intersectoral Impacts of Industrial Water Demand," *Mimeo*.

Kimball, M. (1991): "Precautionary Savings in the Small and the Large." *Econometrica*, 58, 53–73.

Kinnan, C. (2014): "Distinguishing Barries to Insurance in Thai Villages," *Mimeo*.

Knight, F. (1921): *Risk, Uncertainty, and Profit*, Hart, Schaffner, and Marx; Houghton Mifflin Company.

Koren, M. and S. Tenreyro (2007): "Volatility and Development," *Quarterly Journal of Economics*, 122, 243–287.

Krueger, D. (2007): "Coffee production effects on child labour and schooling in rural Brazil." *Journal of Development Economics*, 82, 448–462.

Kudamatsu, M., T. Persson, and D. Stromberg (2014a): "Weather and Infant Mortality in Africa," *Mimeo*.

Kudamatsu, M., T. Persson, and D. Strömberg (2014b): "Weather and Infant Mortality in Africa," *Mimeo*.

Lagakos, D. and M. Waugh (2013): "Selection, Agriculture, and Cross-Country Productivity Differences," *American Economic Review*, 103, 948–980.

Layard, R. (2005): *Happiness: lessons from a new science*, London: Allen Lane.

Lee, J. J., F. Gino, and B. Staats (2012): "Rainmakers: Why Bad Weather Means Good Productivity," Tech. rep., lee2012.

Levinson, A. (2012): "Valuing Public Goods Using Happiness Data: the case of air quality," *Journal of Public Economics*, 96, 869–880.

——— (2013): "Happiness as a Public Policy Tool," *Mimeo*.

Lewis, W. (1954): "Economic Development with Unlimited Supplies of Labor," *Manchester School of Economic and Social Studies*, 22, 139–191.

Lilleør, H. (2015): "Human Capital Diversification among Children," *Mimeo*.

Lorenz, C. and H. Kuntsmann (2012a): "The Hydrological Cycle in Three State-Of-The-Art Reanalyses: intercomparison and performance analysis," *Journal of Hyrdrometeorology*, 13, 1397–1420.

——— (2012b): "The hydrological cycle in three state-of-the-art reanalyses: intercomparison and performance analysis," *Journal of Hyrdrometeorology*, 13, 1397–1420.

Lowenstein, G. (1987): "Anticipation and the Valuation of Delayed Consumption," *Economic Journal*, 97, 666–684.

Lucas, R. (1977): "Understanding Business Cycles," *Carnegie-Rochester Conference Series on Public Policy*.

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

Mackworth, N. (1946): "Effects of Heat on Wireless Telegraphy Operators Hearing and Recording Morse Messages," *British Journal of Industrial Medicine*, 3, 143–158.

——— (1947): "High Incentives versus Hot and Humid Atmospheres in a Physical Effort Task," *British Journal of Psychology: General Section*, 38, 90–102.

Mankiw, G., D. Romer, and D. Weil (1992): "A Contribution to the Empirics of Economic Growth," *Quarterly Journal of Economics*, 107, 407–437.

Marden, S. (2015): "The Agricultural Roots of Industrial Development: Forward Linkages in Reform Era China," *Mimeo*.

Marshall, A. (1891): "Principles of Economics," *Quarterly Journal of Economics*, 5, 319–338.

Matsuyama, K. (1992): "Agricultural Productivity, Comparative Advantage, and Economic Growth," *Journal of Economic Theory*, 58, 317–334.

Mayer, T., M. Melitz, and G. Ottaviano (2014): "Markte Size, Competition, and the Product Mix of Exporters," *American Economic Review*, 104, 495–536.

Metcalfe, R., J. Feddersen, and M. Wooden (2012): "Subjective Well-Being: Weather Matters; Climate Doesn't," Tech. Rep. 627, University of Oxford, Department of Economics.

Mian, A. and A. Sufi (2015): "What Explains the 2007-2009 Drop in Employment?" *Econometrica*, forthcoming.

Miguel, E. and J. Hamory (2009): "Individual Ability and Selection into Migration in Kenya," *UNDP, Human Development Research Paper 2009/45*.

Mobarak, A. and M. Rosenzweig (2014): "Risk, Insurance and Wages in General Equilibrium," *Mimeo*.

Morduch, J. (1995): "Income Smoothing and Consumption Smoothing," *Journal of Economic Perspectives*, 9, 103–114.

Moretti, E. (2011): "Local Labor Markets," in *Handbook of Labor Economics*, ed. by D. Card and O. Ashenfelter, Elsevier.

Morten, M. (2013): "Temporary Migration and Endogenous Risk Sharing in Village India," *Mimeo*.

Mueller, V., C. Gray, and K. Kosec (2014): "Heat stress increases long-term human migration in rural Pakistan," *Nature Climate Change*.

Munshi, K. and M. Rosenzweig (2015): "Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap," *Mimeo*.

Newey, W. and K. West (1987): "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55, 703–708.

Nickell, S. (1978): "Fixed Costs, Employment, and Labour Demand over the Cycle," *Economica*, 45, 329–345.

Oi, W. (1962): "Labour as a Quasi-Fixed Factor," *The Journal of Political Economy*.

Paxson, C. (1992): "Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand," *American Economic Review*, 82, 15–33.

Rabin, M. and J. Schrag (1999): "First Impressions Matter: a model of confirmatory bias," *Quarterly Journal of Economics*, 114, 37–82.

Ramey, G. and V. Ramey (1995): "Cross-Country Evidence on the Link between Volatility and Growth," *American Economic Review*, 85, 1138–1151.

Ravallion, M. and Q. Wodon (2000): "Does child labour displace schooling? Evidence on beahvioural responses to an enrollment subsidy." *The Economic Journal*, 110, 158–175.

Rehdanz, K. and D. Maddison (2011): "The Impact of Climate on Life Satisfaction," *Ecological Economics*, 70, 2437–2445.

Restuccia, D. and R. Rogerson (2008): "Policy Distortions and Aggregate Productivity with Heterogenous Plants," *Review of Economy Dynamics*, 11, 707–720.

Rosenzweig, M. (1995): "Why Are There Returns to Schooling?" *The American Economic Review*, 85, 153–158.

Rosenzweig, M. and H. Binswanger (1993): "Wealth, Weather Risk, and the Profitability of Agricultural Investment," *Economic Journal*, 103, 56–78.

Rosenzweig, M. and C. Udry (2013): "Forecasting Profitability," *NBER Working Paper 19811*.

——— (2014): "Rainfall Forecasts, Weather and Wages over the Agricultural Production Cycle," *American Economic Review Papers and Proceedings*, 104, 278–283.

Rosenzweig, M. and K. Wolpin (1985): "Specific Experience, Household Structure, and Intergenerational Transfers: Farm Family Land and Labor Arrangements in Developing Countries," *Quarterly Journal of Economics*, 100, 961–987.

Roy, A. (1951): "Some Thoughts on the Distribution of Earnings," *Oxford Economic Papers*, 3, 135–146.

Ryan, N. (2014): "The Competitive Effects of Transmission Infrastructure in the Indian Electricity Market," *Mimeo*.

Santangelo, G. (2015): "Firms and Farms: The Impact of Agricultural Productivity on the Local Indian Economy," *Mimeo*.

Santos Silva, J. and S. Tenreyro (2006): "The Log of Gravity," *The Review of Economics and Statistics*, 88, 641–658.

Schady, N. (2004): "Do Macroeconomic Crises Always Slow Human Capital Accumulation?" *World Bank Economic Review*, 18, 131–154.

Schlenker, W. and D. Lobell (2010): "Robust negative impacts of climate change on African agriculture," *Environmental Research Letters*, 5, 1–8.

Schlenker, W. and M. Roberts (2009): "Nonlinear temperature effects indicate severe damages to U.S. crop yields under climate change," *Proceedings of the National Academy of Sciences*, 106, 15594–15598.

Sen, A. (1990): "More than 100 million women are missing." *The New York Review of Books*, 20, 61–66.

Shah, M. and B. Steinberg (2015): "Drought of Opportunities: Contemporaneous and Long Term Impacts of Rainfall Shocks on Human Capital," *Mimeo*.

Simonsohn, U. (2010): "Weather to go to College," *Economic Journal*, 120, 270–280.

Soderbom, M. and B. Rijkers (2013): "The Effects of Risk and Shocks on Non-Farm Enterprise Development in Rural Ethiopia," *World Development*, 45, 119–136.

Somonathan, E., R. Somonathan, A. Sudarshan, and M. Tewari (2015): "The Impact of Temperature on Productivity and Labor Supply: Evidence from Indian Manufacturing," *Mimeo*.

Strauss, J. and D. Thomas (1998): "Health, Nutrition, and Economic Development," *Journal of Economic Literature*, 36, 766–817.

Subak, S., J. Palutikof, M. Agnew, S. Watson, C. Bentham, M. Cannell, M. Hulme, S. McNally, J. Thornes, D. Waughray, and J. Woods (2000): "The Impact of the Anomalous Weather of 1995 on the U.K. Economy," *Climatic Change*, 44, 1–26.

Suri, T. (2011): "Estimating the Extent of Local Risk Sharing Between Households," *Mimeo*.

Taraz, V. (2012): "Adaptation to Climate Change: Historical Evidence from the Indian Monsoon," *Mimeo*.

Thomas, D., K. Beegle, E. Frankenberg, S. Bondan, J. Strauss, and G. Terul (2004): "Education during a Crisis," *Journal of Development Economics*, 74, 53–86.

Townsend, R. (1994): "Risk and Insurance in Village India," *Econometrica*, 62, 539–592.

Udry, C. (1994): "Risk and Insurance in a Rural Credit Market: an empirical investigation in Northern Nigeria," *The Review of Economic Studies*, 61, 495–526.

van den Bos, R., M. Harteveld, and H. Stoop (2009): "Stress and Decision-Making in Humans: performance is related to cortisol reactivity, albeit differently in men and women." *Psychoneuroendocrinology*, 34, 1449–1458.

Vanden Eynde, O. (2011): "Targets of Violence: Evidence from India's Naxalite Conflict," *Mimeo*.

Vollrath, D. (2009): "How Important are Dual Economy Effects for Aggregate Productivity," *Journal of Development Economics*, 88, 325–334.

——— (2014): "The Efficiency of Human Capital Allocations in Developing Countries," *Journal of Development Economics*, 108, 106–118.

Weinstein, N. (1980): "Unrealistic Optimism About Future Life Events," *Journal of Personality and Social Psychology*, 39, 806–820.

Welsch, H. (2002): "Preferences over Prosperity and Pollution: environmental valuation based on happiness surveys," *Kyklos*, 55, 473–494.

——— (2006): "Environment and happiness: evaluation of air pollution using life satisfaction data," *Ecological Economics*, 58, 801–813.

Wooldridge, J. (1999): "Distribution-free Estimation of some Nonlinear Panel Data Models," *Journal of Econometrics*, 90, 77–97.

——— (2010): *Econometric Analysis of Cross Section and Panel Data*, MIT Press: Cambridge (MA).

Yariv, L. (2001): "Believe and Let Believe: axiomatic foundations for belief dependent utility functionals," *Yale University, Cowles Foundation Discussion Papers No. 1344*.

Young, A. (2013): "Inequality, the Urban-Rural Gap and Migration," *Quarterly Journal of Economics*, 128, 1727–1785.

——— (2015): "Structural Transformation, the Mismeasurement of Productivity Growth, and the Cost Disease of Services," *American Economic Review*, forthcoming.

Appendices

Appendix A

Appendices to Chapter 1: Weather, Labour Reallocation, and Industrial Production: Evidence from India

A.1 Theory Appendix

This appendix presents a simple model based on Matsuyama (1992) demonstrating how the direction of labour reallocation in response to a sector-specific productivity shock depends on market integration. Any analysis of labour reallocation across sectors within an economy necessitates a diversified economy and so for simplicity I consider two sectors: agriculture (a) and manufacturing (m).

Preferences

Consider a country composed of a large number of regions i. Each location i is populated by a continuum of workers L_i, which are assumed to be mobile between sectors, immobile between regions, supplied inelastically, and fully employed. Workers earn income $w_{ij}L_{ij}$ and preferences are defined over two types of goods agriculture and manufactured goods. Agricultural consumption is subject to subsistence constraints with a Stone-Geary utility function (Matsuyama, 1992; Caselli and Coleman, 2001; Jayachandran, 2006; Desmet and Parente, 2012).[1] Given prices in sector j, p_{ij} and total income w_iL_i, each worker maximises

$$U_i = (C_{ia} - \bar{a})^{\alpha} {C_{im}}^{1-\alpha} \tag{A.1}$$

which they maximise subject to their budget constraint,

$$p_{ia}C_{ia} + p_{im}C_{im} \leq L_i w_i \tag{A.2}$$

Worker demand for goods in agriculture, $D_{ia} = p_{ia}\bar{a} + \alpha(L_i w_i - p_{ia}\bar{a})$. For manufactured goods $D_{im} = (1-\alpha)(L_i w_i - p_{ia}\bar{a})$. As such, preferences are non-homothetic. Higher food subsistence requirements, higher prices, and lower incomes are associated with an increase in the demand for agricultural goods ($D_{ia}/L_i w_i$).

Production

There are 2 goods that can be produced in each location i, agricultural good a and manufactured goods m.[2] I assume that all regions have access to the same technology and so production functions do not differ across regions within each industry.

[1] Non-homothetic preferences can also be incorporated through a CES utility function where the elasticity of substitution between agricultural goods and other goods is less than one (Ngai and Pissarides, 2007; Desmet and Rossi-Hansberg, 2014).

[2] I will refer to goods and sectors interchangeably.

Different industries may have different production functions. Consequently, I drop the locational subscript unless necessary.

Output of each good j is produced according to the following production function,

$$Y_j = A_j F_j(L_j) \tag{A.3}$$

where A_j is sector-specific productivity and L_j is the set of workers in sector j. I assume that $F_j(0) = 0$, $F_j' > 0$ and $F_j'' < 0$. In addition, I assume that $A_a F^1(1) > \bar{a}L > 0$. This inequality states that agriculture is productive enough to provide the subsistence level of food to all workers. If this condition is violated then workers receive negative infinite utility.

Each firm equates its demand for labour to the value of the marginal product of labour. Consequently, as market clearing requires that $L_a + L_m = L$, the marginal productivity of labour will be equalised across sectors,

$$p_a A_a {F_a}'(L_a) = w = p_m A_m {F_m}'(L_m) \tag{A.4}$$

Equilibrium

Autarky and Equilibrium Prices

Equilibrium is defined as a set of prices, wages, and an allocation of workers across sectors such that goods and labour markets clear. In a state of autarky, the price ensures that the total amount produced is equal to total consumption in each location, so that,

$$\begin{aligned} C_a &= A_a F_a(L_a) \\ C_m &= A_m F_m(L_m) \end{aligned} \tag{A.5}$$

Maximisation of equation A.1 implies that each worker consumes agricultural goods such that,

$$p_a C_a = \bar{a} + \frac{\alpha p_m C_m}{1-\alpha} \tag{A.6}$$

Combining this result with the profit maximisation condition (equation A.4), the labour market clearing condition ($L_m = 1 - L_a$), and the fact that total production must equal total consumption yields,

$$\Omega(L_m) = \frac{\bar{a}}{A_a} \tag{A.7}$$

where,

$$\Omega(L_m) \equiv F_m(L_m) - \frac{F_m{}'(L_m)F_a(1-L_a)}{F_a{}'(1-L_a)} \tag{A.8}$$

In addition, it is the case that $\Omega(0) = F_m(1)$, $\Omega(1) < 0$ and $\Omega'(\cdot) < 0$.

Consequently, in equilibrium a unique interior solution will arise for the employment share in manufacturing L_m,

$$L_m = \Omega^{-1}\left(\frac{\bar{a}}{A_a}\right) \tag{A.9}$$

As preferences are non-homothetic the demand for agricultural goods (food) decreases as income increases (Engel's law). Consequently, an increase (decrease) in agricultural productivity will push (pull) workers into the manufacturing (agricultural) sector. Similarly, a decrease (increase) in the subsistence constraint $\bar{a}$ will push (pull) workers into the manufacturing (agricultural) sector.

Trade and Equilibrium Prices

Without opportunities to trade, consumers must consume even their worst productivity draws. The ability to trade breaks the production-consumption link. In the case of free trade prices, set globally, are taken as given. If the world price for a good j, $\bar{p}_j$, exceeds the autarkic local price p_{ij}, firms and farms will engage in arbitrage and sell to the global market. By contrast, if the world price for a good j is less than the autarkic local price consumers will import the product from outside of the local market. Consequently, local demand does not affect the allocation of labour across sectors, i.e., changes in A_{ij} do not affect prices.

As discussed above the rest of the world differs only in terms of agricultural and manufacturing productivity, $A_{i'a}$ and $A_{i'm}$. Profit maximisation in the rest of the world implies that,

$$p_a A_{i'a} F_{i'a}{}'(L_{i'a}) = p_m A_{i'm} F_{i'm}{}'(L_{i'm}) \tag{A.10}$$

Within industry production functions are assumed to be constant across regions. Under the assumption of free trade and incomplete specialisation manufacturing employment in region i, L_{im}, is now determined jointly by equations A.4 and A.10. Taking the ratio of these equations provides the following equality,

$$\frac{F_{im}{}'(L_{im}}{F_{ia}{}'(L_{ia})} = \frac{A_{ia}A_{i'm}}{A_{i'a}A_{im}} \frac{F_{i'm}{}'(L_{i'm})}{F_{i'a}{}'(L_{i'a})} \tag{A.11}$$

As $\frac{F_{im}{}'(L_{im})}{F_{ia}{}'(L_{ia})}$ is decreasing in L_{im} it follows that,

$$L_{im} \gtreqless L_{ia} \quad iff \quad \frac{A_{i'a}}{A_{i'm}} \gtreqless \frac{A_{ia}}{A_{im}} \tag{A.12}$$

In this case an increase (decrease) in agricultural productivity will pull (push) workers into the agricultural (manufacturing) sector, due to a change in local comparative advantage. This is demonstrated in figure 1.1

Figure 1.1: The Effect of a Reduction in Agricultural Productivity on Equilibrium Employment Shares (Free Trade)

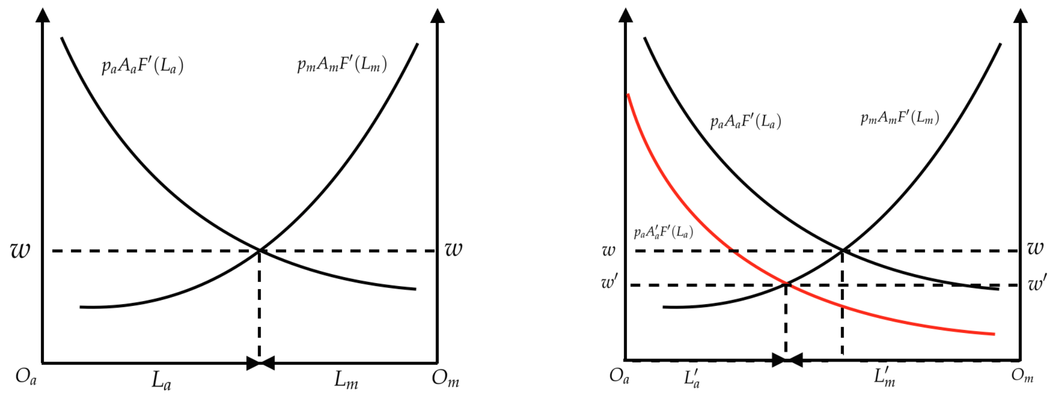

In the case of costly trade, firms (farms) will engage in aribtrage opportunities as before; however, the local price is bounded by a trade cost δ. Consequently, a trader will engage in arbitrage, selling on the global market, as long as the global price is greater than the local price net of trade costs, i.e., $\bar{p}_j/\delta > p_j^A$. Conversely, consumers will import from the global market if the local price is greater than the global price net of trade costs, i.e., $\bar{p}_j < p_j^A/\delta$. Consequently, in the case of homogenous traders where all agents face a constant iceberg trade cost, the local price is bounded by the global price, i.e., $\frac{\bar{p}_j}{\delta} \leq p_j^A \leq \bar{p}_j\delta$.[3]

A.2 Data appendix

A.2.1 Agricultural Data Appendix

This section provides additional details on the Agriculture data used in section III.

As discussed in the main paper, the data is collected from the ICRISAT Village Dynamics in South Asia Macro-Meso Database (henceforth VDSA) which is compiled from a number of official government datasources. Figures 1 provides summary statistics for the 12 crops used.

We observe from the figures that both Rice and Wheat are the most produced crops in terms of cultivated land area and total production (figure 1) and that they also comprise the largest share of production and cultivated land area within-district

[3] See Allen and Atkin (2015) for an analysis of heterogeneous traders.

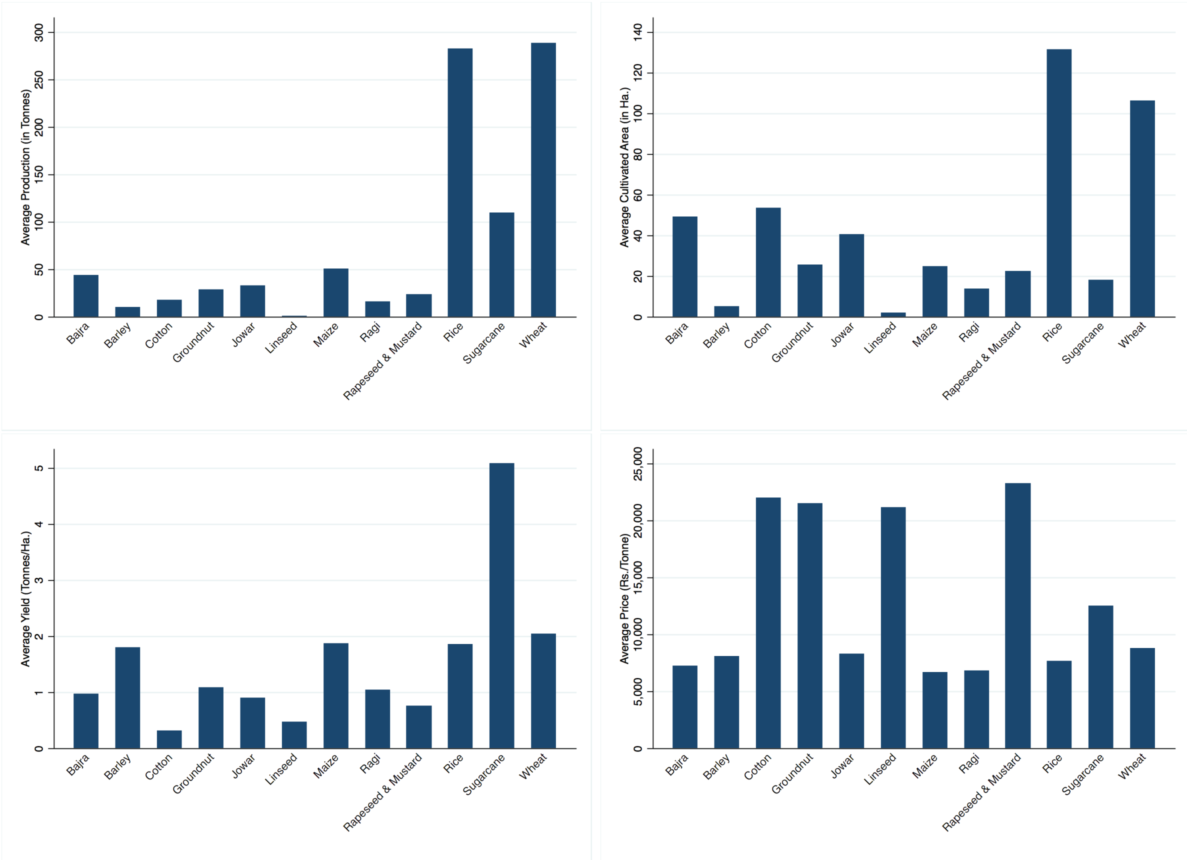

Figure 2.1: Average: (TL) Production; (TR) Cultivated Land Area; (BL) Yield; (BR) Price (2001 Rs.)

(figure 2). However, in terms of yields sugarcane is show to have one of the highest yields and has the largest share of yields within-district (figure 2).

A.2.2 NSS Data Appendix

This section provides additional details on the NSS Employment and Unemployment surveys used in section III. The National Sample Survey Organisation (NSSO) carries out all-India, large sample, household surveys on employment and unemployment every few years. This paper takes advantage of the 60th round (January 2004 – June 2004), the 61st round (July 2004 – June 2005), the 62nd round (July 2005 – June 2006), and the 64th round (July 2007 – June 2008).

Using this data I construct average wage for agricultural labourers (not cultivators), other non-agricultural wages in rural areas, and non-agricultural wages in urban areas. Looking at the breakdown of employment between rural and urban areas it is clear that non-agricultural activities are not restricted to urban areas motivating the differentiation between rural and urban outcomes when examining the effects on wages in section III.

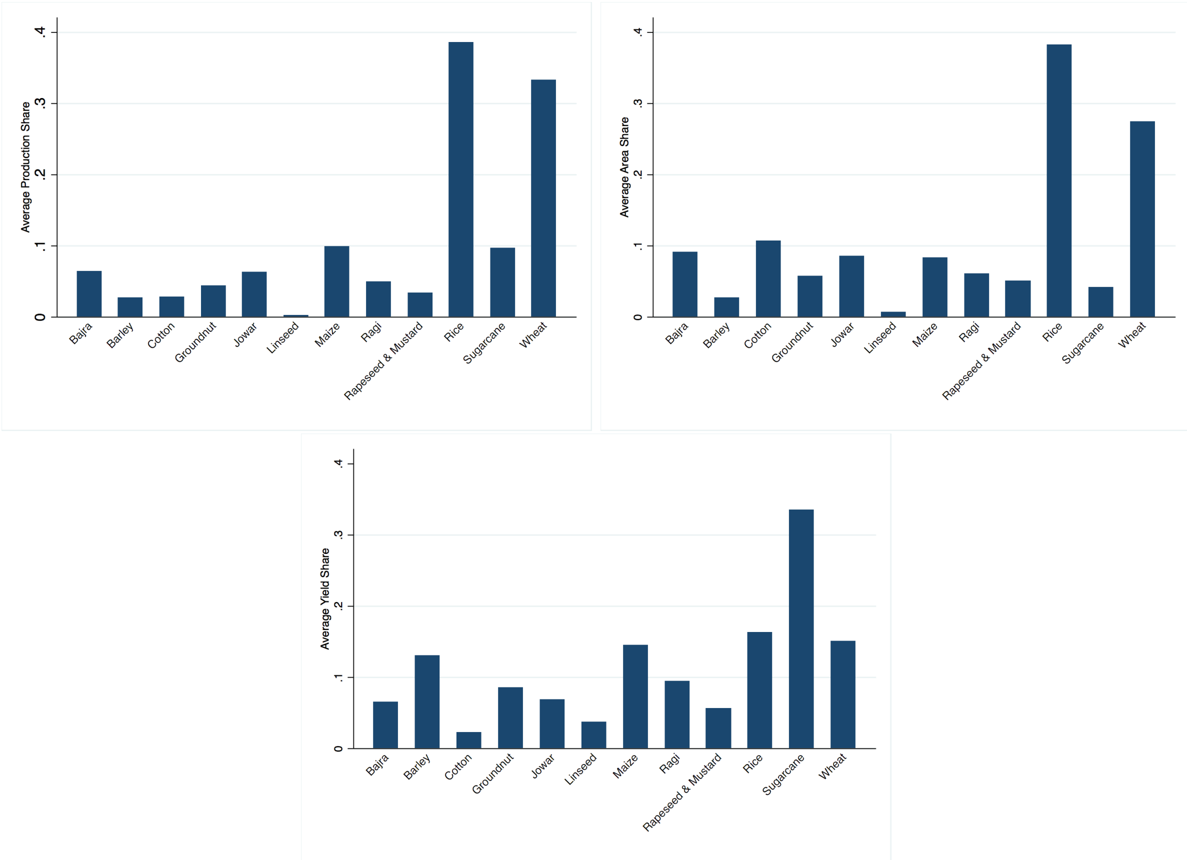

Figure 2.2: District Shares of: (TL) Agricultural Production; (TR) Agricultural Cultivated Land; (B) Agricultural Yields

As one might expect agricultural employment is largely focused in rural areas accounting for an average of 65% of rural employment during this period. However, employment manufacturing and services account collectively for close to 30%. By contrast, in urban areas manufacturing and services account for close to 80% of employment. This is consistent with one of the most striking features of India's recent spatial development, namely the expansion of India's metropolitan areas into rural areas, referred to peri-urbanization (see Colmer (2015) for a more detailed discussion and review of this literature). In the last decade there has been an official increase in urban agglomerations by 25% with populations shifting outwards. Henderson (2010) presents evidence in support of this industrial decentralization for the Republic of Korea and Japan. Desmet et al. (forthcoming) and Ghani et al. (2014) also provide supporting evidence for this process in India. Desmet et al. (forthcoming) show that the services sector has become increasingly concentrated over time, while manufacturing has become less concentrated in districts that were already concentrated and has increased in districts which originally were less concentrated. Ghani et al. (2014) look more specifically at the manufacturing sector and document its movement away from urban to rural areas, comparing the formal and informal sectors. The authors argue that the formal sectors is becoming more rural; however, in practice a lot of

Table 2.1: Employment Shares in India (2001–2007)

	Rural	Urban	Combined
Agriculture	65.3%	9.7%	45.6%
Manufacturing	14.2%	40.3%	23.4%
Services	12.1%	38.1%	21.2%
Construction	6.8%	9.61%	7.8%
Mining	0.6%	0.8%	0.7%
Unemployment	10.1%	18.0%	12.9%

this movement is likely sub-urbanization, rather than ruralisation, in which firms move to the outskirts of urban areas where they can exploit vastly cheaper land and somewhat cheaper labour. Colmer (2015) finds evidence consistent with these papers finding that manufacturing employment growth has become more concentrated in districts which were initially less concentrated, and that this employment growth is significantly higher in less concentrated rural areas compared to less concentrated urban areas.

This process of peri-urbanization also benefits workers reducing the cost of sectoral adjustment and migration costs. Indeed, in many instances it may reduce the need to migrate altogether with workers choosing to commute from home, rather than migrate to urban areas. This is consistent with the non-trivial shares of manufacturing employment and agricultural employment presented in rural and urban areas respectively. Interestingly, we observe that the unemployment share in urban areas is almost twice the size of those in rural areas, suggesting that there is more absorptive capacity in rural areas.

A.2.3 Weather Data Appendix

This section provides additional details on the weather data used throughout this paper.

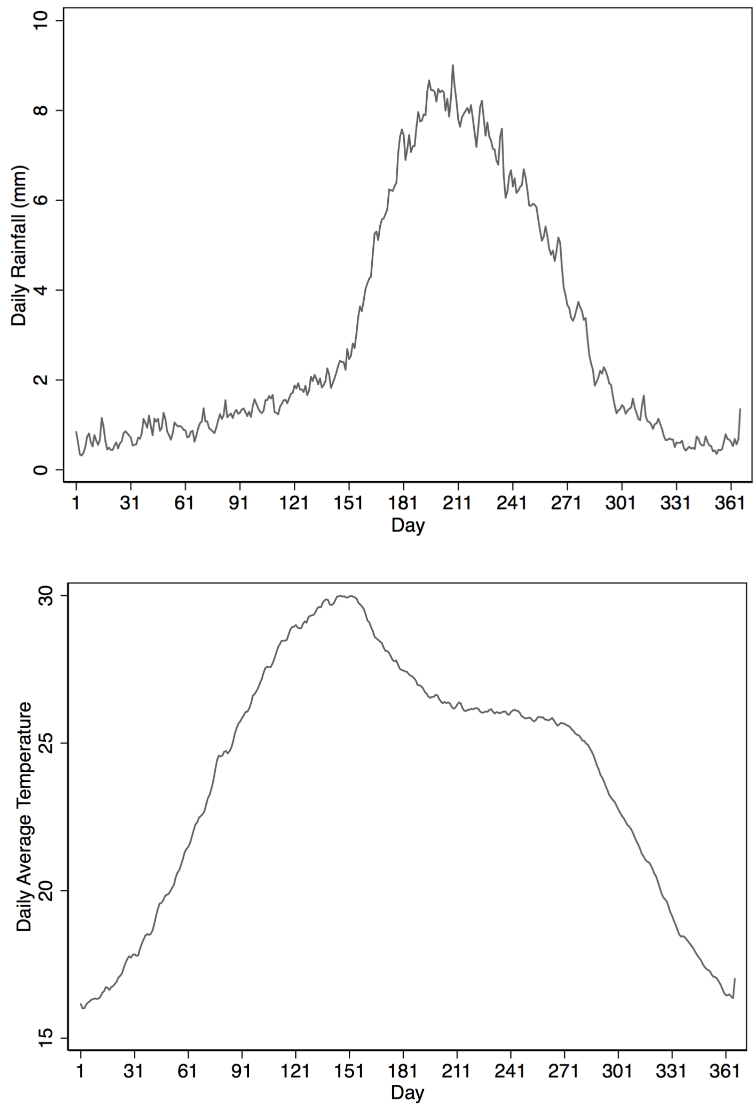

Figure 2.3: Intra-Annual Weather Variation

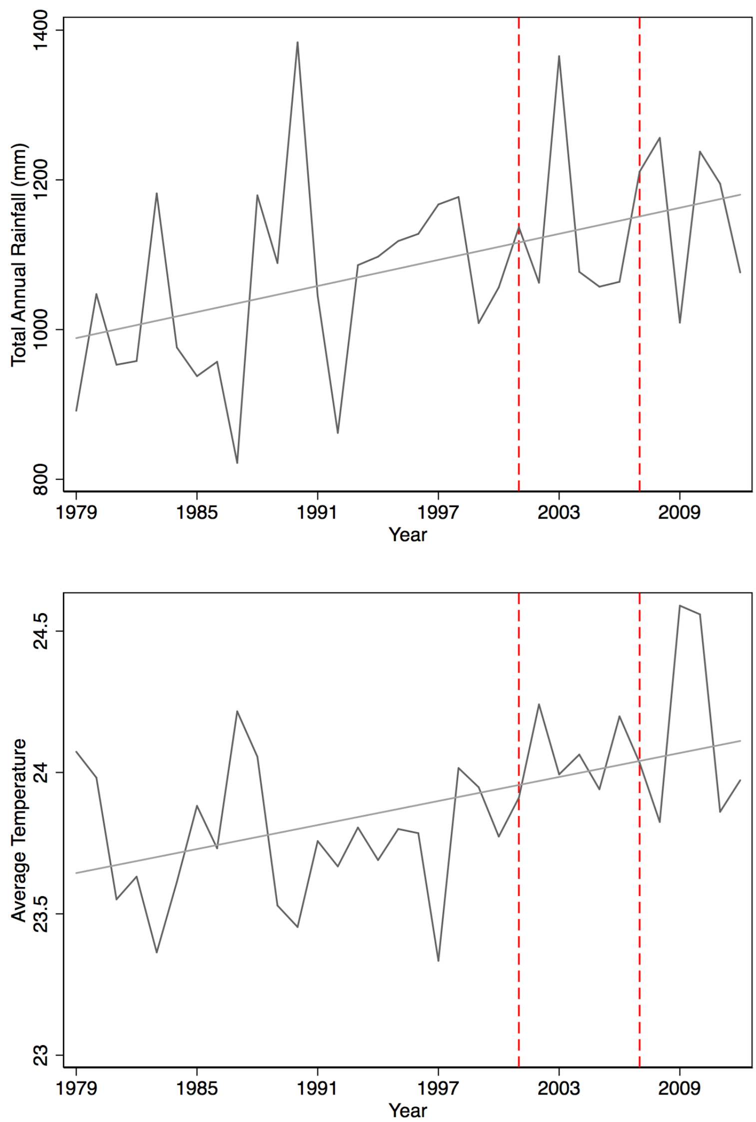

Figure 2.4: Inter-Annual Weather Variation

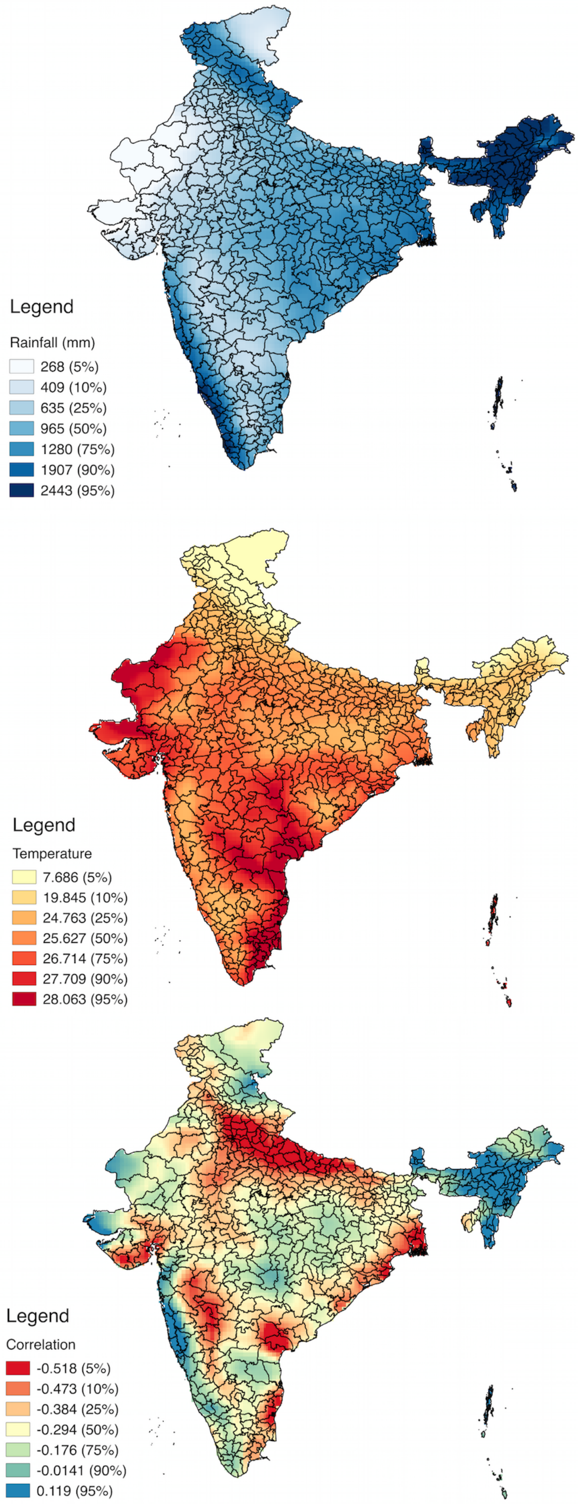

Figure 2.5: Spatial Weather Variation

A.2.4 ASI Data Appendix

This section provides additional details on the Annual Survey of Industries Establishment-level Microdata.

I begin by extracting a subset of variables from the raw data separately for each year and then append each year together before apply the following cleaning processes, summarised in table With this initial sample I begin by dropping all plants that are outside of the manufacturing sector, closed. In addition, I remove all observations with missing or zero total output data due to the importance of the revenue and productivity results. I then combine this data with the weather data taken from the ERA-Interim Reanalysis Data archive. Finally, I drop Union Territories and remove all districts with zero agricultural production. This is due to the focus on agricultural productivity shocks as a driver of labour reallocation.

All financial amounts are deflated to constant 2001-02 Rupees.[4] Revenue (gross sales) is deflated by a three-digit commodity price deflator available from the "Index Numbers of Wholesale Prices in India - By Groups and Sub-Groups (Yearly Averages)" produced by the Office of the Economic Adviser in the Ministry of Commerce & Industry.[5] Material inputs are deflated by constructing the average output deflator for a given industry's supplier industries based on India's 1993-94 input-output table, available from the Central Statistical Organization. Fuel and Electricity costs are deflated by the price index for "Fuel, Power, Light, and Lubricants". Capital is deflated by an implied national deflator calculated from "Table 13: Sector-wise Gross Captial Formation" from the Reserve Bank of India's Handbook of Statistics on the India Economy.[6] Wage costs are deflated using a national GDP deflator.

[4]Thank you to Hunt Allcott, Allan Collard-Wexler, and Stephen O'Connel for publicly providing the data and code to conduct this exercise.

[5]Available from http://www.eaindustry.nic.in/

[6]Available from http://www.rbi.org.in

Table 2.2: ASI Sample Selection

Action Taken	Observations Dropped	Final Sample
Initial Sample	-	371,383
Drop Sectors Outside of Manufacturing	22,644	348,739
Drop Closed Plants	89,437	259,302
Drop Total Output Zero or Missing	32,646	226,656
Merge Weather Data	19,707	206,949
Drop Union Territories	1,430	205,519
Drop Zero Agricultural GDP	18,003	187,516
Merge with Deflators	9	187,507
Drop if Employment ¡ 10	42,861	144,646
Drop if Employment ¡ 20 & No Electricity	864	143,782
Plants Above the Threshold	-	49,112
Plants Below the Threshold	-	94,085

A.2.5 Production Function and Productivity Estimation

In what follows I provide an explicit model of TFPR, in the context of a profit-maximising firm.

Each firm i, in time t, produces output Q_{it} using the following (industry-specific) technology:

$$Q_{it} = A_{it} K_{it}^{\alpha_k} M_{it}^{\alpha_m} E_{it}^{\alpha_E} (L_{it} \bar{\varepsilon}_{it})^{\alpha_L}$$

where K_{it} is the capital input, $L_{it}\bar{\varepsilon}_{it}$ is the heterogeneous labour input, M_{it} is the materials input, and E_{it} is the electricity input. Furthermore, I assume constant returns to scale in production so $\alpha_M + \alpha_E + \alpha_K + \alpha_L = 1$.

The demand curve for the firm's product has a constant elasticity:

$$Q_{it} = B_{it} P_{it}^{-\epsilon}$$

Combining these two equations I obtain an expression for the sales-generating production function:

$$S_{it} = \Omega_{it} K_{it}^{\beta_k} M_{it}^{\beta_M} E_{it}^{\beta_E} (L_{it} \bar{\varepsilon}_{it})^{\beta_L}$$

where $\Omega_{it}(true) = A_{it}^{1-\frac{1}{\epsilon}} B_{it}^{\frac{1}{\epsilon}}$, and $\beta_X = \alpha_X(1 - \frac{1}{\epsilon})$ for $X \in \{K, L, M, E\}$. Within the confines of this paper, I define true productivity as $\omega_{it} \equiv \log(\Omega_{it})$.

To recover a measure of ω_{it}, I compute the value of β_L, β_M, and β_E using median regression for each industry-year cell.

$$\beta_X = median\left(\left\{\frac{P_{it}^X X_{it}}{S_{it}}\right\}\right) \qquad \text{for } X \in \{L, M, E\}$$

To recover the coefficient on capital, β_K, I use the assumption of constant returns to scale in production, i.e., $\sum_X \alpha_X = 1$, such that:

$$\beta_K = \frac{\epsilon - 1}{\epsilon} - \beta_L - \beta_M - \beta_E$$

For ease of measurement I set ϵ to be constant for all firms. Following Bloom (2009) I set $\epsilon = 4$. Using these estimates I compute ω_{it},

$$\begin{aligned} \omega_{it}(est) &= \omega_{it}(true) + \beta_L \bar{\varepsilon}_{it} \\ &= \log(S_{it}) - \beta_K \log(K_{it}) - \beta_M \log(M_{it}) - \beta_E \log(E_{it}) - \beta_L \log_(L_{it}) \end{aligned}$$

This captures the fact that when estimating TFPR the data forces us to treat each worker the same. Consequently, if average worker productivity is decreasing as the number of workers increases we will systematically underestimate $\omega_{it}(true)$ in sectors where the labour force is expanding. By contrast in the presence of adjustment costs where average worker productivity is increasing as the number of workers increases we will systematically overestimate $\omega_{it}(true)$.[7]

[7] This considerations remains even if we had the option of differentiating workers based on observable characteristics (Young, 2015). Within each type the marginal worker is treated as identical to average workers resulting in the same problem identified above. To avoid this problem we would need to know the individual productivity of each worker (accounting for both observable and unobservable characteristics) and calculate TFP based on the summation of each worker's contribution to production.

A.3 The Labour Regulation Environment – Supporting Evidence

This appendix provides supporting evidence for the identification strategy that exploits spatial variation and firm-level exposure to India's labour regulation environment.

A.3.1 Bunching in the Firm-Size Distribution

First I examine the degree to which there is bunching in the firm-size distribution, exploiting differences in the incentives that firms face across different states as well as differences in the regulatory thresholds. I demonstrate that in West Bengal, arguably the state with the most rigid labour regulation environment, that there is a bunching of firms just below the regulatory threshold of 50 workers. However, there is no bunching for the other states around this threshold, in support of the identification strategy.

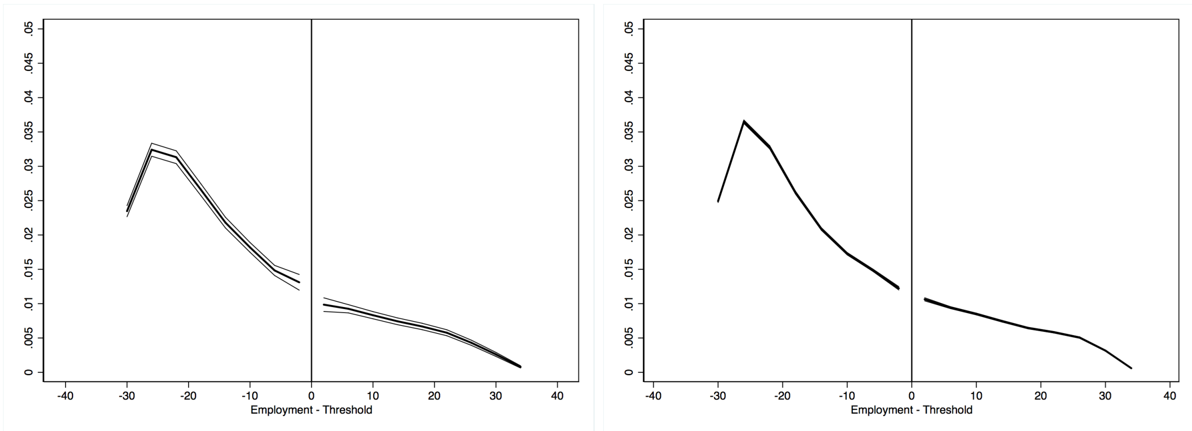

Figure 3.1: Bunching in West Bengal (top) around the regulatory threshold (50 workers) vs. other States (bottom) around the 50 workers threshold.

Identifying bunching around the regulatory threshold of 100 is more challenging as this coincides with the sampling scheme of the ASI, in which there is an oversampling above this threshold. Interestingly, we observe that there is a bunching just above the regulatory threshold but only in flexible states. This is the complete opposite to what one should expect, as there should be no bunching in flexible states. In rigid states we observe that there is no bunching where we should expect there to be. Due to the fact that the sampling scheme of the ASI coincides with the regulatory threshold it is impossible to identify or rule out bunching at this level. However, it is interesting to note that the difference in the distribution between rigid and flexible states around the regulatory threshold of 100 is equal to the difference around the regulatory threshold of 50 where we can identify bunching. This provides support for the conclusion that the sampling scheme of the ASI interferes with the identification; however, as mentioned it is impossible to identify or rule anything out.

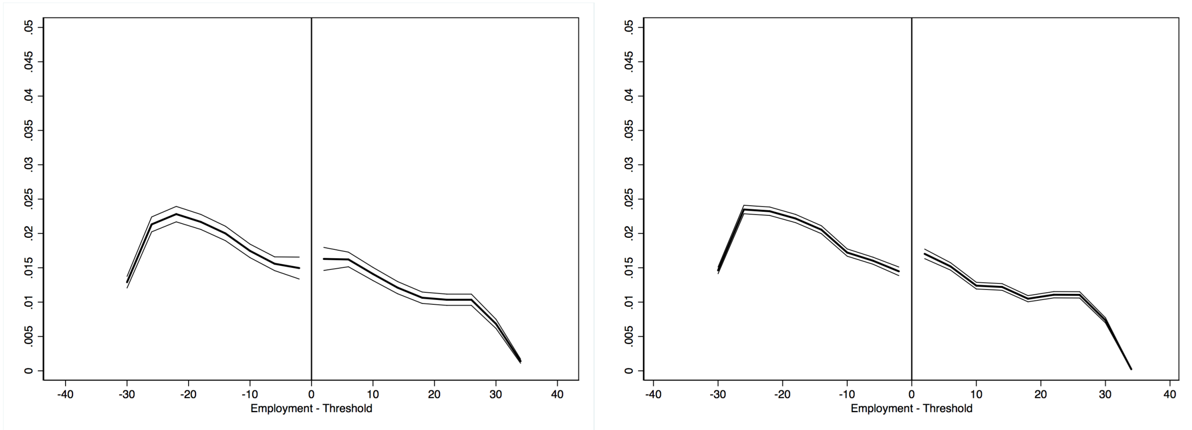

Figure 3.2: Bunching vs. Sampling around the regulatory threshold (100 workers): Rigid (Top); Flexible (Bottom)

In further support of the assumption that the sampling scheme interferes with the identification strategy I don't find any differences around the regulatory threshold of 300 in "flexible" Uttar Pradesh. This suggests that there isn't anything fundamentally related to the labour regulation environment in flexible states that would result in bunching to the right of the regulatory threshold.

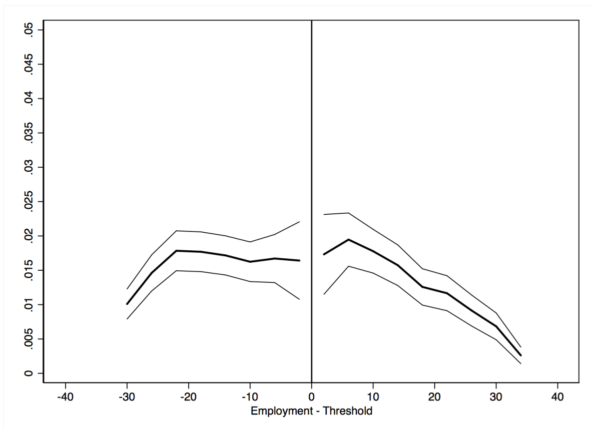

Figure 3.3: No Bunching around the regulatory threshold in "Flexible" Uttar Pradesh (300 workers)

A.3.2 The Effect of the Labour Regulation Environment on Unregulated Sectors

Secondly, I demonstrate that there are no differences in temperature effects across labour regulation environments when examining unregulated sectors such as agriculture, services, construction and mining. In addition there are no differences in temperature effects across labour regulation environments when looking at manufacturing broadly defined to include both the formal sector (those above and below the regulatory threshold) and the informal sector. These results provide further support for the identification strategy as they suggest that there are no additional spatial differences between the collection of states that make up rigid and flexible labour regulation environments, other than the regulatory environment that affects manufacturing establishments above the regulatory threshold.

Table 3.1: The Differential Effect of Temperature on Real GDP - By Sector (2001 – 2012)

	(1) Aggregate	(2) Agriculture	(3) Services	(4) Manufacturing	(5) Construction
Daily Average Temperature (°C)	-0.0375**	-0.157**	0.0107	-0.0981**	0.0236
	(0.0186)	(0.0627)	(0.0190)	(0.0434)	(0.0694)
Temperature × Flexibility	0.0179	0.0678	-0.0188	0.0747	0.00643
	(0.0241)	(0.0881)	(0.0244)	(0.0572)	(0.0856)
Rainfall Controls	Yes	Yes	Yes	Yes	Yes
Fixed Effects	District, Year, and State-Year Time Trends				
Observations	6,792	6,763	6,792	6,264	6,780

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Standard errors are adjusted to reflect spatial dependence (up to 1,480km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

A.3.3 Alternative Definitions of the Labour Regulation Environment

In addition to the evidence in support of the identification strategy, I provide a series of robustness tests in support for the main results.

Table 3.2: Alternative Measures of the Labour Regulation Environment – (Neutral and Flexible Binary Variables)

	Total Output	Items Produced	Employment Contract	Employment Permanent	Output Per Worker	TFPR Contract	Day Wage Permanent	Day Wage
Daily Average Temperature (°C)	-0.0550	-0.0491***	-0.101**	-0.000238	-0.0604	-0.0765***	0.0152	-0.0647***
	(0.0444)	(0.0146)	(0.0468)	(0.0285)	(0.0388)	(0.0244)	(0.0251)	(0.0157)
Temperature × Neutral	0.0869	0.0576***	0.124**	0.0227	0.0821*	0.0581**	-0.0393	0.0734***
	(0.0533)	(0.0183)	(0.0497)	(0.0324)	(0.0460)	(0.0242)	(0.0284)	(0.0190)
Temperature × Flexible	0.0991*	0.0462*	0.0847	0.0477	0.0933*	0.0514*	-0.0631*	0.0615**
	(0.0565)	(0.0240)	(0.0638)	(0.0385)	(0.0534)	(0.0299)	(0.0337)	(0.0242)
Rainfall Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed Effects	Sector × District, Sector × Year, and State-Year Time Trends							
H_0: Neutral = Flexible (p-value)	0.801	0.638	0.424	0.534	0.798	0.748	0.389	0.554
Observations	49,112	49,112	24,557	47,363	49,112	44,305	24,557	47,363

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Standard errors are adjusted to reflect spatial dependence (up to 1,800km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

Table 3.3: Alternative Measures of the Labour Regulation Environment – (Flexible and Neutral Combined Binary Measure)

	Total Output	Items Produced	Employment Contract	Employment Permanent	Output	TFPR Per Worker	Day Wage Contract	Day Wage Permanent
Daily Average Temperature (°C)	-0.0554	-0.0488***	-0.1000**	-0.00104	-0.0607	-0.0763***	0.0161	-0.0643***
	(0.0441)	(0.0147)	(0.0471)	(0.0285)	(0.0387)	(0.0245)	(0.0251)	(0.0157)
Temperature × Labour Regulation	0.0910*	0.0538***	0.113**	0.0311	0.0858*	0.0558**	-0.0464*	0.0695***
	(0.0494)	(0.0169)	(0.0494)	(0.0289)	(0.0440)	(0.0244)	(0.0274)	(0.0186)
Rainfall Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed Effects	Sector × District, Sector × Year, and State-Year Time Trends							
Observations	49,112	49,112	24,557	47,363	49,112	44,305	24,557	47,363

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Standard errors are adjusted to reflect spatial dependence (up to 1,800km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

A.4 Additional Results and Robustness Tests

This appendix provides a series of additional results and robustness tests to support of the main results presented in the paper.

A.4.1 Concerns Relating to the Endogenous Selection of Firms around the Regulatory Threshold

One concern relates to the endogenous selection of firms around the regulatory threshold.

To mitigate these concerns about selection around the firm-size threshold I run the baseline specification applying uniform thresholds to all states. This allows for the analysis to still be applied to regulated firms but abstracts from state-specific thresholds. The results of this exercise can be found in table 4.1. In Panel A and B I restrict the uniform threshold to apply to regulated firms with more than 100 and 150 workers respectively. In addition I drop Uttar Pradesh as firms are only regulated if they have more than 300 workers. In Panel C and D I restrict the uniform threshold to apply to regulated firms with more than 300 and 350 workers respectively. Across all panels the results are broadly robust suggesting that endogeneity concerns relating to selection around the regulatory threshold are not a first-order concern.

A.4.2 Missing Observations

One concern relating to the quality of the data is the large number of missing values associated with the contract worker variables. Missing data can cause biased estimates and reduced efficiency when estimating regression coefficients (Rubin, 1987). The key issue is that it is unclear as to whether these missing values are zeroes or whether they are positive unreported values. The implicit assumption made so far is that these data are missing at random. In practice this is very likely to not be the case and so it is important to test whether the effects differ based on whether the firm reports or fails to report data on contract workers.

A.4.3 Non-Linearities in the Temperature Schedule

To begin I explore the degree to which there are non-linearities in the temperature schedule. A large literature in agricultural science has demonstrated that the relationship between agricultural yields and weather is highly nonlinear (Schlenker and Roberts, 2009; Auffhammer and Schlenker, 2014). To account for these non-linearities

I apply the concept of growing degree days, which measure the amount of time a crop is exposed between a given lower and upper bound with daily exposures summed over the season. Denoting the lower bound as t_l, the upper bound as t_h, and t_d as the daily average temperature on a given day,

$$GDD_{d;t_l;t_h} = \begin{cases} 0 \text{ if } t_d \leq t_l \\ t_d - t_l \text{ if } t_l < t_d < t_h \\ t_h - t_l \text{ if } t_h \leq t_d \end{cases} \tag{A.13}$$

These daily measures are then summed over the period of interest.[8] This approach is appealing for several reasons. First, the existing literature suggests that this simple function delivers results that are very similar to those estimated using more complicated functional forms (Schlenker and Roberts, 2009; Burgess et al. 2014; Burke and Emerick, 2015). Secondly, these other functional forms typically feature higher order terms, which in a panel setting means that the unit-specific mean re-enters the estimation, as is the case with using the quadratic functions (McIntosh and Schlenker, 2006). This raises both omitted variable concerns, as identification in the panel models is no longer limited to location-specific variation over time.

Using the notion of GDD, I model outcomes as a simple piecewise linear function of temperature and precipitation,

$$f(w_{dt}) = \beta_1 GDD_{dt;t_l;t_h} + \beta_2 GDD_{dt;t_h;\infty} + \beta_3 Rain_{dt} \tag{A.14}$$

The lower temperature "piece" is the sum of GDD between the lower bound $t_l = 0$ and kink-point t_h. The upper temperature "piece" has a lower bound of t_h and is unbounded above. The kink-point in the distribution t_h is determined by estimating an agricultural production function, looping over all possible thresholds and selecting the model with the lowest root-mean-square error. This results in a kink-point at 17°C. When the same process is applied to the manufacturing data a kink-point of 19°C is selected. The closeness in kink-points from each exercise is encouraging.

[8] For example, if we set t_l equal to 0°C and t_h equal to 24°C then a given set of observations $\{-1, 0, 8, 12, 27, 30, 33\}$, would provide $GDD_{dt;0;24} = \{0, 0, 8, 12, 24, 24\}$. Similarly if we wanted to construct a piecewise linear function setting t_l equal to 24 and t_h equal to infinity the second "piece" would provide $CDD_{dt;24;\infty} = \{0, 0, 0, 0, 6, 9\}$. These values are then summed over the period of interest, in this case $CDD_{dt;0;24} = 68$ and $CDD_{dt;24;\infty} = 15$. This approach accounts for any differences in the response to this temperature schedule relative to a different schedule with the same daily average temperature.

A.4.4 Alternative Weather Data

In addition to testing the robustness of the results to non-linearities in the temperature schedule I also test the robustness of the results to an alternative weather data set, the University of Delaware Air Temperature and Precipitation dataset, which provides monthly data on a high-resolution grid. The main issue with this data is its reliance on ground station data, which is interpolated onto a $0.5^\circ \times 0.5^\circ$ grid. This interpolation occurs over time and space even in places where there is limited data availability resulting in considerable measurement error that may not necessarily be classical. In addition, I demonstrate that the insignificance of the rainfall results is robust to the use of satellite data from NASA and JAXA's Tropical Rainfall Measuring Mission (TRMM), where non-classical measurement error is avoided.

Table 4.1: Baseline Specification - Uniform Thresholds

	Total Output	Items Produced	Employment Contract	Employment Permanent	Output	TFPR Per Worker	Day Wage Contract	Day Wage Permanent
Panel A: Above 100 UP dropped								
Daily Average Temperature (°C)	-0.138***	-0.0517**	-0.101	-0.0302	-0.104**	-0.0805***	0.0450	-0.0673***
	(0.0530)	(0.0231)	(0.0670)	(0.0378)	(0.0472)	(0.0297)	(0.0344)	(0.0192)
Temperature × Flexibility	0.204**	0.0658*	0.138	0.0840*	0.149**	0.0714*	-0.103**	0.0793**
	(0.0817)	(0.0347)	(0.0958)	(0.0501)	(0.0750)	(0.0396)	(0.0496)	(0.0310)
Observations	46,224	46,224	23,072	44,534	46,224	41,951	23,072	44,534
Panel B: Above 150 UP dropped								
Daily Average Temperature (°C)	-0.0737	-0.0203	-0.0610	-0.00807	-0.0441	-0.0790**	0.0499	-0.0688***
	(0.0637)	(0.0243)	(0.0691)	(0.0477)	(0.0554)	(0.0318)	(0.0377)	(0.0250)
Temperature × Flexibility	0.221**	0.0189	0.0922	0.0588	0.166**	0.0913**	-0.103*	0.106***
	(0.0920)	(0.0370)	(0.101)	(0.0673)	(0.0806)	(0.0421)	(0.0567)	(0.0397)
Observations	35,249	35,249	18,457	34,488	35,249	31,991	18,457	34,488
Panel C: Above 300								
Daily Average Temperature (°C)	-0.110	-0.0398	-0.229**	-0.0588	-0.0479	-0.0745**	0.118***	-0.0543**
	(0.0791)	(0.0351)	(0.0905)	(0.0489)	(0.0774)	(0.0321)	(0.0429)	(0.0268)
Temperature × Flexibility	0.250**	0.0535	0.330**	0.140*	0.138	0.0972**	-0.214***	0.0924**
	(0.116)	(0.0539)	(0.130)	(0.0725)	(0.109)	(0.0477)	(0.0653)	(0.0410)
Observations	18,774	18,774	10,626	18,489	18,774	16,972	10,626	18,489
Panel D: Above 350								
Daily Average Temperature (°C)	-0.0940	-0.0453	-0.200**	-0.0418	-0.0611	-0.0515	0.131***	-0.0407
	(0.0864)	(0.0368)	(0.102)	(0.0542)	(0.0836)	(0.0334)	(0.0505)	(0.0293)
Temperature × Flexibility	0.223*	0.0672	0.283*	0.121	0.166	0.0655	-0.225***	0.0857*
	(0.125)	(0.0552)	(0.147)	(0.0809)	(0.118)	(0.0483)	(0.0759)	(0.0449)
Observations	15,598	15,598	8,868	15,374	15,598	14,108	8,868	15,374

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. UP dropped = Uttar Pradesh dropped. This is because the threshold for UP is 300. Standard errors are adjusted to reflect spatial dependence (up to 1,800km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

Table 4.2: Baseline Specification - Regulated Firms (Missing Data)

	Total Output	Items Produced	Employment Contract	Employment Permanent	Output	TFPR Per Worker	Day Wage Contract	Day Wage Permanent
Daily Average Temperature (°C)	-0.0621	-0.0623***	-0.121*	0.00457	-0.0889*	-0.0846***	0.0401	-0.0897***
	(0.0478)	(0.0212)	(0.0644)	(0.0358)	(0.0459)	(0.0276)	(0.0316)	(0.0210)
Temperature × Flexibility	0.140*	0.0891***	0.171*	0.00867	0.143*	0.0769**	-0.0963**	0.122***
	(0.0779)	(0.0331)	(0.0933)	(0.0512)	(0.0734)	(0.0372)	(0.0462)	(0.0323)
Temperature × Missing	-0.00509	0.0115	–	0.00433	0.00538	-0.00456	–	0.0381***
	(0.0176)	(0.00845)		(0.0224)	(0.0156)	(0.00811)		(0.00920)
Temperature × Missing × Flexibility	0.0235	-0.0247*	–	0.0179	0.00432	0.0151	–	-0.0563***
	(0.0274)	(0.0133)	–	(0.0346)	(0.0230)	(0.0116)		(0.0141)
Missing × Flexibility	-0.603	0.563	–	-0.497	-0.0713	-0.416		1.378***
	(0.717)	(0.343)		(0.895)	(0.605)	(0.301)		(0.362)
Missing	0.0235	-0.207	–	0.578	-0.112	0.132		-0.954***
	(0.450)	(0.217)	–	(0.572)	(0.399)	(0.206)		(0.233)
Rainfall Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed Effects	Sector × District, Sector × Year, and State-Year Time Trends							
Observations	49,112	49,112	24,557	47,363	49,112	44,305	24,557	47,363

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Standard errors are adjusted to reflect spatial dependence (up to 1,800km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

Table 4.3: The Non-Linear Effects of Temperature on Agricultural Outcomes

	Agricultural Outcomes					
	Log Value (All Crops)	Log Yield (All Crops)	Log Price (All Crops)	Log Value (Main Crop)	Log Yield (Main Crop)	Log Price (Main Crop)
Degree Days (10 days)	-0.00467***	-0.00599***	0.00103**	-0.00773***	-0.00882***	-0.00000367
$t_L = 17, t_H = \infty$	(0.00114)	(0.00140)	(0.000445)	(0.00221)	(0.00277)	(0.00100)
Degree Days (10 days)	0.00344	-0.000159	-0.00388*	-0.00277	-0.00814	-0.00319
$t_L = 0, t_H = 17$	(0.00233)	(0.00333)	(0.00232)	(0.00496)	(0.00643)	(0.00295)
Monsoon Rainfall (100mm)	0.00523	0.00536	0.000871	0.0144***	0.00907	-0.00434
	(0.00355)	(0.00396)	(0.00202)	(0.00549)	(0.00689)	(0.00411)
Crop × District FE	Yes	Yes	Yes	Yes	Yes	Yes
Crop × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
State-Year Trends	Yes	Yes	Yes	Yes	Yes	Yes
Observations	16,924	9,723	9,774	2069	1519	1519

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Standard errors are adjusted to reflect spatial dependence (up to 1,000km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

Table 4.4: The Non-Linear Effects of Temperature on Manufacturing Outcomes (17°C kink-point)

	Total Output	Items Produced	Employment Contract	Employment Permanent	Output	TFPR Per Worker	Day Wage Contract	Day Wage Permanent
Degree Days (10 days) $t_L = 17, t_H = \infty$	-0.00245 (0.00149)	-0.00172*** (0.000608)	-0.00303 (0.00198)	-0.0000321 (0.000976)	-0.00246* (0.00146)	-0.00206*** (0.000796)	0.00153* (0.000881)	-0.00219*** (0.000564)
DD High × Flexibility	0.00415* (0.00224)	0.00209** (0.000883)	0.00470* (0.00272)	0.00158 (0.00126)	0.00383* (0.00214)	0.00182* (0.00106)	-0.00306** (0.00125)	0.00236*** (0.000904)
Degree Days (10 days) $t_L = 0, t_H = 17$	-0.000596 (0.00776)	0.000699 (0.00257)	-0.00286 (0.0130)	-0.00565 (0.00444)	-0.000558 (0.00757)	-0.000560 (0.00416)	-0.0104** (0.00475)	0.000708 (0.00297)
DD Low × Flexibility	0.00342 (0.0113)	0.000801 (0.00383)	0.000350 (0.0188)	0.00593 (0.00669)	0.00278 (0.0109)	-0.00166 (0.00608)	0.0146** (0.00703)	0.00283 (0.00439)
Rainfall Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed Effects	Sector × District, Sector × Year, and State-Year Time Trends							
Observations	49,112	49,112	24,557	47,363	49,112	44,305	24,557	47,363

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Standard errors are adjusted to reflect spatial dependence (up to 1,800km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

Table 4.5: The Non-Linear Effects of Temperature on Manufacturing Outcomes (19°C kink-point)

	Total Output	Items Produced	Employment Contract	Employment Permanent	Output	TFPR Per Worker	Day Wage Contract	Day Wage Permanent
Degree Days (10 days) $t_L = 19, t_H = \infty$	-0.00288*	-0.00169**	-0.00255	-0.000142	-0.00274*	-0.00182**	0.00185*	-0.00236***
	(0.00166)	(0.000676)	(0.00212)	(0.00108)	(0.00162)	(0.000875)	(0.000983)	(0.000599)
DD High × Flexibility	0.00476**	0.00211**	0.00436	0.00186	0.00425*	0.00165	-0.00363***	0.00247***
	(0.00239)	(0.000953)	(0.00291)	(0.00136)	(0.00228)	(0.00112)	(0.00136)	(0.000949)
Degree Days (10 days) $t_L = 0, t_H = 19$	0.00240	-0.000556	-0.00723	-0.00188	0.00109	-0.00324	-0.00810**	0.000654
	(0.00672)	(0.00204)	(0.0106)	(0.00384)	(0.00652)	(0.00316)	(0.00379)	(0.00235)
DD Low × Flexibility	-0.00215	0.00183	0.00683	0.000351	-0.000629	0.00212	0.0116**	0.00206
	(0.00986)	(0.00308)	(0.0154)	(0.00564)	(0.00942)	(0.00462)	(0.00563)	(0.00348)
Rainfall Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed Effects	Sector × District, Sector × Year, and State-Year Time Trends							
Observations	49,112	49,112	24,557	47,363	49,112	44,305	24,557	47,363

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Standard errors are adjusted to reflect spatial dependence (up to 1,800km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

Table 4.6: The Effects of Weather on Agricultural Outcomes (UDEL Weather Data)

	Agricultural Outcomes					
	Log Value (All Crops)	Log Yield (All Crops)	Log Price (All Crops)	Log Value (Main Crops)	Log Yield (Main Crop)	Log Price (Main Crops)
Daily Average Temperature (°C)	-0.115***	-0.174***	0.000801	-0.249***	-0.314***	-0.0134
	(0.0361)	(0.0442)	(0.0146)	(0.0742)	(0.0860)	(0.0347)
Monsoon Rainfall (100mm)	0.00728**	0.00711**	-0.00118	0.0158***	0.00987	-0.00486
	(0.00338)	(0.00359)	(0.00187)	(0.00584)	(0.00664)	(0.00399)
Crop × District FE	Yes	Yes	Yes	Yes	Yes	Yes
Crop × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
State-Year Trends	Yes	Yes	Yes	Yes	Yes	Yes
Observations	16,924	9,723	9,774	2,069	1,519	1,519

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Standard errors are adjusted to reflect spatial dependence (up to 1,000km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

Table 4.7: The Effects of Weather on Average Wages (UDEL Weather Data)

	Agriculture	Manufacturing	Services	Construction
Daily Average Temperature (°C)	0.00296	-0.0307	0.166***	-0.00259
	(0.0249)	(0.0485)	(0.0427)	(0.0363)
Monsoon Rainfall (100mm)	0.00532	0.0111	0.00500	0.00602
	(0.00362)	(0.00710)	(0.00472)	(0.00580)
Observations	1755	1748	1824	1701

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Standard errors are adjusted to reflect spatial dependence (up to 360km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1–7 years.

Table 4.8: The Effects of Weather on the District Share of Employment - By Sector (UDEL Weather Data)

	District Employment Shares				
	Agriculture	Manufacturing	Construction	Services	Unemployment
Daily Average Temperature (°C)	-0.0582***	0.0374***	0.00254	0.0168*	0.00982
	(0.0217)	(0.0129)	(0.00586)	(0.00970)	(0.00634)
Monsoon Rainfall (100 mm)	-0.00289	0.00219	-0.00113	0.00129	-0.000209
	(0.00281)	(0.00175)	(0.000933)	(0.00130)	(0.000942)
Average Share	0.546	0.200	0.071	0.163	0.112
Observations	1831	1831	1831	1831	1831

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Standard errors are adjusted to reflect spatial dependence (up to 630km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

Table 4.9: Main Results: Manufacturing Firm Outcomes – Regulated Firms (UDEL Weather Data)

	Total Output	Items Produced	Output Per Worker	TFPR	Day Wage Contract	Day Wage Permanent	Employment Contract	Employment Permanent
Daily Average Temperature (°C)	-0.0876	-0.0219	-0.144**	-0.0777*	0.0938**	-0.0902***	-0.0901	0.0136
	(0.0671)	(0.0260)	(0.0606)	(0.0404)	(0.0410)	(0.0255)	(0.0805)	(0.0432)
Temperature × Flexibility	0.147	0.0626*	0.217**	0.0417	-0.141**	0.131***	0.190*	-0.0184
	(0.0971)	(0.0362)	(0.0849)	(0.0551)	(0.0574)	(0.0368)	(0.115)	(0.0560)
Rainfall Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Fixed Effects	Sector × District, Sector × Year, and State-Year Time Trends							
Observations	49,112	49,112	49,112	44,305	24,557	47,363	24,557	47,363

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Standard errors are adjusted to reflect spatial dependence (up to 1,800km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

Table 4.10: Main Results: Manufacturing Firm Outcomes – Regulated Firms (Rainfall Data Sets)

	Total Output	Items Produced	Employment Contract	Employment Permanent	Output Per Worker	TFPR	Day Wage Contract	Day Wage Permanent
Panel A: ERA-Interim Reanalysis								
Monsoon Rainfall	-0.00329	0.00526**	-0.00596	-0.00557**	-0.00990***	-0.000496	0.0129	0.00325
(100 mm)	(0.00523)	(0.00215)	(0.00631)	(0.00276)	(0.00364)	(0.00207)	(0.00903)	(0.00421)
Rainfall	-0.00204	-0.00728**	0.00217	0.00530	0.0105**	-0.00161	-0.0259**	0.00150
× Flexibility	(0.00831)	(0.00358)	(0.00873)	(0.00525)	(0.00516)	(0.00361)	(0.0132)	(0.00680)
Panel B: UDEL								
Monsoon Rainfall	-0.00154	0.00275	-0.00157	-0.00841**	-0.00300	0.00373	0.0250***	-0.000414
(100 mm)	(0.00799)	(0.00281)	(0.00741)	(0.00347)	(0.00374)	(0.00245)	(0.00877)	(0.00531)
Rainfall	-0.00497	-0.00418	-0.00853	0.00561	0.00699	-0.00499	-0.0265**	0.00467
× Flexibility	(0.00983)	(0.00345)	(0.00909)	(0.00466)	(0.00518)	(0.00333)	(0.0121)	(0.00695)
Panel B: TRMM								
Monsoon Rainfall	0.00586	-0.000318	0.00707	-0.00703**	-0.000539	0.00244	0.0176**	-0.00253
(100 mm)	(0.00674)	(0.00230)	(0.00702)	(0.00323)	(0.00354)	(0.00199)	(0.00854)	(0.00501)
Rainfall	-0.0182*	-0.00315	-0.0238**	0.00453	0.00288	-0.00387	-0.0260*	0.00678
× Flexibility	(0.00942)	(0.00300)	(0.00926)	(0.00535)	(0.00487)	(0.00265)	(0.0135)	(0.00676)

Notes: Significance levels are indicated as * 0.10 ** 0.05 *** 0.01. Standard errors are adjusted to reflect spatial dependence (up to 1,800km) as modelled in Conley (1999) and serial correlation (1-year) as modelled in Newey and West (1987). District distances are computed from district centroids. Kernels are selected to provide the most conservative standard errors, looped over all distances between 10 and 2,000km and 1-7 years.

A.5 Adjustment Costs and the Potential Gains from Reallocation

This appendix provides an upper bound estimate of the gains associated with removing any adjustment costs that impede the movement of casual workers into permanent manufacturing positions. As discussed wage gaps could also be explained by skill differences and so do not necessarily imply a misallocation of talent. Consequently, the lower bound associated with this exercise is zero.

A.5.1 Modelling the Potential Gains from Reallocation

To provide some insight into the potential gains from reallocation I introduce some economic structure to the data and explore quantitatively the impact of removing the distortion between casual manufacturing employment and permanent manufacturing employment – a naïve counterfactual in which it is assumed that all differences in wages are driven by misallocation.

To assess the potential gains from reallocation I compare a hypothetical output level in which labour is efficiently allocated across activities to observed output following a similar approach taken in other firm-level and sector-level studies of misallocation (Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009; Vollrath, 2009, 2014; Gollin et al., 2014).

Assuming that each activity operates with a Cobb-Douglas production technology and maximises profits, the wage distortion τ can be identified from the first-order condition,

$$w_j = (1-\alpha)\Lambda_j L_j^{-\alpha}\left[\frac{1}{\tau}\right] \tag{A.15}$$

where $\Lambda_j = p_j A_j$. The presence of adjustment costs, τ, will distort the amount of labour used in activity j compared to the level used in the absence of the distortion. As τ falls labour becomes relatively cheaper for activity j, and so the amount of labour that is utilised rises. In this context misallocation arises as the marginal revenue product of labour is not equalised across activities. To identify aggregate output two additional assumption are required. First, I assume that labour is perfectly substitutable across activities, i.e., there is no activity specific human capital. This implies that the total amount of labour in the economy is simply $L = \sum_j L_j$.[9] The

[9]This assumption implies that the gains from reallocation are an upper bound of the upper bound; however, if one considers sector-specific human capital as a constraint to reallocation then relaxing this constraint is part of the problem.

second assumption is that prices are exogenously fixed, consistent with a small open economy in which all activities produce output that can be traded internationally.[10] Under these assumptions observed output in the economy with adjustment costs can be written,

$$Y = \left(\sum_j \Lambda_j^{1/\alpha} \left[\frac{1}{\tau} \right]^{1/\alpha} \right)^{\alpha} \left(\sum_j L_j \right)^{1-\alpha}$$

which follows from using equation A.15 for each activity to solve simultaneously for the shares L_j/L, and then taking the sum of output across activities. In the presence of adjustment costs Y is below the output-maximising level. Consequently, one can estimate, given the structure imposed above, how much output would rise under the counterfactual in which these adjustment costs are removed. The counterfactual output after removing these adjustment costs is written,

$$Y^* = \left(\sum_j \Lambda_j^{1/\alpha} \right)^{\alpha} L_j^{1-\alpha}$$

With both observed output and counterfactual output levels, the gains from reallocation can be written as,

$$G = \frac{Y^*}{Y} = \frac{\left(\sum_j \Lambda_j^{1/\alpha} \right)^{\alpha}}{\left(\sum_j \frac{\Lambda_j^{1/\alpha}}{\tau^{1/\alpha}} \right)^{\alpha}}$$

providing a measure of the gains in aggregate productivity from eliminating the adjustment costs that impede the movement of labour across activities.

A.5.2 Estimating the Potential Gains from Reallocation

Given the model structure discussed above I estimate the gains from reallocation, $\hat{G}$, for each firm providing the average gains from reallocation, as well as the distribution of gains.

Under the naïve assumption that the only difference in wages across activities is driven by misallocation the average (observable) wage in the destination activity is,

$$\mathbb{E}[w_j] = \mathbb{E}[w_i]\tau,$$

In log-linear terms, the distortion can therefore be estimated as the log-difference in average wages across sectors,

[10]With endogenous prices the gains from reallocation would be smaller as an equivalent movement of workers out of casual activities raises the marginal revenue product of labour by more than if prices are held fixed.

$$\log \tau = \log \mathbb{E}[w_j] - \log \mathbb{E}[w_i]$$

Taking this to the data, I estimate the following moment for each firm,

$$\mathbb{E}[\tau] \quad = \quad \exp(\log \mathbb{E}[w_j] - \log \mathbb{E}[w_i])$$

In addition, I use estimates of the average permanent manufacturing wage, $\mathbb{E}[w_p]$, the average casual manufacturing wage, $\mathbb{E}[w_c]$, and the number of workers in each activity, L_j.

With these estimates, and an assigned value of α, I estimate output in each activity, Λ_j,

$$\hat{\Lambda}_j = \frac{\hat{w}_j \hat{\tau}}{1-\alpha} \hat{L}_j^{\alpha}$$

These values are then used to construct estimates for the observed level of output $\hat{Y}$, the counterfactual level of output $\hat{Y}^*$, and, with these estimates, the estimated gains from reallocation $\hat{G}$.

$$\hat{G} = \frac{\hat{Y}^*}{\hat{Y}} = \frac{\left(\sum_j \hat{\Lambda}_j^{1/\alpha}\right)^{\alpha}}{\left(\sum_j \frac{\hat{\Lambda}_j^{1/\alpha}}{\hat{\tau}^{1/\alpha}}\right)^{\alpha}}$$

A.5.3 Counterfactual Estimates

In considering the gains from reallocation I construct a counterfactual that removes the total wage gap across activities providing an upper bound on the size of adjustment costs. The results of this exercise are presented in table 5.1, and the distribution of gains are presented in figures 5.1.

I estimate that the removal of adjustment costs τ_j would result in a 9.3% increase in the manufacturing output of regulated firms hiring both casual and permanent workers ($\alpha = 0.3$), a non-trivial increase. Under the assumption that the formal manufacturing sector accounts for two-thirds of manufacturing GDP, and that regulated firms hiring both casual and permanent workers account for 55% of the formal sector manufacturing (author's own calculations). This would increase total manufacturing GDP by 2.8% and total GDP by 0.5%.

As emphasised, it is beyond the scope of this exercise to provide inferences about the relative contribution that adjustment costs may play in explaining the wage gap between casual and permanent manufacturing workers. Instead this exercise provides an upper bound on the gains from reallocation, under the assumption that the total wage gap is driven by adjustment costs. The lower bound is zero. Understanding the relative importance that adjustment costs play in impeding the movement of workers out of casual employment and into permanent positions remains an important area for future research.

Table 5.1: The Average Output Gains from Reallocation

	(1)	(2)	(3)	(4)	(5)
Naïve Gains	1.198	1.132	1.093	1.062	1.043
Labour Share ($(1-\alpha)$)	0.9	0.8	0.7	0.6	0.5

Notes: These estimate provide an upper bound of the static gains from reallocation under the assumption that the total wage gap between casual manufacturing workers and permanent manufacturing workers are driven by adjustment costs. The lower bound estimate of the static gains from reallocation are therefore zero.

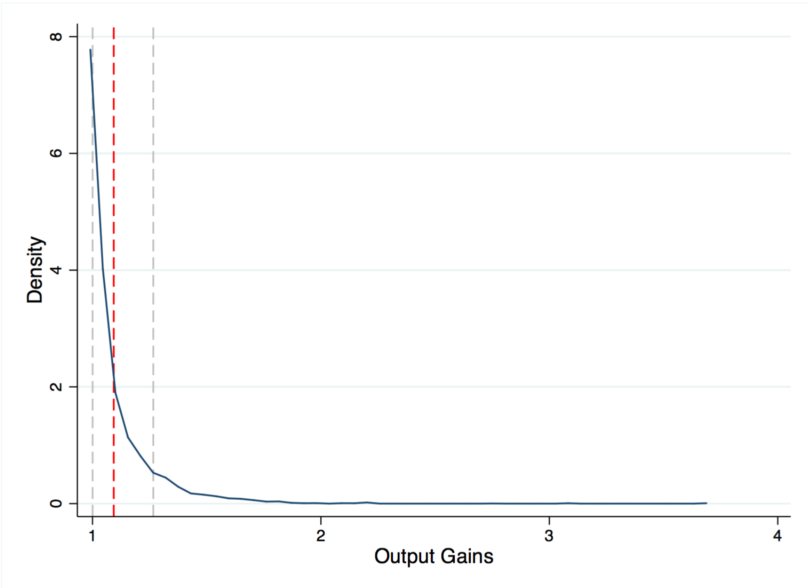

Figure 5.1: The Distribution of Output Gains from Reallocation: α = 0.3