A Weather Eye on Employment and Education:

Essays on Employment Polarization, Technology and Human Capital Formation

Sujaya Sircar

Thesis submitted to the Indian Statistical Institute

in partial fulfilment of the requirements for the degree of

Doctor of Philosophy

To my parents

A Weather Eye on Employment and Education: Essays on Employment Polarization, Technology and Human Capital Formation

Sujaya Sircar

July 2022

Thesis Supervisor: Dr. Tridip Ray

Thesis submitted to the Indian Statistical Institute

in partial fulfilment of the requirements for the degree of Doctor of Philosophy

Acknowledgements

I am extremely grateful to have had the opportunity to work under the guidance of my thesis supervisor, Dr. Tridip Ray. He has been a constant source of motivation and inspiration to me. I have learnt a lot from him over these past few years, both academically and otherwise. His openness to new ideas, his ability to critically analyze and his suggestions for complicated problems have always left me in awe. I am extremely fortunate to have had a supervisor who is a guardian figure to me and who has helped and supported me a lot during my most difficult times. I would like to humbly thank him for his immense contribution to my work.

My work would not have been possible without the help and support of Dr. Arka Roy Chaudhuri. He has been my teacher and has always had the patience and kindness to help me, even with the smallest of queries. I am extremely grateful that he took the time to go through the 'nitty gritties' of basic research work with me. His honest opinions and enthusiasm for my work have always motivated me, even during the most difficult phases. I feel extremely fortunate to have worked with his expertise during these past few years.

I would also like to thank the faculty members of the Economics and Planning Unit, Indian Statistical Institute, Delhi, who taught me and provided invaluable advice for my thesis. I am extremely grateful to Dr. Farzana Afridi, Dr. Chetan Ghate, Dr. Kanishka Kacker and Dr. Abhiroop Mukhopadhyay for their encouragement and interest in my work and for providing me with helpful suggestions. I would also like to extend my gratitude to Mr. Athisii who assisted me with my work. I am also extremely thankful to the ISI staff, mainly Mr. Shukla, Ms. Deepmala and Mr. Rajesh, and the mess and hostel cleaning staff. Without their work and responsibility, my time at ISI would not have been as easy or enjoyable.

I am fortunate to have met my senior and friend Komal di who has always been like an older sister to me. My seniors in ISI have always lent me a helping hand, especially Dhritiman da, Dyotona di, Prachi di, Gursharn di and Digvijay. Life would not have been as enjoyable on campus without their presence. I would like to specially acknowledge my dear friend and fellow researcher Ahana who has made work and campus life memorable. The healthy camaraderie among students and researchers in ISI have made my research experience cherishable. I have found dear friends in Nikita, Ojasvita and Saakshi. This list would not be complete without thanking my life-long friends who have always cheered me on: Adrija, Akanksha, Arya, Harshika, Isha, Prerna and Suvidya, my 'seven wonders'.

This work would have been incomplete without the blessings of my parents, Amitava Sircar and Sharmistha Sircar. Their love and support have made me what I am today and I am extremely grateful to them. My older sister, Sudeshna is not only an inspiration but she has also taught me to follow my dreams. She has been my strength during my most trying times. I am thankful to my brother-in-law, Sanjeeban Majumdar, who has always cheered me on with my work, and has always reminded me to take care of myself when I did not. I would also like to mention 'dimma', my late grandmother, Dr. Lakshmi Ghosh, the only grandparent I was old enough to remember. Memories of her and her life will forever remain an inspiration to me, a reminder to never give up even during the most difficult times.

Finally, I would like to thank my niece, Khushi, who like her name has always made me smile. I am thankful that she has been a constant source of joy for me and hope to be the same for her in the coming years.

Contents

Chapter 1

Introduction

This thesis is related to two widely discussed strands in economics literature: first, the relationship between technology and workforce and second, the impact of environmental factors on economic outcomes. With the rapid pace at which technology has evolved since the late 1990s, it is not surprising that several studies have investigated how the workforce composition has evolved (Katz and Murphy (1992), Acemoglu and Autor (2011), Goos et al. (2014)) and whether machines substitute workers or demand more of them (Acemoglu and Autor (2011), Autor and Salomons (2018), Aghion et al. (2020), Acemoglu and Restrepo (2020)), a hotly debated issue. At the same time there are numerous studies that look at how weather elements such as rainfall and temperature have affected economic outcomes like agricultural productivity, labour productivity, health and education (Jayachandran (2006), Colmer (2021), Somanathan et al. (2021), Park (2022)). These studies are particularly relevant in recent years as climate changes (such as global warming) have significantly increased in pace due to human activities like burning of fossil fuels. This thesis tries to provide an understanding of both these topics focusing in particular on the impact on employment and education.

The World War II saw a phase of rapid technological progress and, by 1970, microprocessors became widely available leading to computational tasks becoming economically feasible to perform. This allowed firms to substitute workers with computer machines in certain tasks (Acemoglu and Autor (2011)). Autor et al. (2003) point out that machines displace workers from occupations which carry out tasks that are relatively easier to code into machines. Middle wage occupations like manufacturing and sales are easier to code into machines unlike low wage occupations like janitors and high wage occupations like doctors. As a result, *job polarization* has occurred in several advanced economies: employment has decreased for middle wage occupations but increased for low and high wage occupations (Autor et al. (2006), Goos et al. (2014)). The first chapter of the thesis focuses on this phenomenon in the context of two developing economies – India and Indonesia.

Skill upgrading in developing economies may not occur in the same manner as in advanced economies as demonstrated by Berman et al. (2005) in the context of India. Similarly, while analyzing employment changes across occupations, studies on developing economies seem to produce a wide range of results (Kupets (2016), Ariza and Raymond Bara (2020), World Bank (2020), Cortes and Morris (2021)). We hypothesize that the presence of non-regular workers, a salient feature of developing economies (ILO (2016)), may be affecting the job polarization analysis in the context of developing economies. Our study accounts for this aspect and examines whether employment polarization has occurred in the developing economies of India and Indonesia.

While the first chapter looks at the evolution of workforce through a macroeconomic lens, the second chapter examines how hiring decisions at the level of firms are affected due to technology investment in the context of Indian manufacturing firms. The relationship between technology and demand for workers is a highly debated issue: do machines have a cost-saving or productivity improving effect leading to an increase in demand for workers, or a labour displacing effect leading to a decrease in demand for workers (Autor and Salomons (2018), Acemoglu et al. (2020))? What makes this question difficult to answer is that technology adoption is neither random nor exogenous since the decision to invest in technology is taken by firms. This has forced researchers to adopt quasi-experimental techniques in their analyses (Bessen et al. (2019), Aghion et al. (2020)). The second chapter of the thesis focuses on this question and exploits weather fluctuations (rainfall shocks) to provide a novel way to circumvent the endogeneity issue. This chapter also brings into focus how environmental factors like rainfall can affect economic outcomes like demand for industrial workers (Adhvaryu et al. (2013), Chaurey (2015)). We examine how hiring decisions of manufacturing firms are affected due to computer capital investments in the face of demand shocks.

The econometric method proposed in the second chapter (rainfall shocks leading to labour demand shocks) provides a natural segue to the next chapter of the thesis. This chapter falls under the broad research area of impact of climate on human capital formation. There have been studies that look at climate affecting economic outcomes indirectly, such as heat affecting crop yield which in turn affects human welfare (Heal and Park (2020)). However, recent studies have emerged focusing on the direct impact of heat on human physiology and cognitive performance, affecting health and labour productivity (Graff Zivin et al. (2020), Park (2022)). Global temperatures are rising, and are predicted to do so over the next century, making these studies increasingly important (Dell et al. (2012)). In the last chapter of the thesis I examine how temperature can affect human capital formation by studying the impact of exam time temperature on the high-stakes secondary exam scores of Indian school students.

The chapters are briefly described below where I discuss the research question, empirical strategy and the results of each study.

1.1 Evolving occupations or occupation-status? Employment polarization in the context of developing countries

Several studies on developed economies have pointed out that the labour market has undergone *job polarization*: a phenomenon in which employment hollows out for occupations in the middle of the wage distribution along with a concurrent rise at the two ends of the wage distribution (Autor et al. (2006), Goos and Manning (2007), Acemoglu and Autor (2011), Goos et al. (2014)). The most cited reason for this phenomenon is technology replacing occupations concentrated in the middle of the wage distribution, as middle-wage occupations tend to consist of more routine tasks which are easier to automate (Autor et al. (2003), Acemoglu and Autor (2011)). This chapter analyzes the phenomenon of employment polarization in the context of two developing countries – India and Indonesia, using data from the NSS Employment and Unemployment Household surveys for India and Indonesia Family Life Surveys for Indonesia. Following the analyses carried out for developed countries, we find, surprisingly, no evidence of job polarization in India and Indonesia for the periods 1983-2004 and 1993-2000, respectively.

The absence of job polarization may be due to organizational and institutional factors specific to developing economies, which could be very different from that of developed economies. A salient feature of developing countries is the significant presence of a large number of non-regular workers (Shyam Sundar (2011), Rothenberg et al. (2016), ILO (2016)), even leading to segmented labour markets within the same economy (Günther and Launov (2012)). To make our analysis reflect the labour market conditions of these developing countries, instead of documenting employment change across occupations, we consider *occupation-status* as the primary unit of analysis. For India we consider the worker statuses of regular, casual and self-employed, while for Indonesia the worker statuses are government, private and self-employed. Interestingly, we do find that employment decreases in the middle of the wage distribution across occupation-worker status categories while increases at the two ends of this distribution, a phenomenon we term as *occupation-status polarization*.

We establish that worker statuses are important not simply in terms of labour market conditions but also in terms of the skills required in an occupation. We apply a mixture model to our data to create heterogeneous skill groups among workers in India in 1983 and workers in Indonesia in 1993. The mixture model is a data-driven endogenous clustering technique that has the advantage of not specifying any cutoffs or terciles but requiring only the data to create heterogeneous classes (Günther and Launov (2012), Maitra (2016), Dell and Querubin (2018)). Following the job polarization literature, we use wages and years of schooling as indicators for skill, and create skill classes in the overall sample as well as within major occupation groups for India and Indonesia.

Finally we examine the association of worker statuses with these heterogeneous skill classes to establish that worker statuses are distinct skill groups within occupations. We find that, in both countries, workers with statuses of non-regular nature are relatively more likely to belong to lower skill classes and less likely to belong to higher skill classes within occupation groups. On the other hand, workers of regular status are relatively more likely to belong to higher skill classes and less likely to belong to lower skill classes within occupation groups. Our results suggest that skills differ across worker statuses even within similar occupations. This provides support that occupation-status is more appropriate as a unit of analysis while documenting employment polarization in developing economies.

1.2 Do clouds have a silicon lining for firms? Contract hiring and computer investment: Evidence from rainfall shocks

The role of technology in driving labour market changes has been well established in the literature (Tinbergen (1974, 1975), Katz and Murphy (1992), Autor et al. (2003), Acemoglu and Autor (2011)). But the literature provides mixed evidence on whether technology adoption leads to a rise or fall in employment. Firm-level studies have documented that technology adoption leads to both decreases in employment (Autor et al. (1998), Autor and Salomons (2018), Bessen et al. (2019), Acemoglu and Restrepo (2020)), as well as increases in employment (Van Reenen (1997), Blanchflower and Burgess (1998), Bessen (2019), Aghion et al. (2020)). In this chapter we look at how Indian manufacturing firms' contract worker hiring decisions are affected due to investment in computer capital in the face of demand shocks for the period 2000-2010. Documenting the relationship between technology and employment is a difficult exercise as technology investments are not random or exogenous shocks.

Our study circumvents this issue by exploiting rainfall shocks as proxies for demand shocks to study how hiring decisions differ across firms with varied computer capital investments.

Following Adhvaryu et al. (2013) and Chaurey (2015), we use data on rainfall shocks as industrial labour demand shocks, as rainfall affects agricultural productivity and yield which in turn leads to higher income and spending for agricultural workers. This in turn results in an increase in demand for industrial goods, leading to an increase in industrial labour demand (Chaurey (2015)). As rainfall shocks are transitory demand shocks, we focus on contract workers as they are not protected by labour regulations like the Industrial Disputes Act, 1947 (IDA) that restrict firms from firing workers. Since rainfall shocks are exogenous demand shocks, identification comes from the interaction of a computer capital investment measure and exogenous demand shocks (Nizalova and Murtazashvili (2016)). For measuring computer capital investment, we create a computer capital share dummy for firms that have above-average computer capital expenditure.

We find that firms with above-average computer capital share tend to hire 2.32 fewer contract workers compared to firms with lower-than-average computer capital share, in response to positive transitory demand shocks. This differential hiring is 11.33% of the sample mean of contract employment. In line with Chaurey (2015), we find that the results are driven by rural firms as the mechanism of rainfall shocks proxying for demand shocks work through changes in income and spending in the agricultural sector. We also find that our results are mostly driven by contract workers who are employed in the main manufacturing processes in the firm and not in peripheral works such as maintenance and security. We conclude that firms investing in computer capital may be shifting towards labour-substituting technologies, and therefore do not need to hire as many workers during transitory demand shocks. Our results are robust to controlling for agricultural inputs, state specific laws, state-industry specific factors such as tariffs, and a host of alternative computer capital investment measures such as the computer capital share of a firm, dummies for above-median computer capital share, industrial average computer capital share, industrial-year computer capital share, and a number of measures based on the US computer capital shares.

1.3 The Heat is on: Temperature and exam scores in India

While there has been an extensive research on how rising temperatures have caused adverse effects on economic outcomes through indirect channels (such as rising sea levels affecting infrastructure), a relatively recent literature has emerged studying the direct impact of temperature on economic outcomes (Heal and Park (2020)). Several studies have analyzed how rising temperatures affect economic outcomes such as health and labour productivity (Deschênes and Greenstone (2011), Somanathan et al. (2021)). This chapter adds to this recent literature investigating the impact of temperature on human capital formation by studying how temperature affects exam scores (Cho (2017), Graff Zivin et al. (2018), Tien Manh (2019), Conte Keivabu et al. (2020), Graff Zivin et al. (2020), Park (2022)).

We use secondary board exam scores from the administrative data of the Central Board of Secondary Education (CBSE) for the period 2012-2015. In India, the secondary board exam is an important exam as the scores potentially affect academic stream choices in higher secondary education, which in turn affect college or university admissions and hence future career paths of students. Also, secondary exam scores are given non-trivial weightage during the admission process of colleges and universities in India.

Our unit of analysis is student-subject standardized exam scores. As a student takes more than one subject exam at their secondary level, this allows us to exploit variation in district temperature across subject exam days for each student (allowing us to control for student specific characteristics). Identification comes from the exogeneity in district temperatures during exam time.

We find that one degree Celsius rise in temperature leads to a fall in standardized scores by 0.003 standard deviations, or a one standard deviation increase in temperature leads to a decrease in exam score by 0.016 standard deviations. We also find that temperature has a non-linear impact on exam scores: the marginal impact of temperature on exam scores increases with temperature. We find further that students located in rural areas or of disadvantaged castes perform relatively better when exposed to high temperatures (compared to urban or general caste students, respectively). Rural and backward caste students may be accustomed to living in uncomfortable temperatures (as they may not be able to afford cooling equipment like air coolers or air conditioners). Hence, our results point to students performing relatively better when they have adapted to heat. In line with the heat adaptation literature, we also find that students in very hot districts (above the 80^{th} percentile) perform relatively better than others (Cho (2017), Alberto et al. (2021)).

Finally, we provide evidence that temperature works through a physiological channel when affecting exam scores by examining whether the impact of temperature differs across male and female students. The literature provides evidence that physical differences between males and females can lead to thermoregulatory differences, with men being relatively more susceptible to thermal discomfort (Deschênes and Greenstone (2011), Mishra and Ramgopal (2013), Bai et al. (2014), Schweiker et al. (2018)). We find that male students perform relatively poorly compared to female students when exposed to higher temperatures. We also establish the impact of heat on cognitive ability working differentially for quantitative vis-a-vis language subjects. Interestingly, the effect of heat on cognitive ability is further amplified by physical heat stress for male students. Our explanation, following the relevant medical and economics literature (Hocking et al. (2001), Deschênes and Greenstone (2011), Bai et al. (2014), Schweiker et al. (2018)), is that female and male students have very different physiological aspects that play a vital role in regulating thermal temperatures of the human body. Thus, our results point to the direction that the physiological channel is at work.

Chapter 2

Evolving Occupations or Occupation-status? Employment Polarization in the Context of Developing Countries

2.1 Introduction

In an era of rapidly evolving technology, questions arise regarding the kind of occupations being replaced by machines. Studies in the context of the USA and other Western countries (Autor et al. (2003); Autor et al. (2006); Goos and Manning (2007); Acemoglu and Autor (2011); Goos et al. (2014)) have analyzed the phenomenon of job polarization. Job polarization entails employment hollowing out in the middle of the wage distribution along with its concurrent rise at the two ends of the wage distribution due to new technology replacing occupations concentrated in the middle of the wage distribution. This is possible because relative to occupations at either end of the wage distribution, middle-wage occupations tend to consist of more routine tasks which are easier to automate (Autor et al. (2003); Acemoglu and Autor (2011)).[1] In this chapter we explore the phenomenon of job polarization in the context of two developing countries – India and Indonesia.

Following the analyses carried out for developed countries, we find, surprisingly, no evidence of job polarization in developing countries like India and Indonesia for the periods 1983-2004 and 1993-2000, respectively. The absence of job polarization in India and Indonesia may stem from the organization of their labour markets, which is different from the labour markets of developed countries. In contrast to advanced economies, a salient feature in developing economies is the overwhelming presence of *non-regular* workers in the labour market.[2] Regular employment refers to all types of employment that have secure employment contracts and offer social security or protection after the termination of a job contract (ILO (2016)), unlike non-regular employment. When we analyze employment changes across occupations, we could fail to take into account that workers may face different labour market conditions or contracts within an occupation.

In the Indian and Indonesian labour market surveys, whether a worker is employed under a regular or non-regular contract is referred to as her employment *status*. To make our analysis reflect the labour market conditions of these developing countries, we consider *occupation-status* as the primary unit of our analysis.

[1] Trade and offshoring have also been analyzed as potential drivers of polarization, where the routine jobs are offshored to developing economies and employment in these jobs contract in the developed economies (Egger et al. (2016); Olsson and Tåg (2017)).

[2] See, for example, Shyam Sundar (2011) for India, Rothenberg et al. (2016) for Indonesia, and ILO (2016) for other developing countries.

We consider three types of worker statuses: regular, casual and self-employed for India, and government employee, private employee and self-employed for Indonesia. In India, regular workers are employed on a contract that does not require frequent renewals. Casual workers, however, do require periodic renewals of their contract with their employers. Self-employed workers, on the other hand, work on their own account and may even employ other workers. In Indonesia, while the government employees generally have a regular contract, the private workers are usually employed under a non-regular contract (Kwon (2013), Rothenberg et al. (2016)). Self-employed workers (with or without help) work on their own account, similar to the Indian self-employed workers. We rank occupation-status categories according to the average wages in these categories. Interestingly, we do find that employment is hollowing in the middle of this distribution while increasing at the two ends of the wage distribution across occupation-worker status categories. We refer to this phenomenon as *occupation-status polarization*.

In what follows we analyze this intriguing phenomenon of occupation-status polarization and the aspects surrounding it. We establish that worker statuses are distinct groups within occupations with the help of NSS Employment and Unemployment household surveys for India and the Indonesia Family Life Survey (IFLS) for Indonesia. We apply a *mixture model* to formally show that worker statuses are different, not just in terms of labour market conditions, but also in terms of skill.[3] The presence of heterogeneity across the worker statuses supports the need for a new unit of analysis for developing countries – *occupation-worker status*. In this study,

[3]Peirce (1884) introduced the latent structure of the mixture model to study the relationship between two observed dichotomous variables in the context of measuring the success of predictions (see, for example, Goodman and Kruskal (1959)). This model has been applied to several contexts like segmenting the informal sector in Ivory Coast (Günther and Launov (2012)), division of income classes in India (Maitra (2016)), grouping hamlets in Vietnam according to security (Dell and Querubin (2018)), among others.

our novel way of looking at skill distribution (across occupation-worker status) highlights the importance of paying heed to organizational and institutional differences in analyzing any labour market phenomenon such as job polarization in the context of developing countries like India and Indonesia.

In terms of worker statuses, labour markets in developed countries are fairly homogeneous as borne out by the relatively small non-regular employment shares in developed economies. Less than 20% of the workforce is self-employed in relatively high-income countries (Schuetze (2000), Kok and Berrios (2019)). Only 10% of the US employees were non-regular workers, while 2.5% of the UK employees were casually employed, a type of non-regular employment, by the end of 2015 (ILO (2016)).

However, there is a high degree of labour market segmentation in terms of worker statuses in developing countries as evinced by the very large share of non-regular employment. Over 50% of workers is self-employed in relatively low-income countries (Kok and Berrios (2019)). Self-employed workers are also known to dominate the unorganized sector in developing economies. Most self-employed workers are not covered by social security or formal retirement planning like regular workers. Low-income countries also have considerably higher shares of casual employment than in high-income nations (Aleksynska et al. (2016)). Casual workers require frequent renewal of contracts and receive little to no security benefits or employment protection (Hussmanns (2004), ILO (2016)).

In India and Indonesia, for the base years 1983 and 1993, respectively, we find that wages are different across worker statuses both in the overall labour market as well as within occupations. In India, within most occupations, wages are the highest for regular workers, followed by the wages of self-employed workers, while casual workers earn the lowest wages, on average. In Indonesia, on an average, government workers earn the highest wages, followed by private workers, and then by self-employed workers, within most occupations. Ranking of worker statuses based on average years of schooling is exactly the same as above in both countries.

To the extent wages and years of schooling reflect skill, the key to the task-based definition of occupations in Autor et al. (2003) and Acemoglu and Autor (2011), we consider occupation-worker status as the primary unit of skill distribution for India and Indonesia. That is, instead of ranking occupations according to the average occupation wages, we rank the occupation-status categories by the occupation-status average wages. And, as mentioned above, we demonstrate occupation-status polarization occurring in both India and Indonesia: employment is hollowing in the middle while increasing at the two ends of the wage distribution across occupation-worker status categories. We also decompose the percentile employment change into the employment changes of the three worker statuses in each percentile to explore how worker statuses are responsible for the employment changes. We find that, for India, occupation-status polarization is driven by a sharp decline in self-employment in the middle wage occupation-status percentiles, while an increase in self-employment and casual employment at lower percentiles and in self-employment and regular employment in higher percentiles. Occupation-status polarization in Indonesia, on the other hand, seems to be driven by an increase in government worker employment and self-employment at higher percentiles, a rise in self-employment at lower percentiles, and a decline in self-employment as well as private worker employment in the middle percentiles.

Research on changing workforce composition in developed countries like the US have focused on the importance of occupations, where occupations differ according to the tasks or activities executed within occupations (irrespective of the industry), hence the skill requirement too differs by occupations (Acemoglu and Autor (2011)). Our findings above indicate that in developing economies, with large shares of non-regular employment, skill may also vary by worker statuses within an occupation. However, unlike the US[4], there are no surveys or measures to show that skill or task requirement differs across occupation-worker status categories in developing countries. To show that the primary unit of the Indian and Indonesian skill distribution is at the level of occupation-worker status, we first formally establish that worker statuses and workers statuses within an occupation are heterogeneous in terms of skill.

We show that there are distinct groups of workers, in terms of skill, in our baseline data, that is, NSS 1983 data for India and IFLS 1993 data for Indonesia. We use wages and years of schooling as measures of skill (Acemoglu and Autor (2011)), to create these groups. We employ a finite mixture model, a probabilistic clustering technique, which allows us to identify latent clusters or groups of observations, where observations are homogeneous within the cluster but are heterogeneous across clusters in terms of some characteristic such as skill in our study. The advantage of the mixture model is that clustering is independent of any boundary criteria proposed by the researcher (such as wage quintiles), but depends solely on the data. For a mixture model, we do not need to define any cutoffs for wages or schooling years to observe if worker statuses are situated differently across these cutoffs to determine worker status heterogeneity. Thus the mixture model avoids any researcher subjectivity since we merely need to specify the skill indicators to identify the heterogeneous skill

[4]The US Department of Labor's Dictionary of Occupational Titles (DOT) and its successor the Occupational Information Network (O*NET) contains numerous task scales for occupations from its survey of US workers.

(latent) clusters.

So our first step is to use education and wages to create clusters in the 1983 Indian and 1993 Indonesian sample of workers using the mixture model as described above. The second step is to examine the relationship between these latent clusters and worker statuses. If each worker status is associated or correlated with a distinct latent class, this will reflect the skill heterogeneity across worker statuses. On the other hand, if all worker statuses are associated with only one latent class, then all the worker statuses are similar since they share the same traits specific to that latent class. We carry out a multinomial logistic regression of latent class on worker status which formally shows how being in a specific worker status might affect the chances of belonging to a latent cluster. We find evidence of significant associations of worker statuses with distinct classes. This result supports our hypothesis that skills are dissimilar across worker statuses, and it would be necessary to segment occupations by the same to study labour market phenomena like job polarization in the context of developing countries like India and Indonesia.

Since we study if occupations should be segmented by worker status, we also estimate the mixture model within five major occupation groups. This method identifies latent clusters within each of the five major occupation groups. This allows us to check if worker statuses are heterogeneous in terms of skill within an occupation. We carry out the same multinomial logistic regression described above and find that worker statuses are distinctly associated with the heterogeneous clusters formed within occupation groups.

This heterogeneity supports the need for segmentation of occupations by worker status. Our study shows that, instead of job polarization, occupation-status polarization is a far more appropriate phenomenon that needs to be studied for developing economies or economies which possess large shares of non-regular employment.

There are several recent studies confirming that developing economies may not exhibit job polarization in the same manner as developed countries, despite technology upgrading. For example, World Bank (2020) reports that China records an inverted U-shaped employment changes across jobs recorded in the Chinese census for the period 2000-2015, where middle-wage occupations see a rise in employment compared to low and high wage jobs. The authors point out that this reflects an expansion in the manufacturing sector due to outsourcing of manufacturing production from high-income economies in the 2000s. On the other hand, Cortes and Morris (2021), investigating whether offshoring from the US to Mexico is a potential cause for the decline in middle-wage employment in the US, find that most of the middle-skill occupations recording a fall in the US also experience a fall in Mexico in the 2000s. While the share of employment in non-routine manual jobs increases in both countries, non-routine cognitive tasks do not record a growth in Mexico, in contrast with the US experience. Ariza and Raymond Bara (2020) analyze the changing task compositions in urban labour markets in Mexico, Columbia and Brazil for the period 2002-2015 and find evidence of declining routine employment and increasing non-routine employment. However, the study needs to categorize routine cognitive jobs into two groups based on the ease of automation in certain jobs, and mainly focuses on results for the routine cognitive subset of jobs that are easier to automate. This categorization within routine cognitive jobs, the authors point out, is required for developing economies that have a large share of self-employed workers and street vendors. It is also worth noting that job polarization similar to a developed country may take place in a developing economy but technology and globalization may not be the only causes. In the case of Ukraine, Kupets (2016) records a decline in employment in middle-wage jobs but attributes this phenomenon to not just technology but also to de-industrialization, expansion of subsistence farming and other factors.

Studies on developing economies seem to produce a wide range of results when analyzing employment changes across occupations. We believe that this diversity in results may be due to the presence of different types of workers within occupations who possess distinct skills such as casual workers or self-employed workers or regular workers. Hence, change in employment for a job as a whole may not reveal information regarding the displacement of workers into different contracts or self-employment, possibly in the unorganized sector. Studies finding the absence of job polarization in developing economies may ignore such displacement and the consequent effects of changing worker status. This motivates our study to look at not just change in job composition but also the worker status composition within the labour market.

The rest of the chapter is organized as follows. In section 2.2 we describe the data sources for our analysis. Section 2.3 explores whether job polarization occurs in India and Indonesia and describes the type of employment polarization experienced by these countries. Section 2.4 describes the methodology of finite mixture models and how we apply it to our data. Section 2.5 presents evidence on the association of worker statuses and the (latent) skill clusters identified by the finite mixture model. Finally, section 2.6 discusses the results and concludes.

2.2 Data

India.— This study uses the NSSO (National Sample Survey Office) employment and unemployment surveys, which are quinquennial surveys that collect socio-economic information of individuals and follow a stratified multi-stage design. The surveys collect individual information such as educational qualification, age, religion, details of their household such as the number of household members, religion, migration details and so on. Other information such as whether the individuals are employed or not and the status of their employment are also recorded. Occupations are recorded according to the National Classification of Occupation (NCO) 1968. We use the 38th survey round (1983), the 50th survey round (1993-1994) and the 61st survey round (2004-2005) for this study. Details regarding employment such as the occupation and industry workers are employed in, the wages they have earned in the current week prior to the day of the survey, if they have any subsidiary employment, information regarding subsidiary employment and so on are collected.

Our sample is restricted to employed individuals. We look at individuals who have worked for at least 5 days a week; part-time workers are excluded from this study.[5] We consider workers who are between the ages of 15 and 65, and engaged in non-farm occupations. We obtain real daily wages by deflating nominal wages by the rural and urban state-level poverty lines (Eswaran et al. (2009), Hnatkovska et al. (2012)).[6]

Following Kijima (2006), we impute wages for self-employed workers. We obtain real daily wages by deflating nominal wages by the rural and urban state poverty lines. We try to improve upon Kijima's method by carrying out a stochastic regression imputation procedure (Little and Rubin (2019)) instead of the regression imputation procedure carried out in Kijima (2006). In case of a regression imputation procedure, wages for non-missing observations are regressed on a set of regressors. The coefficients obtained from the regression are then used to predict the missing

[5]The proportion of part-time workers in the labour market is roughly constant over the time period of our study, with 28% in 1983 and 25% in 2004.

[6]Real daily wages are reported in terms of rural Maharashtra's 1983 state poverty line.

wages, given the value of the regressors for the missing observations. In this method, the distribution of imputed earnings has a smaller variance than the distribution of actual values. A solution to this is to add variability to the predicted wages by carrying out a stochastic regression imputation procedure. In such a procedure, missing wages are predicted by the regression imputation plus a residual randomly drawn from a normal distribution with mean zero and variance equal to the residual variance in the regression (West et al. (1990); Little and Rubin (2019)). This helps in increasing the variability of the distribution, which is otherwise underestimated in a deterministic regression. In Appendix A2.1 we show that the distribution of wages is closer to the true distribution when employing a stochastic regression imputation procedure as opposed to a regression imputation procedure.

Table 2.1 reports the employment shares across occupation categories and employment shares across worker status categories in India for the years 1983 and 2004. For our analysis involving latent classes or skill classes, we use information from the 1983 survey round. We use information such as industry, gender, whether the worker was working in the rural sector and years of schooling. We follow Hnatkovska et al. (2012) to map education categories to years of schooling. Summary statistics of these variables are reported in Table 2.2.[7]

Indonesia.— For Indonesia, this study uses the Indonesia Family Life Survey (IFLS) which provides longitudinal data at the individual and family level on fertility, health, education, migration, and employment, assigning sampling weights to each observation. Occupations are recorded according to the International Standard Classification of Occupations (ISCO).[8] We use the first survey round (1993) and the

[7]The education categories are mapped in the following manner: illiterate: 0 years, pre-primary: 2 years, primary: 5 years, middle: 8 years, secondary and higher secondary: 10 years, above higher secondary: 15 years. All our results hold when we consider education categories instead of years of schooling.

[8]The ISCO version 68, which the IFLS follows, is very similar to NCO-68, since NCO is based

third survey round (2000) for this study. Details regarding employment such as occupation and wages are collected. Unlike the NSSO data, IFLS records the income of self-employed workers, hence no imputation is required for wages of self-employed workers. Real daily wages are reported in terms of Jakarta's 1993 consumer price index.

The sample is restricted to employed individuals. Part-time workers are excluded from this study, and we look at individuals who have worked for at least 5 days a week.[9] We consider workers who are between the ages of 15 and 65, and engaged in non-farm occupations (ISCO 1-digit occupation division '6', agriculture, is not considered.

Table 2.3 reports the employment shares across occupation categories and employment shares across worker status categories in Indonesia for the years 1993 and 2000. We use data from the 1993 IFLS round for our analysis with latent classes, mainly data on gender, whether a worker was employed in the rural sector and years of schooling.[10] Summary statistics of these variables are provided in Table 2.4.

2.3 Polarization in India and Indonesia

2.3.1 Job polarization

In this subsection we will check if the phenomenon of job polarization, as seen in developed countries (Autor et al. (2006), Goos and Manning (2007), Goos et al. (2014)), has occurred across India and Indonesia for the period 1983-2004 and 1993-

on ISCO descriptions.

[9] The proportion of part-time workers is 18% in 1993 and 25% in 2000.

[10] In Indonesian education system primary schooling includes 6 years of education, junior secondary includes 9 years of schooling, senior secondary includes 12 years of schooling and 15 years for education above senior secondary schooling (Zulfikar (2009), Kurniawati et al. (2018)).

2000, respectively. For these two countries, we first create country-specific skill percentiles based on the occupational mean wages in the base years of the study – 1983 for India and in 1993 for Indonesia (Autor et al. (2006)).

Following the studies on developed economies (Autor et al. (2006), Autor and Dorn (2013)), we next look at how employment shares across occupations have changed. To check for job polarization, we record the employment changes for the period 1983-2004 in India across the 1983 Indian skill percentiles in Figure 2.1 and for the period 1993-2000 in Indonesia across the 1993 Indonesian skill percentiles in Figure 2.2.[11] The skill percentile is measured on the x-axis which has occupations ranked by the average wages of workers in an occupation and the y-axis records employment changes. In Figure 2.1, we have also divided the changes in employment for the period 1983-2004 into employment changes for the periods 1983-1993 and 1993-2004. We find no evidence of job polarization in either country. Neither India, nor Indonesia experiences any hollowing out of employment across occupations which has been observed in developed countries. Even when we consider employment changes for the periods 1983-1993 and 1993-2004 separately for India, we find no evidence of job polarization.

India and Indonesia are developing economies and their experience may be quite different from the advanced economies. This result, the absence of job polarization, motivates us to look at how the organization of the labour market of these countries differs from the developed ones. One prominent aspect, as mentioned in the introduction, is the overwhelming presence of *non-regular* workers in the Indian and Indonesian labour markets, similar to many other developing countries. In what

[11] For a detailed description for the construction of skill percentiles, please refer to the Appendix A2.2.1. Employment change here refers to the percentile weighted employment change in each percentile. Please refer to the Appendix A2.2.2, for details on construction.

follows, we explore whether the presence of distinct worker statuses within the same occupations might shed some light on our results.

2.3.2 Heterogeneity across status in employment

For advanced economies, to check for job polarization, occupations are ranked by average occupation wages as wages reflect the skill of workers engaged in a profession (Goos and Manning (2007), Autor and Dorn (2013), Goos et al. (2014)). However, if an occupation employs very distinct classes of workers and employment shares of these distinct classes are not small, it would be more appropriate to compute the average wages of each worker class within the occupation. In relatively low-income countries, over 50% of the workforce is self-employed (Kok and Berrios (2019)). Shares of casual labour employment are also much higher in these countries in comparison with high-income countries (Aleksynska et al. (2016)). In this subsection we show that wages are different across employment statuses, not just in the overall Indian and Indonesian labour markets but also within occupations. We record our observations for the base year (1983 for India, 1993 for Indonesia), as the skill percentiles are constructed from the base year observations.

We study three worker statuses in India: regular, self-employed and casual. Regular workers are employed in farm or non-farm enterprises and receive wages or salaries on a permanent contract basis (the contract does not require periodic renewals). Casual workers, on the other hand, are paid on a temporary contract basis which requires renewals, periodically. Casual workers are also known to work in government-funded poverty alleviation programmes such as road construction and pond digging. Some workers may even engage in household production and receive wages accordingly. Self-employed workers are engaged in their own enterprises, either working on their own account, with a few partners or with hired workers. Self-employed workers form a major part of a developing economy labour market and are known to work largely in the informal sector (ILO (2016)).

IFLS 1993 and 2000 report 5 types of workers under employment status in Indonesia. These are self-employed without help, self-employed with temporary help, self-employed with regular workers, private workers and government workers. We have grouped together all the self-employed worker categories under self-employment status. We treat private workers, who are private company employees, as another worker status and government employees as a separate worker status. In Indonesia, the majority of non-agriculture enterprises in the private sector are small, medium and micro enterprises. The majority of employment in these private sector enterprises is non-regular in nature (Rothenberg et al. (2016), Chang et al. (2019)). On the other hand, government sector or state-owned enterprise workers are regular workers (Tambunan and Purwoko (2002), Kwon (2013)).

Given the nature of construction of the skill percentiles (ranking occupations by their average wage) in the existing literature on job polarization in the context of developed countries (Autor et al. (2006), Acemoglu and Autor (2011)) which we follow in Figures 2.1 and 2.2, we lose out information on possible change in employment share of workers across employment statuses within an occupation. There may be considerable dispersion of earnings within an occupation due to some worker statuses being paid more than others. In Table 2.5, we record the average log daily real wages in 1983 for India and in 1993 for Indonesia across 15 categories – 5 major occupation categories[12] and 3 worker statuses. It is interesting to note from Panel A of Table

[12]The 5 major occupation categories are created by grouping NCO/ISCO 1-digit categories in the following manner: 0-2: professional, technical and administrative, 3: clerks, 4: sales, 5: services, 7-9: production, operatives and labourers.

2.5 that for India, within most occupations, regular workers are the highest earners, followed by self-employed workers (except sales and services). Those engaged in casual employment seem to earn the least in all occupation categories in 1983. The average wage of a casual worker employed even under the category of professional, technical and administrative occupations (the highest paid occupations, on average) are less than the average wage of a regular worker employed under the category of low paid production, operatives and labourers. In our Indian data, on average, years of schooling attained by a regular worker is 7.21, while it is 3.96 for a self-employed worker and 2.22 for a casual worker.

From Panel B of Table 2.5 it seems that for Indonesia, private workers rank second in terms of wages among the employment statuses in every occupation category. The government workers seem to be the highest paid workers in all occupations, except for professional, technical and administrative workers. On the other hand, except for professional occupations, self-employed workers are the lowest paid workers in all occupations. Average wage of a government worker under the category of production, operatives and labourers (a low income occupation) is greater than that of a self-employed clerical worker (where clerks are the second highest paid workers in the economy), reflecting the dispersion of wages across worker statuses. Government workers are also the most educated among the three statuses, attaining, on average, around 12 years of schooling, followed by private workers (7.65) and self-employed workers (5.92).

In another way to demonstrate that wages are dissimilar across worker statuses, we consider the distribution of different worker statuses across the (log real daily) wage percentiles in 1983 for India and in 1993 for Indonesia. From the kernel density of wage percentiles for each employment status in India in Figure 2.3, we find that while most of the casual workers occupy the lower half of the wage distribution and regular workers the upper half, the self-employed workers are present throughout. The scenario is very similar for Indonesia (Figure 2.4): government workers dominate the top percentiles, self-employed workers have a relatively larger density towards the bottom percentiles, and the private workers dominate the middle and are distributed throughout. Thus the evidence shown so far point to a labour market structure in which the 1983 Indian workers and 1993 Indonesian workers form heterogeneous or distinct groups based on their worker status.

2.3.3 Occupation-status polarization in India and Indonesia

In this subsection we examine how employment changes occur across worker statuses within occupations. For India, instead of ranking the NCO 3-digit occupations according to the average 1983 wage of each occupation, we divide each NCO 3-digit occupation into the three worker statuses – regular, casual and self-employed. We record the average 1983 log daily real wages of these occupation-status categories and rank them to construct the skill percentiles based on the average wages of these categories. The skill percentiles thus constructed will reflect the dispersion of wages across worker statuses within an occupation. For Indonesia also, instead of ranking occupations, we divide the ISCO 2-digit occupation categories into self-employed, private worker and government worker statuses and rank the occupation-status categories by the 1993 occupation-status average wages.

After ranking occupation-status categories, we look at the fraction of employment of each worker status across the skill percentiles for both countries. This exercise will let us observe if specific parts of the occupation-status skill percentiles employ more of a certain worker status. We record the employment shares of worker statuses within each percentile in Figure 2.5 for India and Figure 2.6 for Indonesia.[13]

For India, regular workers seem to dominate the higher percentiles implying that they receive relatively higher wages than self-employed and casual workers. Casual workers, on the other hand, seem to be located across the lower half of the skill percentiles, indicating that casual workers are the lowest earners in most occupations in 1983. Self-employed workers seem to be distributed throughout but mostly in the lower and middle percentiles in 1983. For Indonesia, government workers are mostly high earners as seen from their position in the highest tercile suggesting government workers receive higher wages among the three statuses. Self-employed workers dominate below the 50th percentile, with some private workers at the lowest percentiles, implying that most self-employed workers are low earners. Private workers are also present at the 60th to 78th percentiles, which are the lower earning percentiles of the highest tercile.

Next we carry out the same exercise as in Section 2.3.1, but record employment change across the occupation-status skill percentiles, instead of only across occupation skill percentiles.[14] The employment change at each percentile is recorded for the period 1983-2004 in India in Figure 2.7 and for 1993-2000 in Indonesia in Figure 2.8. For India, we have again divided the 1983-2004 employment changes into employment changes for the periods 1983-1993 and 1993-2004. Interestingly, for India there is a significant decline in the middle percentiles with an increase at the two ends, implying a U-shaped curvature for employment change for the period of 1983-2004. It is

[13]In each percentile, employment shares of occupation-status categories is summed up for each status. For example, in India, the employment share of casual workers is calculated by summing all the occupation-casual employment shares within a percentile. Similar calculations are carried out for obtaining the regular and self-employment shares within each percentile. All the shares are expressed in fractions within a percentile. The 3 worker status employment shares thus computed will add up to 1 within a percentile. The same exercise is carried out for the Indonesian worker statuses.

[14]For a detailed technical explanation, refer to the Appendix A2.2.3.

also interesting to note that the U-shaped curvature is seen for the period 1993-2004 and not for 1983-1993. This would imply that the U-shaped employment change seen in India for the period 1983-2004 is mainly being driven by the employment changes for the period 1993-2004. Indonesia, too exhibits a similar phenomena. We see an increase in employment at the two ends of the skill percentiles (a larger increase at higher percentiles) and a decline in the middle, exhibiting a U-shaped curvature. This result confirms that the Indian and Indonesian experiences are very different from the advanced countries. There is no job polarization, as seen in Figures 2.1 and 2.2 in Section 2.3.1. Instead, polarization occurs across *occupation-status*: low and high wage occupation-worker status categories experience a rise in employment along with a decline in middle wage occupation-worker status employment.

Next we record how worker statuses are responsible for the employment changes for India and Indonesia across the 1983 and 1993 skill percentiles, respectively. For this we decompose the percentile employment change into the employment changes of the 3 worker statuses in a percentile. This allows us to inspect which status (or statuses) is responsible for the overall employment change in a percentile for the observed period. It also allows us to see where each worker status expands or contracts over time across the skill percentiles. We record the employment change for each worker status separately across the percentiles.[15] The decomposition results are presented in Figures 2.9 and 2.10 for India, and in Figure 2.11 for Indonesia.

It is interesting to observe from Figure 2.9 that non-regular employment seems to be playing a major role in employment polarization for India. The growth in casual employment is in line with the casualisation of Indian workforce pointed out in the literature (Pais (2002), Goldar and Aggarwal (2012)). Occupation-status

[15]For a detailed technical explanation, refer to Appendix A2.2.3.

polarization in India is characterized by an increase in self-employment and casual employment at low wage occupation-status percentiles, a decline in self-employment in the middle percentiles, and a rise in self-employment and regular employment in higher percentiles.

For Indonesia, Figure 2.11 demonstrates that there is a large increase in government workers which is resulting in the increase in employment at higher percentiles. Occupation-status polarization in Indonesia seems to be driven by an increase in government worker employment and self-employment at higher percentiles, a rise in self-employment at lower percentiles, and a decline in self-employment as well as private worker employment in the middle percentiles. Interestingly, self-employed workers in Indonesia themselves experience polarization, increasing especially at the lower end.

The occurrence of occupation-status polarization is intriguing and brings up a new question: as the significant presence of non-regular employment in developing economies is captured by worker status, should we consider skill at the level of occupation *and* worker status and not just at the level of occupation for these economies? If skills do differ across not just occupations but also across the employment status of workers, it would be prudent to consider employment status for polarization research in developing economies. While the above results are interesting, considering occupation-worker status as the relevant unit of analysis for employment polarization in developing economies need further investigation. In the next sections we explore if skill is indeed heterogeneous across worker statuses and even within occupations, necessitating the need for occupation-status as the unit of analysis.

2.4 Finite mixture model

As shown in the previous section, we are able to observe employment polarization in India and Indonesia at the level of occupation-status categories and not at the level of occupations. This may be because occupation is the relevant unit of skill for developed countries when creating skill percentiles. For developing countries, however, workers with different employment statuses may have distinct skills, even when they are employed within the same occupation. Hence, when the skill percentiles are constructed with occupation-status units, we are able to observe employment polarization (and not when we construct the skill percentiles at the level of occupations). It may be the case that we are able to record movements of workers with heterogeneous skills correctly when we consider this nuanced categorization. Hence, our goal is to show that employment statuses are indeed heterogeneous in terms of skill, even within occupations. To do this we employ a finite mixture model.

Mixture models identify clusters or sub-populations with respect to observed variable in the data, such that the density of the variable will be equal to the weighted sum of the densities of these sub-populations (McLachlan et al. (2019)). In mixture models, clusters are termed as latent classes. These models rely only on observed data with minimal assumptions to create clusters and do not require any pre-specified boundaries.[16]

The ability to create latent classes without *a priori* boundaries has useful applications. Dell and Querubin (2018) applies a mixture model to cluster hamlets in Vietnam into security classes. Maitra (2016) creates three latent classes of con-

[16]Explicit cutoffs or divisions include grouping workers by specific criteria. For example, workers divided into three groups where the first group lies below the 33rd wage percentile, the second group between 33rd and 66th percentiles and the third group forms above the 66th wage percentile. The econometric technique of mixture models require no such percentile limits but only wage observations.

sumption in India based on the consumption of durables, which she identifies as low-, middle- and high-income classes and confirms rising poverty in the 1990s. She requires no income cutoff or poverty lines to distinguish the three income classes but only data on durable consumption. Using a finite mixture model, Günther and Launov (2012) divides the informal labour market of Ivory Coast into two latent groups. Again, no wage cutoff is needed to create the two informal sector groups. Paap et al. (2005) clusters countries based on GDP per capita and not according to any geographical location to analyze if sub-Saharan African countries have low growth rates compared to countries located in other continents. The authors use national income to cluster countries and check which African nations fall in the low-income group. Similar to these papers, we form clusters of workers based on observed variables, and do not specify any wage cutoffs to divide either the 1983 Indian sample or the 1993 Indonesian sample into groups. We create clusters or latent classes without any specific classification such as worker status categories, nor do we need to define any pre-specified cutoffs for our mixture model.

In Section 2.3.2, we compared average wages for different employment statuses within occupations. One can do this in various ways such as comparing worker status employment densities below and above the mean sample wage, or comparing across wage quartiles. However, this brings the researcher's subjectivity to the mix. The mixture model circumvents this problem as we merely need to specify the skill indicators to create heterogeneous groups of workers in our data. Thus, any heterogeneity concluded in our mixture model analysis will depend only on the data.

In this study, we use a two-step process to identify skill classes. The first step includes estimating the parameters of the distribution for each latent class, followed by the second step which involves classifying observations into the latent classes based on Bayesian posterior probability. We use the same variables that are generally used to create the skill percentiles for analyzing job polarization in developed countries – wages and years of schooling (Acemoglu and Autor (2011)) to classify workers from the 1983 Indian and 1993 Indonesian labour market into skill classes. The next section describes the estimation method for the mixture model.

2.4.1 Estimation

In a mixture model, the density of observed variables y (wages and education in our study) is the weighted sum of class specific densities of y. Here weights describe the probability of an observation belonging to class t (t = 1,..,T) which is denoted by π_t and

$$\sum_{t=1}^{T} \pi_t = 1. \tag{2.1}$$

The probability of obtaining observation y in a sample (for a 3-class model or $T = 3$) is given by

$$f(y) = \pi_1 f(y|\theta_1) + \pi_2 f(y|\theta_2) + (1 - \pi_1 - \pi_2) f(y|\theta_3). \tag{2.2}$$

The set of parameters θ_t are the parameters to be estimated for the distributions of each sub-population t. We can repeat this equation for a 2-class model or a 4-class model by removing or adding the appropriate number of classes.

Our *base* specification is described by equation (2.2) where, for India, we consider y[17] as log daily real wages (following a normal distribution) and years of schooling (following a Poisson distribution) to estimate the latent classes. For Indonesia too,

[17]The distribution $f(y)$ here refers to the joint distribution over wages and years of schooling. In mixture models, more than one indicator may be used. For ease of exposition, we simply denote the joint distribution as $f(y)$.

we consider log daily real wages and years of schooling (both following normal distributions) for estimating the latent classes.[18]

We also create latent classes within five major occupation groups using the same specification (equation (2.2)). Thus, we estimate θ_t for all latent classes within major occupation groups for both India and Indonesia. This model considers heterogeneity within occupation groups: equation (2.2) must be estimated across all the five major occupation group subsamples, simultaneously. Estimating equation (2.2) both for the entire sample and also for the major occupation groups separately are important as we are trying to observe if employment statuses are heterogeneous, not just across statuses but also within occupations.

A typical issue of mixture models is observational equivalence. There is no difference between a mixture model that has a parameter vector ordered $(\pi_1, \pi_2, \theta_1, \theta_2)$ and a mixture model with a parameter vector ordered $(\pi_2, \pi_1, \theta_2, \theta_1)$. However, given that our latent classes are skill classes, we can identify the lowest skill class as the latent class with the lowest expected wages and years of schooling and the highest skill class as the latent class with the highest expected wages and years of schooling. Thus, we can order the latent classes by the estimated parameters in each class, which allows us to identify the skill classes.[19]

Now, we need to estimate the parameters θ_t and π_t for class $t = 1, ..., T$. To estimate the parameters for a mixture model, we maximize the log-likelihood function. However, closed-form expressions to estimate parameters for all the latent classes is very difficult with a one-step differentiation. We apply the Expectation

[18]In 1983 India, 34% of the sample report 0 years of education. We consider both normal and Poisson distribution for schooling years and find that the Poisson distribution is a better fit for the data. For selecting the distributions for years of education for India and Indonesia, please refer to Appendix A2.4.

[19]Maitra (2016) uses a similar technique to identify lower, middle and upper consumption classes.

Maximization (E-M) algorithm to resolve this issue.

2.4.1.1 Expectation Maximization (E-M) Algorithm

This section motivates the use of the E-M algorithm. Let us consider a mixture model with variable y_i following a normal distribution (parameters $\mu, \sigma \in \theta$) as an example. Our likelihood function (for the mixture model) is given by

$$L(\theta, \pi | y_1, ..., y_n) = \prod_{i=1}^{n} \sum_{t=1}^{T} \pi_t f(y_i|\theta_t) \tag{2.3}$$

which can be written in log-likelihood form

$$l(\theta, \pi | y_1, ..., y_n) = \sum_{i=1}^{n} log(\sum_{t=1}^{T} \pi_t f(y_i|\theta_t)). \tag{2.4}$$

However, we cannot estimate θ_k with a simple one-step differentiation. One can see this issue in equation (2.4) where there are numerous unknown parameters: π_t in addition to θ_t, and both π_t and θ_t have to be estimated for more than one class. When we maximize the log-likelihood function for θ_k (the parameter θ for class k and $k \in [1, T]$), obtaining a closed form solution is difficult from the first order condition:

$$\sum_{i=1}^{n} \frac{\partial log(f(y_i|\theta)}{\partial \theta_k} \left[\frac{1}{\sum_{t=1}^{T} \pi_t f(y_i|\theta_t)} \pi_k f(y_i|\theta_k) \right]. \tag{2.5}$$

We now look at what our results yield from the one-step differentiation above. Other than the fact that we are not able to estimate θ_k, we obtain a term $\frac{1}{\sum_{t=1}^{T} \pi_t f(y_i|\theta_t)} \pi_k f(y_i|\theta_k)$ in equation (2.5). This is the posterior probability (p_{ik}) of an observation i belonging to latent class k, which is the probability evaluated *after* we gain information on parameters. One can see the obvious dilemma here. If we could estimate the posterior probability term, then we can easily compute the parameters from maximizing the log-likelihood function. Again, if we knew the parameter estimates, we can compute the posterior probability term. The EM algorithm essentially helps overcome this issue. We first start with an initial guess of values for the parameters.

The E-step: We calculate the posterior probability (p_{it}) of each observation i belonging to each class t, given the current set of parameters.

The M-step: Given the posterior probability, we maximize the log-likelihood function and obtain new parameters for π_t and θ_t.

After the M-step we go back to the E-step. The iterative procedure continues until the parameter values converge. After convergence, we obtain estimates for π_t (class weights), θ_t (parameters for each class specific distribution) and p_{it} (posterior probability of observation i belonging to class t).

While we have used the normal distribution as an example here, the issue of closed form solutions in mixture models is persistent for all types of distributions.[20]

2.4.2 Estimates

We first choose the optimal number of latent classes for each specification with the help of information criteria (Akaike Information Criteria (AIC) and Bayesian Information Criteria (BIC)).[21] In Appendix A2.3, we show that 3 classes are optimal for the Indian data as well as for the *base* specification for Indonesian data, whereas 4 classes are optimal for the *within occupation groups* specification for Indonesia.

[20]See for example Maitra (2016) who uses a binomial distribution for constructing latent classes and applies the EM algorithm for a closed form solution.

[21]Leroux (1992) establishes that information criteria such as AIC and BIC consistently estimate the number of classes in a mixture model. Günther and Launov (2012) chooses the number of latent classes in the informal sector of Ivory Coast with the help of information criteria.

Table 2.6 reports the estimates of each latent class constructed in the overall sample (columns 1-3) and within occupations (columns 4-6) for India (Panel A) and Indonesia (Panel B).

For India, we find that the expected log wages and years of schooling increases from latent classes 1 to latent class 3. As the expected wages and years of schooling increase over the three latent classes, we can say that latent class 1 is the low skill class given it has the lowest expected years of education and log daily wages among the three classes, followed by the middle skill class (latent class 2) and finally latent class 3 is the high skill class since it has the highest expected log wages and schooling years among the three. Similarly, for Indonesia too the expected log daily wages and years of schooling increase as we go from the latent class 1 to latent class 3 (or latent class 4 for within occupations latent classes). Thus, the lower the number of the latent class the lower is the skill.

We next calculate the posterior probabilities of the three (or four for *within occupation group* for Indonesia) latent classes for each observation and classify each observation to the latent class which has the highest posterior probability. Posterior probability for observation i and latent class k is given by $p_{ik} = \pi_k f(y_i|\theta_k) / \sum_{t=1}^{T} \pi_t f(y_i|\theta_t)$. The estimated parameters (θ_t) are reported in Table 2.6. We classify observations according to which latent class the observation is most likely to belong to, given the parameters. Thus, for each observation, we are able to create a variable (predicted class) that takes value 1, 2 or 3 (or 4), depending on which latent class records the maximum posterior probability. The shares of the three (or four) predicted classes in the sample across the two specifications are reported in Table 2.7.

As all the observations are now classified into one of the skill classes, we can now check if a worker is more (or less) likely to belong to a specific skill class given his worker status. If worker statuses exhibit heterogeneity in skill, we should expect different worker statuses to have different probabilities of belonging to a skill class. In the next section, we move on to the final step where we estimate the association of worker status with skill classes.

2.5 Association of worker status with latent classes

In this section we check if being a certain worker status affects the probability of belonging to a skill class. For example, in the Indian case, if casual workers are low skilled workers, they will be relatively more likely to belong to latent class 1 or the low skill class, compared to other worker statuses. If regular workers are high skilled workers, they will be relatively more likely to belong to the high skill latent class. Also, if regular workers are less likely to belong to the low skill class, this will confirm that regular and casual workers have different skills.[22]

We employ a multinomial logistic regression and estimate the coefficients of the model (using maximum likelihood) to determine the association of worker status with skill class.[23] The probability that observation i is in skill class t is given by:

$$Pr(Class_i = t) = \begin{cases} \dfrac{1}{1 + \sum\limits_s \exp(X_i \beta_s)} & \text{if t=2} \\ \dfrac{\exp(X_i \beta_t)}{1 + \sum\limits_s \exp(X_i \beta_s)} & \text{if t=1, 3, 4} \end{cases} \tag{2.6}$$

[22]Please refer to Appendix A2.5 for details on a preliminary check.

[23]The multinomial logistic model is typically used in the three step procedure for obtaining estimates for association in latent class and mixture model analysis (Vermunt (2010), Bakk and Kuha (2018)) and also to avoid computation and convergence difficulties (Compton and Pollak (2007)).

where $t = 1, 2, 3, 4$ refers to the skill classes (lower number refers to lower skill and s is casual and regular for India, and self-employed and government for Indonesia.[24] Our reference category is latent class 2 which is also the middle skill class. This allows us to easily interpret our results. X_i includes worker status dummies and controls such as gender, rural dummies for Indonesia and also industry categories for India.[25] For both countries, the base dummy is the worker status group with the largest share of employment among all the worker statuses. For India, the base group for worker status is self-employed. Thus, we check if being a casual (or regular) worker increases or decreases the chance of being part of latent class 1 (or 3) rather than latent class 2, compared to self-employed workers. For Indonesia, the base group of worker status is private worker. Thus, we report the chances of a worker being part of latent class 1 (or 3) versus latent class 2 when she is a self-employed (or government) worker as compared to a private worker.

We have classified observations into the latent classes by selecting the class that reports the highest posterior probability for a worker, as discussed in Section 2.4.2. For ease of interpretation, we report our results in terms of relative risk ratios which is computed by exponentiating the estimated coefficients. The relative risk ratio implies the change in probability that an observation will be in skill class t, relative to the probability of being in skill class 2, when the observation is of a specific worker status compared to the base worker status.

Results are reported in Table 2.8. The *base* and *within occupation* latent classes in Table 2.8 are the same as the latent class construction in Table 2.6: *base* refers to latent classes formed in the entire sample and *within occupation* refers to latent

[24] t takes values 1 to 3 for latent classes in India in the overall sample and within occupations and for Indonesia in the overall sample. t takes values 1 to 4 for Indonesian latent classes within occupation groups.

[25] Industry data is not translated in English for IFLS 1993.

classes formed within 5 major occupation groups in the sample.

For India (Panel A, column 1), we find that being a casual worker increases the probability of being in latent class 1 or the low skill class instead of latent class 2 or the middle class. A casual worker is 1.30 times more likely to belong to the low skill class. Also, a casual worker is less likely to belong to latent class 3 or the high skill class. Results are similar when we construct latent classes within occupation groups. When adding controls, the coefficient becomes smaller for a casual worker for latent class 1 (0.91), but is close to 1 and still greater than that of latent class 3 (0.32). Interestingly, we find that being a regular worker increases the probability of belonging to latent class 3 or the high skill class (2.60 times more likely to belong to latent class 3 than 2). A regular worker is less likely to belong to latent class 1 or a low skill class than class 2 (only 0.58 times likely to belong to latent class 1 than latent class 2). Similar results hold for latent classes constructed within occupations and even with controls (the magnitude of coefficients decrease when adding controls). All results are statistically significant.

Given our results, we can say that casual workers are less likely to belong to a skill class that a regular worker is more likely to belong to. This can be easily seen from how being a casual worker increases the probability of being in a low skill class (which a regular worker is less likely to belong to) and decreases the probability of being in a high skill class (which a regular worker is more likely to belong to), compared to a middle skill class. Opposite associations of regular and casual workers to latent classes or skill classes imply that the worker statuses have very different skills, even within occupations.

For Indonesia (Panel B), being a self-employed worker increases the probability of being part of latent class 1, compared to latent class 2 (although this result is weak with controls). When a worker is self-employed, he is 1.76 times more likely to belong to latent class 1 (column 1). Self-employed workers are less likely to belong to the higher skill classes (for overall and within occupation latent classes). The magnitude of coefficients increases from latent class 4 to 1, implying that self-employed workers are more likely to belong to the lower skill classes in Indonesia and less likely to belong to higher skill classes. In contrast, the coefficients increase in magnitude from latent class 1 to 4 for government workers. For example, being a government worker, a worker is 12.22 times more likely to belong to the highest skill class or latent class 4 within occupations, but only 0.21 times likely to belong to the lowest skill class or latent class 1 (column 4).

As in the Indian case, our results for Indonesia have similar implications: being a government worker decreases the probability of belonging to relatively low skill classes, whereas being a self-employed worker increases the probability of belonging to relatively low skill classes. The converse also holds. Both government and self-employed workers have different associations with different skill classes.

From both the Indian and Indonesian cases we show that the probability of belonging to a specific skill class differs across worker statuses. As worker statuses have different associations to different skill classes, this points to worker statuses sharing unobserved heterogeneity in skill. Our results are obtained both for skill classes in the overall workforce as well as within occupation groups. Thus, distinct associations are present even after considering latent classes within occupation groups, which may be interpreted as worker statuses being heterogeneous groups, even within occupations. This provides support that occupation-status as a category (and not just occupation) should be used to construct skill percentiles, at least for the developing countries.

2.6 Conclusion

Advanced nations have reported the presence of job polarization in the labour market, exhibiting a rise of employment in low- and high-skilled occupations accompanied by a contraction of employment in middle-skilled occupations (Autor et al. (2006); Goos and Manning (2007); Acemoglu and Autor (2011); Goos et al. (2014)). The most cited reason for this occurrence is automation, which displaces workers from middle-wage occupations. The task theory, which introduced the concept that occupations consist of tasks or activities, was proposed by Autor et al. (2003) to explain the polarization phenomenon. Automation replaces middle-skill workers that carry out routine tasks but not workers engaged in non-routine tasks, as programming non-routine tasks is difficult.

Following advanced economy studies, we investigate whether job polarization has occurred in two developing economies – India and Indonesia. We find that for the period 1983-2004 and 1993-2000, no such phenomenon occurred in India and Indonesia, respectively. In the emerging literature, several studies have shown that developing economies do not exhibit job polarization in the same manner as developed economies and provide a wide range of results (Kupets (2016), Ariza and Raymond Bara (2020), World Bank (2020), Cortes and Morris (2021)). We hypothesize that this absence of job polarization in developing countries may be due to organizational factors specific to the developing economies. We take into account the presence of non-regular workers, a prominent part of the workforce in several developing countries, and consider the presence of different worker statuses in an occupation. Worker statuses capture the non-regular nature of work. For India, we analyze self-employment, regular and casual status. Self-employment and casual workers make up more than 50% of the Indian workforce (Shyam Sundar (2011)). The Indonesian surveys in our study report self-employment status and whether a worker is a private or government employee as the other employment statuses.

Instead of employment polarization across occupations, we find evidence of employment polarization across occupation-status categories. In India, non-regular work seems to be driving the polarization and Indonesian occupation-status polarization seems to be driven by self-employment. This throws up questions regarding the unit of analysis for employment polarization in developing economies, as we may be incorrectly considering that non-regular and regular workers possess the same skills within an occupation. Since we find evidence of polarization not across occupations but across a more nuanced category of occupations, we establish that occupation-status is an important unit of analysis.

We apply a finite mixture model to prove heterogeneity in skill across worker statuses. This econometric technique does not need to rely on cutoffs such as wage percentiles or terciles to prove that worker statuses are heterogeneous in terms of skill, as such cutoffs are subjective to some extent. Mixture models create heterogeneous clusters (known as latent classes) in the data based on observed variables that segment the data. We use indicators, mainly wages and years of education which are typically used to represent worker skill (Acemoglu and Autor (2011)), to create heterogeneous skill clusters in the data. We determine the heterogeneous skill classes in our data and then check how worker statuses are associated with these skill classes (Compton and Pollak (2007), Vermunt (2010), Bakk and Kuha (2018)).

If worker statuses are significantly associated with distinct skill classes, this would tell us that worker statuses are very different from each other due to some underlying heterogeneity they share with these classes. The greatest advantage of this method is that we let the data 'speak for itself'. We create skill classes for the overall sample and also within major occupation subsamples (this considers skill heterogeneity within each occupation group). In India, casual workers are relatively more likely to belong to the low skill class and regular workers are relatively more likely to belong to the high skill class. All the results are relative to the middle skill class and self-employed workers. It is interesting to note that casual workers are relatively less likely to belong to high skill classes and regular workers are relatively less likely to belong to low skill classes. Similarly, in Indonesia, self-employed workers are relatively more likely to belong to the low skill class but less likely to belong to the high skill class. Government workers on the other hand are relatively more likely to belong to the high skill class but less likely to belong to the low skill class, the reference group being private workers. Our results point to worker statuses exhibiting skill heterogeneity within major occupation categories, which implies that worker statuses seem to have very different skills, even within the same occupation.

Our results from the analysis with the mixture model shows that worker statuses are heterogeneous in terms of skill, which support the need to segment occupations by worker status when studying employment polarization. The unit *occupation-status* becomes a necessary categorization when analyzing employment polarization in developing countries like India and Indonesia, as this helps to analyze possible employment polarization across a relatively more accurate skill spectrum. This novel unit of analysis takes into account a salient feature of developing countries – the presence of non-regular workers or contracts, and may help to explain the diverse set of results with respect to job polarization observed in studies on developing countries.

Our results motivate further research on how employment polarization occurs for developing countries. Berman and Machin (2000) and Berman et al. (2005) point out that skill upgrading in India is not similar to those of developed countries for the period we examine in this study. The Indonesian workforce experienced a labour shift away from agriculture but over a long period of time (Salam et al. (2018)). For the period we study, there were substantial gains in agriculture labour productivity but not for other sectors (Salam et al. (2018)). More research is required as to how the employment polarization is taking place. Studies such as World Bank (2020) and Cortes and Morris (2021) also point out how factors other than technology (such as trade and offshoring) may play a role in employment polarization. It is important to recognize labour market segmentations, both non-frontier technology and automation, and trade-related aspects when studying how employment polarization may be taking place in developing economies.

Figures and Tables for Chapter 2

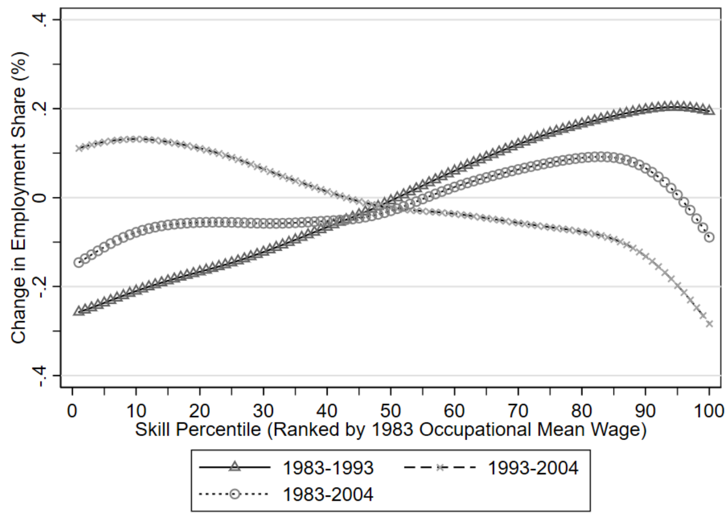

Figure 2.1: *Change in employment by occupation, India 1983-2004*

Note: The lines represent locally weighted scatterplot smoothing of the employment change for the period across occupations. Data sources are NSSO employment and unemployment surveys (1983, 1993-1994, 2004-2005).

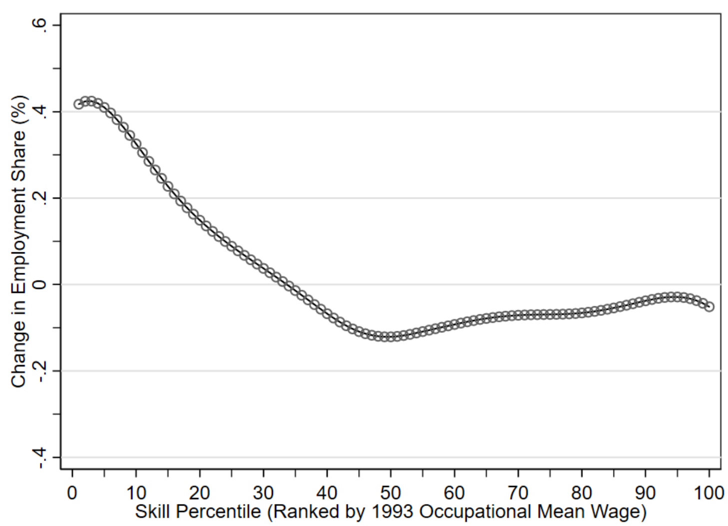

Figure 2.2: *Change in employment by occupation, Indonesia 1993-2000*

Note: The lines represent locally weighted scatterplot smoothing of the employment change for the period across occupations. Data sources are IFLS surveys (1993, 2000).

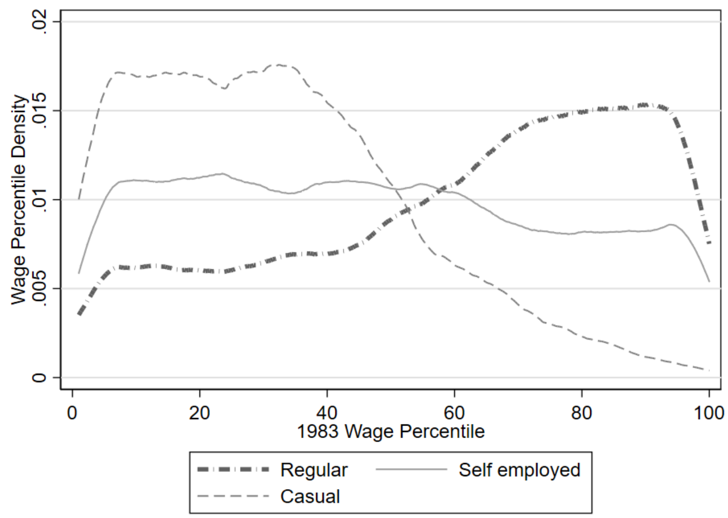

Figure 2.3: *1983 wage density of worker statuses across Indian wage percentiles*

Note: Data source is NSSO employment and unemployment survey (1983).

Figure 2.4: *1993 wage density of worker statuses across Indonesian wage percentiles*

Note: Data sources is IFLS survey (1993).

Figure 2.5: *Employment share of worker status across occupation-status, India 1983*

Note: Employment shares of worker statuses add up to 1 within each percentile. Data source is NSSO employment and unemployment survey (1983).

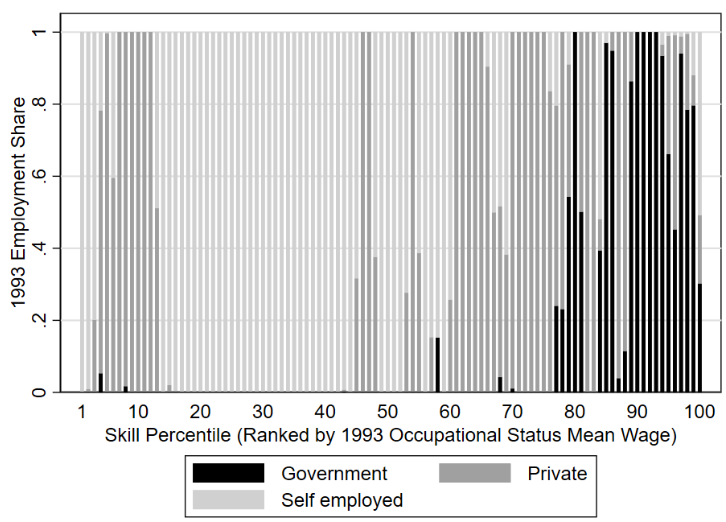

Figure 2.6: *Employment share of worker status across occupation-status, Indonesia 1993*

Note: Employment shares of worker statuses add up to 1 within each percentile. Data source is IFLS survey (1993).

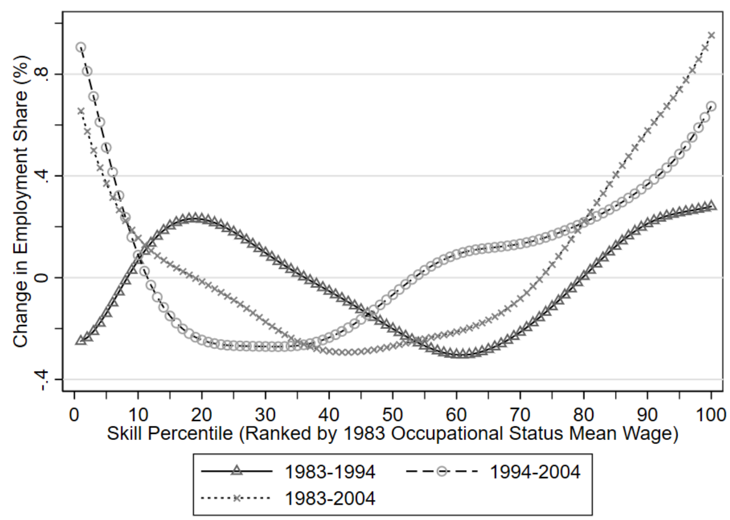

Figure 2.7: *Change in employment by occupation status, India 1983-2004*

Note: The lines represent locally weighted scatterplot smoothing of the employment change for the period across occupation-worker status categories. Data sources are NSSO employment and unemployment surveys (1983, 1993-1994, 2004-2005).

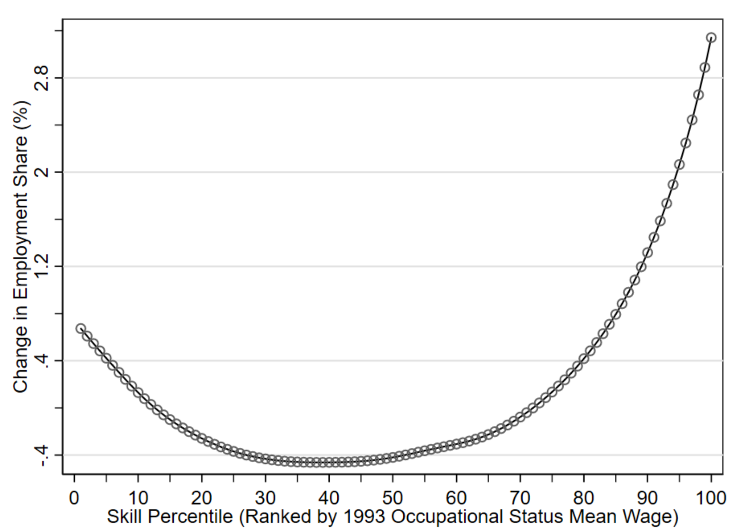

Figure 2.8: *Change in employment by occupation status, Indonesia 1993-2000*

Note: The lines represent locally weighted scatterplot smoothing of the employment change for the period across occupation-worker status categories. Data sources are IFLS surveys (1993, 2000).

Figure 2.9: *Employment change decomposed by status, India 1983-2004*

Note: The bars for each worker status category represent the employment change for that worker status for the period, computed from the locally weighted scatterplot smoothing across occupation-worker status categories. Data sources are NSSO employment and unemployment surveys (1983, 2004-2005).

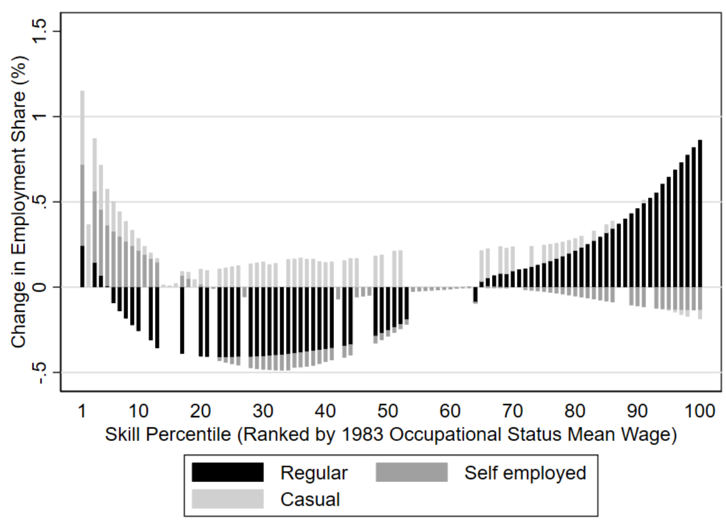

Figure 2.10: *Employment change decomposed by status, India 1993-2004*

Note: The bars for each worker status category represent the employment change for that worker status for the period, computed from the locally weighted scatterplot smoothing across occupation-worker status categories. Data sources are NSSO employment and unemployment surveys (1993-1994, 2004-2005).

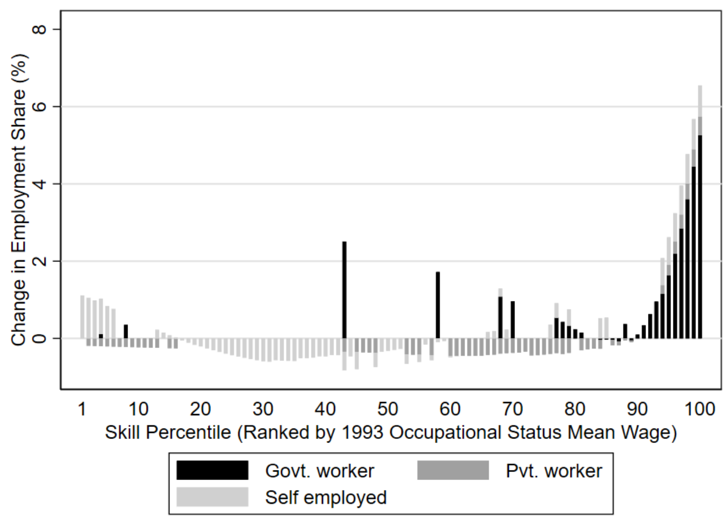

Figure 2.11: *Employment change decomposed by status, Indonesia 1993-2000*

Note: The bars for each worker status category represent the employment change for that worker status for the period, computed from the locally weighted scatterplot smoothing across occupation-worker status categories. Data sources are IFLS surveys (1993, 2000).

Table 2.1: *Indian employment shares*

	(1) 1983	(2) 2004
Occupation categories		
Professional,technical,administrative	0.13	0.18
	(0.34)	(0.39)
Clerk	0.09	0.07
	(0.29)	(0.26)
Sales	0.19	0.19
	(0.39)	(0.40)
Service	0.12	0.10
	(0.32)	(0.30)
Production	0.47	0.45
	(0.50)	(0.50)
Worker status categories		
Regular	0.39	0.37
	(0.49)	(0.48)
Self-employed	0.46	0.47
	(0.50)	(0.50)
Casual	0.14	0.16
	(0.35)	(0.37)
Observations	70845	103465

Note: Data is from NSSO (National Sample Survey Office) employment unemployment rounds 1983, 2004-2005. NSSO sampling weights are applied. Major occupation categories are constructed from the occupation classification provided in NSSO survey rounds (National Classification of Occupation (NCO) 1968 1-digit categories, excluding agricultural occupations). Standard deviation in brackets.

Table 2.2: *Summary statistics India, 1983*

Variables	
Log daily wages	2.38
	(0.80)
Schooling years	4.99
	(0.43)
Male fraction	0.83
	(0.37)
Rural fraction	0.46
	(0.50)
Education categories:	
Illiterate	0.34
	(0.47)
Pre-primary	0.10
	(0.30)
Primary	0.17
	(0.37)
Middle	0.15
	(0.36)
Secondary, higher secondary	0.16
	(0.37)
Above higher secondary	0.07
	(0.26)
Industries:	
Agriculture	0.01
	(0.09)
Mining, quarrying	0.02
	(0.14)
Manufacturing	0.32
	(0.46)
Electricity, gas, water	0.01
	(0.10)
Construction	0.06
	(0.24)
Wholesale, retail trade, restaurants, hotels	0.21
	(0.41)
Transport, storage, communication	0.08
	(0.28)
Financial, insurance, real estate, business services	0.02
	(0.15)
Community, social, personal services	0.27
	(0.44)
Observations	70845

Notes: Data is from NSSO (National Sample Survey Office) employment unemployment round 1983. NSSO sampling weights are applied. Education categories are constructed and mapped to schooling years following Hnatkovska et al. (2012). Industry classifications are mapped from the NIC (National Industrial Classification) 1970 1-digit level. Standard deviation in brackets.

Table 2.3: *Indonesian employment shares*

	(1) 1993	(2) 2000
Occupation categories		
Professional,technical,administrative	0.12	0.08
	(0.33)	(0.28)
Clerk	0.05	0.07
	(0.22)	(0.25)
Sales	0.31	0.23
	(0.46)	(0.42)
Service	0.09	0.21
	(0.28)	(0.41)
Production	0.42	0.40
	(0.49)	(0.49)
Worker status categories		
Government worker	0.15	0.11
	(0.35)	(0.32)
Private worker	0.37	0.50
	(0.48)	(0.50)
Self-employed	0.48	0.38
	(0.50)	(0.49)
Observations	4413	8249

Notes: Data is from IFLS (Indonesia Family Life Survey) rounds 1993, 2000. IFLS sampling weights are applied. Major occupation categories are constructed from the occupation classification provided in IFLS (ISCO or International Standard Classification of Occupations 1-digit categories, excluding agricultural occupations). Standard deviation in brackets.

Table 2.4: *Summary statistics Indonesia, 1993*

Variables	
Log daily real wages	1.45
	(1.29)
Schooling years	7.40
	(3.97)
Male fraction	0.64
	(0.48)
Rural fraction	0.49
	(0.50)
Education categories:	
Illiterate	0.12
	(0.33)
Primary	0.49
	(0.50)
Middle	0.13
	(0.34)
Secondary, higher secondary	0.19
	(0.39)
Above higher secondary	0.07
	(0.25)
Observations	4413

Notes: Data is from IFLS (Indonesia Family Life Survey) round 1993. IFLS sampling weights are applied. Education categories are mapped to schooling years as in Kurniawati et al. (2018). Standard deviation in brackets.

Table 2.5: *Occupation and worker status average wages*

Panel A: India, 1983	Self-employed	Casual	Regular	Total
Professional, technical & administrative	2.57	1.95	3.09	2.89
Clerk	2.70	2.05	2.83	2.81
Sales	2.37	1.77	2.14	2.34
Service	2.17	1.56	2.13	2.08
Production, operatives & labourers	2.23	1.94	2.52	2.24
Total	2.31	1.90	2.63	2.38

Panel B: Indonesia, 1993	Self-employed	Private	Government	Total
Professional, technical & administrative	2.75	2.15	2.37	2.33
Clerk	1.84	2.01	2.15	2.10
Sales	1.17	1.47	2.39	1.20
Service	1.21	1.28	1.87	1.35
Production, operatives & labourers	1.10	1.39	2.11	1.31
Total	1.17	1.48	2.24	1.45

Notes: Average wages are average log daily real wages. NSS sampling weights are applied for India and IFLS sampling weights are applied for Indonesia. Real wages for India are reported in terms of 1983 Maharashtra rural poverty line. Real wages for Indonesia are reported in terms of 1993 Jakarta consumer price index.

Table 2.6: *Latent class segmentation across specifications*

	Base			Within occupation groups		
	(1)	(2)	(3)	(4)	(5)	(6)
	Education	Wage	π	Education	Wage	π
Panel A: India						
Latent Class 1	-48.12	1.99***	0.30	-48.12	2.00***	0.30
		(0.01)			(0.01)	
Latent Class 2	1.51***	2.15***	0.28	1.54***	2.16***	0.29
	(0.01)	(0.01)		(0.01)	(0.01)	
Latent Class 3	2.30***	2.94***	0.42	2.30***	2.97***	0.41
	(0.00)	(0.01)		(0.00)	(0.01)	
		var =0.46				
	log-likelihood=-253044.1			log-likelihood=-252664.3		
Panel B: Indonesia						
Latent Class 1	0.00	0.78***	0.11	0.00	0.77***	0.11
	(0.06)	(0.06)		(0.02)	(0.05)	
Latent Class 2	6.72***	1.32***	0.59	6.00***	1.25***	0.46
	(0.02)	(0.02)		(0.01)	(0.03)	
Latent Class 3	12.74***	2.30***	0.30	9.01***	1.66***	0.15
	(0.04)	(0.03)		(0.02)	(0.05)	
Latent Class 4				12.47***	2.26***	0.28
				(0.02)	(0.04)	
	var=1.57	var=1.50				
	log-likelihood=-18341.71			log-likelihood=-16092.69		

Notes: *p<0.10, **p<0.05, ***p<0.01. Standard errors in parentheses. Expected log years of schooling reported under Education in Panel A. Expected years of schooling reported under Education in Panel B. Variance reported as *var*. For India and Indonesia respectively, the *within occupation group* specification wage variances are 0.45, 1.99 (professional, technical, administrative), 0.29, 1.53 (clerk), 0.46, 1.57 (sales), 0.57, 1.86 (service), 0.47, 1.19 (production, operatives, labourers). For Indonesia, education variances are 3.37, 2.63, 0.22, 0.17, 0.17 respectively for occupations in the same order as above.

Table 2.7: *Share of predicted latent classes across specifications*

Predicted Class	Base	Occupation groups
Panel A: India		
1	30.71	30.71
2	27.23	28.21
3	42.07	41.08
Panel B: Indonesia		
1	10.88	10.88
2	59.78	45.28
3	29.35	14.62
4	-	29.23

Notes: Predicted class refers to the latent class an observation is classified in. Observations are classified into a latent class if the posterior probability of belonging to that class is greater than the other classes. *Base* refers to latent classes being constructed for the 1983 Indian data (Panel A) and 1993 Indonesian data (Panel B). *Occupation groups* refers to latent class constructed within major occupation groups in the 1983 Indian data (Panel A) and 1993 Indonesian data (Panel B). Major occupation groups include Professional, technical and administrative; Clerk; Sales; Service; Production, operatives and labourers.

Table 2.8: *Multinomial logistic regression of latent class on worker status*

	Base		Within occupation	
	(1)	(2)	(3)	(4)
Panel A: India				
Latent Class 1				
Regular	0.58***	0.48***	0.58***	0.48***
	(0.01)	(0.01)	(0.01)	(0.01)
Casual	1.30***	0.90***	1.32***	0.91***
	(0.03)	(0.03)	(0.03)	(0.03)
Latent Class 3				
Regular	2.60***	2.03***	2.61***	2.01***
	(0.05)	(0.05)	(0.05)	(0.05)
Casual	0.32***	0.33***	0.32***	0.32***
	(0.01)	(0.01)	(0.01)	(0.01)
Observations	70845	70845	70845	70845
Pseudo R^2	0.99	0.99	0.99	0.99
Panel B: Indonesia				
Latent Class 1				
Self-employed	1.76***	1.23*	1.65***	1.20
	(0.18)	(0.14)	(0.18)	(0.13)
Govt. worker	0.18***	0.15***	0.24***	0.21***
	(0.09)	(0.08)	(0.13)	(0.11)
Latent Class 3				
Self-employed	0.33***	0.38***	0.73***	0.87
	(0.03)	(0.04)	(0.07)	(0.09)
Govt. worker	7.62***	8.51***	2.73***	3.12***
	(0.82)	(0.95)	(0.48)	(0.56)
Latent Class 4				
Self-employed			0.31***	0.37***
			(0.03)	(0.04)
Govt. worker			10.44***	12.22***
			(1.36)	(1.65)
Observations	4413	4413	4413	4413
Pseudo R^2	0.14	0.15	0.11	0.11
Controls	No	Yes	No	Yes

Notes: Standard errors in parentheses; *p<0.10, **p<0.05, *** p<0.01. Multinomial logistic regression weighted by sample weights. Latent class 2 and self employed (India) and private worker (Indonesia) are base categories. Coefficients are relative risk ratios. Controls include sector and gender for both countries and also industry for India.

Appendix

A2.1 Wage imputation

To graphically show the improvement in imputation, we look at the sample of non-missing earnings in the NSS data. This sample includes only casual and regular workers as earnings data are reported only for them. We artificially create missing values by assigning a number to each observation with the help of a random number generator. Wages above the 50th percentile for the random numbers are considered as missing. We regress log daily real wages on 5 education dummies, 6 experience dummies, 9 occupation dummies and state dummies for rural and urban samples, divided into male and female subsamples for 1983.[26] We carry out the regression imputation procedure followed by Kijima (2006) and a stochastic regression imputation procedure (Little and Rubin (2019)) on this sample to impute earnings for the artificially missing data and compare the imputed data with the recorded or true data. We compare the true wage distribution density plot with the density plot of the dataset with imputed wages. As one can see, the kernel density created from

[26]The 5 education dummies are illiterate, pre-primary, primary, middle, and secondary and above. Years of experience are computed by subtracting 5 and years of schooling from the age of a worker. Following Hnatkovska et al. (2012), we consider years of schooling as 0 for illiterate, 2 for pre-primary, 5 for primary, 8 for middle, 10 for secondary, and 15 for graduate and above. The 6 experience dummies are: below 11 years, 11-20 years, 21-30 years, 31-40 years, 41-50 years and above 50 years. The 9 occupation dummies are formed from the 1-digit NCO codes, excluding agriculture.

the regression imputation (Figure A1) varies significantly from the actual, as compared to the case where stochastic regression imputation is applied (Figure A2). The spikes of the regression imputation kernel density plot show the loss of variability in imputed earnings data. Since the stochastic regression imputation adds a *random* error or residual to predicted wages, carrying out the stochastic regression imputation several times will add different residuals each time to the predictions. One might be concerned if the different stochastic regression imputations yield different earnings distributions. We have carried out several stochastic regression imputations to compare with the recorded earnings data. The kernel density plots do not vary much across different stochastic regression imputations. A simple paired t-test reveals that different stochastic regression imputations have no significant difference with the recorded wage data.

Kijima (2006) carries out a regression imputation procedure to impute earnings for self-employed workers. The stochastic regression imputation model used to impute wage values for artificially missing casual and regular wages may not be appropriate for predicting self-employed wages, as self-employed worker characteristics may be very different from those who work on a contract or on a regular basis and adding a random error may make the imputation worse. To test for this, we look at the Indian Human Development Survey 1 (IHDS-1), 2005 data which has information on regular and salary earnings along with profit obtained by a household in a household run business. After dropping observations with missing occupation data (more than 50%[27]) and restricting the sample for the age group 15-65 and full-time workers, we

[27]After dropping missing occupation data, the proportion of each major occupation category is greater than 9% hence dropping data due to missing occupation data does not remove any major occupation category. IHDS occupation employment fraction for production, operatives and labourers is the largest, followed by sales, then followed by professional, technical and administrative, similar to the NSS sample for 2004-05. The smallest employment share is of services in IHDS unlike NSS, where the smallest employment share is of clerks. Also, the NSS records about 45% of the

carry out the same regression exercise on the regular or salaried workers and impute earnings for the household run business (similar to Kijima (2006)). About 73% of the sample are regular and salaried workers while the rest run household businesses. As before, we then run a stochastic regression imputation procedure and compare the actual earnings with the imputed earnings kernel density plot. The results are similar as before: earnings distribution from the stochastic regression imputation procedure (Figure A3) is closer to predicting the actual earnings distribution than the regression imputation procedure (Figure A4).

One concern that may arise from our cross-validation method is that only 30% of the IHDS earnings required imputation, whereas the NSS dataset has about 50% of self-employed earnings observation missing. Accuracy of imputation methods depends on the percentage of the sample to be predicted. Any imputation method used for the IHDS dataset may seem dubious in terms of applicability to the NSS dataset given the relatively smaller share of data imputed in the IHDS. However, we have shown that for both NSS and IHDS sets, the kernel density of the imputed wages is very similar to the kernel density of the true wages when using the stochastic regression imputation procedure. Hence, the proportion of missing data is not a concern in this case.

In light of the above results, we carry out a stochastic regression imputation procedure where we regress log daily real wages of regular and casual workers for a year on 5 education dummies, 6 experience dummies, occupation dummies (NCO 1-digit) and state dummies. We carry out this regression separately for rural and

employment share in production, operatives and labourers and IHDS reports around 34% for the same after dropping observations. Consequently, a small share of production worker earnings is imputed in the IHDS. However, we can show that the regression imputation does far worse than the stochastic regression imputation despite the differing magnitudes of occupation employment shares across the NSS and IHDS samples.

urban sectors, divided into male and female subsamples for 1983.

A2.2 Non-parametric (LOWESS) strategy

We carry out a LOWESS (Locally Weighted Scatterplot Smoothing) on employment shares across the skill percentiles, following Autor et al. (2006). This provides a non-parametric approach to recording employment changes across the 3-digit NCO occupations for India and the 2-digit ISCO occupations for Indonesia. After sorting the base year data (1983 for India and 1993 for Indonesia) in ascending order with respect to occupation average wages, we create employment weighted percentiles based on the individual level data. We present below computations for the Indian NSS data, but the same method is carried out for the Indonesian IFLS data.

A2.2.1 Percentile construction for LOWESS

In this section we describe how employment weighted percentiles are constructed following (Autor et al. (2006)). We first rank individual observations by occupation average wages. To create weighted percentiles, we utilize employment weights such that occupations with very large employment shares are assigned more than one percentile and occupations with relatively small employment shares are assigned to the same percentile.

The average employment weight ($avlwt$) is generated for each occupation code o:

$$avlwt_o = \sum_{i=1}^{n_o} samplingweight_{io}/n_o, \tag{a1}$$

where n_o is the number of sample observations of occupation o and $samplingweight_{io}$ refers to the NSS (or IFLS) sampling weight of worker i in occupation o. Thus, $avlwt_o$ stands for the average employment weight for an occupation in the base year sample.

The percentile construction is similar to Hazen's percentile (Hazen (1914)) which is computed by subtracting 0.5 from the rank and dividing the result by the total number of observations:

$$hperc_i = 100 * (r_i - 0.5)/n, \tag{a2}$$

where r_i is the rank of observation i and n is the number of observations. Note that for Hazen's percentile, 0.5 is subtracted from the rank in the numerator to make the percentile aligned to the midpoint between observation i and the previous observation $i - 1$.

Our first step is to arrange observations in the ascending order of their occupation average wages. Thus observation 1 will have the lowest occupation average wage, whereas the last observation (observation N) will have the highest occupation wage. Next, for each observation j in this ordered arrangement, we compute:

$$rank_j = \sum_{i=1}^{j} avlwt_{oi}, \tag{a3}$$

where $avlwt_{oi}$ is the average employment weight of observation i (belonging to occupation o). So $rank_j$ is the cumulative employment weight of observation j – the sum of average occupation employment weights of all observations up to observation j. Similar to obtaining the midpoint between observations j and $j - 1$ in Hazen's percentile (equation (a2)), we compute the following:

$$\begin{aligned} hrank_j &= (rank_j + rank_{(j-1)})/2, \\ hrank_1 &= rank_1/2. \end{aligned} \tag{a4}$$

Thus, $hrank_j$ is an employment weighted rank for observation j and will be divided by the total employment in the base year to create percentiles in the next step. Note that we are simply taking the average of two cumulative employment weights, one that is till observation j ($rank_j$) and the other till observation $j-1$ ($rank_{j-1}$). This ensures that $hrank_j$ will lie in between the cumulative employment weights of observation j and the previous observation $j-1$. If the observation j has a relatively larger cumulative employment weight (or $avlwt_o$) than the previous observation $j-1$, the midpoint (or $hrank_j$) will have a much larger value than the case where observation j has a relatively small value of $avlwt_o$. This will result in a relatively larger value of $hrank$ for observation j.

To calculate the percentile for observation j ($prank_j$), we simply divide $hrank_j$ by $rank_N$, where $rank_N$ adds up all the average employment wages ($avlwt_o$) till the last observation N (or is the total employment in the base year), and multiply by 100:

$$prank_j = 100 * hrank_j/(rank_N). \tag{a5}$$

We make the appropriate conversion to integer form of $prank_j$ so that we have percentiles ranging from 1 to 100.

When the fractions in $prank_j$ are converted into integer values, different fractional values are assigned the same integer value. Intuitively, the construction of $hrank$ ensures that observations with small and similar occupation employment shares remain in the same percentile. It is possible that observations with different occupations but small employment shares are assigned to the same percentile. This is because employment shares for these occupations were relatively small, leading to relatively smaller and similar values of $hrank$. It is also possible that an observation with an occupation similar to the previous observation may end up in a new percentile. More than one percentile will be assigned to an occupation if there are many observations and the employment share is relatively large for that occupation.

A2.2.2 Occupation LOWESS

With the employment weighted percentiles being constructed, in this section we describe how we measure the employment change over a period in each percentile. For year t, the share of employment (NSS sample weights are used) is calculated for each occupation code:

$$empshr_{to} = employment_{to}/employment_t, \tag{a6}$$

where $t = 1983$, 2004 and o is the 3-digit NCO code. Employment change in occupation o over the period 1983-2004 is then given by:

$$emp8304_o = (empshr_{2004o} - empshr_{1983o})/empshr_{1983o}, \tag{a7}$$

where $empshr_{1983o}$ and $empshr_{2004o}$ are the employment shares of occupation o in 1983 and 2004, respectively, as defined in equation (a6). Equation (a7) calculates the occupation employment change for the period 1983 to 2004 normalized by the occupation employment share in 1983.

Since there will be many occupations in one percentile in most cases, we must weight the occupations by the relative employment weights within a percentile. A percentile employment weight ($percwt$) is created for each occupation o in percentile p :

$$percwt_{op} = avlwt_{op}/(\sum_o avlwt_{op}). \tag{a8}$$

This weight ensures that observations with relatively larger average occupation employment shares in a percentile are given greater importance or weightage than occupations with relatively smaller average occupation employment shares in the same percentile. Note that the sum of the $percwt_{op}$ over observations in a percentile adds up to 1.

Finally, the average employment change in occupation o over the period 1983-2004 ($pdesh8304$) in a percentile (p) is given by:

$$pdesh8304_p = \sum_o [percwt_{op} * emp8304_o]. \tag{a9}$$

The employment change is in percentage points since we weight all employment changes by $percwt_{op}$ given in equation (a8).

In Section 2.3.1, we compute the average employment change in Indian occupations over the period 1983 to 2004 ($pdesh8304$) in each percentile (p) as in equation (a9), carry out a LOWESS (or a smoothing) on these employment changes, and graphically demonstrate how employment changes have occurred across occupations (Figure 2.1). We repeat the above exercise for Indonesia using ISCO 2-digit occupation categories, IFLS sampling weights and replacing the years 1983 with 1993 (the base period for Indonesia) and 2004 with 2000 (Figure 2.2).

A2.2.3 Occupation-status LOWESS

Employment changes in occupation-statuses in a skill percentile: The employment share of occupation-status categories is expressed as a share in total 1983 employment and denoted by $empshr_{1983os}$, where o refers to occupation and s refers to status. In a percentile, $empshr_{1983os}$ (or employment share of occupation-status categories in 1983) is summed up for each status. For example, the employment share of casual workers is calculated by summing $empshr_{1983os}$ within a percentile for $s = casual$. Similar calculations are carried out for obtaining the regular and self-employed employment shares within each percentile. All the shares are expressed in fractions within a percentile. The 3 worker status employment shares thus computed will add up to 1 within a percentile.

Similar to equation (a7), $emp8304_{os} = (empshr_{2004os} - empshr_{1983os})/empshr_{1983os}$ is the change in employment share for occupation code $o \times$ worker status category s for period 1983-2004 for India. Percentile weights are computed as in equation (a8): $percwt_{osp} = avlwt_{osp}/\sum_{os} avlwt_{osp}$, where $avlwt_{osp}$ is the average employment weight for occupation-status category os in percentile p. The weighted employment change in each percentile p is given by $pdesh8304_p = \sum_{os}[percwt_{osp} * emp8304_{os}]$. Interpretations of all the equations are the same as before, the only distinguishing feature is the level at which they are computed – occupation-status instead of only occupation. We repeat all the above exercises for Indonesia where s = self-employed, private, government. We carry out LOWESS on $pdesh8304_p$ and demonstrate graphically how employment changes have occurred across occupation-statuses in Figure (2.7) for India and Figure (2.8) for Indonesia.

LOWESS decomposition: We compute $pdesh8304_{ps}$ which is the sum of percentile weighted employment changes of occupations belonging to status s in percentile p (similar to equation (a9)):

$$pdesh8304_{ps} = \sum_{o}[percwt_{osp} * emp8304_{os}]. \tag{a10}$$

For each percentile, we obtain percentile weighted employment changes for the 3 statuses such that employment changes for the 3 statuses sum up to the employment change for the percentile:

$$pdesh8304_p = \sum_s pdesh8304_{ps}. \tag{a11}$$

We carry out LOWESS on $pdesh8304_{ps}$ separately for each status s. For India, we plot the regular, casual and self-employed LOWESS values for each percentile in the form of a stacked bar graph in Figures 2.9 and 2.10. Similarly, for Indonesia, we plot the self-employed, private and government LOWESS values for each percentile in the form of a stacked bar graph in Figure 2.11.

A2.3 Optimal number of latent classes

To choose the optimal number of latent segments we use the information criteria (Akaike Information Criteria (AIC) and Bayesian Information criteria (BIC)) and compare models with different number of latent classes.[28] Information criteria are statistical measures that compare the quality of two statistical models. An information criterion number by itself does not reveal any information regarding a model, but it must be compared over models. The smaller the values, the better is the model. The optimal number of classes is selected from the model which sees a great improvement in information criteria, that is, a relatively large drop in the magnitudes of AIC and BIC numbers. We apply the elbow method for choosing the optimal number of classes. If the values are plotted on a graph where the x-axis represents the number of classes and the y-axis records AIC, BIC values, we would see an 'elbow-shaped' line. The x-value at the elbow is considered the optimal num-

[28] AIC $= 2lnL + 2k$, where lnL is the maximized log-likelihood of the model and k is the number of parameters estimated (Akaike (1974)). BIC $= 2lnL + klnN$, where N is the number of observations (Schwarz (1978)). Both criteria penalize overfitting of the model to the sample, BIC more than AIC.

ber of classes (Thorndike (1953)). However, due to identifiability[29] issues, we cannot compare information criteria for a very high number of latent classes.

In Table A1, we have our two specifications – *base* and *within occupation groups* for India (Panel A) and Indonesia (Panel B). The *base* specification is described in equation (2.2) where real wages and years of schooling are used to construct the latent classes. The *within occupation groups* specification carries out the same exercise but within the 5 major occupation subsamples. The column *class* in Table A1 refers to the number of latent classes for a model specification. The table reports the log-likelihood values, degrees of freedom, AIC and BIC values for different models and specifications. For example, for India, when we estimate parameters for the 4-class model, we can see that the improvement when moving from a 2-class model to a 3-class model is much greater than moving from a 3-class model to a 4-class model for all specifications. We can see this from the drop in AIC, BIC values. We can infer that 3 classes are the optimal number of classes for India for both specifications. For Indonesia, the *base* specification exhibits greater improvement when moving from a 2-class to a 3-class model than moving from a 3-class to a 4-class model. However, for the *within occupation groups* specification the improvement is far greater when moving from a 3-class to a 4-class model.

A2.4 Education distribution: Poisson versus Normal

In Table A2, we compare the information criteria, AIC and BIC, across two finite mixture model specifications: (1) wages follow a normal distribution and years of

[29]Identifiability issues occur when we do not have a distinct representation of the model or distinct observational claims (Teicher (1963)). If two or more latent classes have the same set of estimated parameters, these latent classes are equivalent to one latent class. Identifiability issues may arise from collinearity in a linear regression run across latent classes. Convergence may not occur when too many latent classes and consequently, too many parameters need to be estimated.

education follow a Poisson distribution (columns 1 and 2) and (2) both wages and years of education follow a normal distribution (columns 3 and 4). We compare the information criteria for 2, 3 and 4 class models for India in Panel A and for Indonesia in Panel B. Smaller AIC and BIC numbers imply a better model. For India, years of education following a Poisson distribution is better than following a normal distribution, as seen from the smaller AIC and BIC values in columns 1 and 2. For Indonesia, on the other hand, years of education following a normal distribution is better than following a Poisson distribution.

A2.5 Correlation matrix

For a simpler understanding, we check how accurately the skill classes are able to predict the worker status. For this we simply, check the fraction of each skill class within each worker status. One can see from Table A3 that 54% of Indian casual workers are classified as the low skill class and 60% of Indian regular workers are classified as the high skill class. Self-employed workers seem to have a relatively even distribution of all three skill classes. Similarly, for Indonesia, the fraction of middle skill workers is highest for both private and self-employed workers (62%, 71%) and the fraction of high skill workers is highest for government workers (77%).

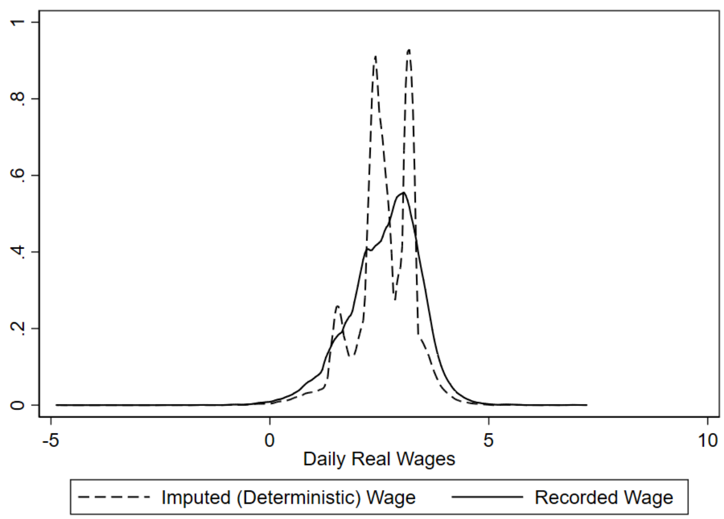

Figure A1: *Density of wages from regression imputation procedure compared with density of recorded wages in NSS*

Note: Data source is NSSO employment and unemployment survey 1983. The data is restricted to regular and casual workers.

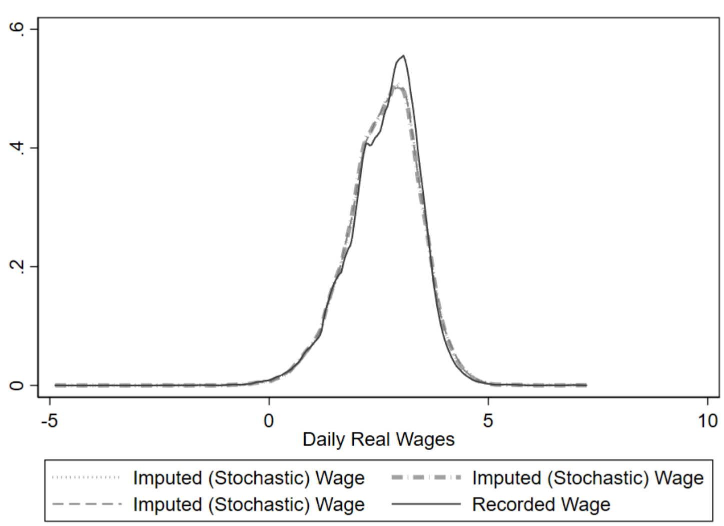

Figure A2: *3 Densities of wages from stochastic regression imputation procedure compared with density of recorded wages in NSS*

Note: Data source is NSSO employment and unemployment survey 1983. The data is restricted to regular and casual workers.

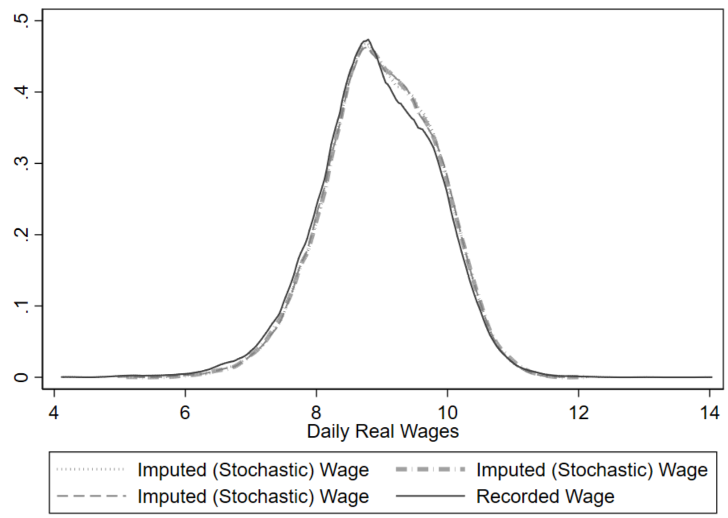

Figure A3: *3 Densities of wages from stochastic regression imputation procedure compared with density of recorded wages in IHDS for self-employed workers*

Note: Data source is IHDS-I. The data is restricted to self-employed workers.

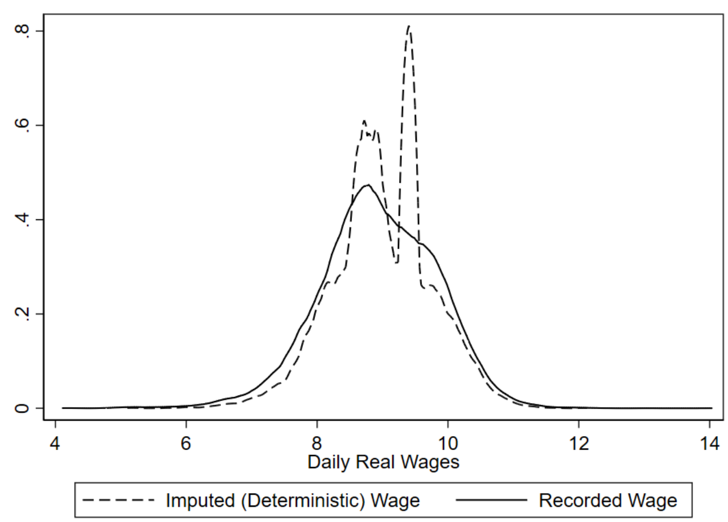

Figure A4: *Density of wages from regression imputation procedure compared with density of recorded wages in IHDS for self-employed workers*

Note: Data source is IHDS-I. The data is restricted to self-employed workers.

Table A1: *Information criteria across mixture model specifications*

Class (1)	log likelihood (2)	df (3)	AIC (4)	BIC (5)
Panel A: India				
Base specification				
2	-262452.6	6	524917.1	524972.1
3	-253044.1	8	506104.3	506177.6
4	-252409.9	12	504843.8	504953.8
Within occupation groups				
2	-262190.2	10	524400.4	524492.1
3	-252664.3	12	505352.7	505462.7
4	-251934.3	15	503898.6	504036.1
Panel B: Indonesia				
Base specification				
2	-19583.7	7	39181.36	39226.11
3	-18341.7	10	36703.43	36767.35
4	-17439.2	13	34904.33	34987.43
Within occupation groups				
2	-19466.6	15	38963.2	39059.08
3	-18229.4	18	36494.85	36609.91
4	-16092.7	21	32227.38	32361.61

Notes: Column 1 reports the number of latent classes constructed in a mixture model. Column 3 reports the degrees of freedom (df) for the respective mixture model. Columns 4 and 5 report the Akaike Information Criteria and Bayesian Information Criteria for the respective mixture model, respectively. *Base specification* refers to latent classes constructed for the 1983 Indian data (Panel A) and 1993 Indonesian data (Panel B). *Within occupation groups* refers to latent classes constructed within major occupation groups in the 1983 Indian data (Panel A) and 1993 Indonesian data (Panel B). Major occupation groups include Professional, technical and administrative; Clerk; Sales; Service; Production, operatives and labourers.

Table A2: *Poisson versus normal for education years*

Latent classes	Poisson		Normal	
	AIC	BIC	AIC	BIC
	(1)	(2)	(3)	(4)
Panel A: India				
2	524917.1	524972.1	565122.3	565186.4
3	506104.3	506177.6	563333.5	563425.1
4	504843.8	504953.8	546472.4	546591.6
Panel B: Indonesia				
2	40665.97	40704.33	39181.36	39226.11
3	37222	37273.14	36703.43	36767.35
4	35624.9	35695.21	34904.33	34987.43

Notes: Column denoted by *Latent classes* reports the number of latent classes constructed in a mixture model for India in Panel A and Indonesia in Panel B. Columns 1 and 2 report the information criteria for the model which considers years of education following a Poisson distribution. Columns 3 and 4 report the information criteria for the model which considers years of education following a Normal distribution.

Table A3: *Fraction of skill classes in a worker status*

Status	Class 1	Class 2	Class 3	Total
(1)	(2)	(3)	(4)	(5)
Panel A: India				
Self-employed	0.36	0.30	0.34	1
Regular	0.17	0.23	0.60	1
Casual	0.54	0.33	0.13	1
Panel B: Indonesia				
Self-employed	0.17	0.71	0.12	1
Government	0.01	0.23	0.77	1
Private	0.08	0.62	0.30	1

Notes: Columns 1-3 denote the skill classes predicted by the mixture model. Fractions add up to 1 for each worker status for India in Panel A and Indonesia in Panel B.

Chapter 3

Do Clouds have a Silicon Lining for Firms? Contract Hiring and Computer Investment: Evidence from Rainfall Shocks

3.1 Introduction

The role of technology in driving labour market changes has been well established in the literature.[1]The literature provides mixed evidence on whether technology adop-

[1]Tinbergen (1974, 1975) linked technology to the relative demand for skills, providing a precursor to the literature on SBTC or skill-biased technological change (Katz and Murphy (1992)). SBTC categorizes workers into high and low skill groups and explores whether improvements in technology increase the demand for high-skilled workers relative to low-skilled workers. The task-based framework (Autor et al. (2003)) describes how machines displaces workers from occupations based on the types of tasks or activities carried out in jobs. Acemoglu and Autor (2011) show how workers are displaced from manufacturing, sales and clerical occupations, given the type of

tion leads to a rise or fall in employment. Firm-level studies have documented that technology adoption leads to both decreases in employment (Autor et al. (1998), Autor and Salomons (2018), Bessen et al. (2019), Acemoglu and Restrepo (2020)), as well as increases in employment (Van Reenen (1997), Blanchflower and Burgess (1998), Bessen (2019), Aghion et al. (2020)). Studying how technology affects hiring decisions is a difficult exercise. Technology is not an exogenous or random shock as the decision to invest in technology is taken by firms. Our study proposes a way to overcome this issue by analyzing how firms with different levels of computer capital expenditure respond differently to exogenous demand shocks. We examine a panel of Indian manufacturing firms for the period 2000-2010 and analyze if hiring decisions in response to exogenous transitory demand shocks vary due to different levels of computer capital expenditure. We use lagged rainfall shocks to proxy for transitory demand shocks experienced by firms (Adhvaryu et al. (2013), Chaurey (2015)). Firm decisions do not affect the occurrence of rainfall shocks, hence demand shocks arising from rainfall shocks are exogenous.

This chapter addresses the following question: Do firms that invest more in computer capital (and hence invest more in computer-aided technology), hire workers differentially in response to transitory demand shocks? Computer capital investment is not an exogenous variable. It is difficult to obtain computer capital investment "shocks" or consider computer capital spending as random across firms, as firms decide how much to invest in computer capital. We create a computer capital share dummy for firms that have above-average computer capital expenditure. Identification comes from the interaction of computer capital share dummy and exogenous demand shocks (Nizalova and Murtazashvili (2016)). For robustness, we

activities or tasks that workers carry out in these occupations.

also construct a host of alternative measures reflecting computer capital investment and check if we obtain similar results.

For measuring exogenous demand shocks, we use data on rainfall shocks following Adhvaryu et al. (2013) and Chaurey (2015) who show that industrial labour demand is sensitive to rainfall fluctuations through weather impact on agricultural productivity. Positive rainfall shocks have a positive impact on agricultural production and agricultural productivity which leads to relatively higher income and higher spending for agricultural workers. Increase in spending suggests an increase in demand for industrial goods. As more production takes place following rainfall shocks, there is a consequent increase in demand for workers in the industrial sector. Thus, rainfall shocks translate to demand shocks for firms.[2] Rainfall shocks are exogenous as firm decisions or unobserved labour market changes will not affect rainfall fluctuations. Hence, the demand shock arising from a rainfall shock is also exogenous to a firm's decisions. When faced with these exogenous transitory demand shocks, firms lay-off or hire more workers to cater to the changes in demand.

We find that firms with above-average computer capital share tend to hire 2.32 fewer contract workers compared to firms with lower-than-average computer capital share, in response to positive transitory demand shocks. This differential hiring is 11.33% of the sample mean of contract employment. Our results are robust to using a number of alternative measures of computer capital expenditure. We find that a 10 percentage point increase in computer capital expenditure share leads to a reduction in average contract hiring by 2.8 workers in the face of demand shocks.

[2]Adhvaryu et al. (2013) and Chaurey (2015) show that positive rainfall shocks lead to an increase in monthly per capita expenditure implying an increase in spending. Chaurey (2015) also shows that both industrial wages and employment increase following rainfall shocks for the period 2000-2010, which are sufficient conditions to conclude a net increase in demand for industrial workers, even if there are supply shocks due to rainfall. We discuss this in detail in Section 3.3.

Firms with computer capital share above the sample median computer capital share reduce contract hiring by 1.7 workers. Our results are also robust to other measures that consider industry and state-industry average computer capital shares. We also confirm that the reduction in hiring is for contract workers who carry out activities directly involved in the manufacturing process and not for activities in peripheral works such as security, cleaning services, etc.

As we are looking at how hiring decisions of firms are affected in the face of rainfall shocks or transitory demand shocks, one may be concerned that labour regulations will protect workers from being laid off in the face of such transitory shocks. The Industrial Disputes Act, 1947 (IDA) is the primary labour regulation in India that restricts firms from laying off workers. However, the IDA does not cover contract workers. Contract workers are temporary workers who are paid for less than 240 days in any 365 days period.[3] Chaurey (2015) finds that firms in states with strict labour regulations tend to hire contract workers in response to transitory shocks and no significant hiring of regular workers is observed following these same shocks. For this reason, we focus on contract workers and analyze how firms' hiring of contract workers vary by computer capital share.

We add to the literature that examines whether technology increases or decreases employment. Empirical studies have found evidence of workers being displaced from automating firms (Autor and Salomons (2018), Bessen et al. (2019), Acemoglu and Restrepo (2020)) and declining employment in manufacturing jobs due to computer technology (Autor et al. (1998), Vashisht (2018)). Machines carry out the work of human labour, essentially causing a labour substitution effect. There is also evidence of

[3]A provision for protection of contract workers exist in the Contract Labor Act (1970) which grants state the authority to ban use of contract labour in any establishment, and makes provisions for proper disbursement and prevent wage payment delays. But it has been found that enforcement of these regulations is weak (Bhandari and Heshmati (2008), Basu et al. (2021)).

technology increasing demand for workers (Blanchflower and Burgess (1998), Bessen (2019), Hjort and Poulsen (2019), Aghion et al. (2020)). Automation can help to reduce production costs and hence prices which leads to an increase in demand for goods, thereby resulting in an increase in production and a consequent increase in demand for labour. Our work is closely related to the recent research that employs quasi-experimental variation to study the relationship between technology and employment. Bessen et al. (2019) studies the impact of automation on employment and circumvents endogeneity issues by reporting firm's employment before and after automation investment spikes. However, this restricts the sample to firms which have made significantly large investments in automation. Aghion et al. (2020) uses a shift-share IV (Instrumental Variable) and estimate employment before and after automation expenditure. However, the analysis is restricted to the subset of firms that imports automation inputs. In our analysis, we look at how firms' hiring decisions vary by computer capital investment in the face of exogenous shocks. Irrespective of the size of computer capital investment or importing decisions, the demand shock is experienced by all firms in our sample. Thus, our strategy has the advantage of not restricting the sample to importing firms or to firms with large investments in technology.

Our study adds to other strands of literature. As our study looks at how heterogeneity in computer capital investment across firms affects hiring decisions in response to demand shocks, we add to studies that focus on firm heterogeneity due to technology. Doms et al. (1995) is one of the first studies that stress on controlling for producer heterogeneity arising from differences in types and level of capital investments. The outcome of interest for the study is firm entry, exit and growth. Yeaple (2005) develops a theoretical model where some firms choose to adopt new technologies whereas others do not. Firms with new (low cost) technologies are then able to produce more and enter the export market. Thus, heterogeneity in technology adoption leads some firms to become exporters while others remain non-exporters. Aboal et al. (2015) investigate the effect of product and process innovation on employment growth and additionally explores the presence of heterogeneous effects of innovation.[4] They find differential effects of innovation on employment across these technologically diverse sub-sectors. Similar to our study, this study highlights the importance of heterogeneous effects on employment arising due to different expenditure on innovation. However, our question is different as we are examining employment decisions in response to exogenous demand shocks.

We also add to the literature which analyzes impact of weather on economic outcomes (Dell et al. (2014)).[5] Studies focusing on the manufacturing sector find that high temperatures have a negative impact on industrial output (Zhang et al. (2018), Somanathan et al. (2021)). Hsiang (2010) records relatively higher output losses in the non-agricultural sector than the agricultural sector due to temperature changes from tropical cyclones. In our study, we focus on how firms are prone to demand shocks due to changes in agriculture productivity from weather fluctuations. Firms respond to demand shocks arising from rainfall shocks by changing output and employment, as discussed earlier. Hence, any impact on hiring decisions arising from heterogeneous computer capital investment during such weather-driven demand shocks becomes important.

The chapter is organized as follows. In the next section we discuss the sources of

[4]Their study analyzes the differences among high-tech and low-tech sub-sectors (high innovation expenditure and low innovation expenditure as share of turnover, respectively). The study points out that these sub-sectors have very different demand profiles for workers leading to dissimilar workforce composition.

[5]Dell et al. (2014) discuss studies that have analyzed the effects of weather variation on industrial output (Hsiang (2010), Jones and Olken (2010), Dell et al. (2012)).

data and how we measure rainfall shocks and computer investment for Indian firms. In section 3.3 we discuss the empirical strategy in detail. Section 3.4 discusses results on how computer capital investment affects hiring decisions during demand shocks and carries out robustness checks. Finally, section 3.5 concludes.

3.2 Data

We use data from several sources for our analysis. We use the Annual Survey of Industries (ASI) of India to obtain firm-level capital expenditure and employment data from 1999-2000 to 2009-2010 for firms in the manufacturing sector.[6] The reference year for the ASI is the accounting financial year from 1st April to 31st March of the following year. All registered industrial units in India, which employ at least 10 workers and use electricity, or employ at least 20 workers but do not use electricity, are covered by the ASI.[7]

ASI surveys provide firm-level annual data such as expenditure on assets and liabilities, employment, labour cost, receipts, expenses, input items (indigenous and

[6] As mentioned above, Chaurey (2015) shows that there is no statistically significant impact of rainfall shock on agricultural wages, but there exists statistically significant positive impacts on industrial wages and monthly per capita expenditure for the period 2000-2010. These findings point to rainfall shocks acting as exogenous labour demand shocks for the industrial sector. We restrict out study to this period as we leverage rainfall shocks as exogenous labour demand shocks in our analysis. This period is also relevant given that the key indicator of technology in our study is computer capital. The 2008 KPMG report (IBEF (2008)) points out that demand for Computer Numerical Control machines (the production machines that require computer capital) have outgrown conventional tools from the early 2000s. The demand for these tools largely arises from the manufacturing sector. Erumban and Das (2020) also document a significant increase in ICT investment for this period in India.

[7] The ASI frame consists of a census sector (units surveyed every year) and a sample sector (sampled every few years). The census sector consists of all firms in five industrially backward states (Manipur, Meghalaya, Nagaland, Tripura and Andaman and Nicobar Islands) and large factories for the rest of the country. The ASI definition of large factory underwent a change. From ASI 2001 onwards, a large factory in the census sector is defined as an industrial unit with 100 or more employees, whereas for years before 2001 the ASI definition of a large factory consists of an unit with 200 or more workers. The firms not covered in the census sector constitute the sampling frame for the sample sector. A third of these firms are randomly selected in the survey each year.

imported), cost and quantity of output produced, etc. We use data on the number of contract workers a firm hires each year. The ASI data also includes information on man-days (the number of workers employed on each day summed over all days in a year) separately for manufacturing and non-manufacturing work. Manufacturing work refers to activities in the main or core production process whereas non-manufacturing work refers to peripheral activities of the production process like repair, maintenance, etc. ASI reports the yearly gross value of computer equipment and software for each firm. We also obtain measures for the total capital value for each firm which includes buildings, plant and machinery, transport equipment, computer equipment including software, and other capital equipment.

We also use the US measures of computer capital share in constructing an alternative measure for our study. We obtain the US industry-wise computer capital expenditure and total capital expenditure for the period 2002-2010 from the Annual Survey of Manufactures (ASM). ASM provides industry-level estimates of key statistics for manufacturing sector[8] establishments with one or more paid employees. It provides measures of US manufacturing activity, products, and location for the public and private sectors. The survey provides yearly data on employment, payroll, hours, cost of materials, receipts, value added, capital expenditures, etc. at the industry level.

For our main analysis, we use monthly rainfall data from the Center for Climatic Research, University of Delaware, which is available from 1900 to 2017, at 0.5 degree by 0.5 degree latitude-longitude grid.[9] We match annual rainfall (in mm) to the

[8] Industries are classified according to North American Industrial Classification Standard (NAICS). We match the 3-digit level NAICS to the 2-digit level Indian National Industrial Classification (NIC) in our study.

[9] The Delaware rainfall data is available after spatial interpolation was carried out with the spherical version of Shepard's algorithm, which employs an enhanced distance-weighting method (Shepard (1968), Willmott et al. (1985), Willmott and Matsuura (2018)).

latitude and longitude nearest to the geographic centroid of each district using district shapefiles.[10]

We define the exogenous demand shock from rainfall data following Jayachandran (2006), Adhvaryu et al. (2013), Chaurey (2015) and Kaur (2019). Exogenous demand shock is a rainfall shock in district d in year t, given by:

$$Rainshock_{dt} = \begin{cases} 1, & \text{if rainfall in district } d \text{ in year } t \text{ is above the 80th percentile,} \\ 0, & \text{if rainfall in district } d \text{ in year } t \text{ is between the 20th and 80th} \\ & \text{percentiles,} \\ -1, & \text{if rainfall in district } d \text{ in year } t \text{ is below the 20th percentile,} \end{cases} \tag{3.1}$$

where the percentiles correspond to the the rainfall distribution for district d for the period 1998-2008. Similar to Chaurey (2015), we look at lagged effect of rainfall. We construct rainfall shock from rainfall measures in the previous calendar year. For example, to correspond to the ASI accounting year from 1st April 2005 to 31st March 2006, rainfall measures are from January 2004 to December 2004. This allows us to account for the time that firms may take to respond to demand shocks (due to the impact of rainfall on the local economy).

Table 3.1 reports the summary statistics of employment and computer capital expenditure data in our sample. On an average, 20.47 contract workers are employed in a firm, compared to 47.55 regular (or permanent) workers. Thus, employment in firms is largely regular. But, at the same time, contract workers employment is

[10]Shapefiles store geographic data such as location, shape and attributes of geographic features in a vector format. Geographic features in a shapefile can be represented by points, lines, or polygon areas which allows one to run analyses with geographic data such as constructing the centroid of a geographic location.

also non-trivial, consisting of around 18.96% of total workforce. Contract workers are mainly involved in the core manufacturing activities of the production process, whereas contract non-manufacturing man-days for peripheral works is very low, on average. The share of computer capital in total capital for a firm is 2.88% on average in our sample. In Table 3.1 we report separately the summary statistics for firms with computer capital share greater than 2.88% (or above-average computer capital share) in column 2 and firms with computer capital share less than or equal to 2.88% (or below-average computer capital share) in column 3. On an average, firms with above-average computer capital share employ relatively fewer workers (both regular as well as contract) and have a lower total capital value compared to firms with below-average computer capital share. The share of observations in our sample experiencing negative rainfall shocks (16%) is similar to the share of observations receiving positive rainfall shocks (20%). We also find that the share of observations receiving positive and negative rainfall shocks are similar for firms with above-average or below-average computer capital share. In Figure 3.1 we demonstrate that the industry distribution also is similar across positive and negative rainfall shocks in our sample. Thus, there should not be any concern that specific industries are facing only one type of rainfall shock in our sample. In Table 3.2, we also provide summary statistics for variables that reflect economic development (such as Gini coefficient and total landholdings) and the political orientation of the state the firm is located in for robustness exercises.[11]

[11] Data on Rural gini coefficient is from the EOPP Indian States Database, also used in Besley and Burgess (2004)). Total landholdings of a state in 2000-01 and percentage variation of holdings in 2000-01 from 1995-96 is from the Agricultural Census 2000-01. Hard-left refers to cumulative years (in 1997) since 1957 that hard left parties were in majority in the state legislature and labor regulation score is a measure of whether the state regulations is pro-worker, neutral or pro-employer (from Aghion et al. (2008)).

3.3 Empirical strategy

Rainfall shocks are exogenous to a firm's decisions – it is unlikely that a particular firm's decision will affect the weather. Hence, the demand shock arising from the rainfall shock is also exogenous to a firm's hiring decisions. Good rainfall (or positive rainfall shocks) will lead to higher agricultural income or monthly per capita expenditure, leading to greater demand for industrial goods and hence for industrial labour (a labour demand shock). However, good rainfall might also lead to a higher demand for agricultural labor, thus resulting in a labour supply shock for industrial sector firms. Therefore, whether rainfall shocks lead to a demand shock or a supply shock for firms depends on whether the demand shock dominates the supply shock. A sufficient condition for demand shock to dominate the supply shock is to observe that, in equilibrium, both industrial wages and employment increase.[12] Figures 3.2 and 3.3 illustrate this point.

In both the figures, labour demand curve shifts to the right, from DD to $D'D'$, representing the labour demand shock. Figure 3.2 illustrates the positive labour supply shock – supply curve shifts to the right from SS to $S'S'$, while Figure 3.3 demonstrates the negative labour supply shock – supply curve shifts to the left from SS to $S'S'$. In Figure 3.2, since both the demand and supply curves shift to the right, equilibrium employment will definitely increase. Whether equilibrium wages will also increase depends on the relative sizes of the shifts. If the shift of the labour

[12]Chaurey (2015) discusses three possible scenarios on the supply side. In the first case, no agricultural worker may move to the industry, resulting in a rise in industrial wages and employment owing to the larger demand for industrial labour. The second case is when agricultural workers move to the industry. In this case, as long as demand exceeds the supply of workers in the industry, there will be an increase in both industrial wages and employment. In the third case, rainfall attracts labour from industries to agriculture so that the industries face a negative supply shock. However, if the industrial wages and employment have increased, then the increase in labour demand must be greater than the decrease in labour supply.

supply curve is lower than the shift of the labour demand curve (distance y between SS and $S'S'$ is less than distance x between DD and $D'D'$), equilibrium wages increase. On the other hand, if the shift of the labour supply curve exceeds the shift of the labour demand curve (distance $x + e$ between SS and $S''S''$ is more than distance x between DD and $D'D'$), equilibrium wages decrease. Clearly, the only way both equilibrium wages and employment can increase is when demand shock dominates the supply shock. In Figure 3.3, since demand increases while supply decreases, equilibrium wages will definitely increase. What happens to equilibrium employment, on the other hand, depends on the relative sizes of the demand and supply shocks. Following a similar analysis as above here also we can conclude that both equilibrium wages and employment can increase only when demand shock dominates the supply shock.

Chaurey (2015) shows that both industrial wages and employment rise for the period 2000-2010 in response to positive rainfall shocks (the converse holds for negative rainfall shocks).[13] Thus, he establishes that rainfall shocks do represent a demand shock for the industrial sector in India for the period 2000-2010. Therefore, for our study, we consider rainfall shocks as exogenous transitory demand shocks for firms.

We analyze how a firm's contract employment varies by computer capital expenditure interacted with whether the district (in which the firm is located) has experienced a demand shock. As a measure of a firm's computer capital expenditure, we construct a dummy ($Compdum_{idt}$) that takes value 1 when the computer

[13] Both industrial wages and monthly per capita expenditure rise for the period 2000-2010 following positive rainfall shocks. The positive impact of rainfall shock on monthly per capita expenditure signifies an increase in spending, and possibly an increase in demand for industrial goods. There is no statistically significant impact of rainfall shocks on agricultural wages for the same period. He shows that industrial employment increases (decreases) in response to positive (negative) rainfall shocks for a subset of the period (1999-2008), confirming that there is, in net, a demand shock for firms when rainfall shocks occur.

capital expenditure share of firm i in district d in year t is above the mean of the sample (2.88%), and 0 otherwise. We run the following regression:

$$\begin{aligned} Contract_{idt} = \gamma_t + \delta_i + \lambda_{kt} + \beta_1 Rainshock_{dt-1} + \beta_2 Compdum_{idt} \\ + \beta_3 Rainshock_{dt-1} * Compdum_{idt} + \epsilon_{idt}, \end{aligned} \tag{3.2}$$

where $Contract_{idt}$ measures the contract employment of firm i in district d in year t. The variable $Rainshock_{dt-1}$ refers to lagged rainfall shock which proxies for transitory demand shock. We include year fixed effects, γ_t, to control for macroeconomic year-specific events in our regressions. We also include firm fixed effects, δ_i, to control for time-invariant firm characteristics. Industry-year fixed effects, λ_{kt}, control for unobserved time-varying factors in an industry (such as changes in tariffs or regulations) which may affect both employment and computer capital investment in an industry.[14] We cluster standard errors at the district-level given that the rainfall shock is measured at the district-level. The impact of lagged rainfall shock on contract employment is captured by β_1. The coefficient β_2 captures the effects of above-average computer capital expenditure on contract employment. Our coefficient of interest β_3 captures the differential effect on contract hiring in response to demand shocks for firms that have above-average computer capital share as compared to firms that have lower-than-average computer capital share.

In an ideal setting, if computer capital expenditure was randomly assigned to each firm then one could easily regress employment on computer capital to obtain causal estimates. However, in reality, the number of workers to hire and how much to invest in computers are both decisions taken by the firm. It is possible that workers may choose to leave when they anticipate that a firm will invest in an automation

[14]Firms are classified into industries according to the National Industrial Classification (NIC). Description of industries is provided in Appendix A3.1.

technology. Alternatively, a firm may start hiring workers who have computer-based skills in anticipation of increased computer capital expenditure. Also, unobserved labour market conditions may affect both computer capital investment and employment leading to endogeneity issues. Thus, simply regressing contract employment on computer capital share will not capture the hiring decisions taken by a firm due to computer capital investment.

In this study, identification comes from the interaction of computer capital dummy and rainfall shocks. Interaction of an endogenous variable (computer capital expenditure) with an exogenous variable (rainfall shock) is exogenous under mild assumptions (Nizalova and Murtazashvili (2016)). In equation 3.2, β_3 gives us causal estimates.

We also check with a host of alternative measures such as computer capital share, a dummy for computer capital share above the median computer capital share, a dummy for computer capital share above the industrial average and a dummy for computer capital share above the average computer capital share in a state in an industry. We also consider variables that classify firms based on the computer capital share in the first sampling year of a firm. As further checks, we construct dummies based on computer capital shares in US industries for the period 2002-2010.

3.4 Results

3.4.1 Main results

We document the relationship between firm-level contract employment and firm-level computer capital expenditure share in response to rainfall shocks. Table 3.3 presents the results of regressing contract employment on rainfall shocks and rainfall shocks interacted with the measure of firms with above-average computer capital share, as specified in equation (3.2). Column 1 runs the regression equation (3.2) with firm fixed effects and year fixed effects. Column 2 presents our preferred specification with firm and year fixed effects, as well as industry-year fixed effects. We find that the coefficient for rainfall shock is statistically significant and positive: positive (negative) rainfall shocks increase (decrease) contract employment, which is in line with previous literature (Adhvaryu et al. (2013), Chaurey (2015)). Following positive rainfall shocks, firms face a positive demand shock and on average hire 1.31 more contract workers. As discussed above, $compdum_{idt}$ takes the value 1 if computer share is above-average (above 2.88%) and 0 otherwise. The interaction term, computer capital dummy interacted with rainfall shock, is statistically significant and negative. Firms on average hire more contract workers in response to a positive demand shock. The negative coefficient on the interaction term implies that compared to firms with lower-than-average computer capital share, firms with above-average computer capital share hire fewer contract workers following positive rainfall shocks. We find that a firm with above-average computer capital share hires 2.32 fewer contract workers compared to firms with lower-than-average computer capital share, in response to a positive rainfall shock. This differential hiring amounts to 11.33% of the sample mean of contract employment (20.47).

As a check, we look at how rainfall shocks affect total sales or output of firms. We mainly check that firms are not facing differential impact on their output when exposed to rainfall shocks. In Table 3.4 column 1, as expected, the value of sales go up when a firm is exposed to a rainfall shock. This reflects the demand shock arising from the rainfall shock. In column 2, we check if the interaction term of computer capital and rainfall shock is also significant, as this would imply that the rainfall shock has differential impact on firms (depending on their computer capital investment). This is not the case: the interaction term is statistically insignificant implying that rainfall shocks have a similar impact on all firms with respect to their output. We also make an additional check with an alternative measure of permanent work versus contract work - intensity of mandays. We find that the results are in line with Chaurey (2015), there is no significant impact on permanent work but there is a significant impact on contract work (discussed in Appendix A3.2).

Table 3.5 presents our regression results separately on rural firms (firms located in the rural sector) and urban firms (firms located in the urban sector). Rainfall shock leads to a statistically significant increase in contract hiring for rural firms whereas there is no statistically significant impact on contract hiring for urban firms. It is not surprising that the impact will be stronger for rural firms since the mechanism of a rainfall shock translating to a demand shock for the industrial sector is through changes in income and spending from the agricultural sector.[15] The coefficient on the interaction term is statistically significant and negative for rural firms: firms with above-average computer capital share hires 4.74 fewer contract workers compared to firms with below-average computer capital share in the face of demand shocks. For urban firms, however, we do not find any statistically significant impact.

We also check our results separately for positive rainfall shocks and negative rainfall shocks. For this exercise, positive rainfall shock takes value 1 if rainfall in district d in year t is above the 80^{th} percentile of rainfall distribution, and 0 otherwise, while negative rainfall shock takes value 1 if rainfall in district d in year t is below the 20^{th} percentile of rainfall distribution, and 0 otherwise. The regression

[15] This result is similar to the findings in Chaurey (2015). He finds a statistically significant impact of rainfall on rural firm employment but no statistically significant impact on urban firm employment. This provides further support that rainfall shocks are demand shocks arising due to the impact of rainfall on the agricultural sector.

results are reported in Table 3.6. In column 1 we observe that firms with above-average computer capital share tend to hire 2.62 fewer contract workers compared to firms with less-than-average computer capital share in response to a positive demand shock. Note that in case of negative rainfall shocks, firms on average tend to layoff contract workers as seen from the statistically significant and negative coefficient of 1.93 in column 2. The interaction term has a statistically significant and positive coefficient of 2.86. As firms on average lay-off workers during negative demand shocks, the positive coefficient on the interaction term implies that firms with relatively above-average computer capital share tend to lay off fewer contract workers following negative demand shocks, compared to firms with less-than-average computer capital share.

We also check how contract hiring differs between machinery production industries and non-machinery production industries.[16] In Table 3.7 we report results of our main regression for contract workers employed in non-machinery industries in column 1 and contract workers employed in machinery industries in column 2. It is interesting to note that the coefficient of the interaction term is negative and statistically significant for machinery industry contract workers (and not statistically significant for non-machinery industry contract workers). Machinery industries with above-average computer capital investment hire 4.02 contract workers less in the face of demand shocks. This points to our results being driven by machinery producing industries compared to non-machinery industries.

An important question may arise following our results: are firms adopting labour-

[16]Machinery production industries include, for example, industries producing office accounting machinery, communication devices, electrical machinery and transport equipment, whereas non machinery production industries include industries producing food and beverages, tobacco products, textiles, rubber products and chemical products. Non-machinery industries include industries up to classification number 28 in Appendix A3.1, whereas classification number 29 onward are considered machinery industries.

saving technologies in their production process when investing in computer capital? One way that we check this is by looking at the impact of rainshock on contract work in the core production processes and the non-core production processes. The ASI data records information on man-days (the number of workers in a day summed over all days in a year) separately for manufacturing and non-manufacturing works. Manufacturing work includes activities that are directly relevant to the main production process (considered as "core" activities), whereas non-manufacturing work (considered as "non-core" activities) includes peripheral works such as security, cleaning services, etc.[17] Contract workers on average are mostly involved in the core manufacturing process (Table 3.1). We check if the interaction of our computer capital dummy and rainshock has different or similar impacts across the two types of contract work.

In Table 3.8 we find that the coefficient on the interaction term is statistically significant and negative for manufacturing contract work (column 1). Firms on average increase the number of contract manufacturing man-days by 408.47 when faced with a demand shock. However, firms with above-average computer capital share carry out contract work by 715.02 fewer man-days than other firms facing a demand shock. Impact on manufacturing man-days would imply that contract employment is mainly affected for the core production processes, essentially the contract workers on the factory floor. This could possibly point to an automation impact of computer capital investment, as contract hiring for the main production process goes down. On the other hand, we do not find any statistically significant impact of the interaction term on non-manufacturing man-days. Thus, it is not the case that contract hiring for peripheral activities are being affected due to computer capital investment

[17] The terminology "core" and "non-core" was used by Chaurey and Soundararajan (2019).

in response to demand shocks. This confirms that contract hiring by firms following demand shocks is relatively lower in the main production process, possibly pointing to a labour-saving effect of computer capital in the main production process.

Why do firms differ in hiring decisions in response to transitory demand shocks depending on their level of computer capital expenditure? It is quite possible that more technologically advanced firms do not require as many contract workers in their manufacturing process.[18] A firm investing in above-average computer capital share may require relatively fewer contract workers for production compared to firms with lower-than-average computer capital share. Following positive demand shocks, firms hire more contract workers to cater to an increase in demand for industrial products. However, firms investing in computer capital may require relatively fewer contract workers to produce industrial products. Therefore, these firms will also demand relatively fewer contract workers to produce additional output facing demand shocks. Thus, even when there is a demand surge, these firms do not need to hire as many workers to cater to the increase in demand.

We also carry out robustness checks and present them in Table 3.9. Firms of certain industries may choose to locate in certain states, leading to different industry distributions between states. We add state-industry fixed effects to control for non-random location decisions of firms and find that our results are similar (column 1). Firms may be subject to changes in state specific labour laws. We account for this by adding state-year fixed effects and our results remain unchanged (column 2). As a robustness check, we also control for the share of agricultural inputs a firm purchases in a given year. Rainfall shocks affect agriculture yield and agricultural productivity. As a result, supply and price of agricultural products are affected during rainfall

[18]We are referring to the labour-substitution effect of technology which removes workers from the manufacturing process (Acemoglu and Autor (2011), Aghion et al. (2020)).

shocks. Some firms are relatively more dependent on inputs from the agriculture sector. For example, relatively more inputs are required from the agricultural sector for firms in the food and beverages industry and textile industry. These firms will face a change in the supply and price of most of their inputs during rainfall shocks, compared to firms not dependent on the agriculture sector for inputs. Firms that use agricultural inputs may account for these changes and optimize their costs during rainfall shocks and then hire workers as required. We control for such price and quantity changes on the input side for a firm. In column 3 we control for the share of agricultural inputs in a firm's total quantity of inputs in a given year, while in column 4 the share of expenditure on agricultural inputs in the total expenditure on inputs is controlled for.[19] Our results are not sensitive to adding agricultural input controls. It is also possible that baseline characteristics such as political control and income inequality may be correlated with computer capital investment decisions of firms in the face of shocks. To account for this we consider measures such as the rural Gini coefficient (1994), a dummy for left parties that were in majority in the state legislature till 1997, landholdings in 2000-01 for each firm and interact with rainfall shocks. In Table 3.10, we show that our results are not sensitive to adding baseline characteristics interacted with rainfall shocks.[20]

[19] The ASI records quantity and purchase value of inputs. Inputs are classified according to the Annual Survey of Industries Commodity Classification (ASICC). We consider ASICC 2-digit codes from 11 to 16 as agricultural inputs.

[20] The controls include rural Gini coefficient (from the EOPP Indian States Database, also used in Besley and Burgess (2004)), total landholdings of a state in 2000-01, percentage variation of holdings in 2000-01 from 1995-96 (Agricultural Census 2000-01), cumulative years (in 1997) since 1957 that hard left parties were in majority in the state legislature (from Aghion et al. (2008)).

3.4.2 Alternative measures

Next we construct a host of alternative measures of a firm's computer capital expenditure share, and run the main regression with these alternative measures. These regression results are reported in Table 3.11. In column 1, $Compshare_{idt}$ is the computer capital share in total capital of firm i in district d in year t. It is a continuous variable as opposed to the dummy variable $Compdum_{idt}$ that we have used so far. The interaction term (computer capital share and rainfall shock) is again statistically significant and negative. This implies that compared to firms that invest relatively less in computer capital, firms that invest more in computer capital are less likely to hire as many contract workers in response to positive rainfall shocks. An increase of one standard deviation in the share of computer capital expenditure of a firm leads to a decline in contract hiring by 1.24 (0.28*4.41) following positive demand shocks. Our result also implies that a 10 percentage point increase in computer capital share leads to a reduction in average contract hiring by 2.8 workers in the face of demand shocks. In column 2, we consider a dummy $Compmed_{idt}$ that takes value 1 if the computer capital share of a firm in a year is above the median computer capital share (1.31%) of the sample, and 0 otherwise. That is, we use the sample median instead of the sample mean to define our computer capital dummy. Again, the coefficient is negative and statistically significant. A firm with computer capital share above the median hires 1.7 fewer contract workers compared to firms with lower-than-median computer capital share in the face of demand shocks.

We also construct another dummy ($Compind_{idt}$) that takes value 1 if computer capital share of a firm in a year is above the industrial average computer capital share in the sample, and 0 otherwise. This measure considers that the average computer capital share may vary across industries. Some industries invests more in computer capital depending on the nature of goods produced by these industries compared to others.[21] In column 3 of Table 3.11 we find that the interaction term is statistically significant and negative. A firm that has computer capital share above its industry average hires 2 fewer contract workers in the face of a positive rainfall shock. We also consider the average computer capital share within each state and industry. We construct a dummy $Compindstt_{idt}$ that takes value 1 if a firm's computer capital share in a year is above the average computer capital share of its state and industry, and 0 otherwise. This alternative measure considers that average computer capital expenditure could vary across states also (which may be due to various reasons such as state laws and resources). Once again our results are similar to the previous measures. Column 4 of Table 3.11 shows that the coefficient on the interaction term with this measure is statistically significant and negative. Contract hiring is lower by 1.55 when a firm has above average computer capital expenditure in its state and industry.

We now use the US data on average computer capital for the period 2002-2010 to construct another set of alternative measures of computer capital share. The US based measures may be far more stringent since the US computer capital investment in an industry is relatively more than India as the US is far more advanced technologically.[22] The results are reported in Table 3.12. Column 1 considers the dummy

[21] For example, firms in food and beverage industry have lower computer capital share than firms in the publishing and printing industry. For a list of average computer capital shares by industry, please refer to Appendix A3.3.

[22] For example, the average US computer capital share for the period 2002-2010 is 6.02%, whereas the average computer capital share for the same period is only 2.84% for Indian firms. In Appendix A3.4, we compare five industries with the highest computer capital shares in year 2002 between US and India. We make a similar comparison for 2010. We find that the magnitude of computer capital shares for the Indian industries is much lower than that of the US for both years. We also find that while the industries are relatively different for India and US in 2002, but they have become relatively similar in 2010. This may suggest that India's computer-aided technology pace is slow and is becoming similar to the US over time.

$CompdumUS_{idt}$ that takes value 1 if the computer capital share of a firm in a year is above the average US computer capital share for the period 2002-2010, and 0 otherwise. In column 2 we consider a dummy $CompyrUS_{idt}$ that takes value 1 if the computer capital share of a firm in a year is above the US average computer capital share for that year, and 0 otherwise. In column 3 we run our regression with the dummy $CompindUS_{idt}$, which takes value 1 if the computer capital share of a firm in a year is above the corresponding US industrial average for the period 2002-10, and 0 otherwise. Finally, column 4 considers variation at the industry and year level: $CompindyrUS_{idt}$ takes value 1 if the computer capital share of a firm in a year is above the corresponding US industrial average for that year, and 0 otherwise. In all cases, irrespective of heterogeneity based on the period, year or industry, we find that the coefficient of the interaction term is statistically significant and negative. Firms that are technologically advanced (in terms of computer capital investment) similar to levels of the US, do not hire as many contract workers when exposed to transitory demand shocks. Thus, our results are robust when we consider exogenous US based measures for creating a dummy capturing computer capital heterogeneity across firms. Our corresponding results based on Indian computer capital measures also hold for the period 2002-2010 and are reported in Appendix A3.5.

The above described measures are time variant since the dummy is allowed to change in any year if the firm invests more (or less) than the sample average in that year. We also consider a set of measures that classify firms based on the firm's computer capital expenditure at the starting year. These measures do not account for any temporal changes in computer capital expenditure of a firm. We construct an alternative measure $compalt_{id}$ that takes value 1 if the computer capital share of the firm in the first sampling year is above the average computer capital share in the first sampling year, and 0 otherwise. Column 1 of Table 3.13 reports the coefficient on the interaction term for this dummy and the results are similar as before: statistically significant and negative. We also classify firms based on the industrial average computer capital expenditure in the first sampling year of the firm: $compaltind_{id}$ takes value 1 if the firm's computer capital share in the first sampling year is above its industry average computer capital share in the first sampling year, and 0 otherwise. As column 2 of Table 3.13 demonstrates, our results are robust to this measure too.

From the above analysis we can conclude that our results are not dependent on the type of computer capital measure we use. Our results are similar and consistent across a host of alternative computer capital measures. Given our estimates, we can conclude that computer capital investment allows a firm to reduce contract worker hiring during transitory demand shocks. It is possible that firms investing in computer capital are less dependent on contract workers when producing output. Therefore, when a demand shock arises, firms investing in computer capital do not need to employ as many additional contract workers to produce the additional output. Hence, these firms hire fewer contract workers to cater to short term changes in output in response to demand shocks, compared to firms that invest relatively less in computer capital. Firms investing in computer capital seem to be relatively protected against demand shocks, positive as well as negative, and can therefore avoid flexible hiring or firing of contract workers (in the sense that they do not need to rely on availability of contract workers in response to demand fluctuations).

3.5 Conclusion

The literature has focused on the impact of different types of technology on employment – from computers (Autor et al. (1998), Autor et al. (2003)) to robotics and artificial intelligence (Acemoglu and Restrepo (2020), Acemoglu et al. (2020)). There is no doubt that establishments are adopting technologies to secure a leaner manufacturing process, as automation has already displaced workers from their occupations on a global scale (ILO-IOE (2019)). However, in developing countries or emerging economies, technology adoption might not be homogeneous across all firms given political, organizational or other factors. Even in developed countries, adopting certain technologies is more prevalent in specific sectors (for example, robotics in automobile sector (Acemoglu and Restrepo (2020))). Thus, differences in firm decisions may arise due to heterogeneous technology adoption.

We study how a firm's hiring decisions in response to demand shocks are affected by investment in computer capital in India for the period 2000-2010. Transitory demand shocks, business cycles or other economic disturbances affect a firm's decisions on employment, output and prices. Given that firms are innovating, examining how hiring decisions in response to demand shocks have changed due to heterogeneous technology adoption is important. We focus on computer capital in our study, as the use of computers and computer-aided technology has risen significantly post World War II (Acemoglu and Autor (2011)). In India, analyzing computer technology (instead of say, robotics) is more relevant given India's technological pace (Berman et al. (2005)) and growth in demand for computer-aided technologies (Erumban and Das (2020), IBEF (2008)).

As we are looking at transitory demand shocks, we focus on contract workers. Contract workers are not covered by India's primary labour regulation, the Industrial

Disputes Act (1947), unlike permanent or regular workers. Hence, firms tend to hire or fire contract workers during transitory demand shocks to cater to any short-term demand changes (Chaurey (2015)). Following the previous literature (Adhvaryu et al. (2013), Chaurey (2015)), we construct lagged rainfall shocks to proxy for demand shocks. We find that firms with above average computer share tend to hire or fire less in response to demand shocks. The reduction is by 2.32 workers which is 11.33% of the sample mean. Our main results are not sensitive to controlling for non-random location of industries, state specific labour laws, expenditure on agricultural inputs and the source of rainfall data.

As robustness checks, we construct a host of alternate computer capital investment measures. We find that contract employment following rainfall shocks is negatively associated with computer capital expenditure share in total expenditure. A 10 percentage point increase in computer capital expenditure share is associated with a decline in hiring and firing in response to demand shocks for firms by 2.8 contract workers. We also construct alternate measures such as a dummy for above-median computer capital share, industrial average computer capital share, industry-state average computer capital share and a number of US based alternative measures. We also classify firms based on their computer capital in the first sampling year. Our results are not sensitive to the type of measure we use.

Our study shows that heterogeneity in firm decisions of contract employment arises due to differences in computer capital expenditure. Reduction in hiring in response to positive demand shocks seem to imply that firms are investing in labour-saving technologies. As firms do not need to hire as many contract workers facing positive demand shocks, it may imply that computer capital investing firms do not require contract workers in parts of their manufacturing process. Our results point to firms possibly automating segments of their manufacturing process when investing in computer capital. Even though our results are for the period 2000-2010, they seem to be in line with recent reports about automation displacing low-skilled workers (ILO-IOE (2019)). Thus, Indian firms with computer capital investments seem to be adopting leaner manufacturing processes in terms of reducing their dependence on contract employment in response to transitory demand shocks.

Figures and Tables for Chapter 3

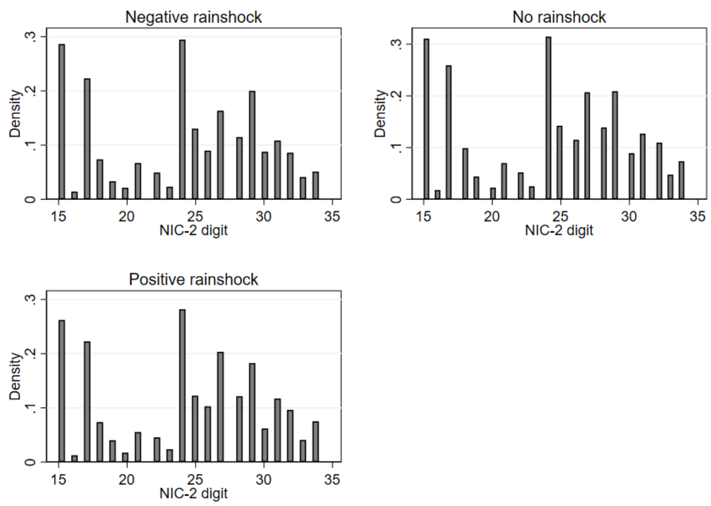

Figure 3.1: *2-digit industries distribution across different rain shock years*

Note: Data sources are the Annual Survey of Industries (ASI), rounds 1999-2000 to 2009-2010. The 2-digit industry classification is provided in Appendix A3.1.

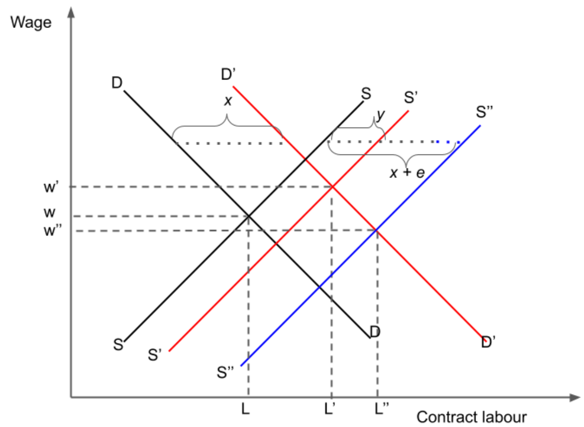

Figure 3.2: *Positive labour demand and positive labour supply*
Note: Supply shift y is less than the demand shift x. Supply shift $x + e$ is more than the demand shift x.

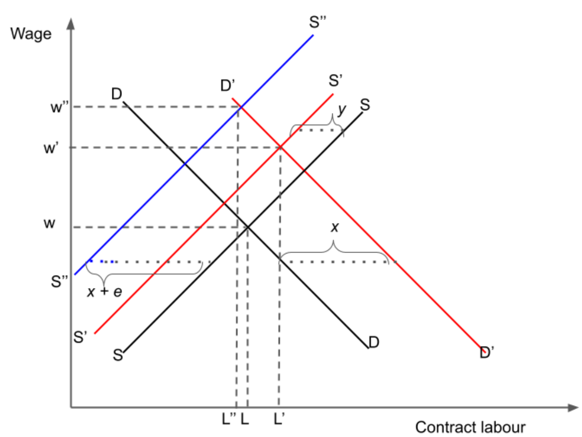

Figure 3.3: *Positive labour demand and negative labour supply*
Note: Supply shift y is less than the demand shift x. Supply shift $x + e$ is more than the demand shift x.

Table 3.1: *Summary statistics 1*

	(1) All firms	(2) Firms with above avg computer share	(3) Firms with below avg computer share
Total capital (mil Rs.)	85.59	49.45	99.53
	(232.52)	(160.16)	(253.64)
Computer capital (mil Rs.)	1.37	2.78	0.83
	(4.18)	(6.25)	(2.84)
Computer capital share	2.88	7.77	1.00
in total capital (%)	(4.41)	(5.94)	(0.76)
Regular workers (number)	47.55	43.14	49.25
	(86.66)	(84.30)	(87.50)
Contract workers (number)	20.47	16.34	22.06
	(59.39)	(52.55)	(61.75)
Contract workers share	18.96	16.88	19.76
in total workers (%)	(32.41)	(31.31)	(32.80)
Contract manufacturing	6317.21	4979.61	6833.39
man-days	(18625.07)	(16239.97)	(19443.03)
Contract non-manufacturing	0.67	0.23	0.84
man-days	(9.23)	(5.56)	(10.30)
Rural firms (fraction)	0.40	0.23	0.40
	(0.49)	(0.42)	(0.49)
Negative rainshock (fraction)	0.16	0.16	0.16
	(0.36)	(0.36)	(0.36)
Positive rainshock (fraction)	0.20	0.19	0.20
	(0.40)	(0.40)	(0.40)
Observations	81635	21780	59855

Note: Data is from the Annual survey of Industries (ASI) from round 1999-2000 to 2009-2010. Observations are weighted by ASI sample weights. Standard deviation in parentheses. All employment and capital values above the 99th percentile are equated to the 99th percentile value to remove influence of outliers. Column 1 includes all firms. Column 2 includes firms with computer capital share greater than average computer capital share of all firms. Column 3 includes firms with computer capital share less than or equal to average computer capital share of all firms. Total workers include contract workers and regular workers.

Table 3.2: *Summary statistics 2*

	(1) Mean
Rural gini coefficient (1994)	28.16
	(2.81)
Total landholding (2000-01)	12333.38
	(6394.74)
Percentage variation in landholdings in 2000-01 over 1995-96	-0.91
	(2.51)
Hard-left (1997)	0.49
	(0.50)
Labour regulation score	0.40
	(1.69)
Observations	80745

Note: Data on Rural gini coefficient is from the EOPP Indian States Database, also used in Besley and Burgess (2004)). Total landholdings of a state in 2000-01 and percentage variation of holdings in 2000-01 from 1995-96 is from the Agricultural Census 2000-01. Hard-left refers to cumulative years (in 1997) since 1957 that hard left parties were in majority in the state legislature and labor regulation score is a measure of whether the state regulations is pro-worker, neutral or pro-employer (from Aghion et al. (2008)).

Table 3.3: *Firm-level computer measures*

	(1)	(2)
Outcome: Contract workers		
$Rainshock_{dt-1}$	1.26**	1.31***
	(0.49)	(0.50)
$Compdum_{idt}$	-0.32	-0.47
	(1.19)	(1.17)
$Compdum_{idt} * Rainshock_{dt-1}$	-1.98***	-2.32***
	(0.74)	(0.82)
Observations	64340	64340
R^2	0.770	0.772
Firm FE	Yes	Yes
Year FE	Yes	Yes
Industry year FE	No	Yes
Age controls	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Observations are weighted by sample weights. Firm-level data is from the Annual Survey of Industries (2000-2010). $Compdum_{idt}$ takes value 1 if the computer capital share of a firm in year t is above the average computer capital share of the sample (2000-2010). Rainfall data is from the Center for Climatic Research University of Delaware. $Rainshock_{dt-1}$ takes value 1 if annual rainfall in district d in year $t-1$ is above the 80^{th} percentile, -1 if rainfall in a district in year $t-1$ is below the 20^{th} percentile of the rainfall distribution, and 0 otherwise. Age controls includes the age and age squared of a firm.

Table 3.4: *Rainfall shock and output sold*

	(1)	(2)
$Rainshock_{dt-1}$	114*	124
	(64.3)	(98.5)
$Compdum_{idt}$		241
		(296)
$Compdum_{idt} * Rainshock_{dt-1}$		-44.8
		(263)
Observations	59536	59536
R^2	0.138	0.138
Firm FE	Yes	Yes
Year FE	Yes	Yes
Industry year FE	Yes	Yes
Age controls	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Observations are weighted by sample weights. Age controls includes the age and age squared of a firm. Sales value refers to the value of the output (product of the price per unit of the good sold and quantity of the good sold). Figures reported in Rs. 10 crore.

Table 3.5: *Rural and urban sectors*

	(1)	(2)
Outcome: Contract workers	Rural	Urban
$Rainshock_{dt-1}$	1.91**	0.90
	(0.81)	(0.55)
$Compdum_{idt}$	-1.58	0.01
	(2.45)	(1.42)
$Compdum_{idt} * Rainshock_{dt-1}$	-4.74***	-1.34
	(1.70)	(0.83)
Observations	26685	33474
R^2	0.791	0.772
Firm FE	Yes	Yes
Year FE	Yes	Yes
Industry year FE	Yes	Yes
Age controls	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Observations are weighted by sample weights. Column 1 restricts the sample to rural sector firms. Column 2 restricts the sample to urban sector firms.

Table 3.6: *Asymmetric rainfall shocks*

Outcome: Contract workers	(1) Positive shock	(2) Negative shock
$PositiveRainshock_{dt-1}$	1.12	
	(0.75)	
$Compdum_{idt} * PositiveRainshock_{dt-1}$	-2.62**	
	(1.09)	
$NegativeRainshock_{dt-1}$		-1.93***
		(0.63)
$Compdum_{idt} * NegativeRainshock_{dt-1}$		2.86**
		(1.19)
$Compdum_{idt}$	0.00	-0.97
	(1.20)	(1.14)
Observations	64340	64340
R^2	0.772	0.772
Firm FE	Yes	Yes
Year FE	Yes	Yes
Industry year FE	Yes	Yes
Age controls	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Observations are weighted by sample weights. $PositiveRainshock_{dt-1}$ takes value 1 if rainfall in district d in year $t-1$ is above the 80^{th} percentile of the distribution, and 0 otherwise. $NegativeRainshock_{dt-1}$ takes value 1 if rainfall in district d in year $t-1$ is below the 20^{th} percentile of the distribution, and 0 otherwise.

Table 3.7: *Machinery industries*

Outcome: Contract workers	(1) Non-machinery	(2) Machinery
$Rainshock_{dt-1}$	0.94**	2.63**
	(0.47)	(1.28)
$Compdum_{idt}$	-0.50	-0.20
	(1.50)	(1.87)
$Compdum_{idt} * Rainshock_{dt-1}$	-1.72	-4.02**
	(1.07)	(1.63)
Observations	46591	16911
R^2	0.778	0.764
Firm FE	Yes	Yes
Year FE	Yes	Yes
Industry year FE	Yes	Yes
Age controls	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Observations are weighted by sample weights. Non-machinery industries include food, beverage, textile, wood, coke, basic metals (up to 2-digit classification 28 in Appendix A3.1). Machinery industries include office, accounting, computing, communication, medical, electrical machinery, transport, manufacturing not elsewhere classified (from 2-digit classification 29 in Appendix A3.1).

Table 3.8: *Manufacturing versus non-manufacturing contract work*

Outcome: Contract Man-days	(1) Manufacturing Man-days	(2) Non-manufacturing Man-days
$Rainshock_{dt-1}$	408.47**	0.08
	(165.62)	(0.18)
$Compdum_{idt}$	-157.45	-0.18
	(362.27)	(0.14)
$Compdum_{idt} * Rainshock_{dt-1}$	-715.02***	-0.16
	(258.36)	(0.17)
Observations	64340	64340
R^2	0.769	0.520
Firm FE	Yes	Yes
Year FE	Yes	Yes
Industry year FE	Yes	Yes
Age controls	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Observations are weighted by sample weights. Man-days represent the total number of days worked and paid for during the year. It is computed by summing-up the number of workers per day over all days. Manufacturing refers to activities involved in the main production process. Non-manufacturing refers to activities that are peripheral to the production process.

Table 3.9: *Robustness checks 1*

Outcome: Contract workers	(1) State-industry	(2) State-year	(3) Agri quantity	(4) Agri price
$Rainshock_{dt-1}$	1.26**	1.55**	1.34***	1.33***
	(0.49)	(0.63)	(0.50)	(0.50)
$Compdum_{idt}$	-0.60	-0.51	-0.50	-0.50
	(1.18)	(1.18)	(1.18)	(1.17)
$Compdum_{idt} * Rainshock_{dt-1}$	-2.27***	-2.42***	-2.30***	-2.30***
	(0.82)	(0.81)	(0.81)	(0.82)
Observations	64330	64339	64154	64288
R^2	0.774	0.774	0.772	0.772
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Industry year FE	Yes	Yes	Yes	Yes
State industry FE	Yes	No	No	No
State year FE	No	Yes	No	No
Age controls	Yes	Yes	Yes	Yes
Agriculture quantity controls	No	No	Yes	No
Agriculture price controls	No	No	No	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Observations are weighted by sample weights. Agriculture quantity controls refers to share of agriculture inputs in total quantity of indigenous inputs of a firm in a year. Agriculture price controls refers to share of firm expenditure on agricultural inputs in total indigenous inputs of a firm in a year. ASI records quantity and purchase value of inputs. Inputs are classified according to the Annual Survey of Industries Commodity Classification (ASICC). ASICC 2-digit codes 11-16 are considered as agricultural inputs.

Table 3.10: *Robustness checks 2*

	(1)	(2)
Outcome: Contract workers		
$Rainshock_{dt-1}$	1.81	2.06
	(5.32)	(5.40)
$Compdum_{idt}$	-0.68	-1.26
	(1.23)	(1.27)
$Compdum_{idt} * Rainshock_{dt-1}$	-2.41***	-2.49***
	(0.90)	(0.94)
Observations	58146	50637
R^2	0.772	0.777
Firm FE	Yes	Yes
Year FE	Yes	Yes
Industry year FE	Yes	Yes
Age controls	Yes	Yes
Controls 1 × Rainshock	Yes	No
Controls 2 × Rainshock	No	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Observations are weighted by sample weights. In column 1, rainshock is interacted with controls 1 which include rural Gini coefficient (1994), total landholding (2000-01) and measure of cumulative years (in 1997) since 1957 that hard left parties in majority in the state legislature for a firm. In column 2, rainshock is interacted with controls 2 which include rural Gini coefficient (1994), percentage variation in landholdings in 2000-01 over 1995-96 and labour regulation score for a firm.

Table 3.11: *Alternative measures 1*

Outcome: Contract workers	(1) Share	(2) >Median	(3) >Industry average	(4) >Industry-state average
$Rainshock_{dt-1}$	1.48***	1.54**	1.28**	1.16**
	(0.52)	(0.63)	(0.52)	(0.52)
$Compshare_{idt}$	-0.17			
	(0.15)			
$Compshare_{idt} * Rainshock_{dt-1}$	-0.28***			
	(0.08)			
$Compmed_{idt}$		0.46		
		(1.06)		
$Compmed_{idt} * Rainshock_{dt-1}$		-1.70**		
		(0.73)		
$Compind_{idt}$			-0.79	
			(1.03)	
$Compind_{idt} * Rainshock_{dt-1}$			-2.00**	
			(0.82)	
$Compindstt_{idt}$				-0.84
				(0.89)
$Compindstt_{idt} * Rainshock_{dt-1}$				-1.55*
				(0.82)
Observations	64340	64340	64340	64340
R^2	0.772	0.772	0.772	0.772
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Industry year FE	Yes	Yes	Yes	Yes
Age controls	Yes	Yes	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Observations are weighted by sample weights. In column 1, $Compshare_{idt}$ refers to computer capital share for a firm in a year. Median computer capital share, industry average and industry-state average computer capital share are calculated for the sample (2000-2010). In column 2, $Compmed_{idt}$ is a dummy that takes value 1 if the firm's computer capital share in a year is above the median computer capital share, and 0 otherwise. In column 3, $Compind_{idt}$ takes value 1 if computer capital share of a firm in a year is above the average computer capital share in the industry of the firm, and 0 otherwise. In column 4, $Compindstt_{idt}$ takes value 1 if computer capital share of a firm in a year is above the average computer capital share in the industry and state of the firm, and 0 otherwise. Results are reported for the period 2000-2010.

Table 3.12: *Alternative measures 2*

Outcome: Contract workers	(1) >Period average	(2) >Year average	(3) >Industry average	(4) >Industry-year average
$Rainshock_{dt-1}$	1.06**	1.05**	1.10**	0.99*
	(0.50)	(0.50)	(0.49)	(0.51)
$CompdumUS_{idt}$	0.62			
	(1.42)			
$CompdumUS_{idt} * Rainshock_{dt-1}$	-2.60**			
	(1.10)			
$CompyrUS_{idt}$		1.16		
		(1.17)		
$CompyrUS_{idt} * Rainshock_{dt-1}$		-2.38**		
		(1.13)		
$CompindUS_{idt}$			-1.99**	
			(0.99)	
$CompindUS_{idt} * Rainshock_{dt-1}$			-2.76***	
			(0.91)	
$CompindyrUS$				-3.30***
				(0.96)
$CompindyrUS_{idt} * Rainshock_{dt-1}$				-1.70*
				(1.03)
Observations	59945	59945	59945	59945
R^2	0.780	0.780	0.780	0.780
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Industry year FE	Yes	Yes	Yes	Yes
Age controls	Yes	Yes	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Observations are weighted by sample weights. Data on average US computer capital share is from the Annual Survey of Manufactures. In column 1, $CompdumUS_{idt}$ is a dummy that takes value 1 if the firm's computer capital share in a year is above the average US computer capital share for the period 2002-2010, and 0 otherwise. In column 2, $CompyrUS_{idt}$ is a dummy that takes value 1 if the firm's computer capital share in a year is above the average US computer capital share for that year, and 0 otherwise. In column 3, $CompindUS_{idt}$ is a dummy that takes value 1 if the firm's computer capital share in a year is above the corresponding US industry average for the period 2002-2010, and 0 otherwise. In column 4, $CompindyrUS_{idt}$ is a dummy that takes value 1 if the firm's computer capital share in a year is above the corresponding US industry average for that year, and 0 otherwise. Results are for the period 2002-2010.

Table 3.13: *Alternative measures 3*

	(1)	(2)
Outcome: Contract workers	Sample average	Industry average
$Rainshock_{dt-1}$	1.10**	1.14**
	(0.50)	(0.50)
$Compalt_{id} * Rainshock_{dt-1}$	-1.67*	
	(0.89)	
$Compaltind_{id} * Rainshock_{dt-1}$		-1.73**
		(0.83)
Observations	55983	55983
R^2	0.870	0.870
Firm FE	Yes	Yes
Year FE	Yes	Yes
Industry year FE	Yes	Yes
Age controls	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Observations are weighted by sample weights. In column 1, $Compalt_{id}$ takes value 1 for a firm throughout the period if the computer capital share in the starting year of the firm in the sample is above the average computer capital share in that year. In column 2, $Compaltind_{id}$ takes value 1 for a firm throughout the period if the computer capital share in the starting year of the firm in the sample is above its industry average computer capital share. Results are reported for the period 2000-2010.

Appendix

A3.1 Industry categories

Firms are classified into industries according to the National Industrial Classification (NIC). In the ASI, firms are classified by the National Industrial Classification (NIC) which is available at the 5-digit level for most years in our sample. The higher the number of digits used to describe the industry the more detailed is the industry description. NIC has undergone some changes in its classification system over the years. Data for years before 2004 follow the NIC-98 classification whereas the data for 2004 to 2007 follows NIC-04. Year 2008 onward, firms are classified by NIC-08. To maintain consistency of industrial classification, we match industries at the 2-digit level, following the detailed industry descriptions provided by MoSPI (Ministry of Statistics and Programme Implementation). Table A1 lists the industry categories.

A3.2 Permanent and contract work intensity

We additionally check that the intensity of permanent work is not significantly different for firms with above average computer capital investment exposed to rainfall shocks. In Table A2, in column 1 we check for contract worker intensity. We can see that the interaction term is negative and statistically significant, implying that the fraction of contact mandays is reduced for firms with above average computer capital investment when exposed to rainfall shocks. In column 2, we check for permanent workers. We find that there is no statistically significant impact on permanent work intensity which is in line with Chaurey (2015).

A3.3 Average computer capital shares by industry

Table A3 reports the average computer share of a firm within its industry. Manufacture of office, accounting equipment, radio, television, communication equipment and medical and optical instruments record the highest average computer capital share (7.26). Manufacture of basic metals records the lowest average computer capital share (1.27).

A3.4 US and India average computer capital share

Table A4 reports the top five industries with the highest computer capital shares in the year 2002 for India (Panel A) and US (Panel B). Computer capital shares for the top five industries for India is lower than that of the US. Table A5 reports the top five industries with the highest computer capital shares in the year 2010 for India (Panel A) and US (Panel B). Again, the computer capital shares are much lower for Indian industries than for the US industries. It is interesting to note that the top five industries are similar across the two countries in the year 2010 as compared to the year 2002. This may suggest that investments in computer capital and possibly use of computer-aided technology used by Indian industries have become similar to the US over time, even though the share is much less than the US. The correlation between US and Indian yearly industrial average computer capital share for the period 2002-2010 is 0.63. If we consider a breakup of the sample period into three-year periods the correlation is 0.54 (2002-2004), 0.67 (2005-2007) and 0.69 (2008-2010) suggesting that India is "catching up" to US computer capital investment.

A3.5 Indian computer capital measures, 2002-2010

In Table A6, we report our results for the period 2002-2010 to check that our results do not change when we change the period of analysis from 2000-2010. Our results hold for all the measures. The measures are constructed in a manner similar to those constructed for the US measures in section 3.4.2.

Table A1: *2-digit Industry*

S. No.	Industry
15	Manufacture of food products and beverages
16	Manufacture of tobacco products
17	Manufacture of textiles
18	Manufacture of wearing apparel; dressing and dyeing of fur
19	Tanning and dressing of leather;
	manufacture of luggage, handbags, saddlery,harness and footwear
20	Manufacture of wood and of products of wood and cork, except furniture;
	manufacture of articles of straw and plaiting materials
21	Manufacture of paper and paper products
22	Publishing, printing and reproduction of recorded media
23	Manufacture of coke, refined petroleum products and nuclear fuel
24	Manufacture of chemicals and chemical products
25	Manufacture of rubber and plastics products
26	Manufacture of other non-metallic mineral products
27	Manufacture of basic metals
28	Manufacture of fabricated metal products, except machinery and equipment
29	Manufacture of machinery and equipment n.e.c.
30	Manufacture of office, accounting and computing machinery
	manufacture of radio, television and communication equipment and apparatus
	manufacture of medical, precision and optical instruments, watches and clocks
31	Manufacture of electrical machinery and apparatus n.e.c.
32	Manufacture of motor vehicles, trailers and semi-trailers
33	Manufacture of other transport equipment
34	Manufacture of furniture; manufacturing n.e.c.

Note: Industry at 2-digit level, concordance has been carried out between NIC-04 and NIC-08 2-digit categories.

Table A2: *Permanent and contract work intensity*

	(1)	(2)
Outcome:	Contract intensity	Permanent intensity
$Rainshock_{dt-1}$	0.00	-0.00
	(0.00)	(0.00)
$Compdum_{idt}$	-0.01	0.01
	(0.01)	(0.01)
$Compdum_{idt} * Rainshock_{dt-1}$	-0.01*	0.01
	(0.00)	(0.00)
Observations	64242	64242
R^2	0.805	0.807
Firm FE	Yes	Yes
Year FE	Yes	Yes
Industry year FE	Yes	Yes
Age controls	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Observations are weighted by sample weights. Age controls includes the age and age squared of a firm. Man-days represent the total number of days worked and paid for during the year. It is computed by summing-up the number of workers per day over all days. Contract intensity refers to contract worker mandays divided by total mandays. Permanent intensity refers to permenent worker mandays divided by total mandays.

Table A3: *Average computer capital share*

Industry	Average computer share
Manufacture of food products and beverages	1.62
Manufacture of tobacco products	3.79
Manufacture of textiles	1.42
Manufacture of wearing apparel; dressing and dyeing of fur	3.90
Tanning and dressing of leather; manufacture of luggage, handbags, saddlery,harness and footwear	2.89
Manufacture of wood and of products of wood and cork, except furniture; manufacture of articles of straw and plaiting materials	1.87
Manufacture of paper and paper products	2.10
Publishing, printing and reproduction of recorded media	5.84
Manufacture of coke, refined petroleum products and nuclear fuel	1.96
Manufacture of chemicals and chemical products	2.48
Manufacture of rubber and plastics products	1.98
Manufacture of other non-metallic mineral products	1.64
Manufacture of basic metals	1.27
Manufacture of fabricated metal products, except machinery and equipment	3.22
Manufacture of machinery and equipment n.e.c.	5.40
Manufacture of office, accounting and computing machinery, manufacture of radio, television and communication equipment and apparatus, manufacture of medical, precision and optical instruments, watches and clocks	7.26
Manufacture of electrical machinery and apparatus n.e.c.	5.16
Manufacture of motor vehicles, trailers and semi-trailers	2.91
Manufacture of other transport equipment	2.83
Manufacture of furniture; manufacturing n.e.c.	4.51

Note: Industry at 2-digit level, concordance has been carried out between NIC-04 and NIC-08 2-digit categories. Average computer capital share is the average computer capital share of firms within each industry (sample weights applied).

Table A4: *Top 5 industries by computer capital share, India and US 2002*

Industry	Computer capital share
Panel A: India	
Manufacture of office, accounting and computing machinery, manufacture of radio, television and communication equipment and apparatus, manufacture of medical, precision and optical instruments, watches and clocks	7.63
Publishing, printing and reproduction of recorded media	7.54
Manufacture of tobacco products	6.42
Manufacture of machinery and equipment n.e.c.	5.24
Manufacture of electrical machinery and apparatus n.e.c.	4.70
Panel B: US	
Manufacture of wearing apparel; dressing and dyeing of fur	13.93
Manufacture of furniture; manufacturing n.e.c.	12.59
Tanning and dressing of leather; manufacture of luggage, handbags, saddlery,harness and footwear	11.50
Publishing, printing and reproduction of recorded media	10.79
Manufacture of office, accounting and computing machinery, manufacture of radio, television and communication equipment and apparatus, manufacture of medical, precision and optical instruments, watches and clocks	9.72

Note: US computer capital shares by industry computed from Annual Survey of Manufactures data (2002). Indian computer capital shares constructed from Annual Survey of Industries data (2002).

Table A5: *Top 5 industries by computer capital share, India and US 2010*

Industry	Computer capital share
Panel A: India	
Manufacture of office, accounting and computing machinery, manufacture of radio, television and communication equipment and apparatus, manufacture of medical, precision and optical instruments, watches and clocks	7.69
Publishing, printing and reproduction of recorded media	5.46
Manufacture of machinery and equipment n.e.c.	5.34
Manufacture of furniture; manufacturing n.e.c.	5.02
Manufacture of electrical machinery and apparatus n.e.c.	4.68
Panel B: US	
Manufacture of wearing apparel; dressing and dyeing of fur	12.90
Manufacture of office, accounting and computing machinery, manufacture of radio, television and communication equipment and apparatus, manufacture of medical, precision and optical instruments, watches and clocks	12.17
Manufacture of furniture; manufacturing n.e.c.	11.15
Publishing, printing and reproduction of recorded media	10.80
Manufacture of machinery and equipment n.e.c.	7.29

Note: US computer capital shares by industry computed from Annual Survey of Manufactures data (2010). Indian computer capital shares constructed from Annual Survey of Industries data (2010).

Table A6: *India-based measures*

	(1)	(2)	(3)	(4)
Outcome: Contract workers				
$Rainshock_{dt-1}$	1.47***	1.46***	1.30**	1.26**
	(0.56)	(0.56)	(0.59)	(0.61)
$Compdum_{idt}$	-0.57			
	(1.14)			
$Compdum_{idt} * Rainshock_{dt-1}$	-2.73***			
	(0.92)			
$Compyr_{idt}$		-1.53		
		(1.12)		
$Compyr_{idt} * Rainshock_{dt-1}$		-2.70***		
		(0.94)		
$Compind_{idt}$			-1.25	
			(1.19)	
$Compind_{idt} * Rainshock_{dt-1}$			-1.88**	
			(0.90)	
$Compindyr_{idt}$				-1.00
				(1.08)
$Compindyr_{idt} * Rainshock_{dt-1}$				-1.72*
				(0.98)
Observations	59945	59945	59945	59945
R^2	0.780	0.780	0.780	0.780
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Industry year FE	Yes	Yes	Yes	Yes
Age controls	Yes	Yes	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Results are for the period 2002-2010. Data on computer capital share is from the Annual Survey of Industries. Observations are weighted by sample weights. In column 1, $Compdum_{idt}$ is a dummy that takes value 1 if the firm's computer capital share in a year is above the period average computer capital share, and 0 otherwise. In column 2, $Compyr_{idt}$ is a dummy that takes value 1 if the firm's computer capital share in a year is above the average computer capital share for that year, and 0 otherwise. In column 3, $Compind_{idt}$ is a dummy that takes value 1 if the firm's computer capital share in a year is above the industry average computer capital share for the period, and 0 otherwise. In column 4, $Compindyr_{idt}$ is a dummy that takes value 1 if the firm's computer capital share in a year is above the industry average computer capital share for that year, and 0 otherwise.

Chapter 4

The Heat is on: Temperature and Exam Scores in India

4.1 Introduction

Rising temperatures across the globe is an old but an extremely important concern since climate change is known to cause adverse effects to ecosystems and life. Heal and Park (2020) point out that there is extensive research on how climate affects economic outcomes through indirect channels such as rising sea levels affecting infrastructure or heat affecting crop yield. However, a relatively recent literature has emerged focusing on direct channels of heat affecting human physiology and cognitive ability (Graff Zivin et al. (2020), Park (2022)). Such channels can also have important consequences such as negatively affecting health and labour productivity, which in turn affects economic output (Deschênes and Greenstone (2011), Somanathan et al. (2021)). Dell et al. (2014) point out that understanding the relationship between temperature and economic performance becomes increasingly important given that the earth is recording warmer temperatures over time. In this context our study focuses on the effect of temperature on human capital formation by analyzing the impact of heat on exam scores. We explore how secondary exam scores (Class X) are affected by exam-time temperature in India for the period 2012 to 2015.

Secondary board exams in India is a particularly relevant context to understand the relationship between temperature and test scores. The exam scores used in this study are important for determining the future labour outcomes of Indian students. Our exam scores come from a unique administrative data set from the Central Board of Secondary Education (CBSE) which is the single largest board of education with an all-India presence. Our study focuses on secondary (or class X) exam scores. These exams are conducted nation-wide and externally[1] at specified dates and time in March. Scores from these exams allow a student to judge their performance in a standardized public exam conducted simultaneously across the nation. This exam serves as the first stepping stone for students in their career choice, as students judge their performance during their secondary exams and decide on their preferred stream choice for higher-secondary education. The secondary exam scores signal the aptitude or ability of a student, based on which the student can decide which subjects they should opt for in their higher-secondary education (Class XI and XII). Depending on the choice of higher-secondary subjects, the student is admitted into a college or university.[2] These scores also become important if a student wishes to shift to another

[1]Students appearing for board-based exams are assigned examination centres away from their school. The examination centre is another school (different from the school the student is enrolled in) and typically falls within 10 kms of the school (CBSE (2013)).

[2]While scores play an important role, stream or combination of subjects a student chooses primarily affects the future career path of the student. For example, a student who opts for a subject combination of English, Hindi, Economics, Accountancy and Business studies in higher-secondary may be admitted into a business administration course. However, the student will not be granted admission into a basic science course in many colleges as she does not have a combination of science-oriented subjects (such as Physics, Chemistry and Math) in her higher-secondary curriculum.

school (or another education board) for her higher-secondary education. The desire to change schools or boards can be due to many reasons such as changing location due to a parent's transferable job or if the other school has reputable teachers for the student's preferred subjects in higher-secondary education. Secondary exam scores play an important role in such cases: a student is permitted admission in another school for higher-secondary education, subject to her secondary exam scores.[3] More importantly, many colleges and universities not only require higher-secondary scores for admission, but also rely heavily on secondary exam scores during the admission process. Hence, these scores significantly affect the future career path of a student.

This chapter investigates the effect of temperature on exam scores during secondary exams for the period 2012-2015 in India. We regress exam scores for each subject of each student on exam-time temperature after adding a set of fixed effects such as student fixed effects, subject fixed effects and year fixed effects to control for student specific, subject specific and year specific unobserved characteristics. Identification comes from exogeneity of the temperature variable as it is highly unlikely that student exam scores will affect temperature. Also, since the CBSE sets the exam schedule, the exam dates and hence the exam-time temperatures are as good as randomly assigned for students across exams. We also carry out a host of robustness checks and find that our results are not sensitive to removing weather controls, including all subjects, removing outliers and clustering standard errors at the school level (instead of at the district level).

We find that a one degree Celsius rise in temperature leads to a fall in standardized exam scores by 0.003 standard deviations. This translates to a fall of 0.016 standard

[3]Details for eligibility of admission to another school in Class XI are provided in the CBSE Curriculum/Syllabus - Senior School Certificate Examination section available at: `https://www.cbse.gov.in/currisyllabus.htm`.

deviations in exam score for each standard deviation rise in temperature. Thus, high temperatures have a negative impact on exam scores. This is in line with the previous literature that finds test-takers exposed to heat perform relatively poorly and obtain lower scores (Garg et al. (2020), Graff Zivin et al. (2020), Park (2022)).

We also find that temperature has a non-linear impact on exam scores. We consider the reference temperature category as temperatures upto the 10^{th} percentile (or 22°C and below). The literature finds evidence of optimal temperature in the range of 18°C to 22°C (Park (2022)). We find that, compared to the optimal temperatures below 22°C, higher temperatures (measured as dummies of 2°C temperature bins) have a negative impact on scores. Also, increasingly higher temperatures have increasingly larger negative impact. Thus, the marginal impact on exam performance increases with temperature or exam performance deteriorates as temperature increases further away from the optimal temperatures. Compared to optimal temperatures below the 10^{th} percentile or 22°C, exam scores fall by 0.05 standard deviations for temperatures above the median (around 28°C).

We check how temperature effects differ by student location and caste. We find that all students report a negative impact of temperature on exam scores irrespective of their castes or whether they are located in the rural or urban sector. The magnitude of impact is lower for rural students (compared to urban students) and disadvantaged castes (compared to general caste). This may point to evidence of students who are adapted to heat performing relatively better during exams in high temperatures. Students located in rural areas or belonging to disadvantaged caste are generally from lower income families and may find it difficult to adopt cooling mechanisms (such as air coolers or air-conditioners) at home as they are expensive. These students may adapt themselves to living in uncomfortable temperatures during the summers and therefore are able to perform relatively better during exams in high temperatures. We also report if very hot districts (with long run March temperatures above the 80^{th} percentile) have a different impact compared to other districts. Students in very hot districts may adopt techniques to avoid heat exposure or may become physiologically adapted to high temperatures over time (Cho (2017), Alberto et al. (2021)). We find that students in very hot districts tend to do relatively better when exposed to heat compared to other districts which, once again, points to the heat adaptation theory in the literature.

As a potential mechanism, we explore whether temperature affects exam scores through a physiological channel by checking if the impact of temperature on scores differs by gender. Schweiker et al. (2018) points out several studies documenting that female bodies prefer higher temperatures than male bodies on account of various physiological differences (such as metabolic rate and area to volume ratio). There is also evidence that females are better at heat adaptation in the short-run (Tien Manh (2019)). Men also tend to have relatively more heat related illnesses due to heat exposure (Deschênes and Greenstone (2011), Bai et al. (2014)). Interestingly, we find that female students tend to perform relatively better than male students when exposed to heat. While both male and female students report negative impact of heat on exam performance, male students seem to record a larger negative impact. Female students record a fall in 0.002 standard deviations in score compared to a fall in 0.004 standard deviations for male students for each degree Celsius rise in temperature. This is interesting as this suggests that temperature may be causing physical stress to students which in turn affects their exam performance.

We next establish that heat has a relatively greater effect on quantitative subject scores but there also exists a statistically significant impact for language subject scores. For both male and female students, there exists a statistically significant and negative impact on quantitative subjects. However, there exists a statistically significant and negative impact for language subjects only for male students and not for female students. Our results point to a cognitive impact of heat on test scores for all students, as implied by the impact on quantitative scores. The additional impact on language subjects (which do not suffer a neurological impact of heat) for male students suggests that there is another channel of heat stress. High temperatures may additionally cause physical stress to male students which affects their language test performance. Given that we find evidence of both physical and cognitive impacts of heat, our results thus suggest that temperature is working on test scores through a physiological channel.

We add to the recent and growing literature on the impact of temperature on exam scores. Cho (2017), Tien Manh (2019), Conte Keivabu et al. (2020), Graff Zivin et al. (2020) and Park (2022) look at the impact of temperature on high-stakes exams for South Korea, Vietnam, Italy, China and the US, respectively. Graff Zivin et al. (2018) is one of the first papers that provides causal estimates of short-run (on test-day) temperature impact on test scores for children taking retests for math and reading tests, which allows the use of child fixed effects. Most of these papers study how exam-day temperature affects high-stakes exam scores (except Cho (2017)[4]). Our study differs from these papers due to several advantages of our data. Unlike Park (2022), our study analyzes scores from the CBSE secondary exams which are not graded locally by the school, nor are they graded on the evening of the associated exam. The CBSE officially appoints graders to evaluate answer scripts at the end

[4]The question addressed by Cho (2017) is different as the study analyzes the impact of summer heat (the summer prior to the exam) on the Korean college entrance exam scores, which takes place in November. Hence, the focus is on learning-time temperature and not exam day temperature.

of the secondary exams, and maintains the anonymity of the answer scripts they are appointed to grade.[5] Hence, our study does not have any concerns regarding grade manipulation or leniency, based on factors such as district temperatures.[6] Our data has subject specific scores available, unlike the data used in Graff Zivin et al. (2020). Graff Zivin et al. (2020) conducts a similar study of effect of temperature on scores from the National College Entrance Examination (NCEE) in China, taking place on two exam days. However, the study relies on the cumulative score of exams taken on each day for a student, and cannot add subject fixed effects for each subject.[7] As our data has subject specific scores available, we are able to analyze the impact of heat over a wider range of exam scores for each student, given subject fixed effects. Thus, our data has more variation in temperature and exam scores for each student.

In a recent paper Garg et al. (2020) examine the relationship between test scores and temperature in India. However, the question differs as the paper looks at the long run impact of temperature on test scores, unlike our study which focuses on test-time or the short-term impact of temperature on test scores. Garg et al. (2020) find that Indian school-age children report a fall in math and reading scores due to hot days in the calendar year prior to the year of the test. Our study, on the other hand, finds evidence that temperatures on the time of the exam have a negative impact on exam scores for Indian students. As the questions differ, the empirical strategy adopted is also different.[8] Also, Garg et al. (2020) use data on primary

[5]CBSE offers students the option to appear in school-based or board-based exams (as proposed in circular number 43/2010). School-based exam answers are graded by the school teachers. However, the board-based exams are graded externally. Our analysis considers only the board-based exams. CBSE (2013) specifies rules regarding receipt and evaluation of answer books or scripts by external examiners.

[6]Park (2022) finds evidence of grade manipulation and points out that there may be a concern that graders are lenient while grading if they feel that the exam-time temperature was too high.

[7]3 compulsory subject exams take place on one day each accounting for 150 marks, whereas optional subject exams take place on another day accounting for 300 marks.

[8]Our empirical strategy requires exam-time temperature: we regress exam scores on exam-time

school children's math and reading test scores, whereas we use data on exam scores for a high-stakes board exam for secondary school children, which includes more advanced subjects than just math and reading at the primary level. Garg et al. (2020) provide interesting evidence that agricultural productivity is the main channel through which temperature affects scores.[9] Our study, on the other hand, points to a physiological channel as we are focusing on exam-time temperatures.[10] Differential impacts for male and female students in our study confirm that physiological heat stress plays an important role for students (as pointed out first by Graff Zivin et al. (2018)). Therefore, our study adds to the literature on impact of temperature on human capital formation in India in terms of different context, data and mechanism in which heat affects exam scores.

Our work also adds to the studies that have reported high temperatures having a relatively larger negative impact on men than women. Studies such as Deschênes and Greenstone (2011) and Bai et al. (2014) find that men are more subject to heat-related illnesses due to extreme temperatures. Karjalainen (2007) points out from studies in homes, offices and a university in Finland that women generally exhibit dissatisfaction with their surrounding temperature and seem to have a narrow range of thermal comfort. Schweiker et al. (2018) provide a review of laboratory and field studies on impact of temperature on physiology and point out that several physiological factors between men and women leads to thermoregulatory differences. Men in general tend to have higher metabolic rates (which affects resting temperature) than

temperature. Garg et al. (2020) regress test scores on different temperature bins for the year before the exam where the temperature bins provide counts of days with the respective temperature.

[9]Garg et al. (2020) point out that heat negatively affects agricultural productivity and provide evidence that this in turn leads to a fall in agricultural income and nutrition and a rise in agricultural labour demand. This leads to a fall in children's human capital investment.

[10]Also, the exam-time temperatures in our study are for the month of March, which does not fall in the growing seasons of Kharif (June to November) and Rabi (October to February).

women. Women also tend to have higher surface-to-volume ratio, smaller average body size, less muscle mass and higher surface-to-mass ratio which gives women a greater advantage in dissipating body heat (Mishra and Ramgopal (2013)). Rupp et al. (2015) review several laboratory studies and find that women generally have lower skin temperatures than men and prefer slightly warmer conditions when they are subject to cooler temperatures. Our work is in line with these studies as we find that temperature has a relatively larger negative effect on male than female students.

In the next section we describe our data sources for exam scores and temperature followed by section 4.3 where we discuss the empirical strategy. In section 4.4 we discuss our main results, robustness checks, alternative specifications, heterogeneity and mechanism. Finally, we conclude in section 4.5.

4.2 Data

Our data on exam scores comes from the administrative data of the Central Board of Secondary Examination (CBSE) which is the single largest board of education with an all-India presence. The primary objectives of the board include maintaining and developing quality and standards of education, giving affiliation to schools both in and outside of India and conducting secondary and higher-secondary examinations and other examinations as may be assigned to the board by the central government.[11] As per the recommendation of the National Policy on Education (1968),[12] the Indian education system follows an education structure of "10+2+3" years: 10 years

[11] The CBSE falls under the purview of the Ministry of Human Resource Development of India. Its objects, functions and powers, along with details on Board committees, duties and obligations of affiliated schools are specified in the Central Board of Secondary Education Bill, 2012. Chapter 3 of the bill specifies the board's primary objectives.

[12] The report of the National Policy on Education (1968) provides recommendations on various aspects of the education system.

comprise of primary, middle and secondary education in schools, 2 years include higher-secondary education in schools or intermediate colleges, followed by 3 years of undergraduate education in colleges or universities. The secondary examination takes place at the end of 10 years (Class X).

Under the CBSE, the areas of learning at the secondary level include two language subjects, Social Science, Mathematics and Science as compulsory subjects. In addition to the compulsory subjects, students may also study other academic elective and skill subjects. Scores for each subject is available for each student in our data, along with the school the student attends and the district in which the school is located. We standardize scores at the level of year and subject to ensure comparability of scores across subjects in our regression results.

Starting in 2011, the CBSE provides Class X students the option of appearing for board-based exams that are evaluated by external examiners or school-based exams which are graded by the student's own school teachers. In order to maintain uniformity and comparability, our sample consists of only those students who took the board-based exams. Our study analyzes scores reported for 1,719,366 students in 7,842 schools across 594 districts in India, over the span of the 4 years – 2012-2015. We drop students who are reported absent, cheating or repeating, and restrict our sample to students who have been awarded a pass or fail grade by the board.[13]

Our unit of analysis is at the level of student-subject observations. Each student is reported more than once in our data as each student takes more than one subject exam. Our data consists of 8,594,268 student-subject observations. Table 4.1 lists the subjects that constitute 95.55% of our data.[14] We report the fraction of subjects

[13]We only drop approximately 3% of our data in any given year when we drop students who are not awarded a pass or fail grade. Also, our data does not include handicapped students and students who are given a compartment status (status for students who were not able to obtain a passing grade in the main exam).

[14]We list the subjects each of which consists of at least 2% of the data at the student-subject

in our data. The core compulsory subjects – Science, Social Science and Mathematics – constitute 60.03% of our data, whereas the language subjects that constitute the largest share in our data is English and Hindi (16.7% and 9.31%, respectively).

The Class X board exams of the CBSE typically takes place in March. Students are assigned an examination centre that is different from the school they attend. While we have information on a student's school location, we do not have location data on the student's assigned examination centre. However, the examination centre is typically within the district where the school is located (CBSE (2013)). Therefore, each student is exposed to their district temperature during an exam. We obtain district temperature data from the European Centre for Medium-Range Weather Forecasts (ECMWF), mainly the database ERA5, which has hourly data from 1979 to present. The ECMWF provides global reanalysis data[15] from 1979 to present time on a number of atmospheric, ocean-wave and land-surface measures. The temperature data is available at 0.25×0.25 degree resolution for every hour of a day. We obtain temperature data for the dates in March on which secondary exams took place for the years 2012 to 2015. Our temperature data approximately captures the temperature a student is exposed to during an exam.[16] We match temperature data for each exam date and time to each district centroid for the years 2012 to 2015.[17]

level and cumulatively make up 95.55% of our data.

[15]Reanalysis data combine information from ground stations, satellites, weather balloons, and other inputs with a climate model to estimate weather variables across a grid (Dell et al. (2014)). For more details on the ERA5 database please visit the official EMCWF website and climate data store where data and information is publicly available: `https://www.ecmwf.int/en/forecasts/datasets/reanalysis-datasets/era5`.

[16]We obtain the exam dates for each subject for the years 2012 to 2015 from the official CBSE website (`www.cbse.gov.in`). The students are given 15 minutes time to read the question paper carefully before they are permitted to write. The reading time starts from 10:15 am. The exam writing time starts from 10:30 am and the duration of an exam is 3 hours. We take the average of hourly temperatures from 10:00 am to 1:00 pm to approximately capture the temperature a student is exposed to during an exam.

[17]We use shapefiles to locate the district centroid for each Indian district. We use Inverse Distance Weighting to match ERA5 temperature data available at 0.25×0.25 degree resolution to the exact

We also match other weather variables to district centroids, mainly humidity and precipitation, which are also from the ERA5 database. We use these variables as additional weather controls during the exam duration. The summary statistics of our data is provided in Table 4.2. The average temperature in our sample is 27°C. The fraction of female students, students located in rural areas and students of disadvantaged castes (scheduled caste, scheduled tribe and other backward classes) in our sample is non-trivial.

We essentially leverage the variation in temperature across districts and across days in our data. As an example, we look at the district temperatures for the middle of the month (15th March). Figure 4.1 shows how temperature varies across districts for the one day (15th March) in year 2012. Students taking an exam in a relatively warmer district will be more exposed to heat during the exam as compared to the students in a relatively cooler district. Figure 4.2 shows how temperature distribution across districts differs across the first exam date (2nd) and the last exam date (26th) in March 2012. The median temperature on 2nd March (27.01°C) is much lower than the median temperature on 26th March (32.53°C). As exams take place on different days of March in a year, a student is exposed to different exam-time temperatures. Each student experiences more than one exam-time temperature since she is required to appear for 5 compulsory subjects and exams take place on different days with different temperatures. Figure 4.3 provides a preliminary look at the relationship between temperature and exam scores: there exists a negative relationship between the two variables.[18]

latitude and longitude of a district centroid. Temperature values at latitude-longitudes nearest to the district centroid are given relatively large weights, whereas the temperature values of latitude-longitudes farther away are given lower weights. The weight assigned is the inverse of the distance (between the district centroid and latitude-longitude) raised to the power 2.

[18]The binned scatter plot, similar to a standard scatter plot, displays the relationship between the x and y variables. The x variable is grouped into equal sized bins. Unlike a standard scatter plot, each scatter point in a binned scatter plot depicts the average x and y values for all points

4.3 Empirical strategy

To analyze the impact of temperature on standardized exam scores we consider the following regression equation:

$$S_{isdty} = \gamma_i + \eta_y + \delta_s + X_{dty} + \beta_1 temperature_{dty} + \epsilon_{isdty}. \tag{4.1}$$

We regress standardized exam scores (S_{isdty}) of student i for subject s taking the exam in district d on date t in year y on temperature of district d on date t in year y. The temperature variable captures the temperature during the exam duration. Our regression considers variation in temperature across different districts on a particular exam day. Also, each student is required to study at least 5 subjects in their secondary curriculum for which they must take exams. Therefore, there also exists variation in temperature for each student over different exam days. Thus, we are able to check how temperature variation across exam days affects a student's performance on exams. We add student fixed effects (γ_i) to capture student specific characteristics such as unobserved student ability. We add year fixed effects (η_y) to capture macroeconomic shocks or year specific unobserved characteristics that may affect a student's exam performance. We also add subject fixed effects (δ_s) which capture subject specific characteristics.[19] We also add controls (X_{dty}) which are humidity and precipitation in a district during an exam. We cluster standard errors at the district level given that temperature is measured at the district level. Our main coefficient of interest is β_1 which captures how a degree Celsius rise in temperature on an exam

within the bin. In Figure 4.3, each point represents approximately 80,000 observations. The line is constructed from a linear regression of standardized exam scores on district temperatures, with student fixed effects.

[19]While we have standardized the scores of subjects, we also include subject fixed effects to control for any unobserved characteristics such as whether the subject is quantitative or subjective in nature.

day affects the standardized score of an exam. If heat exposure negatively affects a student's exam performance, we should observe a negative β_1.

Our identification relies on the assumption that the temperature variable is exogenous to exam scores. Exam-time temperature may affect a student's performance, however, it is unlikely that student scores will affect temperature. Therefore, reverse causality is not a concern in our regression exercise. Students are also not able to choose the date of their exam as the exam date of each subject is set by the CBSE. Hence, exam dates and therefore exam-time temperatures are as good as randomly assigned for students across exams.[20]

In our setting, we observe student scores for different subject exams on different exam days. In an ideal setting, we should observe how a student scores on the same exam (same subject) taken repeatedly across different days, as student ability or aptitude may differ across different subjects.[21] However, it is highly unlikely that temperature would systematically vary with student-subject ability. Since the CBSE sets the exam dates, all students must take the exam on a given date and are therefore exposed to the same temperature within a district, irrespective of their relative aptitude in the subject.

We also consider if temperature has a non-linear impact on exam performance. Heat exposure may not have a linear impact on exam performance. It is possible that the marginal impact on scores may be greater for relatively higher temperatures.

[20]We find that the dates of each subject exam do not follow a systematic pattern. We also find that the average weekly scores do not drop over the four weeks of March (which would have signified that exams of relatively difficult subjects are taking place during the relatively warmer weeks of March). The details are provided in Appendix A4.1.

[21]Also, in the ideal setting the student should not accumulate expertise when taking the exam repeatedly (student ability should be the same for the last exam as well as the first exam). Obtaining such an ideal scenario is very difficult as examinees may gain expertise from repeatedly taking the same exam (for example, a student may become better in time management when retaking the exam).

Since it is difficult to choose the order of a polynomial for our regression (as we are not sure what the nature of non-linearity is), we run an alternative specification with temperature bins. The temperature bins are 2°C bins, mainly exam-time temperatures starting from 22°C. With the temperature bins, we consider the following non-linear specification:

$$S_{isdty} = \gamma_i + \eta_y + \delta_s + X_{dty} + \beta_1 temp(22°C - 24°C)_{dty} + \beta_2 temp(24°C - 26°C)_{dty} + \beta_3 temp(26°C - 28°C)_{dty} + \beta_4 temp(28°C - 30°C)_{dty} + \beta_5 temp(> 30°C) + \epsilon_{isdty}. \tag{4.2}$$

In equation (4.2), scores are regressed on temperature categorical variables that take value 1 if it falls within the temperature range, and 0 otherwise. The reference temperature category is 22°C (which is approximately the 10^{th} percentile) or below. The highest temperature bin is above 30°C which is the 75^{th} percentile of the temperature distribution.[22] Park (2022) provides a review of evidence that suggests that the optimal thermal comfort for physical and cognitive functioning falls in the range of 18°C to 22°C. As our base temperature is 22°C or below, we can check how temperatures above the level considered to be the optimal affects exam performance.[23]

Next we check how temperature affects exam scores by subsets of students based on whether they are located in the rural or urban sector and whether they belong to a disadvantaged caste group (scheduled caste, scheduled tribe, other backward

[22]We have also carried out the analysis for temperature bins upto the 90^{th} percentile (32°C) and have obtained similar results. We break the bins above 30°C into 30°C-32°C and >32°C (32°C is approximately the 90^{th}percentile). We find that the marginal impact is similar for these additional bins above 28°C. The results are provided in Appendix A4.2.1.

[23]We also consider the middle bin 26°C-28°C as the base category (the median temperature of the sample is 27.6°C) and find similar results. Interestingly, the positive and statistically significant coefficient is the largest in magnitude for the < 22°C temperature bin. This, in a sense, reconfirms the optimal thermal comfort condition stated in Park (2022). The results are reported in Appendix A4.2.2.

classes). We also check if female students perform differently than male students when exposed to heat. Finally, we examine how students in outlier or very hot districts perform compared to other students. We check if districts that lie above the 80^{th} percentile of long run district temperatures are driving our results. We take the average March temperatures from 2007 to 2011 for a district to compute the long run *March* temperature for the district. As an additional check, we also take the average annual temperatures from 2007 to 2011 for a district to obtain the long run *annual* temperature for the district. We then examine how exam-time temperature in these very hot districts affects exam scores compared to other districts. This exercise also removes any concern that the very hot districts might largely be driving our results.

As a potential mechanism, we explore whether heat affects exam performance through a physiological channel from our results on female and male students. Physical discomfort from heat can deter focus and attention from exams leading to lower scores. As mentioned in the introduction, physical differences such as muscle mass, metabolic rates, surface-to-volume ratio, or surface-to-mass ratio can lead to thermoregulatory differences between men and women (Mishra and Ramgopal (2013), Schweiker et al. (2018)). Thus, the discomfort experienced due to higher temperatures may be different across men and women due to physiological differences. We look at how temperature affects exam scores separately for male and female students subsamples. If the impact of temperature on exam scores differ by sex, we can confirm that temperature affects exam scores through the physiological channel.

4.4 Results

4.4.1 Main results

We report our main results from running regression equation (4.1) in Table 4.3. Column 1 considers the specification that includes only student and subject fixed effects whereas column 2 also includes year fixed effects. The coefficient of the temperature variable is negative and statistically significant at 1%. That is, an increase in temperature during exams leads to a decrease in standardized scores for a student. We find that the exam score decreases by 0.003 standard deviations for a unit increase in degree Celsius exam-time temperature. This points to a fall of 0.016 standard deviations in performance for one standard deviation rise in temperature (0.003×5.36). This implies that compared to a student writing an exam exposed to the 10^{th} percentile temperature (21.79°C), a student exposed to the 75^{th} percentile temperature (30.25°C) will score 0.025 standard deviations lower in an exam (0.003×8.46). Thus, students with examination centres at relatively warmer places perform relatively poorly simply due to the weather.

Using cross-country data, Evans and Yuan (2019) demonstrates that in low and middle income countries, one standard deviation increase in literacy skill scores requires an additional 4.7 to 6.8 years of schooling.[24] We consider that years of schooling have to increase by 5.75 (midpoint of the highest and lowest estimates, 6.8 and 4.7, respectively) to result in one standard deviation increase in scores for our calculations below. Combining these estimates with the impact of temperature on scores

[24]Evans and Yuan (2019) attempts to quantify standard deviations in learning into meaningful representations for policymakers. Using individual schooling years and wage data across low and middle income countries from surveys such as Skills Towards Employability and Productivity program (STEP), the authors convert test score gains into additional years of schooling and increased wages.

mentioned in the previous paragraph, our results imply that one standard deviation increase in temperature leads to a decline in 0.092 years of schooling (0.016×5.75). Alternatively, compared to a student writing an exam exposed to the 10^{th} percentile temperature (21.79°C), exposure to the 75^{th} percentile temperature (30.25°C) will reduce 0.144 years of schooling (0.025×5.75). These estimates are concerning when compared to the loss in years of schooling for German students during World War II (Ichino and Winter-Ebmer (2004)). German individuals who were born in the nineteen thirties (and were therefore 10 years old during or after the conflict) suffered a loss in 0.21 years of schooling, compared to other cohorts. Years of schooling also translate to future earnings. We consider that one standard deviation increase in test scores leads to a 51% increase in wages of the sample mean (Evans and Yuan (2019), Garg et al. (2020)).[25] A standard deviation increase in temperature would then lead to a loss in earnings by 0.8% (0.016×0.51×100). This loss is quite substantial when compared with average earnings loss (around 3%) faced by individuals born in the nineteen thirties in Germany who were potentially affected by the war, compared to other cohorts (Ichino and Winter-Ebmer (2004)). Note that we consider these magnitudes as a lower bound since the calculations on years of schooling and wage earnings is with respect to literacy scores. Unlike the reading tests that are considered in Evans and Yuan (2019), the secondary exams that we consider are relatively more advanced, and would therefore require relatively more years of schooling for a standard deviation increase in scores.

Our results imply that heat exposure negatively affects exam performance for students. As temperature increases, students suffer from heat exposure which nega-

[25] Evans and Yuan (2019) computes the impact of learning scores on wages under the assumptions that an adult works for 40 years from the age of 20 and at a social discount rate of 3% (following the literature on public finance.)

tively affects their exam performance. This decline in exam performance is reflected in low exam scores. Our results are in line with the previous literature which discusses the negative impact of heat on physiological and cognitive ability of people (Garg et al. (2020), Graff Zivin et al. (2020), Park (2022)). Examinees require focus and concentration during exams as well as physical comfort or ease in writing exams. Heat may cause general discomfiture for examinees during exams which hinders their writing ability. If temperature affects cognitive ability, then heat will also impede the thinking ability of examinees. As a result, with increase in temperature, exam score of a student decreases as the performance of the student deteriorates.[26]

We implement a host of robustness checks and report them in Table 4.4. In column 1, we include all subjects offered by the CBSE (in addition to the subjects listed in Table 4.1) and find that the coefficient is statistically significant and negative and of a similar magnitude as that of our main results. In column 2, we check that our results are not sensitive to excluding other weather controls (humidity and precipitation).[27] In column 3, we drop temperature outliers above the 99^{th} percentile and below the 1^{st} percentile (35.85°C and 8.57°C, respectively). We find that our results are similar as before. Our results are also not sensitive to clustering standard errors at the school level instead of at the district level (column 4).

[26] Park (2022) provides a theoretical discussion where a student gains utility from marks but disutility from heat exposure. A student chooses to balance marginal utility from marks and disutility from physical discomfort, subject to economic stakes. For high stakes exams, a negative impact will be a conservative estimate as the student will choose to put effort into the exam at the cost of physical discomfort.

[27] Inclusion of weather variables such as humidity which are associated with temperature may significantly affect the coefficient of temperature in a regression (Roberts et al. (2013)).

4.4.2 Alternative specifications

4.4.2.1 Non-linear specification

We next run regression equation (4.2) and document the non-linear impact of temperature on exam scores in Table 4.5. We find that compared to exam-time temperatures below 22°C, students facing temperatures between 22°C and 24°C report a fall in score by 0.008 standard deviations, although this result is statistically insignificant. It is interesting to note that the magnitude of coefficients increase progressively for each succeeding temperature bin. For 24°C-26°C exam-time temperature, the magnitude of the coefficient is 0.021, whereas for 26°C-28°C exam-time temperature, the magnitude is 0.035. The magnitude of coefficients are the largest for temperatures above 28°C: 0.052 for temperatures between 28°C and 30°C, and 0.050 above 30°C. Compared to exam-time temperatures below 22°C, exam-time temperatures above 28°C leads to a fall in score by 0.05 standard deviations. All coefficients above 24°C are statistically significant.

Not only does a rise in temperature lead to a fall in exam scores as seen from our main results, but the marginal impact increases as temperature increases. Temperatures of 26°C-28°C, compared to the 10^{th} percentile temperature 22°C, leads to a decline by 0.035 standard deviations in score when we consider the non-linear specification. Our main specification yields a much lower impact of temperature on score: compared to 22°C, at temperatures between 26°C to 28°C scores fall by only 0.012 (0.003×4) to 0.018 (0.003×6) standard deviations. This is mainly due to the specification being linear and assuming that the marginal impact is the same at all temperatures. Thus, the results from the non-linear specification points out that the marginal impact is greater for higher temperatures. Compared to exam temperatures below 22°C, the impact on scores is the highest for temperatures between 28°C and 30°C – scores fall by 0.052 standard deviations.

4.4.2.2 Heterogeneity

We next document the impact of temperature on exam scores for students based on their locations in urban and rural areas in Table 4.6. In column 1, we find that the impact of temperature on exam scores of urban students is negative and statistically significant. A unit degree Celsius rise in temperature leads to a fall in exam score by 0.004 standard deviations. In column 2, we find that rural students also record a statistically significant and negative impact of temperature on exam score. However, the magnitude of impact is relatively lower: a unit degree Celsius rise in temperature leads to a fall in exam score by 0.002 standard deviations. It is quite possible that rural students are more adapted to studying or living in warmer temperatures than urban students. Students in urban areas may be using air coolers or air conditioners at home and are more adapted to cooler temperatures during summers.[28] If students in urban areas are adapted to living in cooler temperatures on average, they are more likely to perform poorly when exposed to high temperatures during exams. Rural students on the other hand may be adapted to living in high temperatures with little or no cooling mechanisms at home. When faced with higher temperatures during exams, they may be able to withstand it better. While the difference in coefficients is not statistically significant, we further explore this angle and check for heterogeneity based on caste.

In Table 4.7 we report the impact of temperature for students separately by caste categories: general category students in column 1 and historically disadvantaged

[28] Rural students may not be able to install air coolers or air-conditioners due to the associated financial costs (such as installation charges, maintenance and electricity costs). In India, rural households are generally composed of low-income households compared to urban households (Jaikumar and Sarin (2015)) and almost 93% of all houses without electricity are located in rural areas (Malakar (2018)). Hence, the use of cooling mechanisms will be relatively lower in rural households.

caste category students from scheduled caste, scheduled tribe and other backward classes in column 2. The impact of temperature is negative and statistically significant for both groups, but, interestingly, the magnitude and power is relatively lower for the disadvantaged caste categories. General caste students report a decrease of 0.004 standard deviations on scores due to one degree Celsius rise in temperature. However, scheduled caste, scheduled tribe and other backward class students report a decrease of 0.002 standard deviations on scores due to one degree Celsius rise in temperature. Again, as in the case of rural and urban students, these results may be due to better adaptation to heat by economically disadvantaged students. In this case, the difference in coefficients is statistically significant at 10%. Wages and consumption levels are relatively lower for disadvantaged castes compared to general castes. Hence, students who cannot afford cooling mechanisms at home are more adapted to warmer temperatures.[29] If general caste students on average adopt more cooling mechanisms at home, this may be one reason why general caste students on average perform poorly when exposed to higher temperatures compared to disadvantaged caste students as they are adapted to living in cooler and comfortable temperatures at home.

Students in very hot districts may also be physiologically adapted to high temperature. This is in line with previous evidence on heat adaptation of students located in relatively warmer regions (Cho (2017), Alberto et al. (2021)), where exam performance is relatively better in very hot regions. In Table 4.8, we include an interaction term of temperature and a dummy that takes value 1 if the observation is from a district that records temperatures above the 80^{th} percentile long run district temper-

[29] Hnatkovska et al. (2012) shows that wage, consumption, education and occupation choices of scheduled caste, scheduled tribe (SC/ST) individuals converge to the levels of non-SC/ST over the period of 1983-2005 in India. However, there remains a gap between the two categories with respect to wages across all wage percentiles.

ature.[30] In column 1, we document the results for districts that record above 80^{th} percentile of long run *annual* temperatures. In column 2, we report the results for districts that have above 80^{th} percentile of long run *March* temperatures. In both columns, we find that the coefficient for temperature is negative and statistically significant at 1%, with similar magnitudes. The interaction term in column 1 is statistically insignificant. However, when we consider past March temperatures in column 2, the interaction term is statistically significant and positive at 10%. As the coefficient of the interaction terms is either insignificant or positive, we can say that very hot districts are not driving the results. This also removes any concerns that outliers or districts that record very high temperatures during March may be driving the results in our analysis. The positive coefficient in column 2 implies that students in very hot districts tend to perform relatively better under heat exposure than students in other districts. These results signal, once again, possible heat adaptation by students residing in very hot districts.

4.4.3 Physiological mechanism

In this section we explore our hypothesis that temperature is affecting exam scores mainly through a physiological channel. Several studies have tried to prove that heat stress can negatively affect cognitive ability. Hocking et al. (2001) examines how thermal stress can affect performance on different cognitive tasks such as memory, attention, verbal learning and information processing, and finds that heat can have a detrimental impact on cognitive ability. Hancock and Vasmatzidis (2003) discusses the literature which finds evidence of heat affecting cognitive tasks differentially,

[30]The 80^{th} percentile long run (past 5 years average, 2007-2011) *March* district temperature is 27.55°C. The 80^{th} percentile long run (past 5 years average, 2007-2011) *annual* district temperature is 26.55°C. The fraction of observations in these very hot districts are 16.28% and 17.22% based on past March and annual temperatures, respectively.

depending on the type of cognitive task. Hygge and Knez (2001), Gaoua (2010) and Mazlomi et al. (2017) are some of the several studies that find field based or laboratory based evidence of heat affecting cognitive ability through a physiological channel.

As we are studying temperatures during exam-time, it is highly unlikely that an agricultural channel as proposed in Garg et al. (2020) is working in our case. We are not considering temperature in the long run or past year which affects agricultural productivity and therefore human capital investment or nutrition of exam-takers. Temperature on exam day is more likely to affect students through a physiological and cognitive channel. Also, CBSE exams take place in the month of March which does not coincide with any of the major agriculture seasons in India, mainly Kharif (June to November) and Rabi (October to February). Therefore, it is unlikely that students will be affected due to any agriculture mechanism arising from exam-day temperatures, if any.

If temperature affects exam scores through a physiological channel, it is quite possible that female students, while recording a negative impact on scores, are slightly better off than their male counterparts, owing to their physiological differences. As discussed in the introduction, studies such as Deschênes and Greenstone (2011), Mishra and Ramgopal (2013) and Bai et al. (2014) have pointed out that men are relatively more affected by high temperatures than women. There are a number of physiological differences between men and women that may cause thermal stress differently between the two sexes (Schweiker et al. (2018)). For example, female bodies have lower muscle mass and lower metabolic rate compared to male bodies. Heat dissipation from greater surface to volume ratio and lower metabolic rates may prove to be an advantage for female students when writing exams in high temperatures and this may be reflected in our results. We check for such evidence in our data.

We report how the impact of temperature on exam scores differs for male and female students in Table 4.9. Columns 1 and 2 report the results for equation (4.1) for male and female subsamples, respectively. We can see that both male and female students record a negative impact of temperature on exam scores.[31] However, the magnitude of impact is different between the two subsamples. For a 1°C rise in temperature, exam scores fall by 0.002 standard deviations for female students, but for male students the fall is found to be 0.004 standard deviations. Thus, the effect of heat on female exam scores is lower compared to male exam scores (the level of significance also falls). Our results show that while temperature has a negative impact on exam scores irrespective of gender, the magnitude of impact does differ by gender. Female students are likely to perform better during exams than male students when faced with higher temperatures. In line with previous laboratory and field studies, heat exposure is affecting male students relatively more negatively than female students. These differences in magnitude by gender confirm that temperature has a significant physiological impact. Heat exposure leads to discomfort for students resulting in possible lower focus and attention during exams leading to poor performance and exam scores.

Heat is also known to affect specific neurological processes as different parts of the brain vary in sensitivity to heat. Hancock and Vasmatzidis (2003) discuss the literature which finds evidence of heat affecting cognitive tasks differentially, depending on the type of cognitive task. Hygge and Knez (2001), Gaoua (2010) and Mazlomi et al. (2017) are some of the several studies that find field based or laboratory based evidence of heat affecting cognition through a physiological channel.

[31] The difference in coefficients are statistically significant at 1%.

Graff Zivin et al. (2018) find that heat stress significantly affects math but not reading scores. The differential impact of heat on subjects (where heat significantly affects reasoning or quantitative subjects but has little to no impact on comprehensive or language-based subjects) suggests that heat affects cognition. To test for this mechanism, we restrict our sample to science and math (considered as quantitative subjects) and language subjects in our next regression.[32] We interact the temperature variable with a dummy which takes value 1 if the subject is quantitative and 0 if the subject is a language.

Column 1 of Table 4.10 reports the results for the whole sample. We observe that while the impact of heat is negative and statistically significant for both language and quantitative subjects, the impact is greater for quantitative or reasoning subjects. We next restrict the sample to female and male students in columns 2 and 3, respectively. In column 2, for female students, the findings are in line with the literature. There is no statistically significant impact on language subjects but there exists a statistically significant and negative impact on quantitative subjects. In column 3, for male students, there is a statistically significant and negative impact for both quantitative as well as language subjects.

The results for male students are quite interesting. While both male and female students record a negative impact on quantitative subjects which points to a cognitive impact of heat, heat also affects language test scores for male students. This suggests that there is potentially an additional channel of heat stress on males other than the cognitive channel. As discussed earlier, males experience greater physical discomfort and stress when exposed to heat than females. Thus, our results again imply that male students may be experiencing physical discomfort or stress in addition

[32] We drop the subjects Social Science and Foundation of IT.

to impact on cognitive ability, which leads to not only a decrease in their quantitative subject scores but also language subject scores. The physiological mechanism operates through temperature affecting cognitive ability and, additionally, causing more physical stress for males as compared to females. Heat exposure leads to a cognitive decline and also physical discomfort for students resulting in possible lower focus and attention during exams. This leads to poor performance and lower test scores.

4.5 Conclusion

Recent literature has emerged on how high temperatures can deter human capital formation. Studies such as Garg et al. (2020), Graff Zivin et al. (2020) and Park (2022) have found evidence of high temperatures negatively affecting test scores of test-takers. Heat stress can have a negative impact on physiological and cognitive ability of people (Park et al. (2020)), and can also reduce human capital investments for children by decreasing agricultural productivity (Garg et al. (2020)). Tests or exams are a vital aspect in determining a person's career path, as admission into colleges, universities and even entry into the workforce are highly dependent on exam scores. Scores act as a signal of ability for the examinee or interviewee. Given that there is scientific evidence that high temperatures negatively affect the physical and cognitive ability of humans and temperature is rising across the globe, it becomes important to study if relatively high temperatures can affect an examinee's performance during exams. Temperature affecting test scores can in turn affect a test-taker's future labour market outcomes when he or she performs poorly during high-stakes exams when exposed to high temperatures.

In our study, we analyze how high temperatures can affect secondary exam test scores in India, administered by a national board, the Central Board of secondary Examination (CBSE). These exams are high stakes exams as the subject scores determine the subject choice in higher secondary education. Subject choice in higher secondary education is an important step, as it determines which field the student will specialize in. Based on the subjects a student selects in higher secondary education, a student will be allowed admission into a college or university department. While higher secondary scores are important for admission into colleges or universities, secondary exam scores are also given non-trivial weightage during admission. Secondary exams therefore tend to be a very important exam for Indian students and can be considered as a high-stakes exam.

We find that temperature negatively affects exam scores. One standard deviation increase in temperature leads to a fall in 0.016 standard deviations in score. When we consider that temperature may have a non-linear impact on scores, we find that the magnitude of marginal impact on scores, moving from the optimal temperature (below 22°C) to a higher temperature, increases as temperatures increase. The magnitude of the impact is also sizable: compared to exam-time temperatures below 22°C, temperatures above 28°C leads to a fall in exam score by 0.05 standard deviations. We carry out a variety of heterogeneity checks such as dividing our sample by caste, rural or urban sector, and checking the impact of heat on exam scores for extremely warm districts (above the 80^{th} percentiles). Our results point to the heat adaptation theory. In line with literature on heat adaptation (Cho (2017), Alberto et al. (2021)), we find that students in very hot districts tend to do better compared to students in other districts (the negative impact of temperature is relatively lower). We carry out a host of robustness checks and confirm that our results are not driven by outliers or very warm districts.

As a potential mechanism, we explore whether heat affects exam performance through a physiological channel from our analysis of the impact of temperature on exam scores of female and male students separately. Our results are interesting as the effect of heat exposure on female test scores is lower compared to male test scores, while both male and female students report negative impact of temperature on exam scores. We also establish the impact of heat on cognitive ability working differentially for quantitative vis-a-vis language subjects. Interestingly, the effect of heat on cognitive ability is further amplified by physical heat stress for male students. Our explanation, following the relevant medical and economics literature (Hocking et al. (2001), Deschênes and Greenstone (2011), Bai et al. (2014), Schweiker et al. (2018)), is that female and male students have very different physiological aspects that play a vital role in regulating thermal temperatures of the human body. Thus, our results point to the direction that the physiological channel is at work.

Thus, we find robust and meaningful evidence that students perform poorly in exams when exposed to higher temperatures. Our results seem to point to physiological channels as seen from the difference in impact between male and female students. As exam scores are vital for entering the labour market and making career choices, significant negative impact of temperature on exam scores becomes a vital issue in human capital formation. Adopting cooling strategies for students during exams or shifting high-stakes exams to cooler months of the year may help to reduce the negative impact of temperature on exam scores. It becomes important to adopt such cooling strategies as students residing in relatively warmer places or facing heat waves and temperature shocks during exams will be severely disadvantaged when entering the workforce. This would in turn also affect their future earning potential. Our results on rural versus urban, general versus disadvantaged caste, and male versus female students are very interesting as it would seem that the marginalized sections of the society are relatively less adversely affected from higher temperature. This is very different from the usual evidence in the literature that marginalized sections of the society face relatively larger adverse impacts of climate changes (UN (2016)).

Figures and Tables for Chapter 4

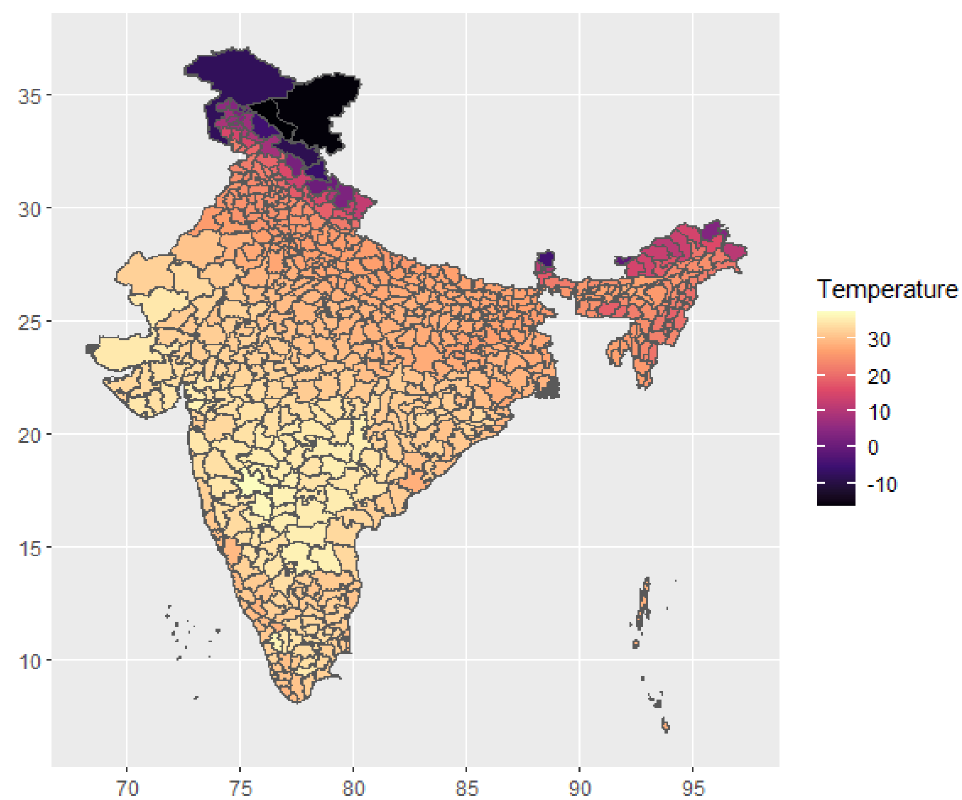

Figure 4.1: *Temperature distribution across districts on March 15th, 2012*

Note: The data for exam-time temperature for each district is from the European Centre for Medium-Range Weather Forecasts (ECMWF) ERA5 database.

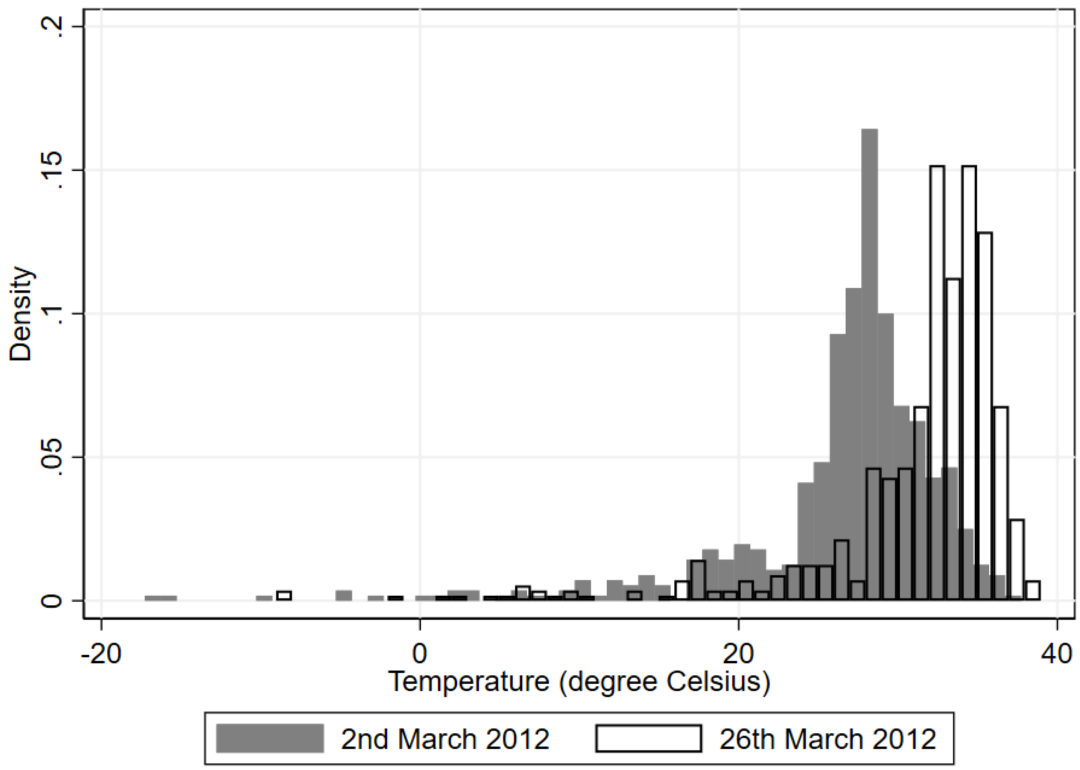

Figure 4.2: *Temperature distribution across two days in March*

Note: The data for exam-time temperature for each district is from the European Centre for Medium-Range Weather Forecasts (ECMWF) ERA5 database.

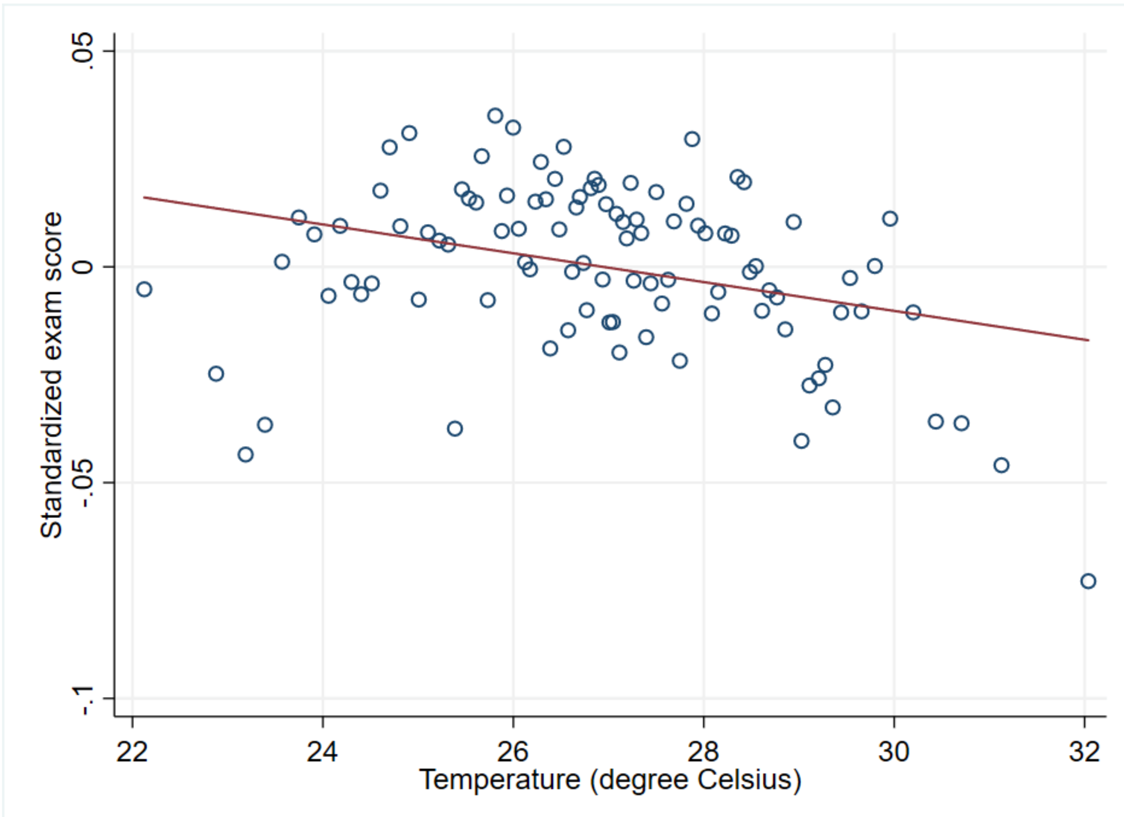

Figure 4.3: *Binned scatter*

Note: The line is constructed from a linear regression of standardized exam scores on district temperatures, with student fixed effects. Each point represents approximately 80,000 observations.

Table 4.1: *Major Subjects*

Subject	Percent
Science	20.01
Social Science	20.01
Mathematics	20.01
English Comm.	16.70
Hindi Course A	9.31
Hindi Course B	5.12
English Lng & Lit.	3.30
Foundation of IT	2.84
Comm. Sanskrit	2.71
Total	100.00

Note: Data is at the student-subject level (each student writes more than one subject exam). Total number of student-subject observations is 8,594,268. We restrict the data to the listed subjects: 95.55% of the total subject-student data consists of the above subjects, cumulatively. Comm. refers to Communicative, Lng. refers to Language, Lit. refers to Literature. The subject names are as reported in the secondary exam schedules of CBSE.

Table 4.2: *Summary statistics*

Variable	Mean
Temperature (°C)	26.94
	(5.36)
Exam scores	69.88
	(16.19)
Female fraction	0.39
	(0.49)
Rural fraction	0.41
	(0.49)
SC,ST,OBC fraction	0.34
	(0.47)
Controls:	
Precipitation (mm)	0.07
	(0.24)
Humidity (°C)	13.05
	(5.44)

Note: Summary statistics of student-subject data for 2012-2015 is reported. Subjects include those listed in Table 4.1. SC,ST,OBC refers to scheduled caste, scheduled tribe and other backward classes, respectively. Data on weather variables are from ECMWF. Temperature is temperature of air, 2 metres above the surface of land, sea or in-land waters. Humidity or 2 metre dewpoint temperature (°C) or humidity of air is the temperature to which the air, at 2 metres above the surface of the Earth, would have to be cooled for saturation to occur. Precipitation refers to total water accumulated and measured as the depth of water (reported here in millimetres) over the grid area (0.25x0.25 degree latitude longitude). Standard deviation is reported in parantheses.

Table 4.3: *Linear impact of temperature on scores*

Dependent variable: Score	(1)	(2)
$temperature_{dty}$	-0.003***	-0.003***
	(0.001)	(0.001)
Observations	8594268	8594268
R^2	0.882	0.882
Controls	Yes	Yes
Student FE	Yes	Yes
Subject FE	Yes	Yes
Year FE	No	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Controls include precipitation and humidity during exam-time. Column 1 includes student and subject fixed effects. Column 2 includes student, subject and year fixed effects. Exam scores are standardized at the subject year level. Data is at the level of student-subject for years 2012-2015. Subjects include those listed in Table 4.1.

Table 4.4: *Robustness checks*

Dependent variable: Score	(1) All subjects	(2) Exclude controls	(3) Exclude outliers	(4) SE clustered at school level
$temperature_{dty}$	-0.003***	-0.003***	-0.003***	-0.003***
	(0.001)	(0.001)	(0.001)	(0.000)
Observations	8994433	8594268	8420391	8594268
R^2	0.873	0.882	0.883	0.882
Controls	Yes	No	Yes	Yes
Student FE	Yes	Yes	Yes	Yes
Subject FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level for columns 1 to 3. Controls include precipitation and humidity during exam-time. Column 1 also includes subjects other than those listed in Table 4.1. Column 2 does not include controls. Column 3 excludes outliers (outliers are temperatures below 8.57°C, which is the 1st percentile and above 35.85°C, which is the 99th percentile temperature). Column 4 clusters standard errors at the school level.

Table 4.5: *Non-linear impact of temperature on scores*

	(1)	(2)
Dependent variable: Score		
Temperature bin (°C):		
22-24	-0.008	-0.008
	(0.007)	(0.007)
24-26	-0.021**	-0.021**
	(0.009)	(0.009)
26-28	-0.035***	-0.035***
	(0.009)	(0.009)
28-30	-0.052***	-0.052***
	(0.009)	(0.009)
>30	-0.050***	-0.050***
	(0.011)	(0.011)
Observations	8594268	8594268
R^2	0.882	0.882
Student FE	Yes	Yes
Subject FE	Yes	Yes
Year FE	No	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Controls include precipitation and humidity during exam-time. Temperature bin takes value 1 if the temperature falls within the respective temperature bin range, and 0 otherwise. Column 1 includes only student and subject fixed effects. Column 2 includes student, subject and year fixed effects. The base category includes temperatures at most 22°C.

Table 4.6: *Impact of temperature on urban and rural scores*

	(1)	(2)
Dependent variable: Score	Urban	Rural
$temperature_{dty}$	-0.004***	-0.002**
	(0.001)	(0.001)
Observations	5074154	3520114
R^2	0.881	0.885
Controls	Yes	Yes
Student FE	Yes	Yes
Subject FE	Yes	Yes
Year FE	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Controls include precipitation and humidity during exam-time. Column 1 restricts the sample to urban students, column 2 restricts the sample to rural students.

Table 4.7: *Impact of temperature on general and SC-ST-OBC caste scores*

	(1)	(2)
Dependent variable: Score	General	SC-ST-OBC
$temperature_{dty}$	-0.004***	-0.002*
	(0.001)	(0.001)
Observations	5647337	2946907
R^2	0.880	0.885
Controls	Yes	Yes
Student FE	Yes	Yes
Subject FE	Yes	Yes
Year FE	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Controls include precipitation and humidity during exam-time. SC, ST and OBC refers to the disadvantaged castes of Scheduled Caste, Scheduled Tribe and Other Backward Classes, respectively. General refers to students who are not of these disadvantaged castes. Column 1 restricts the sample to General caste students, column 2 restricts the sample to disadvantaged caste students.

Table 4.8: *Districts above 80th percentile of district annual and monthly temperatures (past 5 year average)*

Dependent variable: Score	(1) Past annual	(2) Past monthly
$temperature_{dty}$	-0.003***	-0.004***
	(0.001)	(0.001)
$temperature * above80perc_d$	0.002	0.003*
	(0.002)	(0.002)
Observations	8594268	8594268
R^2	0.882	0.882
Controls	Yes	Yes
Student FE	Yes	Yes
Subject FE	Yes	Yes
Year FE	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Controls include precipitation and humidity during exam-time. District annual and March temperatures of past 5 years (2007 to 2011) are averaged for a district to obtain the long run district temperature. The 80^{th} percentile is calculated for district long run temperatures. Dummy variable $above80pec_d$ takes value 1 if long run district temperature is above the 80^{th} percentile, and 0 otherwise. Column 1 interaction term is with respect to long run annual district temperature, column 2 interaction term is with respect to long run March district temperature.

Table 4.9: *Impact of temperature on male and female scores*

Dependent variable: Score	(1) Male	(2) Female
$temperature_{dty}$	-0.004***	-0.002**
	(0.001)	(0.001)
Observations	5239543	3354725
R^2	0.882	0.884
Controls	Yes	Yes
Student FE	Yes	Yes
Subject FE	Yes	Yes
Year FE	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Controls include precipitation and humidity during exam-time. Column 1 restricts the sample to male students, column 2 restricts the sample to female students.

Table 4.10: *Male versus female, Quantitative versus Language Subjects*

	(1) Overall	(2) Female	(3) Male
$temperature_{dty}$	-0.002**	-0.000	-0.004***
	(0.001)	(0.001)	(0.001)
$temperature_{dty} \times Quant$	-0.003***	-0.003***	-0.003***
	(0.001)	(0.001)	(0.001)
Observations	6630486	2589598	4040888
R^2	0.889	0.889	0.890
Student FE	Yes	Yes	Yes
Subject FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors are clustered at the Pincode level. Controls include precipitation and humidity during test time. The dummy Quant takes value 1 for math and science and takes value 0 for language subjects. Columns 2 and 3 restricts the sample to female and male students, respectively.

Appendix

A4.1 Exam dates and weekly averaged scores

Figure A1 reports the number of times an exam of a subject took place on a particular date. The x-axis represents the dates of March and the y-axis represents the density (since the dates are for 4 years, 0.25 implies that the exam took place once on that date during the 4 years and 0.5 implies that the exam took place twice on that date during the 4 years.) There seems to be no systematic pattern for the exam dates for each subject.

We also check how weekly scores (averaged over all subjects and all years) vary by each week. As weekly average temperatures increase with each week in March (Figure A2), it becomes a concern if relatively difficult subjects are towards the end of March. The decline in scores may then simply be because of the difficulty of the subject and not due to temperatures. However, we find that weekly average scores are not decreasing over the weeks, implying that this is not a concern (Figure A3). Together with exam dates not following a pattern, we conclude that the exam-time temperature is as good as randomly assigned across the subjects for each student.

A4.2 Alternative temperature bin specifications

A4.2.1 Temperature bins till 90th percentile

We consider additional temperature bins above 28°C (28-30, 30-32 and >32°C) till 32°C which is approximately the 90thpercentile. We document the results in Table A1 and find that for temperature bins above 28°C the coefficients are statistically significant, negative and relatively similar.

A4.2.2 Alternative reference category

We consider 26°C-28°C as the base temperature bin as an alternative base temperature bin which consists of the median temperature of the sample (27.6°C). Signs on the coefficients are as expected: relatively lower temperature bins have a positive statistically significant coefficient and relatively higher temperature bins have a negative statistically significant coefficient. We report the results in Table A2.

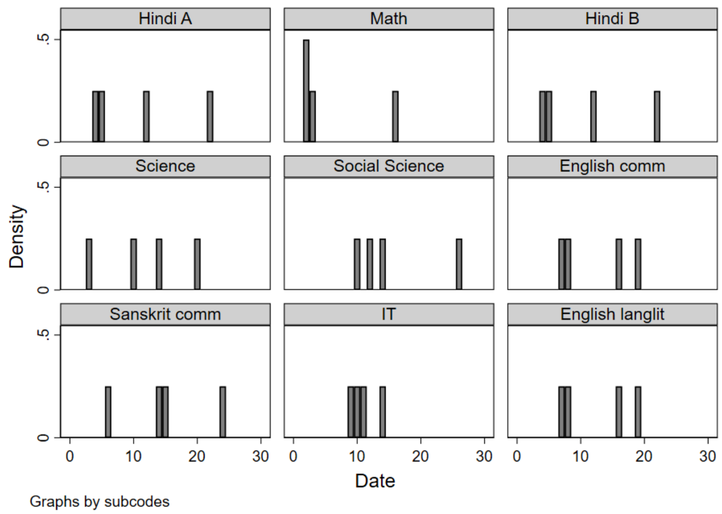

Figure A1: *Exam date distribution for each subject*

Note: The y-axis represents the number of times an exam has taken place on a date for a subject: 0.25 implies the exam has taken place once on that date, 0.50 implies the exam has taken place twice on that date during the period 2012-2015. The x-axis represents the dates of March.

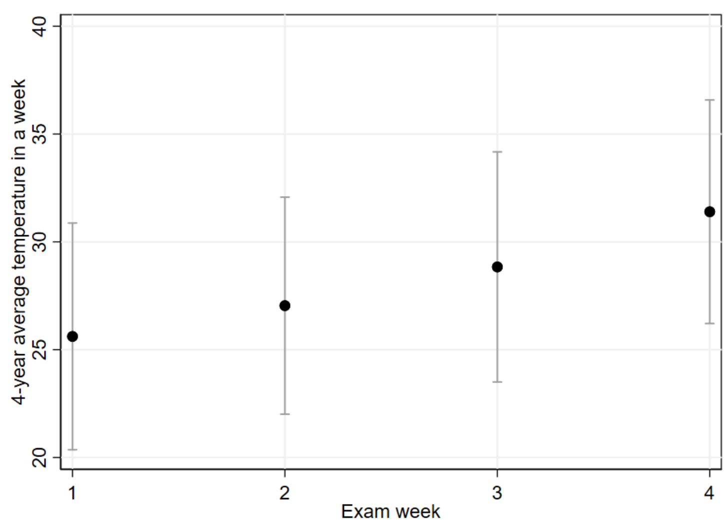

Figure A2: *Weekly average exam-time temperature, 2012-2015*

Note: The y-axis represents the average of exam-time temperatures in our data for a specific week of March for the period 2012-2015. The x-axis represents the 4 exam weeks of March.

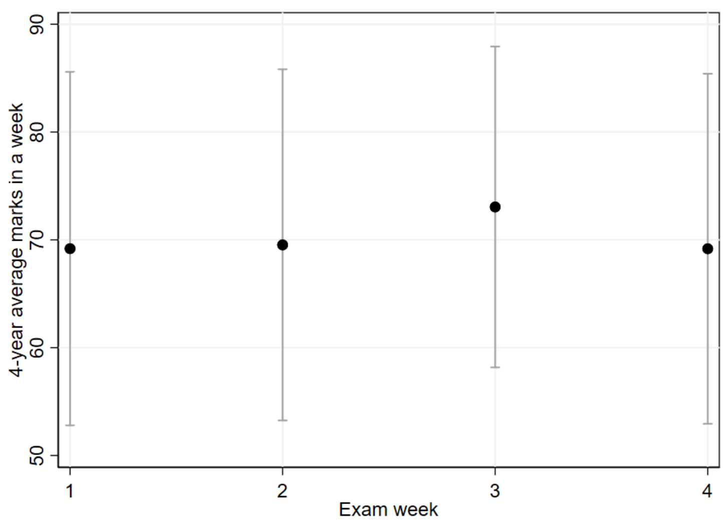

Figure A3: *Weekly average exam scores, 2012-2015*

Note: The y-axis represents the average of all exam scores in our data for a specific week of March for the period 2012-2015. The x-axis represents the 4 exam weeks of March.

Table A1: *Alternative non-linear specification*

	(1)
Dependent variable: Score	
Temperature bin (°C):	
22-24	-0.007
	(0.007)
24-26	-0.020**
	(0.009)
26-28	-0.035***
	(0.009)
28-30	-0.051***
	(0.009)
30-32	-0.053***
	(0.010)
>32	-0.043***
	(0.011)
Observations	8594268
R^2	0.882
Controls	Yes
Student FE	Yes
Subject FE	Yes
Year FE	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Temperature bin takes value 1 if the temperature falls within the respective temperature bin range, and 0 otherwise. The base category includes temperatures at most 22°C. Controls include precipitation and humidity during exam-time.

Table A2: *Alternative base temperature bin*

	(1)
Dependent variable: Score	
Temperature bin (°C):	
<22	0.035***
	(0.009)
22-24	0.028***
	(0.007)
24-26	0.015***
	(0.005)
28-30	-0.017***
	(0.006)
>30	-0.015**
	(0.006)
Observations	8594268
R^2	0.882
Controls	Yes
Student FE	Yes
Subject FE	Yes
Year FE	Yes

Note: *p<0.1, **p<0.05, ***p<0.01. Standard errors clustered at the district level. Temperature bin takes value 1 if the temperature falls within the respective temperature bin range, and 0 otherwise. The base category includes temperatures 26°C-28°C. Controls include precipitation and humidity during exam-time.

Bibliography

Aboal, D., Garda, P., Lanzilotta, B., and Perera, M. (2015). Innovation, firm size, technology intensity, and employment generation: Evidence from the Uruguayan manufacturing sector. *Emerging Markets Finance and Trade*, 51(1):3–26.

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier.

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2020). AI and jobs: Evidence from online vacancies. Working Paper 28257, National Bureau of Economic Research.

Acemoglu, D. and Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6):2188–2244.

Adhvaryu, A., Chari, A. V., and Sharma, S. (2013). Firing costs and flexibility: Evidence from firms' employment responses to shocks in India. *The Review of Economics and Statistics*, 95(3):725–740.

Aghion, P., Antonin, C., Bunel, S., and Jaravel, X. (2020). What are the labor and product market effects of automation? New evidence from France. *CEPR Discussion Paper No. DP14443*.

Aghion, P., Burgess, R., Redding, S. J., and Zilibotti, F. (2008). The unequal effects of liberalization: Evidence from dismantling the license raj in India. *American Economic Review*, 98(4):1397–1412.

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723.

Alberto, I. C., Jiao, Y., and Zhang, X. (2021). Too hot or too cold to study? The effect of temperature on student time allocation. *Economics of Education Review*, 84:102152.

Aleksynska, M., Berg, J., et al. (2016). Firms' demand for temporary labour in developing countries: Necessity or strategy. *Conditions of Work and Employment Series*, (77).

Ariza, J. and Raymond Bara, J. L. (2020). Technical change and employment in Brazil, Colombia, and Mexico. Who are the most affected workers? *International Labour Review*, 159(2):137–159.

Autor, D. and Salomons, A. (2018). Is automation labor-displacing? Productivity growth, employment, and the labor share. Working Paper 24871, National Bureau of Economic Research.

Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, 103(5):1553–97.

Autor, D. H., Katz, L. F., and Kearney, M. S. (2006). The polarization of the US labor market. *American Economic Review*, 96(2):189–194.

Autor, D. H., Katz, L. F., and Krueger, A. B. (1998). Computing inequality: Have computers changed the labor market? *The Quarterly Journal of Economics*, 113(4):1169–1213.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333.

Bai, L., Ding, G., Gu, S., Bi, P., Su, B., Qin, D., Xu, G., and Liu, Q. (2014). The effects of summer temperature and heat waves on heat-related illness in a coastal city of China, 2011–2013. *Environmental Research*, 132:212–219.

Bakk, Z. and Kuha, J. (2018). Two-step estimation of models between latent classes and external variables. *Psychometrika*, 83(4):871–892.

Basu, A. K., Chau, N. H., and Soundararajan, V. (2021). Contract employment as a worker discipline device. *Journal of Development Economics*, 149:102601.

Berman, E. and Machin, S. (2000). Skill-biased technology transfer around the world. *Oxford Review of Economic Policy*, 16(3):12–22.

Berman, E., Somanathan, R., and Tan, H. W. (2005). Is skill-biased technological change here yet? Evidence from Indian manufacturing in the 1990's. *Annals of Economics and Statistics*, No.79/80:299–321.

Besley, T. and Burgess, R. (2004). Can labor regulation hinder economic performance? evidence from india. *The Quarterly journal of economics*, 119(1):91–134.

Bessen, J. (2019). Automation and jobs: When technology boosts employment. *Economic Policy*, 34(100):589–626.

Bessen, J. E., Goos, M., Salomons, A., and Van den Berge, W. (2019). Automatic reaction - What happens to workers at firms that automate? *Boston University School of Law, Law and Economics Research Paper*, (19-2).

Bhandari, A. K. and Heshmati, A. (2008). Wage inequality and job insecurity among permanent and contract workers in India: Evidence from organized manufacturing industries. *ICFAI Journal of Applied Economics*, 7(1):80–111.

Blanchflower, D. G. and Burgess, S. M. (1998). New technology and jobs: Comparative evidence from a two country study. *Economics of Innovation and New Technology*, 5(2-4):109–138.

CBSE (2013). Central board of secondary education - examination bye-laws 1995 (updated to 2013). `https://www.cbse.gov.in/examin~1/CBSE-Examination%20by%20Law%20Book_2013.pdf`.

Chang, P. H., Daga, R., Mehta, G. V., and Cali, M. (2019). Creating markets in Indonesia: Unlocking the dynamism of the Indonesian private sector. *The World Bank*, (146343):1–71.

Chaurey, R. (2015). Labor regulations and contract labor use: Evidence from Indian firms. *Journal of Development Economics*, 114:224–232.

Chaurey, R. and Soundararajan, V. (2019). Banning contract work: Implications for input choices and firm performance. *Indian Institute of Management Bangalore Repository 2010-2019P*. `https://repository.iimb.ac.in/handle/2074/13956`.

Cho, H. (2017). The effects of summer heat on academic achievement: A cohort analysis. *Journal of Environmental Economics and Management*, 83:185–196.

Colmer, J. (2021). Temperature, labor reallocation, and industrial production: Evidence from India. *American Economic Journal: Applied Economics*, 13(4):101–24.

Compton, J. and Pollak, R. A. (2007). Why are power couples increasingly concentrated in large metropolitan areas? *Journal of Labor Economics*, 25(3):475–512.

Conte Keivabu, R. et al. (2020). Too hot to study? Gender and SES differences in the effect of temperature on school performance. *SocArXiv*. `https://doi.org/10.31235/osf.io/whtf5`.

Cortes, G. M. and Morris, D. M. (2021). Are routine jobs moving south? Evidence from changes in the occupational structure of employment in the USA and Mexico. `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3943875`.

Dell, M., Jones, B. F., and Olken, B. A. (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics*, 4(3):66–95.

Dell, M., Jones, B. F., and Olken, B. A. (2014). What do we learn from the weather? The new climate-economy literature. *Journal of Economic Literature*, 52(3):740–98.

Dell, M. and Querubin, P. (2018). Nation building through foreign intervention: Evidence from discontinuities in military strategies. *The Quarterly Journal of Economics*, 133(2):701–764.

Deschênes, O. and Greenstone, M. (2011). Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US. *American Economic Journal: Applied Economics*, 3(4):152–85.

Doms, M., Dunne, T., and Roberts, M. J. (1995). The role of technology use in the survival and growth of manufacturing plants. *International Journal of Industrial Organization*, 13(4):523–542.

Egger, H., Kreickemeier, U., Moser, C., and Wrona, J. (2016). Offshoring and job polarisation between firms. *CESifo Working Paper Series No. 6142*. `https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2877095`.

Erumban, A. A. and Das, D. K. (2020). ICT investment and economic growth in India: An industry perspective. In *Digitalisation and Development*, pages 89–117. Springer, Singapore.

Eswaran, M., Kotwal, A., Ramaswami, B., and Wadhwa, W. (2009). Sectoral labour flows and agricultural wages in India, 1983-2004: Has growth trickled down? *Economic and Political Weekly*, 44(2):46–55.

Evans, D. and Yuan, F. (2019). Equivalent years of schooling: A metric to communicate learning gains in concrete terms. *World Bank Policy Research Working Paper No. 8752*.

Gaoua, N. (2010). Cognitive function in hot environments: A question of methodology. *Scandinavian Journal of Medicine & Science in Sports*, 20(s3):60–70.

Garg, T., Jagnani, M., and Taraz, V. (2020). Temperature and human capital in India. *Journal of the Association of Environmental and Resource Economists*, 7(6):1113–1150.

Goldar, B. and Aggarwal, S. C. (2012). Informalization of industrial labor in India: Effects of labor market rigidities and import competition. *The Developing Economies*, 50(2):141–169.

Goodman, L. A. and Kruskal, W. H. (1959). Measures of association for cross classifications. II: Further discussion and references. *Journal of the American Statistical Association*, 54(285):123–163.

Goos, M. and Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in Britain. *The Review of Economics and Statistics*, 89(1):118–133.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8):2509–26.

Graff Zivin, J., Hsiang, S. M., and Neidell, M. (2018). Temperature and human capital in the short and long run. *Journal of the Association of Environmental and Resource Economists*, 5(1):77–105.

Graff Zivin, J., Song, Y., Tang, Q., and Zhang, P. (2020). Temperature and high-stakes cognitive performance: Evidence from the national college entrance examination in China. *Journal of Environmental Economics and Management*, 104:102365.

Günther, I. and Launov, A. (2012). Informal employment in developing countries: Opportunity or last resort? *Journal of Development Economics*, 97(1):88–98.

Hancock, P. A. and Vasmatzidis, I. (2003). Effects of heat stress on cognitive performance: The current state of knowledge. *International Journal of Hyperthermia*, 19(3):355–372.

Hazen, A. (1914). Storage to be provided in impounding municipal water supply. *Transactions of the American Society of Civil Engineers*, 77(1):1539–1640.

Heal, G. and Park, J. (2020). Reflections — Temperature stress and the direct impact of climate change: A review of an emerging literature. *Review of Environmental Economics and Policy*, 10(2):347–362.

Hjort, J. and Poulsen, J. (2019). The arrival of fast internet and employment in Africa. *American Economic Review*, 109(3):1032–79.

Hnatkovska, V., Lahiri, A., and Paul, S. (2012). Castes and labor mobility. *American Economic Journal: Applied Economics*, 4(2):274–307.

Hocking, C., Silberstein, R. B., Lau, W. M., Stough, C., and Roberts, W. (2001). Evaluation of cognitive performance in the heat by functional brain imaging and psychometric testing. *Comparative Biochemistry and Physiology Part A: Molecular & Integrative Physiology*, 128(4):719–734.

Hsiang, S. M. (2010). Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America. *Proceedings of the National Academy of Sciences*, 107(35):15367–15372.

Hussmanns, R. (2004). Defining and measuring informal employment. *International Labour Office, Bureau of Statistics, Geneva*.

Hygge, S. and Knez, I. (2001). Effects of noise, heat and indoor lighting on cognitive performance and self-reported affect. *Journal of Environmental Psychology*, 21(3):291–299.

IBEF (2008). Machine tools market and opportunities, a report by Klynveld Peat Marwick Goerdeler (KPMG) for Indian Brand Equity Foundation (IBEF). `https://www.ibef.org/download/Machine_Tools_170708.pdf`.

Ichino, A. and Winter-Ebmer, R. (2004). The long-run educational cost of World War II. *Journal of Labor Economics*, 22(1):57–87.

ILO (2016). Non-standard employment around the world: Understanding challenges, shaping prospects. *International Labour Organization, Geneva*. http://www.ilo.org/global/publications/books/WCMS_534326/lang--en/index.htm.

ILO-IOE (2019). Changing business and opportunities for employer and business organizations. *International Labour Office and International Organisation of Employers – Geneva*. https://www.ilo.org/actemp/areas-of-work/WCMS_679582/lang--en/index.htm.

Jaikumar, S. and Sarin, A. (2015). Conspicuous consumption and income inequality in an emerging economy: Evidence from India. *Marketing Letters*, 26(3):279–292.

Jayachandran, S. (2006). Selling labor low: Wage responses to productivity shocks in developing countries. *Journal of Political Economy*, 114(3):538–575.

Jones, B. F. and Olken, B. A. (2010). Climate shocks and exports. *American Economic Review*, 100(2):454–59.

Karjalainen, S. (2007). Gender differences in thermal comfort and use of thermostats in everyday thermal environments. *Building and Environment*, 42(4):1594–1603.

Katz, L. F. and Murphy, K. M. (1992). Changes in relative wages, 1963–1987: Supply and demand factors. *The Quarterly Journal of Economics*, 107(1):35–78.

Kaur, S. (2019). Nominal wage rigidity in village labor markets. *American Economic Review*, 109(10):3585–3616.

Kijima, Y. (2006). Why did wage inequality increase? Evidence from urban India 1983–99. *Journal of Development Economics*, 81(1):97–117.

Kok, J. and Berrios, M. (2019). Small matters: Global evidence on the contribution to employment by the self-employed, micro-enterprises and SMEs. *International Labour Organization*. `https://www.ilo.org/global/publications/books/forthcoming-publications/WCMS_723387/lang--en/index.htm`.

Kupets, O. (2016). Education-job mismatch in Ukraine: Too many people with tertiary education or too many jobs for low-skilled? *Journal of Comparative Economics*, 44(1):125–147.

Kurniawati, S., Suryadarma, D., Bima, L., and Yusrina, A. (2018). Education in Indonesia: A white elephant? *Journal of Southeast Asian Economies*, 35(2):185–199.

Kwon, H. J. (2013). Global economic downturn and social protection in East Asia: Coping with crisis and reducing poverty. In *A Handbook of Comparative Social Policy, Second Edition*. Edward Elgar Publishing.

Leroux, B. G. (1992). Consistent estimation of a mixing distribution. *The Annals of Statistics*, 20(3):1350–1360.

Little, R. J. and Rubin, D. B. (2019). *Statistical Analysis with Missing Data*, volume 793. John Wiley & Sons.

Maitra, S. (2016). The poor get poorer: Tracking relative poverty in India using a durables-based mixture model. *Journal of Development Economics*, 119:110–120.

Malakar, Y. (2018). Evaluating the role of rural electrification in expanding people's capabilities in India. *Energy Policy*, 114:492–498.

Mazlomi, A., Golbabaei, F., Farhang Dehghan, S., Abbasinia, M., Mahmoud Khani, S., Ansari, M., and Hosseini, M. (2017). The influence of occupational heat exposure on cognitive performance and blood level of stress hormones: A field study report. *International Journal of Occupational Safety and Ergonomics*, 23(3):431–439.

McLachlan, G. J., Lee, S. X., and Rathnayake, S. I. (2019). Finite mixture models. *Annual Review of Statistics and its Application*, 6:355–378.

Mishra, A. K. and Ramgopal, M. (2013). Field studies on human thermal comfort — An overview. *Building and Environment*, 64:94–106.

Nizalova, O. Y. and Murtazashvili, I. (2016). Exogenous treatment and endogenous factors: Vanishing of omitted variable bias on the interaction term. *Journal of Econometric Methods*, 5(1):71–77.

Olsson, M. and Tåg, J. (2017). Private equity, layoffs, and job polarization. *Journal of Labor Economics*, 35(3):697–754.

Paap, R., Franses, P. H., and Van Dijk, D. (2005). Does Africa grow slower than Asia, Latin America and the Middle East? Evidence from a new data-based classification method. *Journal of Development Economics*, 77(2):553–570.

Pais, J. (2002). Casualisation of urban labour force: Analysis of recent trends in manufacturing. *Economic and Political Weekly*, 37(7):631–652.

Park, R. J. (2022). Hot temperature and high stakes performance. *Journal of Human Resources*, 57(2):400–434.

Park, R. J., Goodman, J., Hurwitz, M., and Smith, J. (2020). Heat and learning. *American Economic Journal: Economic Policy*, 12(2):306–39.

Peirce, C. S. (1884). The numerical measure of the success of predictions. *Science*, (93):453–454.

Roberts, M. J., Schlenker, W., and Eyer, J. (2013). Agronomic weather measures in econometric models of crop yield with implications for climate change. *American Journal of Agricultural Economics*, 95(2):236–243.

Rothenberg, A. D., Gaduh, A., Burger, N. E., Chazali, C., Tjandraningsih, I., Radikun, R., Sutera, C., and Weilant, S. (2016). Rethinking Indonesia's informal sector. *World Development*, 80:96–113.

Rupp, R. F., Vásquez, N. G., and Lamberts, R. (2015). A review of human thermal comfort in the built environment. *Energy and Buildings*, 105:178–205.

Salam, U., Lee, S., Fullerton, V., Yusuf, Y., Krantz, S., and Henstridge, M. (2018). Indonesia case study: Rapid technological change–challenges and opportunities. *Pathways for Prosperity Commission Background Paper Series, Oxford Policy Management*.

Schuetze, H. J. (2000). Taxes, economic conditions and recent trends in male self-employment: A Canada–US comparison. *Labour Economics*, 7(5):507–544.

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464.

Schweiker, M., Huebner, G. M., Kingma, B. R., Kramer, R., and Pallubinsky, H. (2018). Drivers of diversity in human thermal perception – A review for holistic comfort models. *Temperature*, 5(4):308–342.

Shepard, D. (1968). A two-dimensional interpolation function for irregularly-spaced data. In *Proceedings of the 1968 23rd ACM National Conference*, pages 517–524.

Shyam Sundar, K. (2011). Non-regular workers in India: Social dialogue and organizational and bargaining strategies and practices. ILO working papers, International Labour Organization. `https://EconPapers.repec.org/RePEc:ilo:ilowps:994665563402676`.

Somanathan, E., Somanathan, R., Sudarshan, A., and Tewari, M. (2021). The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing. *Journal of Political Economy*, 129(6):1797–1827.

Tambunan, T. T. and Purwoko, B. (2002). Social protection in Indonesia. *Social Protection in Southeast and East Asia–Towards a Comprehensive Picture*, 1(1.66):1–35.

Teicher, H. (1963). Identifiability of finite mixtures. *The Annals of Mathematical Statistics*, pages 1265–1269.

Thorndike, R. L. (1953). Who belongs in the family? *Psychometrika*, 18(4):267–276.

Tien Manh, V. (2019). Effects of weather on human capital in Vietnam. *Asian Growth Research Institute Working Paper Series*, (2019-01):1–36.

Tinbergen, J. (1974). Substitution of graduate by other labour. *Kyklos: International Review for Social Sciences*, 27(2):217–226.

Tinbergen, J. (1975). *Income Differences: Recent Research*. North-Holland Publishing Company, Amsterdam.

UN (2016). New un report: Inequalities cause and exacerbate climate impacts on poor and vulnerable people. `https://news.un.org/en/story/2016/10/541742#:~:text=Evidence%20is%20increasing%20that%20climate,by%20the%20United%20Nations%20today`.

Van Reenen, J. (1997). Employment and technological innovation: Evidence from UK manufacturing firms. *Journal of Labor Economics*, 15(2):255–284.

Vashisht, P. (2018). Destruction or polarization: Estimating the impact of technology on jobs in Indian manufacturing. *The Indian Journal of Labour Economics*, 61(2):227–250.

Vermunt, J. K. (2010). Latent class modeling with covariates: Two improved three-step approaches. *Political Analysis*, 18(4):450–469.

West, S. A., Butani, S., Witt, M., et al. (1990). Alternative imputation methods for wage data. In *Proceedings of the Section on Survey Research Methods, American Statistical Association*, pages 254–259. Citeseer.

Willmott, C. J. and Matsuura, K. (2018). Terrestrial precipitation: 1900-2017 gridded monthly time series. `http://climate.geog.udel.edu/~climate/html_pages/Global2017/README.GlobalTsP2017.html`.

Willmott, C. J., Rowe, C. M., and Philpot, W. D. (1985). Small-scale climate maps: A sensitivity analysis of some common assumptions associated with grid-point interpolation and contouring. *The American Cartographer*, 12(1):5–16.

World Bank (2020). *Innovative China: New Drivers of Growth*. World Bank Group, and the Development Research Center of the State Council, P. R. China, `https://elibrary.worldbank.org/doi/abs/10.1596/978-1-4648-1335-1`.

Yeaple, S. R. (2005). A simple model of firm heterogeneity, international trade, and wages. *Journal of International Economics*, 65(1):1–20.

Zhang, P., Deschenes, O., Meng, K., and Zhang, J. (2018). Temperature effects on productivity and factor reallocation: Evidence from a half million Chinese manufacturing plants. *Journal of Environmental Economics and Management*, 88:1–17.

Zulfikar, T. (2009). The making of Indonesian education: An overview on empowering Indonesian teachers. *Journal of Indonesian Social Sciences and Humanities*, 2:13–39.