Railroads, Economic Development and Famine Prevention: Theory and Evidence from India, 1861-2000

David Donaldson

Dissertation submitted for the degree of

Doctor of Philosophy in Economics at

The London School of Economics and Political Science

August 2009

Declaration

I declare that the work presented in this thesis is my own except where the collaboration with coauthors is explicitly acknowledged.

David Donaldson, 7 August 2009

I agree with the above statement.

Prof. Robin Burgess (Supervisor), 7 August 2009

The London School of Economics and Political Science

Declaration

I certify that chapter 3 of this thesis, "Weather and Death in India: Mechanisms and Implications of Climate Change", was coauthored with Robin Burgess, Olivier Deschenes and Michael Greenstone. David Donaldson contributed 50 percent to the genesis of the project, 50 percent to the empirical work on the data, and 50 percent to the writing of the text.

Robin Burgess

Prof. Robin Burgess (Supervisor), 7 August 2009

The London School of Economics and Political Science

Abstract

The first chapter of this thesis draws on newly collected archival data from colonial India to estimate the impact of India's vast railroad network on real income there from 1870 to 1930. Guided by four predictions from a general equilibrium trade model I find that railroads: (1) reduced trade costs and interregional price gaps; (2) expanded interregional and international trade; (3), raised real income levels (but harmed neighbouring regions without railroad access); and (4), that a sufficient statistic for the impact of railroads on welfare in the model accounts for virtually all of the estimated reduced-form impact of railroads on real income.

The second chapter examines the role that openness to trade, as brought about by railroads in colonial India, affected the extent to which real incomes in the agricultural sector were exposed to variation in rainfall (an input to agricultural production). This relationship is theoretically ambiguous, but I find strong empirical support for trade openness reducing the exposure of real income to rainfall (prices become less responsive and nominal incomes more responsive to local rainfall, but real incomes are on net less responsive). I find analogous effects on mortality rates, implying that the responsiveness of consumption to rainfall may have fallen as well. These results suggest that openness can reduce real income volatility, and that railroads did much to rid India of famine.

The final chapter extends the above investigation of the role that weather fluctuations play in Indian citizen's lives into the modern period, from 1957-2000. A first set of results document large effects of excess temperature and surfeit rainfall on both adult and infant mortality—further, these effects are only present in the growing periods of the year and are entirely absent among urban citizens. A second set of results

documents an analogous pattern of deleterious effects of adverse weather on agricultural yields, wages and prices.

Acknowledgements

This work would not have been possible without the help of an enormous number of people. My deepest intellectual debt is to my PhD supervisors—Tim Besley, Robin Burgess and Stephen Redding. They have been involved at all stages of this project, from settling on a research topic, to assisting with grant applications, to countless hours of discussions about how to approach the research questions I have chosen to tackle here. Along the way, an uncountable number of friends and colleagues have taken the time to talk about this work with me, and to read drafts. My most frequent encounters were with faculty members and fellow students at EOPP and the CEP's Globalisation Group. I was also fortunate enough to be able to meet with the visiting seminar speakers at EOPP and the CEP, and these were an incomparable twice-weekly lesson on how some of the profession's best economists think about research. Finally, with the backing of Tim, Robin and Steve, as well as that of Michael Greenstone, I was very fortunate to be able to present much of the work in this thesis to tens of seminar audiences and conferences and to benefit from the always original, perceptive and helpful comments I received from colleagues at these events.

A number of people were instrumental in helping me to assemble the dataset that lies beind this thesis. Qianzi Zeng and Yifan Wang took thousands of photographs of the required pages of original documents. Ray Iyer and his team at Structured Concepts worked to convert these photographs into digitized versions of the data they contain. And finally, Sinduja Srinivasa, Rashmi Harimohan, Erasmus Ermgassen, Mike Oliver, Jin Wang, and Ben Faber assisted me in converting these digital documents into coherent datasets. Behind the scenes of this process, the support staff at the LSE's Research Laboratory offered a level of extraordinary support for this project that a researcher could only dream of. In particular, Leila Alberici, Sue Coles, Jane Dickson, and Angela Swain helped with my endless requests to keep the logistics running, Nic Warner and Joe Johannes created and supported a computing architecture that was always one step ahead of anything I could ask for, and Tanvi Desai helped me in my initial searches for funding. This funding was essential for a project of this scale, and I am extremely grateful to the Royal Economic Society, the British Academy, the Nuffield Foundation, the ESRC, DfID and the LSE's Bagri Fellowship for their financial support.

My family—of Donaldsons and Ermgassens—have been uniformly and extraordinarily tolerant of my seemingly eternal student status, and of my unusual railway-based dinner conversation for longer than I can bear to remember. Without their support and sincere interest I am certain that I wouldn't have had the energy to see this project through. My parents and Stine have been supporting my student status for longer than *they* care to remember, but their love and support has been entirely unwavering and unequestioning. They have given me the luxury of patience in finding a career path that I love. My mom's commitment to this thesis work was particularly heroic, when, at an early stage of this project, she spent weeks in the University of Pennsylvania library completing work that I had started but didn't have time to finish: photocopying volumes of the Indian Census and FedExing them to India, without giving off even the slightest impression that she thought there must have been a better way.

Finally, my deepest debt and thanks are to Stine, who, time and time and time again, has made personal sacrifices so that I could see this

project to completion—and, together with Adele and Jesper, brought more joy into my life than I ever dreamed possible.

Preface

This thesis investigates determinants of economic welfare in India from 1861 to the present. A central theme is the potential for transportation infrastructure improvements to improve economic welfare in the agricultural sector, both through raising real incomes and by mitigating the dependency of real incomes on the supply of rainfall.

Chapter 1 focuses on an analysis of the effect of India's colonial-era railroad system on real agricultural income growth. I start by estimating the extent to which railroads actually reduced the costs of trading, and then examine how much new trade was generated by these trade cost reductions. The elasticity of trade flows to trade costs says a great deal in the theoretical model that I work with about how large one could expect the static gains from trade, brought about due to lower trade costs, to be. I estimate this elasticity and compute the model's predicted real income gains due to new static gains from trade. These predicted gains are highly correlated with the observed real income gains attributable to railroads in the data. This is suggestive evidence of an important role for transportation infrastructure in promoting development by allowing regions to trade with one another at lower costs.

Chapter 2 then extends Chapter 1 by looking at a second potential welfare improvement due to railroads, their role in reducing real income volatility. I look at the extent to which agricultural productivity shocks, due to rainfall extremes, affect agricultural incomes in both the pre-railroad and post-railroad eras. When railroads arrive in an average district of India they significantly reduce the responsiveness of real agricultural income—and the mortality rate—to the vagaries of India's monsoon rains. Chapter 2 presents further results that aim to better understand the genesis of this result, beginning with price volatility and working up. The findings with respect to dampened mortality responses to rainfall attributable to railroads speak to the question of whether railroads were responsible for ridding India of peacetime famine from 1906 onwards.

Chapter 3 continues the investigation of the relationship between weather and death taken up in Chapter 2, by extending this work into the post-independence era in India. I also extend the analysis to include both temperature and rainfall variation and show that temperature extremes have a particularly pernicious effect on mortality in India, but that this effect is concentrated entirely in rural areas, and is confined to extremes that occur only during the agricultural growing months of the year. The remainder of Chapter 3 attempts to better understand the mechanisms behind this finding.

A final contribution of this thesis is the creation of a new dataset on the evolution of India's economy, and the health of its citizens, from 1861 to 2000. India is unique among developing countries—and even developed ones—in presenting researchers with a legacy of high quality data on a variety of variables of interest, for fine geographic units, over the extremely long-run. While this has not been a feature of any of the analysis in this thesis, these geographic units can be traced through time as their borders are largely invariant. It is hoped that the creation of this new resource will generate further research on questions of economic development—and long-run development in particular—that India is uniquely positioned to answer.

Contents

Chapter 1

Railroads of the Raj: Estimating the Impact of Transportation Infrastructure

1.1 Introduction

In 2007, almost 20 percent of World Bank lending was allocated to transportation infrastructure projects, a larger share than that of education, health and social services combined (World Bank 2007). These projects aim to reduce trade costs. In prominent models of international and interregional trade, reductions in trade costs will increase the level of real income in trading regions.[1] Unfortunately, despite an emphasis on reducing trade costs in both economic theory and contemporary aid efforts, we lack a rigorous empirical understanding of the extent to which trans-

[1] Workhorse trade models featuring trade costs include Dornbusch, Fischer, and Samuelson (1977), Krugman (1980), Eaton and Kortum (2002) and Melitz (2003). In all of these theories, two trading regions will both gain when the (iceberg) cost of trading goods between them falls symmetrically. However, there are theoretical settings in which symmetric trade cost reductions can harm one of two trading regions; for example, if increasing returns to scale production technologies mix with factor mobility, as in Krugman (1991), or traded intermediate goods, as in Krugman and Venables (1995), then one of two trading regions can experience a welfare loss.

portation infrastructure projects actually reduce the costs of trading, and how the resulting trade cost reductions affect welfare.

In this chapter I exploit one of history's great transportation infrastructure projects—the vast network of railroads built in colonial India (India, Pakistan and Bangladesh; henceforth, simply 'India')—to make three contributions to our understanding of transportation infrastructure improvements. In doing so I draw on a comprehensive new dataset on the colonial Indian economy that I have constructed. First, I estimate the extent to which railroads improved India's trading environment (ie reduced trade costs, reduced interregional price gaps, and increased trade flows). Second, I estimate the reduced-form welfare gains (higher real income levels) that the railroads brought about. Finally, I assess, in the context of a general equilibrium trade model, how much of these reduced-form welfare gains could be plausibly interpreted as newly exploited gains from trade.

The railroad network designed and built by the British government in India (then referred to as 'the Raj') brought dramatic change to the technology of trading there. Prior to the railroad age, bullocks carried most of India's commodity trade on their backs, traveling no more than 30 km per day along India's sparse network of dirt roads (Deloche 1994). By contrast, railroads could transport these same commodities 600 km in a day, and at much lower per unit distance freight rates. As the 67,247 km long railroad network expanded (from 1853 to 1930), it penetrated inland districts (local administrative regions), bringing them out of near-autarky and connecting them with the rest of India and the world. I use the arrival of the railroad network in each district to investigate the economic impact of this striking improvement in transportation infrastructure.

This setting is unique because the British government collected detailed records of economic activity throughout India in this time period—remarkably, however, these records have never been systematically digitized and organized by researchers. I use these records to construct a new dataset with almost seven million observations on district-level prices, output, daily rainfall and interregional and international trade in India, as well as a digital map of India's railroad network in which each 20 km segment is coded with its year of opening. This dataset allows me to track the evolution of India's district economies before, during and after the expansion of the railroad network. The records on *interregional* trade are particularly unique and important here. Information on trade flows within a country is rarely available to researchers, yet the response of these trade flows to a transportation infrastructure improvement says a great deal about the potential for gains from trade (as I describe explicitly below).

To guide my empirical analysis I develop a Ricardian trade model with many regions, many commodities, and where trade occurs at a cost. Because of geographical heterogeneity, regions have differing productivity levels across commodities, which creates incentives to trade to exploit comparative advantage. A new railroad link between two districts lowers their bilateral trade cost, allowing consumers to buy goods from the cheapest district, and producers to sell more of what they are best at producing. There are thousands of interacting product and factor markets in the model. But the analysis of this complex general equilibrium problem is tractable if production heterogeneity takes a convenient but plausible functional form, as shown by Eaton and Kortum (2002).

I use this model to assess empirically the importance of one particular mechanism linking railroads to welfare improvements—that railroads reduce trade costs and thereby allow regions to gain from trade. The model makes four predictions that drive my four-step empirical analysis:

1. *Inter-district price differences are equal to trade costs (in special cases):* That is, if a commodity can be made in only one district (the 'origin') but is consumed in other districts ('destinations'), then that commodity's origin-destination price difference is equal to its origin-destination trade cost. I use this result to infer trade costs (which researchers never fully observe) by exploiting widely-traded commodities that could only be made in one district. Using inter-district price differentials, along with a graph theory algorithm embedded in a non-linear least squares routine, I estimate the trade cost parameters governing traders' endogenous route decisions on a network of roads, rivers, coasts and railroads. This is a novel method for inferring trade costs in networked settings. My resulting parameter estimates reveal that railroads significantly reduced the cost of trading in India.

2. *Bilateral trade flows take the 'gravity equation' form:* That is, holding constant exporter- and importer-specific effects, bilateral trade costs reduce bilateral trade flows. I find that railroad-driven reductions in trade costs (estimated in Step 1) increase bilateral trade flows, and show that the parameters estimated from the gravity equation identify my model.

3. *Railroads increase real income levels:* That is, when a district is connected to the railroad network its real income rises; however, improvements in the railroad network that by-pass a district reduce the district's real income (a negative spillover effect). Empirically, I find that own-railroad access raises real income by 18 percent, but a neighbor's access reduces real income by 4 percent. However,

these are reduced-form estimates that could be due to a number of mechanisms. A key goal of Step 4 below is to assess how much of the reduced-form effect of railroads can be attributed to gains from trade due to the trade cost reductions found in Step 1.

4. *There exists a sufficient statistic for the welfare gains from railroads:* That is, despite the complexity of the model's general equilibrium relationships, the impact of the railroad network on welfare in a district is captured by one variable: the share of that district's expenditure that it sources from itself. A prediction similar to this appears in a wide range of trade models but has not, to my knowledge, been tested before.[2] I test this prediction by regressing real income on this sufficient statistic (as calculated using the model estimated in Steps 1 and 2) alongside the regressors from Step 3 (which capture the reduced-form impact of railroads).[3] When I do this, the reduced-form coefficients on railroad access estimated in Step 3 fall to a level that is close to zero. This finding provides support for prediction 4 of the model and suggests that decreased trade costs account for virtually all of the real income impacts of the Indian railroad network.

These four results demonstrate that India's railroad network improved the trading environment (Steps 1 and 2), generated welfare gains (Steps 3), and that these welfare gains arose predominantly because railroads allowed regions to exploit gains from trade (Step 4).

[2]Arkolakis, Klenow, Demidova, and Rodriguez-Clare (2008) show that this prediction applies to the Krugman (1980), Eaton and Kortum (2002), Melitz (2003), and Chaney (2008) models of trade, but these authors do not test this prediction in their empirical application.

[3]This procedure is similar in spirit to the "sufficient statistic approach" proposed by Chetty (2008) as a compromise between reduced-form and structural methods of welfare analysis.

Because railroads were not randomly assigned to districts, I pursue three strategies to mitigate concerns of bias due to a potential correlation between railroad placement and unobserved changes in the local economic environment. First, I estimate *four placebo specifications* using over 40,000 km of railroad lines that reached advanced stages of costly surveying but were never actually built, but find no spurious effects from these unbuilt lines. Second, I estimate *instrumental variable* specifications in which I instrument for railroad construction post-1884 with rainfall shortages in the 1876-78 agricultural years (because the 1880 Indian Famine Commission recommended that railroad lines be built in regions that experienced drought in the 1876-78 famine), and find IV results that are very close to my OLS results. Finally, in a *bounds check*, I find similar results among railroad lines whose estimates are likely to be biased upwards and lines whose estimates are likely to be biased downwards.

This chapter contributes to a growing literature on estimating the economic effects of large infrastructure projects,[4] as well as a literature on estimating the 'social savings' of railroad projects.[5] A distinguishing feature of my approach is that, in addition to estimating reduced-form relationships between infrastructure and welfare as in the existing literature, I fully specify and estimate a general equilibrium model of how railroads affect welfare. The model makes auxiliary predictions and suggests a sufficient statistic for the role played by railroads in raising welfare—all of which shed light on the economic mechanisms that could

[4]For example, Dinkelman (2007) estimates the effect of electrification on labor force participation in South Africa, Duflo and Pande (2007) estimate the effect of dam construction in India on agriculture, Jensen (2007) evaluates how the construction of cellular phone towers in South India improved efficiency in the fish market, and Michaels (2008) estimates the effect of the US Interstate Highway system on the skilled wage premium. An older literature, beginning with Aschauer (1989), pioneered the use of econometric methods in estimating the benefits of infrastructure projects.

[5]Fogel (1964) first applied the social savings methodology to railroads in the United States, and Hurd (1983) performed a similar exercise for India. In section 1.6.7 I compare my estimates to those from using a social savings approach.

explain my reduced-form estimates. Using a model also improves the external validity of my estimates because the primitive in my model—the cost of trading—is specified explicitly, and is portable to a range of settings in which the welfare benefits of trade cost-reducing polices might be sought.[6] By contrast, my reduced-form estimates are more likely to be specific to the context of railroads in colonial India. Finally, the model suggests a general equilibrium treatment externality of railroads that, if ignored, would bias estimates of the effects of this infrastructure project by almost 20 percent.[7] This point has not, to my knowledge, been incorporated before in the infrastructure literature.

This chapter also contributes to a rich literature concerned with estimating the welfare effects of openness to trade, because the reduction in trade costs brought about by India's railroad network rapidly increased each district's exposure to trade.[8] My chapter adds to this literature in two ways. First, I explicitly connect my empirical approach to an estimable, general equilibrium model of trade. One advantage of doing so is that the model suggests an appropriate way to measure 'openness'. This has been an object of significant debate in this literature but the measurement of openness suggested by my model (the share of a region's expenditure it sources from abroad) has not been used before in an esti-

[6]For example, Raballand and Macchi (2008) find in surveys of African trucking firms that transportation costs are relatively high in Africa because of a number of policy-relevant features (e.g. poor roads, expensive inputs, and underutilized payload capacity). Similarly, Djankov, Freund, and Pham (2006) survey freight forwarding firms in 126 countries to measure wider, policy-relevant costs of trading (e.g. inspections, technical clearance, mandatory storage, port handling, and customs clearance).

[7]Most of the policy evaluation literature assumes that policy treatments received by one unit of observation do not affect outcomes for any other units (the "stable unit treatment value assumption," in the language of Rubin (1978).) Heckman and Abbring (2007) survey the recent literature on general equilibrium policy evaluation.

[8]Frankel and Romer (1999), Rodriguez and Rodrik (2001), Irwin and Tervio (2002) and Alcala and Ciccone (2004) use cross-country regressions of real GDP levels on 'openness' (defined in various ways) to estimate the effect of openness on welfare. Topalova (2005), Trefler (2004) and Pavcnik (2002) instead analyze trade liberalizations within one country by exploiting cross-sectional variation in the extent of liberalization across either industries or regions.

mation approach.[9] Second, I pursue an approach that treats openness to both internal and external trade as similar phenomena, while previous work has focused only on the gains from external openness.[10]

The next section describes the historical setting in which the Indian railroad network was constructed and the new data that I have collected from that setting. In section 1.3, I outline a model of trade in colonial India and the model's six predictions. Sections 1.4 through 1.7 present six empirical steps that test the model's six predictions qualitatively and quantitatively. Section 1.8 concludes.

1.2 Historical Background and Data

In this section I discuss some essential features of the economy in colonial India and the data that I have collected in order to analyze how this economy changed with the advent of railroad travel. I go on to describe the transportation system in India before and after the railroad era, and the institutional details that determined when and where railroads were built.

1.2.1 New Data on the Indian Economy, 1870-1930

In order to evaluate the impact of the railroad network on economic welfare in colonial India I have constructed a new panel dataset on 239 Indian districts. The dataset tracks these districts annually from 1870-1930, a period during which 98 percent of British India's current railroad lines were opened. Table 1.1 contains descriptive statistics for the variables

[9]Waugh (2007), however, uses Eaton and Kortum (2002), and hence the definition of 'openness' employed here, to inform a development accounting approach to estimating the contribution of trade to growth.

[10]An exception is Poncet and Wei (2005), who estimate the effect of trade costs on growth at the provincial level in China.

that I use in this chapter and describe throughout this section. Appendix A contains more detail on the construction of these variables.

During the colonial period, India's economy was predominantly agricultural, with agriculture constituting an estimated 66 percent of GDP in 1900 (Heston 1983).[11] For this reason, district-level output data was only collected systematically in the agricultural sector. Data on agricultural output was recorded for each of 17 principal crops (which comprised 93 percent of the cropped area of India in 1900). Retail prices for these 17 crops were also recorded at the district-level. I use these price figures to construct a nominal agricultural GDP series for each district and year and then a real agricultural income per acre figure by dividing by a consumer price index and district land area.[12] The resulting real agricultural income per acre variable provides the best available measure of district-level economic welfare in this time period.

Real incomes were low during my sample period, but there was 22 percent growth between 1870 and 1930.[13] Real incomes were low because crop yields were low, both by contemporaneous international standards and by Indian standards today.[14] One explanation for low yields, which featured heavily in Indian agricultural textbooks of the day (such as

[11]Factory-based industry—which Chandler (1977) and Atack, Haines, and Margo (2008) argue benefited from access to railroads in the United States—contributed only 1.6 percent of India's GDP in 1900.

[12]The use of real income, rather than real GDP, in open-economy settings is advocated by Diewert and Morrison (1986), Feenstra (2003) and Kehoe and Ruhl (2008). As the latter authors argue, real income captures the gains from trade in a wide range of trade models, but real GDP does not.

[13]For comparison, Heston (1983) estimates that in 1869, on the basis of purchasing power exchange rates, per capita income in the United States was four times that in India. This income disparity rises to ten if market exchange rates are used instead of PPP rates.

[14]For example, the yield of wheat in India's 'breadbasket', the province of Punjab, was 748 lbs/acre in 1896. By contrast, for similar types of wheat, yields in Nevada (the highest state yields in the United States) in 1900 were almost twice as high (see plate 15 of United States Census Office (1902)) and yields in (Indian) Punjab in 2005 were over five times higher than those in 1896 (as calculated from the *Indian District-wise Crop Production Statistics Portal*, `http://dacnet.nic.in/apy/cps.aspx`).

	Number of Observations	Beginning of Available Data	End of Available Data
Value of agricultural output per acre, all crops (current rupees)	14,340	27.3 (10.4)	111.3 (40.8)
Agricultural prices (average over all crops, current rupees per maund)	14,340	2.37 (1.37)	5.07 (2.35)
Real agricultural income per acre (1870 rupees)	14,340	27.3 (10.4)	38.0 (13.8)
Price of salt, all sources (current rupees per maund)	7,329	5.19 (1.96)	3.45 (0.465)
Total annual rainfall (meters)	14,340	1.011 (0.798)	1.145 (1.302)
Crop-specific rainfall shock (meters)	73,000	0.638 (0.614)	0.662 (0.602)
Exports per trade block (millions of 1870 rupees)	6,581,327	0.711 (0.649)	3.581 (2.444)

Table 1.1: Summary Statistics Notes: Values are sample means over all observations for the year and question, with standard deviations in parentheses. Beginning and end of available data are: 1870 and 1930 for agricultural output and real agricultural income; 1861 and 1930 for agricultural prices and salt prices; 1867 and 1930 for all rainfall variables; and 1880 and 1920 for trade data. A 'maund' is equal to 37.3 kg and was the standardized unit of weight in colonial India. Data sources and construction are described in full in Appendix A.

Leake (1923), Mollison (1901) and Wallace (1892)), was inadequate water supply. Only 12 percent of cultivated land was irrigated in 1885; while this figure had risen to 19 percent in 1930, the vast majority of agriculture maintained its dependence on rainfall.[15]

Because rainfall was important for agricultural production, 3614 meteorological stations (plotted in Figure 1.1) were built throughout the country to record the amount of rainfall at each station on every day of the year. Daily rainfall data was recorded and published because the distribution of rainfall throughout the year was far more important to farmers and traders than total annual or monthly amounts. In particular, the intra-annual distribution of rainfall governed how different crops (which were grown in distinct stretches of the year) were affected by a given year's rainfall. In sections 1.5 through 1.7, I use daily rainfall data from each of India's 3614 meteorological stations to construct crop-specific measures of rainfall and use these as a source of exogenous variation in crop-specific productivity.

Commensurate with the increase in the level of real agricultural income in India was a dramatic rise in interregional and international trade. The final component of the dataset that I have constructed on colonial India is data on these internal and external trades.

A great deal of change in internal trade volumes appears to have occurred in India, though this is difficult to measure accurately due to only limited amounts of available data on road trade (data on trade by rail, river and coastal shipping, by contrast, are well recorded throughout the period). The average Indian trade block (a geographical unit containing approximately 5 districts) saw its real exports along all non-road means

[15] These figures encompass a wide definition of irrigation, including the use of tanks, cisterns, and reservoirs as well as canals. See the *Agricultural Statistics of India*, described in Appendix A. 1885 is the first year in which comprehensive irrigation statistics were collected.

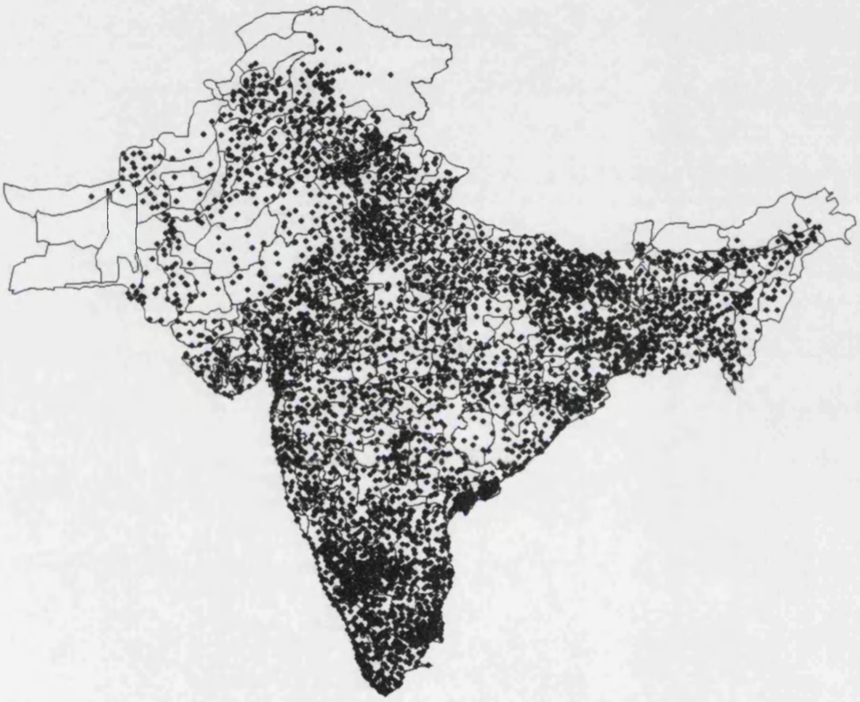

Figure 1.1: Meteorological stations in colonial India: Dots represent the 3914 meteorological stations with rain gauges collecting daily rainfall in the period from 1891-1930. District borders are also shown. Source: *Global Historical Climatological Network* an author's calculations; see Appendix A for full details.

of transportation (ie along rail, river, and coastal routes) grow by a factor of 5 between 1870 and 1920. In particular two regions, however, it is possible to say more about the nature of trade along all available modes of transportation. First, in the Bengal Presidency, for example, trade to and from the port of Calcutta rarely took place on roads, because river transport was so plentiful. Total real exports (along both rivers and railroads) in Bengal, from inland regions to Calcutta, grew by a factor of 6.2 from 1876 to 1920. A second example comes from the Chhattisgarh region of the Central Provinces (a landlocked region with no navigable rivers), where data were collected in 1870 on the region's imports/exports via road, before the railroad network entered the region. This region's real exports grew by at least a factor of 7 between 1870 and 1920 (though this figure is an underestimate of total trade growth because the 1920 trade data excludes road trade).

India's economy became increasingly engaged in international trade from 1870 to 1930, with international trade flows roughly doubling (as a share of national income) over the period (Chaudhuri 1983). India's dominant exports in 1860 were opium, raw cotton and food grains, but by 1930 the opium trade had collapsed, and jute and tea had taken its place. Almost 80 percent of these exports went to Great Britain and China in 1860, but by 1930, while Britain was still predominant (with 25 percent of exports), Japan and the United States also important. While exports, both in 1860 and 1930, were primarily agricultural and other primary commodities, processed textiles (cotton piecegoods) and metals/machinery made up the bulk of imports. These imports came almost entirely (85 percent) from Britain in 1860, though by 1930 this share had fallen to 23 percent with the United States and Japan contributing approximately 10 percent of India's imports each.

1.2.2 Transportation in Colonial India

Prior to the railroad era, goods transport within India took place on roads, rivers, and coastal shipping routes. The bulk of inland travel was carried by bullocks, along the road network.[16] Bullocks were employed either as 'pack bullocks' (which carried goods strapped to their backs and usually traveled directly over pasture land), or 'cart bullocks' (which pulled a cart containing goods and traveled along improved roads). On the best road surfaces and during optimal weather conditions, cart bullocks could cover 20-30 km per day. However, high-quality roads were extremely sparse and the roads that did exists were virtually impassable in the monsoon season (Deloche 1994). Pack bullocks were more versatile than cart bullocks, but their freight rates were three times higher per unit distance and weight (Derbyshire 1985).

Water transport was far superior to road transport, but it was only feasible on the Brahmaputra, Ganges and Indus river systems.[17] In optimal conditions, downstream river traffic (with additional oar power[18]) could cover 65 km per day; upstream traffic needed to be towed from the banks and struggled to cover 15 km per day. Extensive river travel was impossible in the rainy monsoon months, or the dry summer months, and piracy was a serious hazard (Deloche 1995). Coastal shipping, however, was perennially available along India's long coastline. This form of shipping was increasingly steam-powered post-1840. Steamships were fast, covering over 100 km per day, but they could only service major ports. The bulk of this trade, both before and after the railroads, therefore

[16] Camels were also used in sandy areas. Horses, ponies, donkeys, mules and elephants were less common forms of animal-powered transportation.

[17] Navigable canals either ran parallel to sections of these three rivers or were extremely localized in a small number of coastal deltas (Stone 1984, Whitcombe 1983).

[18] Steamboats had periods of success in the colonial era, but were severely limited in scope by India's seasonal and shifting rivers (Derbyshire 1985).

consisted of shipments between the major ports (Naidu 1936).

Against this backdrop of costly and slow internal transportation, the appealing prospect of railroad transportation in India was discussed as early as 1832 (Sanyal 1930)—though it was not until 1853 that the first track was actually laid. From the outset, railroad transport proved to be far superior to road, river or coastal transport (Banerjee 1966). Railroads were capable of traveling up to 600 km per day and they offered this superior speed on predictable timetables, throughout all months of the year, without any risk of piracy (Johnson 1963). Railroad freight rates were also considerably cheaper: 4-5, 2-4, and 1.5-3 times cheaper than road, river and coastal travel, respectively (Deloche 1994, Deloche 1995, Derbyshire 1985, Hurd 1975).

1.2.3 Railroad Line Placement Decisions

Throughout the history of India's railroads, all railroad line placement decisions were made by the Government of India. It is widely accepted that the Government had three motives for building railroads: military, commercial, and humanitarian—in that order of priority (Thorner 1950, Macpherson 1955, Headrick 1988). In 1853, Lord Dalhousie (head of the Government of India) wrote an internal document to the East India Company's Court of Directors that sketched railroad policy in India for decades to come.[19] Military motivations for railroad-building appeared on virtually every page of this document,[20] and these motivations gained

[19]The Government's policy on how best to operate railroads, however, changed from decade to decade. In times of private ownership, companies were enticed to build the Government's desired lines by a guarantee scheme. In exchange for building a line of their specification (and allowing free military and postal traffic), the Government would guarantee the company a minimum five percent return on capital (and half of any returns over five percent).

[20]For example, from the introduction: "A single glance...will suffice to show how immeasurable are the political advantages to be derived from the system of internal communication, which would admit of full intelligence of every event being transmit-

new momentum when the 1857 'mutiny' highlighted the importance of military communications (Headrick 1988). Dalhousie's minute described five 'trunk lines' that would connect India's five major provincial capitals along direct routes and maximize the "political advantages" of a railroad network.

Between 1853 and 1869, all of Dalhousie's trunk lines were built—but not without significant debate over how best to connect the provincial capitals. Dalhousie and Major Kennedy, India's Chief Engineer, spent over a decade discussing and surveying (at great cost) their competing, but very different, proposals for a pan-Indian network (Davidson 1868, Settar 1999). This debate indicates the vicissitudes of railroad planning in India and it was repeated many times by different actors in Indian railroad history. I have collected planning documents from a number of railroad expansion proposals that, like Kennedy's proposal, were debated and surveyed at length, but were never actually built. As discussed in section 1.6.4, I use these plans in a placebo strategy to check that unbuilt lines display no spurious 'impact' on the district economies in which they were nearly built.

By 1876, railroad expansion had slowed significantly in India. But railroads benefited from new enthusiasm in the wake of the 1880 Famine Commission, which recommended railroads as a means for future famine prevention. (Chapter 2 of this thesis investigates the merits of this recommendation by evaluating the extent to which railroads reduced the volatility of prices, real incomes, and the incidence of death due to drought.) The Commission's recommendations for specific railroad lines formed the bedrock on which more detailed plans over the ensuing 15

ted to the Government...and would enable the Government to bring the main bulk of its military strength to bear upon any given point in as many days as it would now require months, and to an extent which at present is physically impossible." (House of Commons Papers 1853).

years were built. In section 1.6.5 I describe how this motivates an instrumental variable for railroad construction. A second consequence of the 1880 Famine Commission report, from the perspective of my identification strategy in section 1.6.6, is that all railroad proposals from 1883 to 1904 were required to be designated according to their intended purpose. I use this feature to motivate a set of bounds on my estimates of the average effect of railroads.

As is clear from Figure 1.2, the railroad network in place in 1930 (by and large, the same network that is open today) had completely transformed the transportation system in India. 67,247 km of track were open for traffic, constituting the fourth-largest network in the world. From their inception in 1853 to their zenith in 1930, railroads were the dominant form of public investment in British India.[21] But influential observers were highly critical of this public investment priority—the Nationalist historian, Romesh Dutt, argued that they did little to promote agricultural development,[22] and Mahatma Gandhi argued simply that "there can be little doubt that they [railroads] promote evil."[23] In the remainder of this chapter I use new data to assess quantitatively the effect of railroads on India's trading environment and agricultural economy.

[21]Kumar (1983) summarizes government revenues and expenditure in India. Public investment (which included railroads, roads, irrigation, buildings, health and education) accounted for around 20 percent of total expenditure (the largest category behind defense) and railroads accounted for over 40 percent of this category.

[22]For example, on page 174 of his textbook on Indian economic history: "Railways...did not add to the produce of the land." (Dutt 1904)

[23]From Chapter IX of the 1938 English translation of Gandhi's 1909 *Hind Swaraj* [Indian Home Rule], his influential newspaper columns. Other passages are equally polemic: "...but for the railways, the English could not have such a hold on India as they have. The railways, too, have spread the bubonic plague...Railways have also increased the frequency of famines, because, owing to facility of means of locomotion, people sell out their grain, and it is sent to the dearest markets....They [railways] accentuate the evil nature of man. Bad men fulfill their evil designs with greater rapidity." (Gandhi 1938)

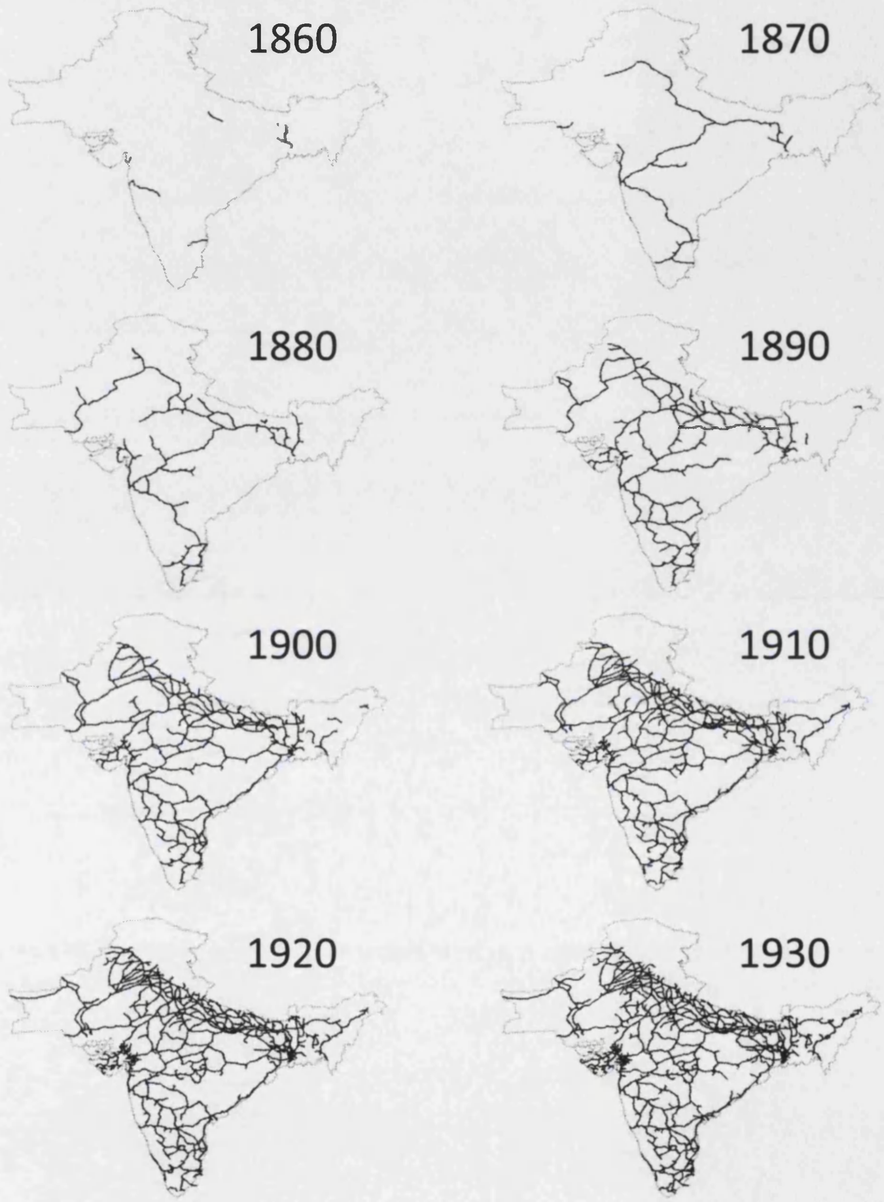

Figure 1.2: The evolution of India's railroad network, 1860-190: These figures display the decadal evolution of the railroad network (railroads illustrated with thick lines) in colonial India (the outline of which is illustrated with thin lines). Source: Author's calculations based on official publications. See Appendix A for full details.

1.3 A Model of Railroads and Trade in Colonial India

In this section I develop a general equilibrium model of trade among many regions in the presence of trade costs. The model is based on Eaton and Kortum (2002), extended to a setting with more than one commodity; this extension allows me to generate cross-commodity predictions that exploit the full richness of my commodity-level data.[24] The model serves two purposes. First, it delivers four predictions about the response of observables to trade cost reductions. Second, I estimate the unknown parameters of the model and use the estimated model to assess whether the observed reduction in trade costs due to the railroads can account, via the mechanism stressed in this model, for the observed increase in welfare due to railroads. Both of these features inform our understanding of *how* transportation infrastructure projects can raise welfare.

1.3.1 Model Environment

The economy consists of D regions (indexed by either o or d). There are K commodities (indexed by k), each available in a continuum (with mass normalized to one) of horizontally differentiated varieties (indexed by j). In my empirical application I work with data on prices, output and trade flows that refer to commodities, not individual varieties. While my empirical setting will consider 70 years of annual observations, for simplicity the model is static; I therefore suppress time subscripts until they are necessary.

[24]While Eaton and Kortum (2002) use a continuum of varieties (j in the notation I use below), these are all varieties from one commodity (k in the notation I use below). The main predictions of their model, and the model that I present here, are at the level of commodities rather than varieties.

Consumer Preferences:

Each region o is home to a mass (normalized to one) of identical agents, each of whom owns L_o units of land. Land is geographically immobile and supplied inelastically. Agents have Cobb-Douglas preferences over commodities (k) and constant elasticity of substitution preferences over varieties (j) within each commodity; that is, their (log) utility function is

$$\ln U_o = \sum_{k=1}^{K} \left(\frac{\mu_k}{\varepsilon_k} \right) \ln \int_0^1 (C_o^k(j))^{\varepsilon_k} dj, \tag{1.1}$$

where $C_o^k(j)$ is consumption, $\varepsilon_k = \frac{\sigma_k - 1}{\sigma_k}$ (where σ_k is the (constant) elasticity of substitution), and $\sum_k \mu_k = 1$. Agents rent out their land at the rate of r_o per unit and use their income $r_o L_o$ to maximize utility from consumption.

Production and Market Structure:

Each variety j of the commodity k can be produced using a constant returns to scale production technology in which land is the only factor of production.[25] Let $z_o^k(j)$ denote the amount variety j of commodity k that can be produced with one unit of land in region o. I follow Eaton and Kortum (2002) in modeling $z_o^k(j)$ as the realization of a stochastic variable Z_o^k drawn from a Type-II extreme value distribution whose parameters vary across regions and commodities; that is,

$$F_o^k(z) \doteq \Pr(Z_o^k \leq z) = \exp(-A_o^k z^{-\theta_k}) \tag{1.2}$$

[25]It is straightforward to extend this setting to an arbitrary number of immobile factors (provided that the stochastic productivity term remains Hicks-neutral) by replacing the land rental rate r_o with a general unit cost function $c^k(\mathbf{w}_o)$, where $\mathbf{w}_o$ is the vector of factor payments in region o.

where $A_o^k \geq 0$ and $\theta_k > 0$. These random variables are drawn independently for each variety, commodity and region.[26] The exogenous parameter A_o^k increases the probability of high productivity draws and the exogenous parameter θ_k captures (inversely) how variable the productivity of commodity k in any region is around its average.

There are many competitive firms in region o with access to the above technology; consequently, firms make zero profits.[27] These firms will therefore charge a pre-trade costs price of $p_{oo}^k(j) = r_o/z_o^k(j)$, where r_o is the land rental rate in region o.

Opportunities to Trade:

Without opportunities to trade, consumers in region d must consume even their region's worst draws from the productivity distribution in equation (1.2). The ability to trade breaks this production-consumption link. This allows consumers to import varieties from other regions in or-

[26]The assumption of within-sector heterogeneity characterized by a continuous stochastic distribution of productivities is a standard feature in the literature on trade with heterogeneous firms (eg Melitz (2003)). It is common in that literature to assume that the productivity distribution is Pareto (to which the upper tail of a Type-II extreme value distribution converges) and that productivities are drawn independently across varieties (firms), commodities, and countries (eg Melitz and Ottaviano (2007), Chaney (2008) and Helpman, Melitz, and Rubinstein (2008).) An attraction of the Type-II extreme value distribution is its plausible micro-foundations: Kortum (1997) applies the extremal types theorem to show that the distribution of productivities among producers who use only the highest draws from any iid process of 'ideas' will converge to an extreme value distributional form. Nevertheless, Costinot and Komunjer (2008) show that the key features of the Eaton and Kortum (2002) model hold locally around a symmetric distribution of exogenous productivity terms A_o^k for *any* continuous productivity distribution.

[27]My empirical application is primarily to the agricultural sector. This sector was characterized by millions of small-holding farmers who were likely to be price-taking producers of undifferentiated products (varieties j in my model). For example, in the 1901 census in the province of Madras, workers in the agricultural sector (67.9 percent of the almost 20 million strong workforce) were separately enumerated by their ownership status, and 35.7 percent of these workers were owner-cultivators (extremely small-scale farms) (Risley and Gait 1903). Nevertheless, Bernard, Eaton, Jensen, and Kortum (2003) and Eaton, Kortum, and Kramarz (2005) extend the Eaton and Kortum (2002) framework to allow for Bertrand and monopolistic competition, respectively. While in principle it is possible to estimate these alternative models, the most natural way to do so uses firm-level trade data, which is unavailable in my setting.

der to take advantage of the favorable productivity draws available there, and allows producers to produce more of the varieties for which they received the best productivity draws. These two mechanisms constitute the gains from trade in this model.

However, there is a limit to trade because the movement of goods is subject to trade costs (which include transport costs and other barriers to trade). These trade costs take the convenient and commonly used Samuelson (1954) 'iceberg' form. That is, in order for one unit of commodity k to arrive in region d, $T^k_{od} \geq 1$ units of the commodity must be produced and shipped in region o; trade is free when $T^k_{od} = 1$. (Throughout this chapter I refer to trade flows between an *origin* region o and a *destination* region d; all bilateral variables, such as T^k_{od}, refer to quantities *from o to d*.) Trade costs are assumed to satisfy the property that it is always (weakly) cheaper to ship directly from region o to region d, rather than via some third region m: that is, $T^k_{od} \leq T^k_{om}T^k_{md}$. Finally, I normalize $T^k_{oo} = 1$. In my empirical setting I proxy for T^k_{od} with measures calculated from the observed transportation network, which incorporates all possible modes of transport between region o and region d. Railroads enter this transportation network gradually over time, reducing T^k_{od} and creating more gains from trade.

Trade costs drive a wedge between the price of an identical variety in two different regions. Let $p^k_{od}(j)$ denote the price of variety j of commodity k produced in region o, but shipped to region d for consumption there. The iceberg formulation of trade costs implies that any variety in region d will cost T^k_{od} times more than in region o; that is, $p^k_{od}(j) = T^k_{od}\, p^k_{oo}(j) = r_o\, T^k_{od}/z^k_o(j)$.

<u>Equilibrium Prices and Allocations:</u>

Consumers have preferences for all varieties j along the continuum of varieties of commodity k. But they are are indifferent about where a given variety is made—they simply buy from the region that can provide the variety at the lowest cost. I therefore solve for the equilibrium prices that consumers in a region d actually face, given that they will only buy a given variety from the cheapest source region (including their own).

The price of a variety sent from region o to region d, denoted by $p_{od}^k(j)$, is stochastic because it depends on the stochastic variable $z_o^k(j)$. Since $z_o^k(j)$ is drawn from the CDF in equation (1.2), $p_{od}^k(j)$ is the realization of a random variable P_{od}^k drawn from the CDF

$$G_{od}^k(p) \doteq \Pr(P_{od}^k \leq p) = 1 - \exp[-A_o^k(r_o T_{od}^k)^{-\theta_k} p^{\theta_k}]. \tag{1.3}$$

This is the price distribution for varieties (of commodity k) made in region o that could *potentially* be bought in region d. The price distribution for the varieties that consumers in d will *actually* consume (whose CDF is denoted by $G_d^k(p)$) is the distribution of prices that are the lowest among all D regions of the world:

$$\begin{aligned} G_d^k(p) &= 1 - \prod_{o=1}^{D}[1 - G_{od}^k(p)], \\ &= 1 - \exp\left(-\left[\sum_{o=1}^{D} A_o^k(r_o T_{od}^k)^{-\theta_k}\right] p^{\theta_k}\right). \end{aligned}$$

Given this distribution of the actual prices paid by consumers in region d, it is straightforward to calculate any moment of the prices of interest. The price moment that is important for my empirical analysis is the expected value of the equilibrium price of any variety j of commodity k found in region d, which is given by

$$E[p_d^k(j)] \doteq p_d^k = \lambda_1^k \left[\sum_{o=1}^{D} A_o^k (r_o T_{od}^k)^{-\theta_k} \right]^{-1/\theta_k}, \tag{1.4}$$

where $\lambda_1^k \doteq \Gamma(1 + \frac{1}{\theta_k})$.[28] In section 2.4.1 I treat these expected prices as equal to the observed prices collected by statistical agencies.[29]

Given the price distribution in equation (1.3), Eaton and Kortum (2002) derive two important properties of the trading equilibrium that carry over to the model here. First, the price distribution of the varieties that any given origin actually sends to destination d (ie the distribution of prices for which this origin is region d's cheapest supplier) is the same for all origin regions. This implies that the share of expenditure that consumers in region d allocate to varieties from region o must be equal to the probability that region o supplies a variety to region d (because the price per variety, conditional on the variety being supplied to d, does not depend on the origin). That is $X_{od}^k / X_d^k = \pi_{od}^k$, where X_{od}^k is total expenditure in region d on commodities of type k from region o, $X_d^k \doteq \sum_o X_{od}^k$ is total expenditure in region d on commodities of type k, and π_{od}^k is the probability that region d sources any variety of commodity k from region o. Second, this probability π_{od}^k is given by

$$\pi_{od}^k = \frac{X_{od}^k}{X_d^k} = \lambda_3^k \, A_o^k \, (r_o T_{od}^k)^{-\theta_k} \, (p_d^k)^{\theta_k}, \tag{1.5}$$

[28] $\Gamma(.)$ is the Gamma function defined by $\Gamma(z) = \int_0^\infty t^{z-1} e^{-t} dt$.

[29] A second price moment that is of interest for welfare analysis is the exact price index over all varieties of commodity k for consumers in region d. Given CES preferences, this is $\widetilde{p}_d^k \doteq \left[\int_0^1 (p_d^k(j))^{1-\sigma_k} dj \right]^{1/1-\sigma_k}$, which is only well defined here for $\sigma_k < 1 + \theta_k$ (a condition I assume throughout). The exact price index is given by $\widetilde{p}_d^k = \lambda_2^k p_d^k$, where $\lambda_2^k \doteq \frac{\gamma^k}{\lambda_1^k}$ and $\gamma^k \doteq [\Gamma(\frac{\theta_k + 1 - \sigma_k}{\theta_k})]^{1/(1-\sigma_k)}$. That is, if statistical agencies sampled varieties in proportion to their weights in the exact price index, as opposed to randomly as in the expected price formulation of equation (1.4), then this will not jeopardize my empirical procedure because the exact price index is proportional to expected prices.

where $\lambda_3^k = (\lambda_1^k)^{-\theta_k}$, and this equation makes use of the definition of the expected value of prices (ie p_d^k) in equation (1.4).

Equation (1.5) characterizes trade flows conditional on the endogenous land rental rate, r_o (and all regions' land rental rates, which appear in p_d^k). It remains to solve for these land rents in equilibrium, by imposing the condition that each region's trade is balanced. Region o's trade balance equation requires that the total income received by land owners in region o ($r_o L_o$) must equal the total value of all commodities made in region o and sent to every other region (including region o itself). That is:

$$r_o L_o = \sum_d \sum_k X_{od}^k = \sum_d \sum_k \pi_{od}^k \, \mu_k \, r_d \, L_d, \tag{1.6}$$

where the last equality uses the fact that (with Cobb-Douglas preferences) expenditure in region d on commodity k (X_d^k) will be a fixed share μ_k of the total income in region d ($r_d L_d$). Each of the D regions has its own trade balance equation of this form. I take the rental rate in the first region (r_1) as the numeraire good, so the equilibrium of the model is the set of D-1 unknown rental rates r_d that solves this system of D-1 (non-linear) independent equations.

1.3.2 Four Predictions

In this section I state explicitly four of the model's predictions. These predictions are presented in the order in which I test for them in my empirical analysis (ie Steps 1-4).

<u>Prediction 1: Price Differences Measure Trade Costs (in Special Cases):</u> In the presence of trade costs, the price of identical commodities will differ across regions. In general, the cost of trading a commodity between

two regions places only an upper bound on their price differential.[30] However, in the special case of a homogeneous commodity that can only be produced in one origin region, equation (1.4) predicts that the (log) price differential between the origin o of this commodity and any other region d will be equal to the (log) cost of trading the commodity between them. That is:

$$\ln p_d^o - \ln p_o^o = \ln T_{od}^o, \tag{1.7}$$

where the commodity label k is replaced by o to indicate that this equation is only true for commodities that can only be made in region o. This prediction is important for my empirical work because it allows trade costs (T_{od}^o), which are typically unobserved, to be inferred.[31] But it is important to note that this prediction—essentially just free arbitrage over space, net of trade costs—is common to many models of spatial equilibrium.[32]

[30]This can be easily seen in two simple settings where bilateral trade costs are infinite, but bilateral inter-regional price differences are zero: (i) two identical autarkic economies will have a price differential of zero, but infinite trade costs vis-a-vis each other; (ii) two regions that have infinite trade costs vis-a-vis each other could both buy a commodity from some common third region from which they are both separated by the same trade cost (meaning that they face the same price for this commodity and therefore have a price difference of zero).

[31]There are two obstacles to using inter-regional price differentials to infer trade costs in wider settings than that employed here. First, the commodity whose price is being compared over space must be identical in the two regions—for example, Broda and Weinstein (2007) use barcode data to illustrate the misleading inferences that have been drawn from comparing prices of commodities that are similar, but not identical, across the Canada-US border. Second, even with a homogeneous commodity, only if two regions actually trade the commodity will their inter-regional price difference be equal to their bilateral trade cost. Restricting attention to a commodity that is only made in one region but is consumed elsewhere, as I do in this chapter, helps to ensure that the commodity is homogeneous and guarantees that the commodity was actually traded between regions.

[32]A class of exceptions is those with some forms of imperfect competition and in which producers can charge separate prices in separate markets, as in Brander and Krugman (1983) or Melitz and Ottaviano (2007). However, my empirical application of this prediction will be to salt, which was produced under strict government license at a small number of locations and then had to be sold (under conditions of the license) to an unrestricted trading community at the 'factory' gate (United Provinces of Agra and Oudh 1868). That is, in this setting, producers only charged one factory gate price.

Prediction 2: Bilateral Trade Flows Take the 'Gravity Equation' Form:

Equation (1.5) describes bilateral trade flows explicitly, but I re-state it here in logarithms for reference: (log) bilateral trade of any commodity k from any region o to any other region d is given by

$$\ln X_{od}^k = \ln \lambda_k + \ln A_o^k - \theta_k \ln r_o - \theta_k \ln T_{od}^k + \theta_k \ln p_d^k + \ln X_d^k. \quad (1.8)$$

This is the gravity equation form for bilateral trade flows: bilateral trade costs reduce bilateral trade flows, conditional on importer- and exporter-specific terms.[33]

Prediction 3: Railroads Increase Real Income Levels:

My focus here is on real income (which, as prediction 4 shows explicitly, is equal to welfare in this model). To simplify notation, let W_o represent real income per unit land area (ie $W_o \doteq \frac{r_o}{\widetilde{P}_o}$, where $\widetilde{P}_o$ is the aggregate price index in region o, defined explicitly in Prediction 4).

Unfortunately, the multiple general equilibrium interactions in the model are too complex to admit a closed-form solution for the effect of reduced trade costs on agricultural prices. To make progress in generating qualitative predictions (to guide my empirical analysis) I therefore assume a much simpler environment for the purpose of obtaining prediction 3 only. I assume: there are only three regions (called X, Y and Z); there is only one commodity (so I will dispense with the k superscripts

[33] A number of theoretical trade frameworks also predict a gravity equation for trade flows. Examples include Anderson (1979), Deardorff (1998), Helpman, Melitz, and Rubinstein (2008) and Chaney (2008). In a traditional gravity equation, bilateral trade flows (for each commodity) are proportional to the expenditure of the importing region, the output of the exporting region, and inversely proportional to the bilateral cost of trading between the two regions. Equation (1.5) can be easily manipulated to take this form. However, what matters for my empirical procedure is simply that, conditional on importer and exporter fixed effects, bilateral trade costs reduce bilateral trade flows—as in equation (1.5) and in a traditional gravity equation.

on all variables); the regions are symmetric in their exogenous characteristics (ie $L_o = L$ and $A_o = A$ for all regions o); and the three regions have symmetric trade costs with respect to each other.[34] I consider the comparative statics from a local change around this symmetric equilibrium, where it is straightforward to show that:

1. $\frac{dW_X}{dT_{YX}} < 0$: Real income in a region (say, X) rises when the cost of trading between that region and any other region (say, Y) falls.

2. $\frac{dW_X}{dT_{YZ}} > 0$: Real income in a region (say, X) falls when the cost of trading between the two other regions (ie T_{YZ}) falls.

These results suggest that a reduction in trade costs in one part of the network is not good for all regions. A railroad project that reduces trade costs between two regions will raise welfare in these two regions; but this project will reduce welfare in the third, excluded region whose trade costs were unaffected by the project. This negative effect on excluded regions arises because of two effects: first, the excluded region's trading partners' land rental costs have increased (because these partners' own trade costs have fallen), which raises the prices of the commodities that the partners ship to the excluded region; and second, the excluded region loses demand for its exports because its trading partners now have a cheaper supplier (in each other).

Prediction 4: A Sufficient Statistic for Railroad Impact on Welfare:

Given the utility function in equation (1.1), the indirect utility function

[34] An alternative means of obtaining analytical predictions would be to invoke the commonly-used assumption that one commodity can be traded at zero trade cost, and is important enough to be produced in positive quantities everywhere and always. (This equates r_o to the nominal productivity in this zero trade cost sector). It is difficult to imagine a commodity that satisfies these conditions in colonial India.

per unit of land (denoted by W_o as in prediction 3) in region o is

$$W_o = \frac{r_o}{\prod_{k=1}^{K}(\widetilde{p}_o^k)^{\mu_k}} \doteq \frac{r_o}{\widetilde{P}_o}. \tag{1.9}$$

The numerator of this expression is nominal income (per unit land area) in region o relative to the numeraire, and the denominator is the exact consumer price index across all commodities k denoted by $\widetilde{P}_o$. That is, welfare is equal to real income. Using the bilateral trade equation (1.5) evaluated at $d = o$, (log) real income per unit of land (defined as W_o as in prediction 3) can be re-written as

$$\ln W_o = \Omega + \sum_k \tfrac{\mu_k}{\theta_k} \ln A_o^k - \sum_k \tfrac{\mu_k}{\theta_k} \ln \pi_{oo}^k, \tag{1.10}$$

where $\Omega \doteq -\sum_k \mu_k \ln \gamma^k$. This result states that welfare is a function of only two terms: local productivity (A_o^k), and a term that I will refer to as 'autarkiness' (ie π_{oo}^k), the fraction of region o's expenditure that region o buys from itself). Because of the complex general equilibrium relationships in the model, the full vector of trade costs (between every bilateral pair of regions), the full vector of productivity terms in other regions, and the sizes of every region all influence welfare in region o. But these terms (that is, every exogenous variable in the model other than local productivity) affect welfare only through their effect on autarkiness. Put another way, autarkiness (the appropriately weighted sum of π_{oo}^k terms over goods k) is a sufficient statistic for welfare in region o, once local productivity is controlled for. If railroads affected welfare in India through the mechanism in the model (by reducing trade costs, giving rise to gains from trade), then prediction 4 states that they did so only by changing π_{oo}^k.

1.3.3 From Theory to Empirics

To relate the static model in section 1.3 to my dynamic empirical setting (with 70 years of annual data) I take the simplest possible approach and assume that all of the goods in the model cannot be stored, and that inter-regional lending is not possible. Furthermore, I assume that the stochastic production process described in section 1.3.1 is drawn independently in each period. These assumptions imply that the static model simply repeats every period, with independence of all decision-making across time periods. Throughout the remainder of the chapter I therefore add the subscript 't' to all of the variables (both exogenous and endogenous) in the model, but I assume that all of the model parameters (θ_k, σ_k and μ_k) are fixed over time.

The four theoretical predictions outlined in section 1.3.2 take a naturally recursive order, both for estimating the model's parameters, and for tracing through the impact of railroads on welfare in India. I follow this order in the four empirical sections that follow (ie Steps 1-4). In Step 1, I evaluate the extent to which the railroads reduced trade costs within India. To do this I use Prediction 1 to relate the unobserved trade costs term in the model (T^k_{odt}) to observed features of the transportation network. In Step 2, I use Prediction 2 to measure how much the reduced trade costs found in Step 1 increased trade in India. This relationship allows me to estimate the unobserved model parameter θ_k, and to relate the unobserved productivity terms (A^k_{ot})[35] to rainfall, which is an exogenous and observed determinant of agricultural productivity. Steps 1 and 2 therefore deliver estimates of all of the model's parameters.

[35]The productivity terms A^k_{ot} are unobserved because they represent the location parameter on region o's *potential* productivity distribution of commodity k, in equation (1.2). The productivities actually used for production in region o will be a subset of this potential distribution, where the scope for trade endogenously determines how the potential distribution differs from the distribution actually used to produce.

In Step 3, I test Prediction 3 by estimating how the level of a district's real income is affected when the railroad network is extended to that district, and when it is instead extended to other nearby district. However, the empirical findings in Step 3 is reduced-form in nature and could arise through a number of possible mechanisms.[36] Therefore, in Step 4 I use the sufficient statistic suggested by Prediction 4 to compare the reduced-form effects of railroads on the level of real income (found in Step 3) with the effects predicted by the model (as estimated in Steps 1 and 2).

1.4 Empirical Step 1: Railroads and Trade Costs

In the first step of my empirical analysis I estimate the extent to which railroads reduced the cost of trading within India. Because this chapter stresses a trade-based mechanism for the impact of railroads on welfare, it is important to establish that railroads actually reduced trade costs. Further, the relationship between railroads and trade costs, which I estimate in this section, is an important input for the three empirical steps that follow this one.

[36]For example, railroads could have: reduced the cost of technology transfer between regions, or the monitoring costs of multi-regional enterprises (as in the model of Ramondo and Rodriguez-Clare (2008), who construct a model similar to that here in which there is diffusion of technology and multinational production); encouraged factor mobility, potentially giving rise to efficiency gains if factors are heterogeneous, or increasing the elasticity of labor supply (as Jayachandran (2006) found in post-Independence India); increased the size of the market, encouraging innovation (as Sokoloff (1988) found in the case of US canals and patenting behavior) or allowing economies of scale to be exploited (as found in Ades and Glaeser (1999)); or even altered the political environment in favor of a commercial class that favored growth-enhancing institutions (as Acemoglu, Johnson, and Robinson (2005) argue explained the growth of European port cities with access to the new trade opportunity of Atlantic trade post-1500).

1.4.1 Empirical Strategy

Researchers rarely observe trade costs.[37] But Prediction 1 suggests an instance under which trade costs can be *inferred*: If a homogeneous commodity can only be made in one region, then the difference in retail prices (of that commodity) between the origin region and any other consuming region is equal to the cost of trading between the two regions.[38]

Throughout Northern India, several homogeneous types of salt were consumed, but each of these varieties could only be made in one unique location. Traders and consumers would speak about 'Kohat salt' (which could only be produced at the salt mine in the Kohat region) as a different commodity from 'Sambhar salt' (which could only be produced at the Sambhar Salt Lake).[39] I have collected data on salt prices in Northern India, in which the prices of eight regionally-differentiated types of salt are reported in 124 districts. Crucially, because salt is an essential commodity, it was consumed throughout India both before and after the construction of railroads.

I use this salt price data, with the help of Prediction 1, to estimate how Indian railroads reduced trade costs. To do this I estimate equation

[37] Even when shipping receipts are observed, as in Hummels (2007), these may fail to capture other barriers to trade, such as the time goods spend in transit (a focus of Evans and Harrigan (2005)), or the risk of damage or loss in transit (a major concern in colonial India). In lieu of direct measures of trade costs, a large literature, surveyed by Anderson and van Wincoop (2004), uses a proxy variable strategy (similar to that I employ in this section) to estimate trade costs.

[38] Anderson and van Wincoop (2004) suggest (on p. 78) the solution I pursue here: "A natural strategy would be to identify the source [region] for each product. We are not aware of any papers that have attempted to measure trade barriers this way."

[39] The leading (nine-volume) commercial dictionary in colonial India, Watt (1889), describes the market for salt in this manner, as do Aggarwal (1937) and the numerous provincial *Salt Reports* that were brought out each year.

(1.7) of Prediction 1 as follows:[40]

$$\ln p^o_{dt} = \underbrace{\beta^o_{ot}}_{=\ln p^o_{ot}} + \underbrace{\beta^o_{od} + \phi^o_{od}t + \delta \ln TC(\mathbf{R}_t)_{odt} + \varepsilon^o_{odt}}_{=\ln T^o_{odt}}. \tag{1.11}$$

In this equation, p^o_{dt} is the price of type-o salt (that is, salt that can only be made in region o) in destination district d in year t. I estimate this equation with an origin-year fixed effect[41] (β^o_{ot}) to control for the price of type-o salt in its region of origin o (ie p^o_{ot}) because I do not observe salt prices exactly at the point where they leave the source. (My price data is at the district level and was recorded as the average price of goods over 10-15 retail markets in a district.)

The remainder of equation (1.11) describes how I model the relationship between trade costs T^o_{odt}, which are unobservable, and the railroad network (denoted by $\mathbf{R}_t$), which is observable.[42] I use two different proxy variables, denoted by $TC(\mathbf{R}_t)_{odt}$ and explained in detail below, that relate trade costs to the railroad network. This specification includes an origin-destination fixed effect (β^o_{od}) which controls for all of the time-invariant

[40]The model in section 1.3 underpinning prediction 1 assumes that trade costs take an *ad valorem* (that is, per unit *value*) form, which is inconsistent with the evidence in Hummels and Skiba (2004). To test for a non-*ad valorem* trade cost specification I have estimated equation 1.11 with an additional interaction term between $\ln TC(\mathbf{R}_t)$ and the level of an excise tax charged on salt as it left the 'factory' gate. This was a very high tax (in the range of 100-300 percent of the value of salt), that initially varied across provinces, but which fell precipitously in 1874, 1878 and 1883 so that all provinces had the same tax rate (I take the data on excise rates from Aggarwal (1937)). However, the coefficient on this interaction term is never statistically significant. This is consistent with my assumption that, regardless of the factory gate price of salt, trade costs took a form that was proportional to the price of the commodity shipped.

[41]That is, each salt origin o has its own fixed effect in each year t. I use this notation when referring to fixed effects throughout this chapter.

[42]An alternative empirical strategy would be to estimate trade costs directly from equation (1.11), by simply equating trade costs to observed price differences. However, this method faces two drawbacks when compared to the method I follow here. First, it would only uncover trade cost estimates for the *o-d-t* observations for which salt prices are observed separately for each region of origin (that is, no estimates would be available in Southern India). Second, it would be vulnerable to the concern that p^o_{ot} is not measured exactly at the point where type-o salt leaves the source in region o.

determinants of the cost of trading salt between districts o and d (such as the distance from o to d, or caste-based or ethno-linguistic differences between o and d that may hinder trade).[43] The specification also includes a separate trend term ($\phi^o_{od}t$) for each origin-destination pair; these trend terms control for any trade costs between o and d that vary over time in a (log) constant way. Finally, ε^o_{odt} is an error term that captures any remaining unobserved determinants of trade costs (or measurement error in $\ln p^o_{dt}$).[44]

I use two different measures for $TC(\mathbf{R}_t)_{odt}$, the proxy for the unobservable trade costs between the origin o and destination d districts in any year t:

1. *Bilateral railroad dummy variable:* I denote this variable by $RAIL_{odt}$. This dummy variable is equal to one in all years when it is possible to travel from district o to district d by railroad (and zero otherwise).[45] This proxy variable has the advantage of simplicity. But the coefficient on this variable is likely to be biased downwards because the spread of the railroad network will potentially reduce trade costs between o and d in years before they are fully connected by a railroad line. Furthermore, this variable ignores any heterogeneity in the effect of railroads on trade costs between two districts, such as that due to the distance between them, the directness of their railroad connection, or the local non-railroad transportation system.

[43] Rauch (1999) and Anderson and van Wincoop (2004) document a series of findings that are consistent with large communication-based barriers to trade in contemporary international trade data.

[44] In this specification and all others in this chapter I allow this error term to be heteroskedastic and serially correlated within districts (or trade blocks, in section 1.5) in an unspecified manner.

[45] Andrabi and Kuehlwein (2005) and Keller and Shiue (2007b) use this dummy variable approach when studying the effect of railroads on price differences in 19$^{\text{th}}$ Century India and Europe, respectively.

2. *Lowest-cost route distance:* I denote this variable by $LCR(\mathbf{R}_t, \boldsymbol{\alpha})_{odt}$. This measure models the cost of trading goods between any two locations under the assumption that agents take the lowest-cost route—using any mode of transportation—available to them.[46] Two inputs are needed to calculate the lowest-cost route between districts o and d in year t. The first input is the *network* of available transportation routes open in year t, which I denote by $\mathbf{R}_t$. A network is a collection of nodes and arcs. In my application, nodes are finely-spaced points in space, and arcs are available means of transportation between the nodes (hence an arc could be a rail, river, road or coast connection). In modeling this network I allow agents to travel on navigable rivers, the coastline, the road network (which I take to be continuous over space and hence connecting any two nodes along the straight line between them), and the railroad network open in year t. The second input is the relative cost of traveling along each arc, which depends on which mode of transportation the arc represents. I model these costs as being proportional to distance,[47] where the proportionality, the *relative per unit distance cost*, of using each mode is denoted by the vector of parameters $\boldsymbol{\alpha} \doteq (\alpha^{rail}, \alpha^{road}, \alpha^{river}, \alpha^{coast})$. I normalize $\alpha^{rail} = 1$ so the other three elements of $\boldsymbol{\alpha}$ are costs relative to the cost of

[46]To the best of my knowledge, this is a new method for measuring trade costs over multiple modes of transportation over a network, where users are free to choose their route over the network. Houde (2008) is related, as he uses Dijkstra's algorithm to find the likely commuting paths of automobile drivers over a road network (in order to define retail gasoline markets). But unlike my procedure, he treats the parameters that govern these users' path choices as known.

[47]This rules out any fixed costs of switching modes of transportation (such as handling charges), or other economies of scale in the transportation sector either internal to trading firms or external to them (such congestion effects). It is difficult to know whether these features were applicable to non-rail transportation in colonial India, but the simple freight rates charged by the railroads did not feature either a fixed handling charge or a bulk discount for large shipments, the trading sector was characterized by a large mass of small-scale traders (Bayly 1983, Yang 1999), and congestion effects on the railroads were rarely a deterrent to trade (Sanyal 1930).

using railroads. Because of this normalization, $LCR(\mathbf{R}_t, \boldsymbol{\alpha})_{odt}$ is measured in units of railroad-equivalent kilometers; in this sense, a finding that all of the non-rail elements of $\boldsymbol{\alpha}$ are greater than one implies that railroads effectively shrunk distance, as measured in railroad-equivalent units. The parameter $\boldsymbol{\alpha}$ is unknown, so I treat it as a vector of parameters to be estimated.[48] Conditional on a value of $\boldsymbol{\alpha}$, it is possible to calculate $LCR(\mathbf{R}_t, \boldsymbol{\alpha})_{odt}$—a calculation that is made computationally feasible by Dijkstra's shortest-path algorithm (Ahuja, Magnanti, and Orlin 1993). But since $\boldsymbol{\alpha}$ is unknown, I estimate it using non-linear least squares (NLS). That is, I search over all values of $\boldsymbol{\alpha}$, recomputing the lowest-cost routes at each step, to find the value that minimizes the sum of squared residuals in equation (1.11).

1.4.2 Data

I use data on retail prices of 8 types of salt, observed annually from 1861-1930 in 124 districts of Northern India (the region in which salt prices were reported broken down by region of origin). Further details on the data I use in this and other sections of this thesis are provided in Appendix A.

1.4.3 Results

Column 1 of Table 1.2 presents results from estimating equation (1.11) by OLS using the bilateral railroad dummy ($RAIL_{odt}$) as the proxy for trade

[48] As discussed in section 1.2, relative freight rates for each mode of transport are available from a number of historical sources. However, as with overall trade costs, the cost of using railroads relative to another mode of transport may include elements (such as increased certainty or time savings) that are not included in observed freight rates. Consistent with this idea, I find in table 1.2 that my estimates of $\boldsymbol{\alpha}$ are higher (for road, river and coastal travel relative to rail travel), than the relative freight rates observed in these historical sources.

costs. The coefficient on this proxy variable is negative and statistically significant, indicating that when two regions are connected by a railroad line the cost of trading between them falls by approximately 10 percent. However, as argued above, this measure is likely to be biased toward zero and to ignore significant heterogeneity in the effect of railroads on trade costs.

To address these shortcomings of the bilateral railroad dummy variable, columns 2-5 present estimates of equation (1.11) using my alternative proxy variable for trade costs, the lowest-cost route $(LCR(\mathbf{R}_t, \boldsymbol{\alpha})_{odt})$. In column 2 I estimate the effect of the lowest-cost route distance on trade costs when the relative costs of each mode $(\boldsymbol{\alpha})$ are set to observed historical relative freight rates (in 1900). I use the relative per unit distance freight rates described in section 1.2 (at their midpoints): $\alpha^{road} = 4.5$, $\alpha^{river} = 3.0$, and $\alpha^{coast} = 2.25$ (all relative to the freight rate of railroad transport, normalized to 1). Column 2 demonstrates that the elasticity of trade costs with respect to the lowest-cost route distance, calculated at observed freight rates, is 0.135.

However, as argued above, it is possible that these observed relative freight rates do not capture the full benefits of railroad transport relative to alternative modes of transportation. For this reason the NLS specifications in columns 3-5 estimate the relative freight rates (ie the parameters $\boldsymbol{\alpha}$) that minimize the sum of squared residuals in equation (1.11). In column 3 I estimate equation (1.11) without district-source specific trends included and find an elasticity of salt prices with respect to the lowest-cost route distance 0.255.

A potential concern with the specification in column 3 is that the lowest-cost route distance may be trending over time because an approaching railroad line will steadily reduce trade costs. Unobserved de-

Dependent variable: Log salt price at destination	(1)	(2)	(3)	(4)	(5)
Source connected to destination by railroad	-0.112 (0.046)***				0.023 (0.342)
Log distance to source, along lowest-cost route (at historical freight rates)		0.135 (0.038)***			
Log distance to source, along lowest-cost route (at estimated mode costs)			0.255 (0.059)***	0.247 (0.063)***	0.233 (0.074)***
Estimated mode costs:					
Railroad (normalised to 1)			1 N/A	1 N/A	1 N/A
Road			7.341 (1.687)***	7.880 (1.913)***	7.711 (2.032)***
River			3.396 (0.760)***	3.821 (1.034)***	4.118 (1.381)***
Coast			3.213 (1.893)*	3.942 (2.581)	3.612 (2.674)
Salt type x Year fixed effects	YES	YES	YES	YES	YES
Salt type x Destination district fixed effects	YES	YES	YES	YES	YES
Salt type x Destination district trends	YES	YES	NO	YES	YES
Observations	7329	7329	7329	7329	7329
R-squared	0.841	0.960	0.953	0.974	0.975

Table 1.2: Railroads and Trade Costs (Step 1) Notes: Regressions estimating equation (1.11) using data on 8 types of salt (listed in Appendix A), from 124 districts in 5 Northern Indian provinces (listed in Appendix A), annually from 1861 to 1930. Columns 1 and 2 are OLS regressions, and columns 3-5 are NLS regressions (with block-bootstrapped standard errors). 'Source connected to destination by railroad' is a dummy variable equal to one in all years when it is possible to travel by railroad from any point in the district containing the salt source to any point in the destination district. The 'distance to source, along lowest-cost route' variable is a measure of the railroad-equivalent kilometers (because railroad freight rate is normalized to 1) between the salt source and the destination district, along the lowest-cost route given relative mode costs per unit distance. 'Historical freight rates' used are 4.5, 3.0 and 2.25 respectively for road, river and coastal mode costs per unit distance, all relative to rail transport. Data sources and construction are described in full in Appendix A. Heteroskedasticity-robust standard errors corrected for clustering at the destination district level are reported in parentheses. *** indicates statistically significantly different from zero at the 1 % level; ** indicates 5 % level; and * indicates 10 % level.

terminants of trade costs (such as a steadily improving institutional environment conducive to interregional trade) may also be trending over time, and there is a risk that these unobserved determinants may be attributed to the railroad network. I therefore allow each district-salt source pair to have its own (log) linear trend term in column 4. This reduces the coefficient on the lowest-cost route measure by a small amount, but this coefficient is still economically and statistically significant. Column 4 is my preferred specification. Even when controlling for all unobserved, time-constant and trending determinants of trade costs between all salt sources and destinations, reductions in trade costs along lowest-cost routes (estimated from time variation in these routes alone) have a large effect on salt prices.

The non-linear specification in column 4 also estimates the relative trade costs by mode that best explain observed salt price differentials. Each of the three alternative modes of transport is larger than one, implying that these alternative modes are more expensive (per unit distance) than rail travel. Further, each of these non-rail modes has higher estimated costs, relative to railroads, than historically observed freight rates. This suggests that the advantages of railroads to encouraging trade were significant, but not entirely reflected in observed freight rates.

To summarize the results in column 4: the coefficient on the lowest-cost route distance ($\widehat{\delta}$) is positive, which implies that trade costs increase with effective distance (in railroad-equivalent kilometers); and the estimated mode-specific per-unit distance costs ($\widehat{\alpha}$) are all much greater than one, implying that railroads were instrumental in reducing effective distance when compared to alternative modes of transportation (especially when compared to roads, which I find were almost eight times more costly to use per unit distance than railroads).

Finally, in column 5 I include both the bilateral railroad dummy variable, and the lowest-cost route variable. The bilateral railroad dummy is no longer statistically significantly different from zero, and its point estimate is much smaller than in column 1. By contrast, the lowest-cost route trade costs variable is still large and statistically significant, and its magnitude is similar to that in column 4. This suggests that the lowest-cost route measure is explaining genuine features of the railroad network as it impacted on salt prices. I use the estimates in column 4, my preferred specification, in the next stages of my empirical strategy.

1.5 Empirical Step 2: Railroads and Trade Flows

The first step of my empirical strategy demonstrated that railroads reduced trade costs. I now estimate the extent to which the reduction in trade costs brought about by India's railroad network (estimated in Step 1) affected trade flows within India, and trade flows between India and its international trade partners. This step is important for two reasons. First, an expansion of trade volumes as a result of the railroad network is a necessary condition for the mechanism linking railroads to welfare gains in the model. Second, as I show below, estimating the model's gravity equation allows all of the model's parameters to be inferred. Equipped with these parameter estimates I am better able to test prediction 4 below.

1.5.1 Empirical Strategy

Prediction 2 of the model suggests a particular relationship between bilateral trade flows and bilateral trade costs—a gravity equation describing trade between any two regions. Substituting the empirical specification for T^k_{odt} introduced in equation (1.11) into equation (1.8) yields

$$\begin{aligned}\ln X^k_{odt} = \beta^k_{od} + \phi^k_{od}t &+ \ln A^k_{ot} - \theta_k \ln r_{ot} - \theta_k \delta \ln TC(\mathbf{R}_t)_{odt} \\ &+ \theta_k \ln p^k_{dt} + \ln X^k_{dt} + \varepsilon^k_{odt}. \qquad (1.12)\end{aligned}$$

Here, X^k_{odt} refers to the value of exports of commodity k from region o to region d in year t (and the other variables were defined in section 1.3.

I estimate two versions of this bilateral exports equation, each with a different goal in mind. The *first version* investigates whether the construction of India's railroad network increased trade in India. To do this I estimate the equation

$$\ln X^k_{odt} = \beta^k_{ot} + \beta^k_{dt} + \beta^k_{od} + \phi^k_{od}t + \rho \ln TC(\mathbf{R}_t)_{odt} + \varepsilon^k_{odt}. \qquad (1.13)$$

In this specification, the term β^k_{ot} is an origin-year-commodity fixed effect and β^k_{dt} is a destination-year-commodity fixed effect (the inclusion of these two fixed-effects are suggested by the model in equation (1.12)); β^k_{od} is an origin-destination-commodity fixed effect and the term $\phi^k_{od}t$ allows for each origin-destination-commodity to have its own trend term (these two terms were motivated in section 1.4 by the concern that some costs of trading may be unobservable). The coefficient ρ on the trade costs proxy variable $\ln TC(\mathbf{R}_t)_{odt}$ is therefore estimated purely from time variation in the railroad network that affects an exporter differently across its trading destinations. Prediction 2 is that the coefficient ρ will be negative—that, conditional on importer and exporter fixed effects (for each commodity), the lower trade costs brought about by railroads increase trade.[49]

[49]The coefficient ρ is not the general equilibrium elasticity of trade flows with respect to trade costs. This is because equation (1.12) also contains the endogenous variables, land rental rates r_{ot}, goods prices p^k_{dt}, and aggregate expenditure X^k_{dt}. I

In following sections of this chapter I assume that the trade cost parameters for salt (which I estimated in Step 1) apply to other commodities as well. One potential concern with this assumption is that the parameter δ, which relates the lowest-cost route variable ($LCR(\mathbf{R}_t, \boldsymbol{\alpha})_{odt}$) to trade costs, may vary across commodities. To explore this possibility while estimating equation (1.13) (which is pooled across commodities), I use $LCR(\mathbf{R}_t, \boldsymbol{\alpha})_{odt}$ as a proxy for the trade cost term $TC(\mathbf{R}_t)_{odt}$ and also include interaction terms between $LCR(\mathbf{R}_t, \boldsymbol{\alpha})_{odt}$ and commodity-specific characteristics.[50] The characteristics I include are weight per unit value and the 'freight class' in which each commodity was placed by railroad companies—both measured at the start of the period. If the effect of $LCR(\mathbf{R}_t, \boldsymbol{\alpha})_{odt}$ on trade flows does not vary across commodities according to these characteristics (which capture the most obvious reasons why per unit distance trade costs could differ by commodity) then this would be consistent with the assumption that trade costs for salt are representative of trade costs for other commodities. A second concern with this assumption is that the relative per unit distance cost of using each mode of transportation ($\boldsymbol{\alpha}$) may also vary across commodities, so that my parameter estimates of $\boldsymbol{\alpha}$, also obtained from salt, do not carry over to other commodities. I discuss evidence that is inconsistent with this second concern below.

control for these endogenous variables (by the use of fixed effects) in my estimating equation (1.13), so they do not present a risk of bias to the estimation of ρ; however, comparative statics exercises (such as the general equilibrium elasticity of trade costs with respect to trade flows) must allow these endogenous variables to adjust. Following Anderson and van Wincoop (2003), the coefficient ρ is best thought of as a partial equilibrium elasticity, which holds constant all of these necessary adjustments in other goods markets and the factor market. However, as these authors note, the sign of the true (general equilibrium) elasticity will be of the same sign as the partial equilibrium elasticity estimated here.

[50]In this specification $\ln LCR(\widehat{\boldsymbol{\alpha}}, \mathbf{R}_t)_{odt}$ is a generated regressor because the parameter $\boldsymbol{\alpha}$ was estimated in Step 1. I therefore correct the standard errors in this regression to account for the presence of a generated regressor using a two-step bootstrap procedure.

The *second version* of equation (1.12) that I estimate takes the model more seriously in order to estimate unknown parameters of the model. Two sets of parameters remain unknown: the technology parameter θ_k (for each commodity), and the productivity parameters A^k_{ot} (for each district, year and commodity). I first estimate the parameters θ_k by substituting the lowest-cost route distance proxy for trade costs estimated in Step 1 (ie $\widehat{\delta}\ln LCR(\widehat{\boldsymbol{\alpha}}, \mathbf{R}_t)_{odt}$, where $\widehat{\delta}$ and $\widehat{\boldsymbol{\alpha}}$ were estimated in Step 1) into equation (1.12) to obtain

$$\ln X^k_{odt} = \beta^k_{ot} + \beta^k_{dt} + \beta^k_{od} + \phi^k_{od}t - \theta_k\widehat{\delta}\ln LCR(\mathbf{R}_t, \widehat{\boldsymbol{\alpha}})_{odt} + \varepsilon^k_{odt}. \quad (1.14)$$

In this specification, the coefficient on the lowest-cost route distance variable ($\hat{\delta}\ln LCR(\widehat{\boldsymbol{\alpha}}, \mathbf{R}_t)_{odt}$) is exactly θ_k. Intuitively, the scope for comparative advantage (the inverse of θ_k) governs how much a reduction in trade costs translates into an expansion of trade. I therefore estimate this equation separately for each of the 85 commodities in my trade flows dataset, in order to estimate 85 values of θ_k (one for each commodity k).

Armed with the parameter estimates $\widehat{\theta}_k$ it is then possible to estimate the other unknown variable in the model, the unobserved productivity term, A^k_{ot} (though this is only possible for agricultural commodities). I relate A^k_{ot} to observables by assuming that A^k_{ot} is a function of a *crop-specific rainfall shock*, denoted by $RAIN^k_{ot}$. As argued in section 1.2, rainfall was an important determinant of agricultural productivity in India because irrigation was uncommon. However, a given distribution of annual rainfall would affect each crop differently because each crop has its own annual timetable for sowing, growing and harvesting, and these timetables differ from district to district. To shed light on these crop- and district-specific agricultural timetables, I draw on the 1967 publication, the *Indian Crop Calendar* (Directorate of Economics and Statistics 1967), which lists sowing, growing and harvesting windows for each crop and district in my sample.[51] To construct the variable $RAIN_{ot}^{k}$, I use daily rainfall data to calculate the amount of rainfall in year t that fell between the first sowing date and the last harvest date listed for crop k in district o.

It is then possible to estimate the relationship between rainfall and productivity by noting that the exporter-commodity-year fixed effect (β_{ot}^{k}) in equation (1.14) can be interpreted in the model as $\beta_{ot}^{k} = \ln A_{ot}^{k} - \theta_k \ln r_{ot}$, by comparing equations (1.12) and (1.14).[52] I model the relationship between productivity (A_{ot}^{k}) and rainfall ($RAIN_{ot}^{k}$) in the simplest possible semi-log manner: $\ln A_{ot}^{k} = \kappa RAIN_{ot}^{k}$.[53] Guided by this relationship, I estimate the parameter κ in the following estimating equation:

[51]This publication describes the technology of agricultural practice (related to scheduling of activities) in each district. This particular aspect of agricultural technology is unlikely to have changed between my sample period and 1967 because the optimal sowing date for a crop depends on the amount of water in the soil on that date, which is governed by the type of soil and the local climate (Mukerji 1915), both of which are unlikely to change over this time frame. Nevertheless, to test for this stability of agricultural scheduling I use the earliest *Crop Calendar* that I was able to access, which is from 1908. This volume presents data at larger geographic areas than the district. However, when I calculate crop-specific sowing and growing rainfall amounts for these larger geographic areas the correlation between these and the (area-weighted average) 1967 district-level amounts is 0.78. This is a strong correlation, especially given that the district-level data is aggregated, indicating that the timing of agricultural activities has not changed dramatically (from 1908 to 1967 at least).

[52]Intuitively, higher productivity in commodity k (A_{ot}^{k}) will increase region o's propensity to export to any location, leading to a higher exporter fixed effect β_{ot}^{k}; however, higher productivity will also raise the land rental rate (r_{ot}), decreasing the propensity to export, and a lower exporter fixed effect.

[53]It is also possible that the marginal productivity of rainfall is diminishing, becoming detrimental to production at some point. With such effects in mind I have also estimated a specification where I include both $RAIN_{ot}^{k}$ and $(RAIN_{ot}^{k})^2$. However, in this alternative specification, the coefficient on the squared amount of rainfall is actually positive (implying increasing marginal productivity of rainfall), but never statistically significant. This finding appears to contradict the findings of a statistically significant diminishing effect from rainfall found in Chapter 3 (using post-1956 data). An important distinction between the results here and those in Chapter 3, however, is that here I use crop-specific rainfall variation while in Chapter 3 I pool across crops. It is possible that the diminishing effect of rainfall found on the pooled regressions reflects an aggregation issue rather than a true technological relationship.

$$\widehat{\beta}_{ot}^{k} + \widehat{\theta}_k \ln r_{ot} = \beta_o^k + \beta_t^k + \beta_{ot} + \kappa RAIN_{ot}^k + \varepsilon_{odt}^k. \qquad (1.15)$$

In this equation, $\widehat{\beta}_{ot}^{k}$ is the estimated exporter-commodity-year fixed effect, and $\widehat{\theta}_k$ is the estimated technology parameter, both of which are estimated in equation (1.14) above. The terms β_o^k, β_t^k, and β_{ot} represent exporter-commodity, commodity-year and exporter-year fixed effects, respectively. I include these terms to control for unobserved determinants of exporting success that do not vary across regions, commodities *and* time. For example, the exporter-commodity fixed effect (β_o^k) controls for all time-invariant factors that make region o successful at exporting commodity k (such as the region's altitude). As a result, the coefficient κ is estimated purely from the variation in rainfall over space, commodities and time. The final term in equation (1.15) is an error term (ε_{odt}^k) that includes any determinants of exporting success, other than rainfall, that vary across regions, commodities and time.

In summary, the method described in this second version of estimating equation (1.12) estimates the parameter θ_k for each of the 85 goods k for which I have trade data. This method also estimates the relationship between the unobserved productivity terms A_{ot}^k and crop-specific rainfall $RAIN_{ot}^k$ (governed by the parameter κ).

1.5.2 Data

I estimate equations (1.13), (1.14) and (1.15) using over six million observations on Indian trade flows that I have collected. The trade flow data relate to both internal trade data (between 45 trade blocks of India) and external trade data (between each of these 45 internal trade blocks and 23 foreign countries), over rail, river and coastal transport routes, for

85 commodities, annually from 1880 to 1920. When estimating equation (1.15), I use the crop-specific rainfall measure ($RAIN^k_{ot}$) described briefly above (and in more detail in Appendix A) and, lacking reliable data on land rental rates, I use nominal agricultural GDP per acre as a measure of r_{ot} (since in my model these two measures are equivalent).

1.5.3 Results

Table 1.3 presents results from this section. Column 1 contains estimates of equation (1.13) using OLS, where the trade costs proxy used is the bilateral railroad dummy variable.[54] The coefficient on the railroad dummy variable is positive and statistically significant—even though, as argued in section 1.4, the coefficient on this trade costs proxy is likely to be biased downwards. very precisely estimated. This suggests that railroads significantly boosted trade, and provides support for prediction 2.

In column 2 of table 1.3 I estimate equation (1.13) again, this time with the lowest-cost route variable used to proxy for trade costs instead of the railroad dummy variable. The lowest-cost route distance proxy depends on the unknown parameters $\boldsymbol{\alpha}$, the per unit distance trade costs along each mode of transportation, relative to rail transport. In order to compute the lowest-cost route distance in estimating equation (1.13), I use the estimated value of the parameters $\boldsymbol{\alpha}$ presented in column 4 of table 1.2. This requires the maintained assumption that the *relative* cost of transporting any commodity by rail (relative to other modes) is the same as that for salt; that is, the per-unit distance trade cost may differ across commodities, but in a way such that the relative cost of using

[54]Because my trade flow data is available for trade blocks, which are larger than districts, I define the 'bilateral railroad dummy' variable here ($RAIL_{odt}$) as the share of district pairs between trade block o and trade block d that are connected by the railroad network.

Dependent variable: Log value of exports	(1)	(2)	(3)
Fraction of origin and destination districts connected by railroad	1.482 (0.395)***		
Log distance beween origin and destination along lowest-cost route		-1.303 (0.210)***	-1.284 (0.441)***
(Log distance beween origin and destination along lowest-cost route) x (Weight per unit value of commodity in 1880)			-0.054 (0.048)
(Log distance beween origin and destination along lowest-cost route) x (High-value railroad freight class of commodity in 1880)			0.031 (0.056)
Origin trade block x Year x Commodity fixed effects	YES	YES	YES
Destination trade block x Year x Commodity fixed effects	YES	YES	YES
Origin trade block x Destination trade block x Commodity fixed effects	YES	YES	YES
Origin trade block x Destination trade block x Commodity trends	YES	YES	YES
Observations	6,581,327	6,581,327	6,581,327
R-squared	0.943	0.963	0.964

Table 1.3: Railroads and Trade Flows (Step 2) Notes: Regressions estimating equations (1.13) and (1.14), using data on 85 commodities, 45 trade blocks, and 23 foreign countries, annually from 1880 to 1920. 'Fraction of origin and destination districts connected by railroad' is the share of the district pairs between trade block o and traded block d that for which it is possible to travel entirely by railroad from any point in one district to any point in the other district. The 'distance between origin and destination along lowest-cost route' variable is a measure of the railroad-equivalent kilometers (due to the normalized railroad freight rate to 1) between the centroid of the origin and destination trade blocks in question, along the lowest-cost route given relative freight rates for each mode of transport (as estimated in table 1.2). 'Weight per unit value in 1880' is the weight (in maunds) per rupee, as measured by 1880 prices. 'Railroad freight class in 1880' is an indicator variable for all commodities that were classified in the higher (more expensive) freight class in 1880; salt was in the omitted category (low-value commodities). Data sources and construction are described in full in Appendix A. Standard errors are reported in parentheses. In column 1 these are heteroskedasticity robust standard errors adjusted for clustering at the exporting block level. In columns 2-3 these are bootstrapped standard errors (using a two-stage block bootstrap at the exporting block level) to correct for the generated regressor. *** indicates statistically significantly different from zero at the 1 % level; ** indicates 5 % level; and * indicates 10 % level.

non-rail transport relative to rail transport is the same for all commodities.[55] The results in column 2 provide further support for prediction 2, as the lowest-cost route measure is estimated to reduce bilateral trade (conditional on the fixed effects used) with a statistically significant elasticity of (minus) 1.3. This result is in line with a large body of work on estimating gravity equations.[56]

In column 3 of table 1.3 I investigate the possibility that the elasticity of trade flows with respect to lowest-cost route distance routes varies by commodity. I do this by including interaction terms between the lowest-cost route distance variable and two commodity-specific characteristics: weight per unit value (as observed in 1880 prices, averaged over all of India), and 'freight class' (an indicator used by railroad companies in 1880 to distinguish between 'high-value' and 'low-value' goods). The results in column 3 are not supportive of the notion that commodities had trade flow elasticities with respect to trade costs that depend on weight, or freight class; that is, neither of these interaction terms is significantly different from zero (nor are they jointly significantly different from zero). This lends support to the maintained assumption throughout this chapter that trade cost parameters for the shipment of salt can be applied to other

[55]One piece of evidence consistent with this assumption comes from data on district-to-district trade flows (for each of 15 goods, one of which is salt) in Bengal from 1877 to 1881, along each of the three modes of transport available in that area (rail, river and road). I regress log bilateral exports by road relative to exports by rail on exporter-importer-year fixed effects, and a fixed effect for each commodity. The F-test that these commodity-level fixed effects are all equal to each other has a p-value of 0.34, so it cannot be rejected at the 5 percent level. A similar test for a regression with exports by river relative to exports by rail has a p-value of 0.28. These results are consistent with the view that, within an exporter-importer-year cell, goods do not have systematically different trade costs.

[56]Head and Disdier (2008) conduct a meta-study of 103 papers estimating the coefficient on bilateral distance in a gravity equation. They find a mean estimate of this coefficient of 0.9 (with 90 percent of estimates lying between 0.28 and 1.55). That my result is higher than the mean estimate in these 103 papers is unsurprising because they were estimated primarily on post-1960 data. Technological improvements in transportation (ocean shipping and air freight) and telecommunications are likely to have reduced the trade-impeding effects of distance, when compared to railroad transportation and communication in 1880-1920 India.

commodities, without doing injustice to the data.

Finally, I estimate equation (1.14) in a manner that allows me to estimate all of the remaining unknown parameters of the model. I begin by estimating equation (1.14), one commodity at a time (for each of the 85 commodities in the trade flows data), in order to obtain estimates of the comparative advantage parameters θ_k for each commodity. The mean across all of these 85 commodities is 5.2 (with a standard deviation across commodities of 2.1). This is slightly lower than the preferred estimate of 8.28 in Eaton and Kortum (2002) obtained from intra-OECD trade flows in 1995, treating all of the manufacturing sector as one commodity.[57] However, the mean value of θ_k across only the 17 principal agricultural commodities for which I have output and price data (and therefore use heavily below) is 3.8 (with a standard deviation of 1.2). This suggests a greater scope for comparative advantage based gains from trade among agricultural goods than among manufacturing goods, at least in colonial India.

I then estimate equation (1.15) to obtain an estimate of κ, the parameter that relates crop-specific rainfall to (potential) productivity (A^k_{ot} in the model). I estimate a value of 0.441 for κ (with a standard error of 0.082), implying that a one standard deviation (0.605 m) increase in crop-specific rainfall causes a 27 percent increase in agricultural productivity (as defined by A^k_{ot} in the model). This suggests that rainfall has a positive and statistically significant effect on productivity, as expected given the importance of water in crop production and the paucity of irrigated agriculture in colonial India (as discussed in section 1.2).

In summary, the results from this section demonstrate that railroads significantly expanded trade in India. This finding is in line with Pre-

[57] Using two alternative methods Eaton and Kortum (2002) also obtained estimates of 3.60 and 12.86.

diction 2 and suggests that the expansion of trade brought about by the railroad network could have given rise to welfare gains due to increasingly captured gains from trade. A second purpose of this section was to use the empirical relationship between trade costs (estimated in Step 1) and trade flows to estimate the remaining unknown model parameters, θ_k and A^k_{ot}. These parameters are important inputs for Step 4 below.

1.6 Empirical Step 3: Railroads and Real Income Levels

Steps 1 and 2 above have established that Indian railroads significantly reduced trade costs, expanded trade, and reduced price responsiveness. These findings suggest that railroads dramatically changed the trading environment in India. I now go on to investigate the welfare consequences of railroad expansion in India by estimating the effect of railroads on real income levels.

1.6.1 Empirical Strategy

Prediction 3 of the model states that a district's real income will increase when it is connected to the railroad network, but that its real income will fall as one of its neighbors is connected to the railroad network (holding its own access constant). These predictions motivate an estimating equation of the form

$$\ln\left(\frac{r_{ot}}{P_{ot}}\right) = \beta_o + \beta_t + \gamma RAIL_{ot} + \psi(\tfrac{1}{N_o}) \sum_{d \in \mathbf{N}_o} RAIL_{dt} + \varepsilon_{ot}. \quad (1.16)$$

In this estimating equation, $\frac{r_{ot}}{P_{ot}}$ represents real agricultural income per acre in district o and year t. In my model, r_{ot} is the nominal land rental rate, but I have been unable to find systematic data on land rents in this time period. However, in my model, nominal land rents are equal to nominal output per unit area, on which data was collected in the agricultural sector (the dominant sector of the colonial Indian economy), so I use this to measure r_{ot}.[58] The denominator, $\widetilde{P}_{ot}$, is a consumer price index.[59]

The first regressor in equation (1.16) is $RAIL_{ot}$, a dummy variable that is equal to one in all years t in which some part of district o is on the railroad network. The summation term captures the effect of railroads in other, neighboring districts on the level of real income in the district of observation o.[60] Finally, I estimate equation (1.16) using fixed effects at

[58]Real income per acre is equal to welfare (for a representative agent) in my model, but may not be in my empirical setting because output per acre may diverge from output per capita if the population of each district is endogenous, and related to railroad expansion. Population could be endogenous for two reasons. First, fertility and mortality may have been endogenous in the setting of colonial India—in a Malthusian limit, fertility and mortality would adjust to any agricultural productivity improvements (due to railroads) and hold output per capita constant. However, the potential for endogeneous fertility and mortality responses is likely to vary from setting to setting so while an effect of railroads on output per acre is transferable to alternative settings, an effect on output per capita is potentially less so. Second, migration could respond to differential productivity improvements over space. Migration, however, was extremely limited in colonial India when compared to other countries in the same time period (a feature that is still true today, and that Munshi and Rosenzweig (2007) argue is due to informal insurance provided by localized caste networks), and the little migration that occurred was vastly skewed toward women migrating to marry (Davis 1951, Rosenzweig and Stark 1989). Nevertheless, to test the hypothesis that the limited migration was correlated with railroad construction I have collected data on district-to-district bilateral migration as revealed from birthplaces, recorded in the decadal censuses in colonial India. I find (in OLS regressions that control for district pair fixed-effects) that there is no statistically significant net migration into districts receiving a railroad line from neighboring districts left off the railroad network (so migration is unlikely to have been strong enough to act to equalize output per capita), and that bilateral railroad connections between districts do not statistically significantly correlate with bilateral migration between districts (so the railroads do not seem to have facilitated migration). However, I do not observe intra-district migration, which may have been significant.

[59]In the model this price index is given in equation (1.9). However, it would be unsurprising if a price index calculated in the manner suggested by my theory fits my model well. To perform a more powerful test of the model I therefore use a flexible price index (the Fisher ideal price index) of the sort that is commonly used to construct real GDP measures from national income statistics.

[60]As in section 2.4.1, I take the neighborhood of district o (denoted by $\mathbf{N}_o$, and containing N_o districts) to consist of any districts for which any part of the district lies within a 250 km radius of the centroid of district o.

the district (β_o) and year (β_t) levels, so that the effect of railroads is identified entirely from variation within districts over time, after accounting for common macro shocks affecting all districts. The district fixed effect is particularly important because it controls for permanent features of districts that may have made them both agriculturally productive, and attractive places to build railroads.

Prediction 3 states that the coefficient γ on district o's own railway access will be positive, but the coefficient ψ on district o's neighbors' access will be negative. A number of alternative theories (whether stressing trade mechanisms or otherwise) could make similar predictions about the signs of these coefficients (especially about γ). For this reason, in Step 4 I go beyond the qualitative test of my model provided by the signs of γ and ψ and assess the *quantitative* performance of the model in predicting real income changes due to the expansion of the railroad network.

I begin below (in section 1.6.3) by estimating equation (1.16) using OLS. Unbiased OLS estimates require there to be no correlation between the error term (ε_{ot}) and the regressors ($RAIL_{ot}$ and $(\frac{1}{N_o})\sum_{d\in\mathbf{N}_o} RAIL_{dt}$), conditional on the district and year fixed effects. This requirement would fail if railroads were built in districts and years that were expected to experience real agricultural income growth, or if railroads were built in districts that were on differing unobserved trends from non-railroad districts. For this reason I pursue three strategies to assess the potential magnitude of bias in my OLS results due to non-random railroad placement: four placebo specifications (section 1.6.4), instrumental variable estimates (section 1.6.5) and a bounds check (section 1.6.6).

1.6.2 Data

I estimate equation (1.16) using annual data on real agricultural income (per acre of land) in 239 districts, from 1870 to 1930. This variable (calculated as nominal agricultural GDP calculated from 17 crops, deflated by a consumer price index and then divided by the district's land area) was described briefly in section 1.2 and in more detail in Appendix A. The variables $RAIL_{ot}$ and $RAIL_{dt}$ were described in section 1.5.

1.6.3 Results

Column 1 of table 1.4 presents OLS estimates of equation (1.16), with only the own-district regressor included. The coefficient estimate is 0.164, implying that in the average district, the arrival of the railroad network raised real agricultural income by over sixteen percent.

Prediction 3, along with the customs union literature in international trade theory, predicts that a district can suffer from trade diversion when one or more of its trade partners gains improved access to a third region's market. Because the arrival of the railroad network is spatially correlated, the specification in column 1 may confound the positive effects of a district's own access to the railroad network with the negative effect of access by its neighbors. Column 2 of table 1.4 checks for this negative effect of railroads by including as an additional regressor the extent to which a district's neighbors are connected to the railroad network (as in equation (1.16)). The coefficient on this additional variable is negative and statistically significant, indicating that losses from trade diversion are at work when a district's trading partners reduce trade costs between them but not the district of observation. In addition, the coefficient on own-district railroad access is higher in column 2 than in column 1—the point estimate (which is still statistically significant) now

Dependent variable: Log real agricultural income per acre	(1)	(2)
Railroad in district	0.164	0.182
	(0.056)***	(0.071)**
Railroad in neighboring district		-0.042
		(0.020)**
District fixed effects	YES	YES
Year fixed effects	YES	YES
Observations	14,340	14,340
R-squared	0.744	0.758

Table 1.4: Railroads and Real Income Levels (Step 3) Notes: OLS Regressions estimating equation (1.16) using real income constructed from crop-level data on 17 principal agricultural crops (listed in Appendix A), from 239 districts in India, annually from 1870 to 1930. 'Railroad in district' is a dummy variable whose value is one if any part of the district in question is penetrated by a railroad line. 'Railroad in neighboring districts' is the variable 'railroad in district' averaged over all districts within a 250 km radius of the district of observation. Data sources and construction are described in full in Appendix A. Heteroskedasticity-robust standard errors corrected for clustering at the district level are reported in parentheses. *** indicates statistically significantly different from zero at the 1 % level; ** indicates 5 % level; and * indicates 10 % level.

suggests that railroad access increases real agricultural income per acre by over 18 percent.

The results from including, in column 2, neighbors' railroad access highlight that railroad projects had a treatment externality on untreated districts. It is important to control for this treatment externality to prevent bias in estimates of the effect of railroad access. This result also highlights potential distributional consequences of railroad construction, whereby a project that is good for one region may be bad for its neighbors.[61]

The OLS results described here are in line with Prediction 3 and suggest that railroads has a large effect on real income in India. In the following three sections I pursue three strategies to explore the robustness of these findings to concerns over the non-random placement of railroads.

1.6.4 Four Placebo Specifications

Empirical Strategy:

The first strategy I use to mitigate concerns of bias due to non-random railroad placement is to estimate the effects of 'placebo' railroad lines: over 40,000 km of railroad lines that went through various stages of the planning process, but were never actually built.[62] I group these placebo

[61]My estimates suggest that the districts where railroads are built could feasibly make transfers to their neighbors that would compensate neighbors while leaving the constructing district better off. Unfortunately, I have been unable to find any data that would shed light on the extent of such transfers. Duflo and Pande (2007) use data from a household consumption survey to measure the effect of dam construction on consumption and poverty; the use of consumption data, which was recorded after any potential compensating transfers, allows these authors to argue that compensation of districts harmed by dam construction appears to be incomplete (and especially incomplete in districts with a history of relatively more extractive institutions). Unfortunately, such household consumption surveys only began in India in 1950.

[62]This strategy is similar in spirit to that in Greenstone and Moretti (2004), who study the welfare impact of large industrial plants in the United States. They compare economic outcomes in the counties where these plants were built to outcomes in the plant's second-choice county (where the plant was not built).

lines into four categories:

1. *Four-stage planning hierarchy:* From 1870-1947, India's Railways Department used one constant system for the evaluation of new railroad projects.[63] Line proposals received from the Indian and provincial governments would appear as *proposed* in the Department's annual *Railway Report.* This invited further discussion, and if the proposed line survived this criticism it would be *reconnoitered.* Providing this reconnaissance uncovered no major problems, every meter of the proposed line would then be *surveyed*, this time in painstaking and costly detail (usually taking several years to complete).[64] These detailed surveys would provide accurate estimates of expected construction costs, and lines whose surveys revealed modest costs would then be passed on to the Government to be *sanctioned*, or given final approval. The railroad planning process was therefore arranged as a four-stage hierarchy of tests that proposed lines had to pass.[65] I estimate equation (1.16), but additionally include regressors for railroad lines abandoned at each of these four planning stages (with separate coefficients on each). If line placement decisions were driven by unobservable determinants of agricultural income, it is likely that unbuilt lines would exhibit spurious effects (relative to the excluded category, areas in which lines were never even discussed) on agricultural income in OLS regres-

[63]Strachey and Strachey (1882) review the early history of the Railways Department (part of the Department of Public Works until 1878). The Railway Department's annual *Railway Reports* describe the planning system in each year through to 1947.

[64]Reconnaissance was a form of low-cost survey of possible track locations (typically within 100 m of their eventual location), along with a statement of all necessary bridges, tunnels, cuttings and embankments. As Davidson (1868) and Wellington (1877) make clear, surveying was much more detailed, as its end goal was to identify the exact position of the intended lines, and a precise statement of all engineering works (down to the number of bricks required to build each bridge).

[65]This process of sequentially more detailed investigation is echoed in Wellington (1877), the standard textbook for railroad engineers and surveyors, in all English-speaking countries, in its day (which ran to six editions by 1906).

sions. Further, it is likely that the lines that reached later planning stages would exhibit larger spurious effects than the lines abandoned early on (because higher expected benefits would be required to justify the increasingly costly survey process). The absence of such a pattern would cast doubt on the extent to which India's Railways Department was selecting districts for railroad projects on the basis of correlation with the error term in equation (1.16).

2. *Lawrence's proposal:* In 1868, Viceroy John Lawrence (head of the Government of India) proposed and had surveyed a 30-year expansion plan, broken into 5-year segments, that would begin where Dalhousie's trunk lines (described in section 1.2.3) left off.[66] Lawrence consulted widely about the optimal routes for this railroad expansion, and drew upon his twenty-six years of experience as an administrator in India. Upon his retirement (from his fixed, five-year term) in 1869, construction on Lawrence's plan had just begun. But Lawrence's successor, the Earl of Mayo, immediately halted construction and vetoed Lawrence's proposal. Mayo was an outsider (who had never been to India before his appointment) and a fiscal conservative, and he wasted no time in criticizing the high costs of railroad construction in India. Instead, Mayo followed a more cautious approach to railroad expansion and Lawrence's plan was never built. However, Lawrence's plan provides a useful window on the trajectory that he and his Government expected in the districts where they planned to expand the railroad network. If anyone was capable of forecasting developments in each district's trading environment, developments that may be correlated with the error term

[66] These segments appear in the plan (published in 1868) as "to be built over the next 5 years", "to be built between 6 and 10 years from now", etc.

in equation (1.16), it was likely to be Lawrence. To check for this, I estimate equation (1.16) and additionally include lines that were part of Lawrence's proposal. Because Lawrence's proposal was broken into six, five-year segments, I allow for separate coefficients on each of these segments and assume that the stated lines in a given five-year period would have opened at the beginning of the period. This provides an additional check: lines that Lawrence proposed to be built in relatively early time segments were presumably more attractive, higher priority proposals, that in addition were made under a shorter forecast horizon. Therefore, to the extent that Lawrence was able to forecast district-level developments, larger spurious effects should be found on these segments.

3. *Bombay and Madras Chambers of Commerce proposals:* In 1883, the Bombay and Madras Chambers of Commerce (bodies representing commercial interests) were invited to submit railroad expansion proposals. Their proposals recommended railroad expansion into areas with unrealized commercial potential (where the Chambers' interests lay). However, the Chambers' proposals were dismissed for paying too little attention to the potential costs of building these lines (costs that the Chambers would not incur). Because it is plausible that the Chambers possessed a great deal of expertise in the identification of commercial opportunities,[67] the Chambers' expansion proposals provide a unique window on the expected commercial trajectory in the regions where the Chambers recommended construction. I estimate equation (1.16) and additionally include lines that were mentioned in the Bombay and Madras Chambers

[67]The potential for such expertise is clear in histories of the Bengal, Madras, Upper India, and Indian Chambers of Commerce in Tyson (1953), Times of India (1938), Tirumalai (1986), and Namjoshi and Sabade (1967), respectively.

of Commerce proposals. If the expected commercial trajectories identified by these Chambers are correlated with the error term in equation (1.16), then unbuilt lines in the Chambers' proposals should display spurious effects on real agricultural income. If no such effects are observed then this would call into question the ability for less commercially-interested agents, such as the Government of India (which planned India's railroad network) to systematically forecast commercial developments in India's districts.

4. *Kennedy's proposal:* India's early line placement followed the suggestions of Lord Dalhousie (then head of the Government of India), but only after Dalhousie's decade-long debate with Major Kennedy (then India's Chief Engineer, who was charged with planning India's first railroad lines) over optimal route choice. Kennedy was convinced that railroad construction would be extremely expensive in India (Davidson 1868). He therefore sought to connect Dalhousie's chosen provincial capitals with a network of lines that followed the gentlest possible gradients, along river gradients and the coastline wherever possible.[68] Kennedy's proposal is useful for my identification strategy because it identifies districts with low railroad construction costs. Geographical features that favor low construction costs (such as topography, vegetation, and climate) may also favor agricultural production, and may result in differential unobservable trends in the real agricultural income of districts with favorable construction conditions; if favorable construction

[68]The network that was built, by contrast, took straight lines in almost all circumstances, requiring in many cases (such as the Thal and Bhor Ghats) some of the most advanced railroad engineering works the world had ever seen (Andrew 1883). By 1869 it was clear that Kennedy's anticipated construction costs were, if anything, underestimates. These high construction costs were a major factor in Mayo's decision to abort Lawrence's plan, as described in my second placebo variable.

conditions drove railroad placement decisions then OLS estimates of equation (1.16) would erroneously attribute unobserved trends to railroad construction. I therefore estimate equation (1.16) including a variable that is an interaction between an indicator variable that captures districts that would have been penetrated by Kennedy's proposed network and a time trend.[69] If this variable predicts real agricultural income then this would be a concern for my identification strategy as it would suggest that the features that Kennedy found favorable for railroad construction (features that are presumably just as favorable to his successors) are correlated with real agricultural income growth. Because Kennedy's subdivided his proposal into high and low priority lines I also look for differential trends across these designations.

Results:

Table 1.5 presents estimates of the four placebo specifications described above. Column 1 compares the effect of railroad lines that were actually built to unbuilt railroad lines that were abandoned at various stages of the *four-stage planning hierarchy*. The coefficients on unbuilt lines are never statistically significantly different from zero, or of the same order of magnitude as built lines. Importantly, the coefficients on each hierarchical stage of the approval process do not display a tendency to increase as they reach advanced stages of the planning process.

Column 2 looks for spurious effects from lines identified in *Lawrence's proposal*. The coefficients on the lines that he proposed are all close to zero, an order of magnitude smaller than the coefficient on built lines, and never statistically significant. Further, the estimated coefficients on

[69]Since Kennedy's proposal was first submitted in 1848, but my real agricultural income data begins in 1870, I cannot estimate the contemporaneous impact of Kennedy's proposed lines in the same manner as my other three placebo specifications.

Dependent variable: log real agricultural income per acre	(1)	(2)	(3)	(4)
Railroad in district	0.172	0.190	0.167	0.188
	(0.099)	(0.099)*	(0.083)**	(0.075)**
Railroad in neighboring district	-0.031	-0.028	-0.058	-0.047
	(0.039)	(0.031)	(0.029)**	(0.034)
Unbuilt railroad in district, abandoned after proposal stage	0.008			
	(0.021)			
Unbuilt railroad in district, abandoned after reconnoitering stage	-0.002			
	(0.048)			
Unbuilt railroad in district, abandoned after survey stage	0.014			
	(0.038)			
Unbuilt railroad in district, abandoned after sanction stage	0.010			
	(0.082)			
(Unbuilt railroad in district, included in Lawrence Plan 1869-1873)		0.010		
x (post-1869 indicator)		(0.056)		
(Unbuilt railroad in district, included in Lawrence Plan 1874-1878)		-0.054		
x (post-1874 indicator)		(0.067)		
(Unbuilt railroad in district, included in Lawrence Plan 1879-1883)		0.008		
x (post-1879 indicator)		(0.051)		
(Unbuilt railroad in district, included in Lawrence Plan 1884-1888)		0.068		
x (post-1884 indicator)		(0.104)		
(Unbuilt railroad in district, included in Lawrence Plan 1889-1893)		-0.092		
x (post-1889 indicator)		(0.087)		
(Unbuilt railroad in district, included in Lawrence Plan 1894-1898)		0.041		
x (post-1894 indicator)		(0.058)		
(Unbuilt railroad in district, in Bombay Chamber of Commerce plans)			0.003	
x (post-1883 indicator)			(0.033)	
(Unbuilt railroad in district, in Madras Chamber of Commerce plans)			-0.063	
x (post-1883 indicator)			(0.098)	
(Unbuilt railroad in district, included in Kennedy plan, high-priority)				0.0005
x (year-1848)				(0.038)
(Unbuilt railroad in district, included in Kennedy plan, low-priority)				-0.001
x (year-1848)				(0.026)
District fixed effects	YES	YES	YES	YES
Year fixed effects	YES	YES	YES	YES
Observations	14,340	14,340	14,340	14,340
R-squared	0.769	0.769	0.770	0.770

Table 1.5: Railroads and Real Income Levels (Step 3)—Placebo Specifications Notes: OLS regressions similar to those in Table 1.4. 'Railroad in district' and 'Railroad in neighboring districts' are defined in the notes to Table 1.4. 'Unbuilt railroad in district, abandoned after X stage' is a dummy variable whose value is one if a line that was abandoned after 'X' stage penetrates a district, in all years after then line was first mentioned as reaching stage 'X' in official documents. Stages 'X' are: 'proposal', where line was mentioned in official documents; 'reconnoitering', where line route was explored by surveyors in rough detail; 'survey', where the exact route of the line and nature of all engineering works were decided on after detailed survey; and 'sanction', where the surveyed line was given official permission to be built. Heteroskedasticity-robust standard errors corrected for clustering at the district level are reported in parentheses.

Lawrence's early proposals are no larger on average than those on his later proposals.

Column 3 performs a similar exercise using lines chosen by the *Bombay and Madras Chambers of Commerce.* The coefficients on the two Chambers' proposed lines are positive but very close to zero and not statistically significantly different from zero. And, as in column 2, these coefficients are an order of magnitude smaller than the (statistically significant) coefficients on built railroad lines.

Finally, column 4 examines the extent to which districts identified in *Major Kennedy's proposal*, as inexpensive districts in which to construct railroads, display different real agricultural income trends from other districts. The coefficients on Kennedy's two types of identified lines (high and low priority) are both close to zero and not statistically significantly different from zero. Crucially, the inclusion of this variable does not change appreciably the coefficient on built railroads.

These four sets of results display a consistent pattern. Regardless of the expert choosing potential railroad lines (India's public works department, India's most senior administrator at the height of his 26-year career, commercial interest groups, or India's chief engineer), or the motivation for doing so (lines attractive to the government for many potential reasons, commercially attractive lines, or low costs of construction) unbuilt lines identified by these experts are uncorrelated with time-varying unobservable determinants of real agricultural income growth. These results cast doubt on the extent to which the Government of India was willing or able to allocate railroads to districts on the basis of their expected evolution (or factors correlated with this evolution) in real agricultural income.

1.6.5 Instrumental Variable Estimates

Empirical Strategy:

After the 1876-78 famine in India, an official UK parliamentary commission—the 1880 Indian Famine Commission—met in London to inquire into the causes of this famine, and how future famines might be prevented. Of the 11 commissioners, nine were Members of Parliament, and no commissioner possessed particular expertise in Indian railroads (Bhatia 1967). Nevertheless, the Commission was unique among previous and subsequent famine commissions in recommending that railroads could prevent famine. Regions that received inadequate rainfall (and therefore suffered from famine) in the 1876-78 agricultural years were highlighted for railroad construction.[70]

This Commission's recommendation motivates my instrumental variable (IV) approach. I instrument for railroad construction in a district with a variable that is an interaction between the deviation of rainfall in the district in the 1876-78 agricultural years (May 1876 to April 1878) from its long-run (1870-1930) mean (over pairs of agricultural years),[71] and an indicator for the post-1884 period.[72] I demonstrate below that this variable has significant predictive power for railroad construction in a district.

The exclusion restriction required for this instrument to provide consistent estimates of the coefficient γ in equation (1.16) is that rainfall shortages in the 1876-78 agricultural year affect real agricultural income

[70]The 1880 Commission argued that the 1876-78 famine had been exacerbated by slow transportation of food into famine-stricken districts.

[71]For simplicity, I use the total amount of rainfall that fell in this period. However, I obtain similar results when I instead use a weighted average over the 17 crop-specific rainfall variables introduced in Chapter 2, with weights suggested by the model.

[72]I allow four years for the Commission's recommended lines to be constructed because the average length of time between a line first being proposed and being opened for traffic in my sample was 4.3 years.

six years later only because of their effect on railroad construction (due to the 1880 Famine Commission). There are two potential concerns with this exclusion restriction. The first potential concern is that rainfall may affect real agricultural income directly, because rainfall is an important input for rain-fed agriculture (I find direct evidence for this in Chapter 2 below, and indirect evidence for this as revealed in trade flows in Step 2 above). For this reason, I control for rainfall and rainfall lagged up to 10 years in my IV regressions. As I show below, there is no evidence for statistically significant effects of rainfall after a lag of more than one year, which casts doubt on the concern that rainfall shortages in 1876-78 have a direct effect on real agricultural income post-1884.

A second potential concern with this IV strategy is that famines (or the official inquiries that followed them) may have long-lived effects on real income by potentially changing policies, institutions, demographics (through mortality, fertility, or out-migration), or (animal and human) capital stocks. For this reason, I examine whether rainfall deviations (from long-run means) in the ten other years in which famine was officially declared (and official inquiries were conducted) in India appear to affect either railroad construction or real agricultural income six years later.[73] First, I find that rainfall in non-1876-78 famine years does not predict railroad construction six years later; this is an important falsification exercise because, of all the ten non-1880 famine inquiries, it was only the 1880 inquiry that mentioned railroad construction. Second, I find that in no other (ie non-1876-78) famine year do rainfall anomalies affect real agricultural income six years later. To the extent that all famines and their inquiries had the same potential for long-lived effects

[73]There were official parliamentary famine commissions after the 1896-97 and 1899-1900 famines, in addition to that after the 1876-78 famine. Official government inquiries were also commissioned after the 1866-67, 1868-70, 1873-74, 1888-89, 1905-06, 1906-07, 1907-08 and 1911-12 famines.

on agriculture, this suggests that it was not the famine or its inquiry *per se* that caused rainfall shortages in 1876-78 to have long-lived effects on real agricultural income.

In the light of these checks, a remaining concern with my instrumental variable's exclusion restriction is that the 1876-78 famine, or its 1880 Commission, was unique in some way other than its effect on railroad construction.[74] While the exclusion restriction is fundamentally untestable, it is comforting that the most obviously unique feature of the 1880 Commission was its recommendation of railroad construction (Bhatia 1967).

Results:

Table 1.6 presents instrumental variable estimates of equation (1.16), beginning with first-stage estimates for the 1880 Famine Commission instrumental variable in column 1. These estimates demonstrate that the instrumental variable has a strong and statistically significant effect on railroad location (even after controlling for contemporaneous rainfall, lagged rainfall, of up to 3 lags, and district and year fixed effects). As this instrument has a high t-statistic, and the model is just-identified, standard concerns over weak instruments are unlikely to arise here (Stock, Wright, and Yogo 2002).[75] Column 1 also demonstrates that contemporaneous and lagged rainfall variables (up to three years of lags) do not predict railroad construction in general—these four variables are indi-

[74]This seems unlikely. For example, Visaria and Visaria (1983) summarize (in their Appendix Table 5.2) the famines in my sample period in tabular form along four dimensions: the number of people killed, and the geographic regions (ie districts), land area, and number of people "affected". The 1876-78 famine is not an outlier in any of these dimensions. Lengthier treatments of famines in this time period, such as Bhatia (1967), McAlpin (1983), and Maharatna (1996), do not see the 1876-78 famine as particularly unique among India's colonial-era famines, especially when compared to the more severe 1896-97 and 1899-1900 famines.

[75]The (herteroskedasticity and serial correlation robust) F-statistic on the excluded instrument in the first stage is 7.91 (ie the square of this variable's t-statistic).

vidually and jointly insignificant.[76] However, rainfall anomalies in one particular period, the 1876-78 agricultural years that were under the remit of the 1880 Famine Commission, do predict railroad construction after 1884.

Column 2 checks whether rainfall anomalies in ten other famine years that were officially declared as famines by the Government of India, other than 1876-78, predict railroad construction. After each of these famine years, official reports were commissioned to recommend policies for future famine, just as after the 1876-78 famine. To avoid estimating ten coefficients (one for each of the ten famines) I estimate only two different effects from these non-1876-78 famines: one for the five famines that were relatively more extreme (as defined by the number of people "affected"),[77] and one for the five that were relatively less extreme. The report on the 1876-78 famine was unique in strongly recommending railroad construction for famine prevention, and the results in column 2 are consistent with this view: I find that rainfall anomalies in other officially-declared famine years do not predict railroad construction, but anomalies in 1876-78 do.

The second-stage results, using the instrument to predict the railroad dummy variable ('railroad in district'), are presented in columns 3 and 4 of table 1.6. Column 3 includes the own-railroad and neighboring district railroad dummy variables; these IV estimates are statistically significantly different from zero, and of a very similar magnitude to the OLS results presented in table 1.4. This suggests that railroad line placement decisions were unlikely to have been driven by unobservable and

[76] An F-test for their joint significance has a p-value of 0.78 implying that the null hypothesis that their coefficients are all zero cannot be rejected.

[77] I take this measure from Visaria and Visaria (1983), Appendix 5.2. This is a more reliable measure of famine intensity than the number of people killed because of the difficulty of measuring deaths in these instances. The five most severe famines were those in 1868-70, 1873-74, 1896-97, 1899-1900 and 1907-08.

Dependent variable:	Railroad in district (First Stage)	Railroad in district (First Stage)	Log real ag. income (Second Stage)	Log real ag. income (Second Stage)
Estimation method:	OLS	OLS	IV	IV
	(1)	(2)	(3)	(4)
(Rainfall in 1876-78 ag. year minus long-run mean) x (Post-1884 indicator)	-0.051 (0.018)***	-0.047 (0.019)**		
Railroad in district			0.184 (0.104)	0.193 (0.122)
Railroad in neighboring district			-0.056 (0.034)	-0.052 (0.031)
Rainfall in district	0.013 (0.089)		1.123 (0.518)**	1.118 (0.429)**
Rainfall in district (lagged 1 year)	-0.003 (0.048)		0.328 (0.294)	
Rainfall in district (lagged 2 years)	0.009 (0.064)		0.024 (0.182)	
Rainfall in district (lagged 3 years)	-0.001 (0.058)		-0.043 (0.195)	
(Rainfall in severe famine ag. year minus long-run mean) x (Indicator for 6 years after famine year)		0.015 (0.034)		0.010 (0.025)
(Rainfall in mild famine ag. year minus long-run mean) x (Indicator for 6 years after famine year)		-0.003 (0.027)		0.006 (0.021)
District fixed effects	YES	YES	YES	YES
Year fixed effects	YES	YES	YES	YES
Observations	14,340	14,340	14,340	14,340
R-squared	0.651	0.650	0.733	0.743

Table 1.6: Railroads and Real Income Levels (Step 3)—Instrumental Variable Estimates Notes: Regressions estimating equation (1.16) using real income constructed from crop-level data on 17 principal agricultural crops (listed in Appendix A), from 239 districts in India, annually from 1870 to 1930. 'Railroad in district' is a dummy variable whose value is one if any part of the district in question is penetrated by a railroad line. 'Rainfall in 1876-78 agricultural year minus long-run mean' is the amount of rainfall (in metres) in a district from 1 May 1876 to 31 April 1878, minus the district's average annual rainfall in agricultural years from 1870 to 1930. 'Railroad in neighboring districts' is the variable 'railroad in district' averaged over all districts within a 250 km radius of the district of observation. 'Rainfall in district' is a measure (in meters) of the amount of crop-specific rainfall that fell in the district, averaged over all 17 crops using the appropriate weighting given in the text. Rainfall in severe/mild famine agricultural year minus long-run mean' is similar to the variable defined above for the 1876-78 famine, but for five other famines designated as either 'severe' or 'mild' as in the text. Data sources and construction are described in full in Appendix A. Heteroskedasticity-robust standard errors corrected for clustering at the district level are reported in parentheses. *** indicates statistically significantly different from zero at the 1 % level; ** indicates 5 % level; and * indicates 10 % level.

time-varying determinants of real agricultural income, other than those already controlled for. Importantly, in column 3 I find that while contemporaneous rainfall has a large and statistically significant effect on real agricultural income (in line with OLS results in table 1.4), lagged values of rainfall (up to three lags) appear to have no effect.[78] This is reassuring from the perspective of the exclusion restriction for the use of rainfall in 1876-78 as an instrument for railroad construction post-1884. Finally, in column 4 I check whether rainfall anomalies in officially-declared famine years (other than 1876-78) have an effect on real agricultural income. I find no statistically significant coefficients on these variables (individually or jointly), which suggests that there are no long-run effects of rainfall anomalies in famine years, other than in 1876-78. This is likely to be due to the unique feature of the 1880 Famine Commission—that it recommended railroad construction.

1.6.6 Bounds Check

Empirical Strategy

A concern when estimating equation (1.16) is that of indeterminate bias due to either positive or negative selection on time-varying unobservables: some railroad projects may have targeted districts where growth was expected and infrastructure would earn higher returns (which would introduce positive bias due to selection on unobservables); other railroad projects may have targeted lagging districts that were nonetheless politically important (which would introduce negative bias due to selection on unobservables). The final strategy I employ to mitigate concerns about non-random placement explores the empirical relevance of this positive and negative bias due to potential selection on unobservables.

[78] I find that the same is true for rainfall lags up to 10 years.

All lines built between 1883 and 1904 were required to be placed in one of four categories: 'productive' (expected to be commercially remunerative), 'protective' (intended to promote development in poorer, famine-prone regions), 'productive and protective', or 'military' (built for strategic motives). These categories were used for administrative purposes, but did not have any bearing on how a line could be used.[79] I interpret the lines that were categorized as 'productive' as being expected to earn high returns (and lead to positive bias), and the lines categorized as 'protective' as being targeted towards lagging regions (leading to negative bias). Therefore, a comparison of the effects of lines that are designated as 'productive' and 'protective' will reveal bounds on the true effect of railroads.[80] If these bounds are tight then bias due to non-random railroad placement is unlikely to be quantitatively important. As a further check on this procedure, the effect of lines designated as 'protective and productive' should lie in between those from 'protective' and 'productive' lines. Finally, the lines designated as 'military' could be biased in either direction.

To implement this strategy I estimate equation (1.16) with five separate coefficients on the own-railroad regressor ($RAIL_{ot}$): one coefficient on each of the four categories of lines built between 1883 and 1904 inclusive, and a fifth coefficient on lines built before 1883 or after 1904 (during which the categorization of railroad lines did not occur).

Results:

[79] The annual *Railway Reports* reported railroad projects according to these categories. The initial motivation for this categorization scheme first appeared in House of Commons Papers (1884).

[80] An alternative interpretation of differential effects would be that different types of railroad projects have heterogeneous treatment effects. As long as protective lines have lower treatment effects than productive lines, the average treatment effect will still be bounded by the OLS estimates from these two different types of lines. This concern simply widens the bounds.

Table 1.7 contains the results of the bounds check that I use to assess the magnitude of potential bias due to non-random placement. For purposes of comparison with earlier results, column 1 replicates column 2 of Table 1.4.

Column 2 of Table 1.7 presents OLS estimates from five different types of railroad lines used in equation (1.16)—four from the four different categories in which lines were placed between 1883 and 1904, and one for all lines built before 1883 or after 1904. The point estimates on these five types of lines are all similar to each other. As anticipated by the argument above, lines categorized as 'productive' have the highest estimated coefficients, reflecting a potential upward bias on these estimates (due to selection on unobservables). Likewise, the coefficient on lines categorized as 'protective' is the lowest of the five coefficients, reflecting bias due to negative selection on unobservables. However, the difference between these two coefficients is very small in comparison with the magnitude of the effect of an average line. This suggests that the scope for positive or negative bias due to endogenous selection is small. Put another way, the coefficient on 'protective' lines—which is likely to be an underestimate of the effect of railroads on real agricultural income—is still almost 17 percent, suggesting an important role for railroads in increasing real income.

Two further points are of note in table 1.7. First, the lines categorized as 'productive and protective' have a coefficient that lies between those on 'productive' and 'protective' lines. This is sensible, but was in no way preordained, so it provides a check on the logic of the bounds procedure. Second, the lines categorized as 'military' have a coefficient that is similar in magnitude to that on all other types of lines. This coefficient is difficult to interpret without a clear prior on the direction of its bias. Nevertheless,

Dependent variable: log real agricultural income per acre	(1)	(2)
Railroad in district	0.182	
	(0.071)***	
Railroad in neighboring district	-0.042	-0.037
	(0.020)**	(0.031)
Railroad in district, built pre-1883 or post-1904		0.174
		(0.100)*
Railroad in district, line labelled as 'productive'		0.212
		(0.119)
Railroad in district, line labelled as 'protective'		0.168
		(0.144)
Railroad in district, line labelled as 'productive and protective'		0.173
		(0.138)
Railroad in district, line labelled as 'military'		0.204
		(0.197)
District fixed effects	YES	YES
Year fixed effects	YES	YES
Observations	14,340	14,340
R-squared	0.758	0.760

Table 1.7: Railroads and Real Income Levels (Step 3)—Bounds Check Notes: OLS regressions estimating equation (1.16) using real income constructed from crop-level data on 17 principal agricultural crops (listed in Appendix A), from 239 districts in India, annually from 1870 to 1930. 'Railroad in district' is a dummy variable whose value is one if any part of the district in question is penetrated by a railroad line (calculated over all years in the sample). 'Railroad in neighboring district' was defined in Table 1.4. 'Railroad in district, built pre-1883 or post-1904' is a similar variable defined in these time periods only, because in these time periods lines were not designated according to primary intended use. 'Railroad in district, line labelled as X' is a dummy variable whose value is one if any part of the district in question is penetrated by a line that whose primary intended use was designated (between 1883 and 1904, when all lines required such a designation) as 'X'. The intended primary uses 'X' are: 'productive', where line was expected to be commercially remunerative; 'protective', where line was intended to be redistributive towards lagging regions; 'productive and protective', where the line was intended to have both of the previous primary uses; and 'military' if the line was built for military/strategic reasons. Heteroskedasticity-robust standard errors corrected for clustering at the district level are reported in parentheses. *** indicates statistically significantly different from zero at the 1 % level; ** indicates 5 % level: and * indicates 10 % level.

it is reassuring that this coefficient is similar to that on other types of lines (though I cannot rule out the case that military lines had both a different treatment effects from other types of lines, and a countervailing bias due to selection on unobservables).

1.6.7 Summary and Interpretation

Taken together, the results from my placebo, instrumental variable, and bounds procedures suggest that my earlier OLS results in Table 1.4 can be interpreted as close approximations to unbiased estimates of the effect of railroads on real agricultural income in India. The impression left by these three procedures is that administrators in colonial India allocated railroads to districts at times that were not related to unobservable, time-varying determinants of real agricultural income. This is perhaps unsurprising given the strong military motivations for building railroads in India outlined in section 1.2, and the difficulty in forecasting the attractiveness of competing railroad plans (as evidenced by the stark disagreements among top-level Indian administrators described in section 1.6.4).

The results from this section suggest that railroads caused a large (18 percent) increase in real agricultural income in India. This estimate is slightly larger than the estimate I would obtain from using a social savings methodology, of 14.8 percent.[81] The social savings approach is known to suffer from indeterminate bias, so my results here suggest that

[81] The social savings approach (Fogel 1964) seeks to estimate the decrease in national income that would have resulted had railroads not existed, and if the factors of production used in the railroad sector had instead been employed in their next-best substitute (O'Brien (1977) and Fishlow (2000) review this literature). Hurd (1983) performs a social savings calculation for India, which I adapt here. Hurd uses a transportation price reduction of a factor of four due to railroads; my results from Table 1.2 suggest that this was an underestimate, so I instead us a reduction of a factor of 5.3 (the average reduction between any pair of districts in my sample). Using this reduction of 5.3 rather than four leads to a social savings of 9.7 percent of aggregate GDP; expressed as a fraction of real agricultural income this is 14.8 percent.

the net bias to social savings estimates, in the case of India, is negative.[82]

As a final note, both the results in this section and those obtained from a social savings calculation ignore any welfare effects from changes in the volatility of real income due to railroads. In Chapter 2 of this thesis I look for evidence of such additional welfare effects due to railroads.

1.7 Empirical Step 4: A Sufficient Statistic for Railroads

Steps 1 and 2 of this chapter have argued that railroads significantly improved the trading environment in India. Step 3 demonstrated that railroads also raised the level of real agricultural income. These two sets of results are qualitatively consistent with each other, in the context of the model in Section 1.3 above. In this section I explore whether these two sets of results are also *quantitatively* consistent with each other in the context of my model. Because the reduced-form impact could arrive through a number of mechanisms, the exercise in this section can also be thought of as determining the share of the observed reduced form impact of railroads that can be explained by the trade-based mechanism in my model.

[82]Bias arises from two sources. First, because he was arguing against the 'indispensability' of railroads, Fogel (1964) chose to evaluate social savings in a manner (assuming that the demand for transportation was perfectly inelastic) that was deliberately biased upwards. Second, as Fishlow (1965), David (1969), Williamson (1974), and Fogel (1979) have argued, the social savings methodology ignores several potential effects of railroads (and hence arrives at an underestimate). For my analysis, the most relevant of these uncounted effects is that, by reducing trade costs, railroads may have given rise to aggregate efficiency gains due to reallocations in transport-using sectors. This is the mechanism stressed in the trade model that I develop (and find empirical support for) here.

1.7.1 Empirical Strategy

In order to compare the reduced-form impact of the railroad network on each district's real agricultural income (estimated in Step 3 above) to the impact that is predicted by my model, I exploit prediction 4. This prediction is equation (1.10), restated here for convenience:

$$\ln\left(\frac{r_{ot}}{P_{ot}}\right) = \sum_k \frac{\mu_k}{\theta_k} \ln A^k_{ot} - \sum_k \frac{\mu_k}{\theta_k} \ln \pi^k_{oot}. \tag{1.17}$$

Prediction 4 thus states that real agricultural income ($\frac{r_{ot}}{P_{ot}}$) is a function of only two terms: technology (A^k_{ot}) and 'autarkiness' (π^k_{oot}, the share of district o's expenditure that it buys from itself), appropriately summed over all commodities k.

To estimate this equation I need to substitute in observable variables for the unobserved terms A^k_{ot} and π^k_{oot},[83] and the unobserved parameters θ_k and μ_k. I estimated the function $\ln A^k_{ot} = \kappa RAIN^k_{ot}$ (where $RAIN^k_{ot}$ is observable) and the parameters θ_k in Step 2 (using trade data), and the parameter μ_k is simply the consumer's budget share.[84] Finally, I *compute* the autarkiness term that emerges in equilibrium in my model when it is evaluated at the parameters $\widehat{\kappa}$, $\widehat{\theta}_k$, and $\widehat{\mu}_k$, and $\widehat{\delta}$ and $\widehat{\boldsymbol{\alpha}}$ (estimated in Step 1 using salt price data). I refer to the computed autarkiness term as $\pi^k_{oot}(\widehat{\Theta}, \mathbf{RAIN}_t, \mathbf{R}_t, \mathbf{L})$ to denote its dependence on the full vector of estimated model parameters $\widehat{\Theta} \equiv (\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\mu}}, \widehat{\boldsymbol{\alpha}}, \widehat{\delta}, \widehat{\kappa})$, as well as the full vector (across districts, commodities and years) of exogenous variables, rainfall

[83] While it would be possible in principle to use trade data to observe π^k_{oot} in the data, this faces two limitations: first, as the model makes clear, π^k_{oot} is endogenous to the error term in equation (1.17), so an instrumental variables methodology would be necessary; and second, the only internal trade data available from colonial India are presented at a more aggregated level, and begin in a later year, than the data on all other variables in equation (1.17).

[84] I estimate these Cobb-Douglas weights as the average (over trade blocks and years) expenditure share for commodity k, where expenditure is calculated as output minus trade.

($\mathbf{RAIN}_t$), the transportation network ($\mathbf{R}_t$), and land sizes ($\mathbf{L}$).

Prediction 4 (ie equation (1.17)) states that, once rainfall (ie A^k_{ot}) is controlled for (and weighted over commodities k in the manner suggested by this equation), autarkiness (π^k_{oot}) in year t is a sufficient statistic for the impact of the entire railroad network open in year t on real income in year t. To test prediction 4 I estimate equation (2.2), but additionally include the sufficient statistic variable, autarkiness (π^k_{oot}):

$$\begin{aligned}\ln\left(\frac{r_{ot}}{P_{ot}}\right) &= \beta_o + \beta_t + \gamma RAIL_{ot} + \psi_1\left(\frac{1}{N_o}\right)\sum_{d\in\mathbf{N}_o} RAIL_{dt} \\ &+ \psi_2\left[\sum_k \frac{\hat{\mu}_k}{\hat{\theta}_k}\hat{\kappa} RAIN^k_{ot}\right] + \psi_3 RAIL_{ot}\times\left[\sum_k \frac{\hat{\mu}_k}{\hat{\theta}_k}\hat{\kappa} RAIN^k_{ot}\right] \\ &+ \eta\sum_k \frac{\hat{\mu}_k}{\hat{\theta}_k}\ln\pi^k_{oot}(\widehat{\Theta},\mathbf{RAIN}_t,\mathbf{R}_t,\mathbf{L}) + \varepsilon_{ot}. \qquad (1.18)\end{aligned}$$

If autarkiness is truly a sufficient statistic, as predicted by my model, then when autarkiness is included in equation (1.18) all other railroad variables should lose predictive power. That is, prediction 4 states that the coefficients γ, ψ_1 and ψ_3 should be zero in this regression. Further, taking the model equation (1.17) literally, prediction 4 also states that the coefficients ψ_2 and η will equal one and minus one, respectively.[85]

1.7.2 Results

Table 1.8 presents OLS estimates of equation (1.18) in order to test Prediction 4 and shed light on the role for a trade-based mechanism,

[85]The computed autarkiness term, $\pi^k_{oot}(\widehat{\Theta},\mathbf{RAIN}_t,\mathbf{R}_t,\mathbf{L})$ is a generated regressor, so conventional standard errors obtained when using it will be incorrect. In principle, it is possible to obtain correct standard errors by using a bootstrap procedure (as in Step 2), but this is computationally expensive here because this is the third step of an estimation procedure (and hence a three-step bootstrap procedure would be required). Adding to the difficulty, the first step is non-linear and the second step involves 85 separate regressions. I have not calculated bootstrapped standard errors for the regressions in this section because of the computation time required. However, the empirical procedure in this section is concerned primarily with the magnitude of point estimates rather than statistical inference about these estimates.

such as that highlighted in my model, to account for the reduced-form impact of railroads on real income levels and volatility.

Column 1 restates column 2 of table 1.4 (discussed in the previous section) as a point of departure. This specification makes it clear that there is a large reduced-form impact of railroads on both the level and volatility of real agricultural income in the average district in India. While these results could reflect the increased opportunities to trade that railroads brought about (an effect for which I found evidence in Step 1), other possible mechanisms could also be at work.

Following the strategy laid out in equation (1.18), column 2 of table 1.8 adds a variable, 'autarkiness' (which I compute in my model using parameter estimates from Steps 1 and 2), to the regression in column 1. Consistent with prediction 4 of the model, the coefficients on own-railroad access and neighboring districts' railroad access—both of which were statistically and economically significant in column 1—have all fallen to a level that is close to zero (and whose 95 percent confidence intervals include zero). This is consistent with the idea that autarkiness is a sufficient statistic for the impact of railroads on real agricultural income (and its responsiveness to rainfall), as predicted by the model.

In further agreement with prediction 4, the coefficient on the autarkiness term is close to minus one, implying that autarkiness, when measured in a model-consistent manner, is a strong determinant of real agricultural income. Notably, the model parameters that enter the autarkiness term were not estimated using data that enters the current estimating equation, so the impressive fit of the autarkiness term was not preordained. Reassuringly, the last part of prediction 4, that the coefficient on the rainfall measure should be one, is now also corroborated in a statistical sense; the coefficient on rainfall has fallen (when compared

Dependent variable: Log real agricultural income per acre	(1)	(2)
Railroad in district	0.252	0.021
	(0.132)*	(.0096)
Rainfall in district	2.434	1.044
	(0.741)***	(0.476)**
(Railroad in district) x (Rainfall in district)	-1.184	0.042
	(0.482)**	(0.64)
Railroad in neighboring district	-0.022	0.003
	(0.027)	(0.041)
"Openness", as computed in model		-0.942
		(0.152)***
District fixed effects	YES	YES
Year fixed effects	YES	YES
Observations	14,340	14,340
R-squared	0.770	0.788

Table 1.8: A Sufficient Statistic for Railroad Impact (Step 4) Notes: OLS Regressions estimating equation (1.18) in column 1 and equation (21) in column 2, using real income constructed from crop-level data on 17 principal agricultural crops (listed in Appendix A), from 239 districts in India, annualy from 1870 to 1930. 'Railroad in district' is a dummy variable whose value is one if any part of the district in question is penetrated by a railroad line. 'Rainfall in district' is a weighted sum of a district's crop-specific rainfall amounts (in meters), summed over all 17 crops with weights as suggested by my model (where for reasons explained in the text, the weights sum to 4.6). 'Railroad in neighboring districts' was defined in the notes to Table 1.4. 'Openness' is the share of a district's expenditure that it buys from itself; this varible is computed in the equilibrium of the model, where the model parameters are set to those estimated in Steps 1 and 2, and the exogenous variables (the transportation network, rainfall, and district land sizes) are as observed. Data sources and construction are described in full in Appendix A. Heteroskedasticity-robust standard errors corrected for clustering at the district level are reported in parentheses. *** indicates statistically significantly different from zero at the 1 % level; ** indicates 5 % level; and * indicates 10 % level.

to column 1) to a level that is close to one.

Finally, taking the point estimate of 0.021 on own-railroad access ($RAIL_{ot}$) seriously, implies that only 12 percent (ie 0.021 divided by 0.186, expressed as a percentage) of the total impact of the railroads estimated in column 1 *cannot* be explained by the mechanism of enhanced opportunities to trade according to comparative advantage, represented in the model. That is, 88 percent of the total impact of the railroads on real income in an average district can be explained by the model.

The results in table 1.8 establish a firm, quantitative connection between the earlier results in this chapter—that railroads improved the ability to trade within India (Steps 1 and 2) and that railroads raised real incomes (Step 3). These results suggest that the important welfare gains that railroads brought about can be well accounted for by the specific mechanism of comparative advantage-based gains from trade.

1.8 Conclusion

This chapter has made three contributions to our understanding of the effects of large transportation infrastructure projects, in the context of an enormous expansion in transportation infrastructure—the construction of India's railroads. Using new district-level data that I have collected from archival sources, my first contribution is to estimate the effect of India's railroads on the trading environment there. I find that railroads reduced the cost of trading, reduced inter-regional price gaps, and increased trade volumes.

My second contribution is to estimate the effect of India's railroads on welfare in colonial India. I find that that when the railroad network was extended to the average district, real agricultural income in that district rose by approximately 18 percent. While it is possible that railroads were deliberately allocated to districts on the basis of time-varying characteristics unobservable to economists today, I find little evidence for this potential source of bias to my results in placebo, instrumental variable, or bounds checks. These reduced-form findings suggest that railroads brought welfare gains to colonial India, but say very little about the economic mechanisms behind these gains.

Finally, my third contribution is to shed light on the mechanisms at work by relating the observed railroad-driven reduction in trade costs to the observed railroad-driven increase in welfare. To do so requires a calibrated, general equilibrium model of trade with many regions, many goods, and unrestricted trade costs. I extend the work of Eaton and Kortum (2002) to construct such a model and estimate its unknown parameters using auxiliary model equations. The model identifies a sufficient statistic for the effect of trade cost reductions on real income, which, when calibrated, accounts empirically for virtually all of the observed real income effect of railroads. This suggests that railroads raised real income in India primarily because they reduced the cost of trading, and enabled districts to enjoy more of the static gains from trade due to comparative advantage.

The methods used in this chapter could be applied to other settings, contemporary and historical. However, in settings (unlike colonial India) in which the manufacturing sector is large, it may be important to allow for an increasing returns to scale production technology and tradable intermediate inputs. Under such assumptions, models such as Krugman and Venables (1995) predict that the lower trade costs brought about by transportation infrastructure projects can be detrimental to some regions. Quantitative assessment of these additional effects, in settings where the transportation infrastructure improved rapidly, is an important area for future research.

The second chapter of this thesis provides a natural extension to the investigation contained in the present chapter: an inquiry into the extent to which railroads reduced the volatility of agricultural real income in India, and reduced the incidence of famines (as indeed the 1880 Famine Commissioners claimed they would.)

Chapter 2

Can Openness to Trade Reduce Income Volatility? Evidence from Colonial India's Era of Famines

2.1 Introduction

Recent decades have seen an unprecedented integration of countries into the world economy as a result of trade liberalization and reductions in transportation costs. A large empirical literature—including the investigation in Chapter 1 of this thesis—has sought to understand the consequences of these trade cost reductions for the *level* of real incomes. Much less is known, however, about the effect of trade cost reductions on the *volatility* of real incomes. This shortcoming of the empirical literature on the connection between trade openness and economic welfare is especially serious because, as I discuss below, even some of the simplest models of trade yield ambiguous theoretical predictions regarding the relationship between trade cost reductions and the volatility of real incomes.[1]

This chapter investigates the effect of trade openness on the volatility of real incomes by studying the extent to which local real incomes respond to local productivity shocks, and then comparing this responsiveness across regimes of high and low trade costs. I do this in the context of colonial-era India (from 1861 to 1930), where incomes were predominantly agricultural, and agriculture was predominantly rain-fed, so that the real incomes of the majority of citizens depended on the vagaries of India's monsoon rains. As a shifter of the trade costs regime I exploit the arrival of railroads in each district of India, as in Chapter 1.

While the effect of productivity shocks on economic welfare can be significant in a broad range of settings, these shocks take on a special gravitas in the low-income context of colonial-era India. In many years, rainfall shortages in some localities were so severe as to lead to widespread death due to starvation—events commonly referred to as famines. Indeed, the setting for this chapter is one of the worst strings of famine in recorded history, those that occurred in colonial India from 1861 to 1906. Over this time period, India was visited by eleven officially declared famines, with an estimated total death toll in the range of 15-35 million people. Because of the dependence of agriculture on the vagaries of monsoon rainfall, India had suffered from recurrent famines for millennia. But widespread famines largely ceased after 1906.[2]

[1]In addition, Newbery and Stiglitz (1981) present a wide range of models in which openness to trade can either increase or decrease real income volatility.

[2]A major exception to this statement is the 1943 famine in Bengal. But, as argued by O'Grada (2008), it is plausible that the war-time conditions that were commensurate with this calamity were an extreme case. O'Grada (2008) argues, in particular, that because Bengal's railroad lines and rolling stock were being diverted to the war effort, market integration withing Bengal suffered considerably. In addition, Burma, a major rice supplier to Bengal, was captured by Japan in early 1942 and all trade between the regions was suspended—though the plausibility of this claim is disputed by Sen (1981). Finally, the 1943 rice harvest fell victim to a rare infestation of brown spot disease, which Padmanabhan (1973) and Tauger (2003) argue was the extraneous circumstance that gave rise to this anomalous famine. Another potential exception

This dramatic improvement in the incidence of famine occurred simultaneously with the dramatic reduction in trade costs brought about by India's new railroad network (as outlined in Chapter 1). A number of commentators in the colonial era, and since, have argued that the reduction in trade costs brought about by railroads played a decisive role in ridding India of famine. Others are more sceptical and point to other forces, such as changes in the political environment, for an explanation of this sudden change.

In this chapter, in addition to examining the extent to which real incomes became less responsive to rainfall after the arrival of railroads, I investigate the claim that railroads weakened the lethal mapping behind famines in India—in which local rainfall shortages led to local mortality. Because mortality is likely to be a proxy for consumption in this low-income, poor-health setting, this additional step sheds light on the extent to which railroads reduced consumption volatility as well as real income volatility.

To guide the empirical analysis, I outline a simple, general equilibrium model of Ricardian-style trade among three regions based on the model in Chapter 1. Climatic shocks are modeled as productivity shock in one region ('home'). Railroads are modeled as a reduction in trade costs, building on the evidence in Chapter 1 that railroads did just this. A number of predictions emerge from this simple model, which drive the empirical approach followed in this chapter:

1. *Railroads reduce the responsiveness of local prices to local productivity shocks:* I find empirical support for this prediction. Specifically, I find that rainfall has an important effect on agricultural prices

occurred in the state of Maharashtra in 1972-73; O'Grada (2007) argues that this was a famine that killed 130,000 people.

before the arrival of railroads in a region the dependency of local prices on local rainfall falls to almost zero (even when focusing purely on rainfall variation across crops, within a district and year). This implies that railroads brought India's district economies close to the small open economy limit where local conditions have no effect on local prices.

2. *Railroads increase the responsiveness of local nominal incomes to local productivity shocks:* In a closed economy, a negative productivity shock in one sector reduces the physical output from that sector, but local prices are likely to rise and thereby off-set the effect of the shock on nominal incomes. In a perfectly small and open economy, however, there can be no price adjustment. The effect of an equivalent shock, therefore, on nominal incomes is larger in an open economy than in a closed economy. If find empirical support for this logic, in that railroads increase the extent to which local nominal agricultural incomes respond to local rainfall shocks.

3. *Railroads reduce the responsiveness of local real incomes to local productivity shocks:* The net effect of railroads on the responsiveness of real incomes (ie nominal income divided by a consumer price index) to productivity shocks will depend on the strength of the weakened price responsiveness (result 1) relative to that of the heightened nominal income responsiveness (result 2). This is ambiguous in a wider class of models, but in the model I work with real income responsiveness is predicted to fall with trade openness. I find empirical support for this, in that railroads reduce the responsiveness of local real agricultural income to local rainfall.

4. *Railroads reduce the responsiveness of local mortality rates to lo-*

cal productivity shocks: The extent to which reduced real income responsiveness passes through to reduced consumption responsiveness will depend on rural residents' abilities to borrow, save and insure. While these smoothing mechanisms are outside of my static model, what is clear is that if trade openness reduces real income responsiveness (as found in result 3) then it should also reduce consumption responsiveness. I lack data on consumption, so I cannot test this prediction directly. But in this low-income, poor-health environment it is likely that mortality rates are correlated with consumption levels and hence the mortality responsiveness to rainfall shocks should fall as regions open to trade. This is precisely what I find. Indeed, my results suggest that railroads brought the responsiveness of mortality to rainfall shocks from a very high level to one that is virtually non-existent. This is strong evidence consistent with the argument that railroads dramatically mitigated the the scope for famine in India.

Overall these results suggest that trade openness can play an important role in reducing the exposure of rural citizens' livelihoods to the riskiness of their environments. Further, this set of results documents benefits of trade openness and transportation infrastructure that go beyond the traditional effects of higher income levels found in Chapter 1 and work such as Frankel and Romer (1999).

These findings relate to several distinct literatures. One concerns the relationship between openness and real income volatility. The central theoretical mechanisms through which trade openness can alter price, nominal income and real income volatility was explored in depth by Newbery and Stiglitz (1981). Rodrik (1997) echoes these mechanisms—in particular, that relating to nominal income volatility—when stressing how (nominal) wage volatility may be higher in open economies than in closed ones due to higher equilibrium elasticities of labor demand in open economies. Hasan, Mitra, and Ramaswamy (2007) document empirical evidence from post-independence India that is consistent with higher trade openness leading to larger elasticities of labor demand, and Krishna and Senses (2009) find that nominal labor incomes are indeed relatively more volaltile in the United States in sectors that are relatively more open to trade. Finally, di Giovanni and Levchenko (2009) investigate, in a cross-section of countries, the reduced-form relationship between trade openness and real income volatility; in contrast to my results here, they find that more open countries have higher real income volatility.

One advantage of the present setting for investigating income volatility is that the dominant source of the volatility—rainfall fluctuations, in this agricultural environment—is both observable and exogenous. That is, rather than investigating the reduced-form relationship between trade openness and income volatility, as in the existing literature, I go one step further and ask how trade openness changes the mapping between a given stochastic input shock and equilibrium output. A second advantage of the setting I exploit here is that the change in trade costs brought about by railroads was large and differentiated over regions and time.

This chapter also contributes to the empirical literature on famines. Empirical work on famines has been largely concerned with testing the relative empirical importance of the Sen (1981) view of famines (as changes in the distribution of purchasing power over food) against the 'food availability decline' view of famines (in which the aggregate stock of food falls). Lin and Yang (2000) is in this vein. Very little empirical or theoretical work has examined how openness to trade changes the nature of famine, or the economic mechanisms behind this change, as I do in this chapter. McAlpin (1983), however, is a closely related qualitative investigation into the changing nature of famines in Bombay—a region in my sample here—in the late colonial period.

Finally, this chapter documents the potential for improvements in a nation's transportation infrastructure to improve not only the level of real incomes (as shown in Chapter 1) but also the volatility of real incomes. To the best of my knowledge, the magnitude of this channel has not been estimated before in the transportation infrastructure literature.

The remainder of this chapter proceeds as follows. The next section describes the setting of the colonial era in which the empirical exercises of this chapter are conducted; the emphasis here is on the elements of this setting—concerning volatility, mortality and famines—not covered in Chapter 1. Section 2.3 outlines a simple general equilibrium trade theory that delivers four predictions about the responsiveness of observable variables to rainfall shocks in India, and how this responsiveness is predicted to change when railroads arrive in India. Section 2.4 then reports on empirical results that are motivated by these four theoretical predictions. Finally, Section 2.5 concludes.

2.2 Background and Data

In this section I describe some of the essential background features of colonial India's era of famines, and the economic conditions that gave rise to these extreme events. I also describe the data I have collected in order to shed new empirical light on the relationship between trade openness and real income volatility from this unique setting. Additional notes on the data can be found in Chapter 1 and Appendix A.

2.2.1 Rainfall and Agricultural Incomes in India

Rainfall in India was extremely volatile from year to year (a statement that is equally true in the 1956-2001 period considered in Chapter 3). And, as discussed in Chapter 1, only 12 % of cultivated land was irrigated in 1885. The volatility of rainfall and the rain-fed nature of the vast majority of agricultural production gave rise to the common description of colonial Indian agriculture as a "gamble in monsoons."[3]

To shed light on this climatic volatility I use the measure of crop-specific rainfall (the amount of rainfall that fell in a district and year during which a given crop was under the soil) introduced in Chapter 1. Table 2.1 (which contains summary statistics relating to the data used in this chapter) documents the extent of this rainfall volatility, and demonstrates that there was no tendency for this climatic volatility to change significantly over time.

Like rainfall, prices, nominal agricultural incomes, and real agricultural incomes were also volatile over time within districts. This reflected the fact that rainfall was an essential input to agricultural production in this predominantly rain-fed environment. However, as shown in Table 2.1, there is some evidence that the volatility of these equilibrium economic variables—measured using the variables introduced in Chapter 1—changed over time in a way that rainfall did not. That is, there was a tendency for the volatility of prices and real income to decrease over the time period under study in this chapter, but for the volatility of nominal incomes to rise. Below I explore the potential for the arrival of railroads to have changed the mapping between an exogenous and stochastic production input, rainfall, and these economic variables.

[3] See, for example, Gadgil, Rajeevan, and Francis (2007). The phrase is still used to refer to agriculture today—for example, in the state of Orissa's 2005 *Economic Survey*, a state in which only 56 percent of cultivated land is irrigated.

	Number of observations	Beginning of available data	End of available data
Crop-specific rainfall shock, coefficient of variation over past 5 years	68,821	0.520 (0.417)	0.541 (0.480)
Agricultural prices, averaged over all crops, coefficient of variation over past 5 years	13,384	0.108 (0.114)	0.024 (0.031)
Nominal agricultural output, coefficient of variation over past 5 years	13,384	0.110 (0.085)	0.135 (0.106)
Real agricultural income, coefficient of variation over past 5 years	13,384	0.06 (0.04)	0.04 (0.03)
Mortality rate due to all causes of death, per 1,000 population	14,672	31.8 (24.0)	35.4 (21.2)
Mortality rate due to all causes of death, coefficient of variation over past 5 years	13,006	0.221 (0.236)	0.116 (0.127)

Table 2.1: Summary Statistics: Notes: Values are sample means over all observations for the year and question, with standard deviations in parentheses. Beginning and end of available data are: 1870 and 1930 for real and nominal agricultural incomes, and mortality rates; and 1861 and 1930 for agricultural prices. Data sources and construction are described in full in Appendix A.

2.2.2 Mortality in Colonial India

The agricultural income volatility described in the previous section occurred against a background of extremely low average agricultural incomes and subsistence living by many. In addition, the health infrastructure was poor, and only a very small minority of citizens had access to formalized health care of even the best Victorian standards (Arnold 1993).

In such a setting it is natural to expect high mortality rates and low life expectancies. I have collected the best available data on district-level death rates from 1870 onwards in order to examine how the death rate responds to rainfall variation throughout the period from 1870 to 1930, and how railroads changed this responsiveness. The mortality rate estimates stem from a (compulsory) vital events registration system that most of British India's provinces had in place by 1870. This system was the precursor to that which generates the mortality and births data used in the post-independence era analyzed in Chapter 3. Like the vital events data from many registration-based systems, the mortality rates used here are known to suffer from under-reporting, so the average death rate reported over time is almost surely an underestimate (Dyson 1991, Davis 1951). However, as Table 2.1 shows, even this underestimate is very high by both contemporary and historical standards.

Table 2.1 also documents the volatility of the mortality rate and its evolution over time. Under-reporting may have been particularly bad in times of crisis, such as famines, so the volatility is also likely to be understated. Still, there is considerable volatility in these registration data. In this low-income and low-health environment it is natural to expect real income volatility to be particularly damaging to human survival, and to give rise to volatile death rates. But notably, like the series for agricultural prices and real agricultural income, the volatility of the mortality rate is falling over time from 1870 to 1930.

2.2.3 Colonial India's Famine Era

Famines were a recurrent blight on India prior to and throughout most of the colonial period—they were India's "late Victorian Holocaust" in the words of historian Mike Davis. There is little doubt that the principal cause of these famines was a shortage of rainfall, giving rise to crop failure.[4] In the low-income setting of colonial-era India, where almost three-quarters of citizens earned their incomes in the agricultural sector directly,[5] it is not surprising that rainfall shortages can lead to death, given the importance of rainfall in India's rain-fed agricultural environment.

In times of extreme mortality in a district, India's colonial government would officially declare the district to be in a 'famine'. The exact criteria for this declaration are unclear, potentially changing over time, and potentially endogenous to the arrival of railroads. For these reasons I deliberately do not use the official famine declarations in my analysis, and focus on the raw mortality data instead.

However, the official famine declarations, along with anecdotal sources, reveal four striking features of India's 11 officially declared famines between 1860 and 1930. First, they were responsible for enormous amounts of death—15 million people at extremely conservative estimates.[6] Natu-

[4]McAlpin (1979) argues: "There is not much dispute that India has been subject to periodic crop failures for centuries, nor that the proximate cause of these crop failures is the lack of adequate timely rainfall" (p. 143). This is echoed in the standard references on Indian famines such as Srivastava (1968) and Bhatia (1967).

[5] Among the employed adult population, 69 % earned their incomes in the agricultural sector in the 1901 census (and this share fell from only 73 % in 1872, the first census year in my sample, to 68 % in 1931). Of course, the rest of rural India (which comprised almost 85 % of Indian citizens in 1900) may also have been exposed to climatic variation through interlinkages between agriculture and non-agricultural activities.

[6]I calculate this figure from Appendix Table 5.2 of Visaria and Visaria (1983).

rally, rainfall shortages may have been responsible for considerable loss of life during the many years that lay beneath the extreme 'famine' threshold.

A second important feature of the officially declared famine districts and events, and even those that were widely agreed to be suffering from 'scarcity', was their regional concentration in any given year. Srivastava (1968) plots India's famines and 'lesser scarcities' (the latter according to his definition) from 1858 to 1918. While the location of the distressed regions occurred in different areas of the subcontinent from year to year, rarely did a given years's famine or scarcity distress an area larger than the size of a single British Province—or approximately one fifteenth the area of colonial India. This limited spatial extent of famines and crop failures suggests a potential role for other regions, not suffering from crop failure, to ship food to famine-stricken districts, if trade costs are low enough to allow for these trades to occur.

A third feature of the officially declared famine events is their diminishing frequency over time during the late colonial era (and the period under study here, 1861-1930, in particular). Only one, in 1906, occurred after 1899—and as discussed in the introduction, to most observers the famine in 1906 was the last peace-time famine that India has seen.

As a final point about famines in colonial India it is important to consider how food shortages might map into death in times of famine. The exact trajectory of death during a famine, whether those in colonial India or modern Africa, is an area of uncertainty and active research. Famines are often times of widespread starvation, accidental poisoning, epidemic disease, unsanitary temporary living conditions, dislocation from fami-

Davis (2001) argues that this is a gross underestimate and suggests an upper estimate of 34 million deaths. In particular, Visaria and Visaria (1983) do not include any deaths from the severe famine of 1899-1900, for want of data.

lies and support networks, and violence, and all of these are contributing factors to the aggregate death toll of a famine. In India, epidemic disease seems to have been particularly important, while dislocation and migration was comparatively less so (Dyson 1991). But the fundamental cause of death due to epidemic disease during a famine is still regarded as inadequate nutrition, which leaves an individual's immune system compromised and vulnerable to disease (Maharatna 1996). An implication of this is that the aggregate consumption-mortality elasticity is likely to be larger than the individual equivalent, because of the externalities of mortality due to disease.

2.2.4 Famine Prevention in India and the Role of Railroads

Despite the enormous human tragedy posed by famine and 'ordinary' rainfall variation, the British colonial government in India was reluctant to intervene. Stokes (1959) and McAlpin (1979) argue that the Indian civil service, educated in the *laissez-faire*tradition of Adam Smith, preferred to help market to prevent famine. First and foremost in this endeavour was the 1880 Famine Commission's recommendation discussed in Chapter 1—to build new railroad lines in a manner that might alleviate famine. In Chapter 1 I demonstrated evidence consistent with a significant contribution of railroads to the reduction of trade costs in India. With this in mind, it is likely that food was able to move from famine-free districts to famine-stricken districts at low cost and great speed. Indeed, contemporary observers quoted in Johnson (1963), argued as much (p. 123):

"The contribution of railways in moving supplies and meeting

> the critical shortages in areas affected by famine was indeed considerable. Quite apart from the economy in the costs of carriage, the knowledge that in a few days at most, more supplies would be arriving by railway trains, helped often to keep down the expensive rise in prices which increased demands produced. When, for example, a famine raged in northern India in 1884, railway stations in Coimbatore, hundreds of miles away, were crammed with grains for transportation and sale in north India. Even a hint of scarcity sufficed to attract movement of goods by railways from areas of relative abundance."

In the remainder of this chapter I explore the potential for this movement of food brought about by railroads to mitigate the effect of rainfall shortages on death in colonial India. This investigation sheds light on the relationship between trade openness and real income volatility, and on the role played by railroads in ridding India of famine.

2.3 Model

In this section I outline a theoretical framework in which multiple regions are able to trade with one another, at a cost, and explore the implications for real income volatility of reducing this trade cost. To simplify the presentation, and to ease connection with my empirical approach below, I consider a purely static model. Of course, it is hard to consider issues of volatility in a static model. But in this static model I consider how the level of an equilibrium variable (for example, the price of a good) responds to changes in the supply of a given productive input. If this productive input is stochastic (and uncorrelated with other shocks), then studying comparative statics relating to the responsiveness of any equilibrium economic variable to this productive input is isomorphic to a dynamic extension of the model in which the volatility of prices emerges as an equilibrium variable.

To shed light on the scope for trade openness to mitigate real income volatility—or equivalently, for railroads to mitigate the effects of productivity shocks on real incomes—I draw on the same theoretical model introduced in Chapter 1 of this thesis. Unfortunately, the multiple general equilibrium interactions in that model are too complex to admit a closed-form solution for the effect of reduced trade costs on agricultural prices and their responsiveness to productivity shocks. To make progress in generating qualitative predictions to guide my empirical analysis I therefore work in this chapter with a simplified version of the model environment outlined in Chapter 1.

Specifically, throughout this section I assume that there are only three regions (called X, Y and Z), that there is only one commodity (so I will dispense with the k superscripts on all variables), that the regions are symmetric in their exogenous characteristics (ie $L_o = L$ and $A_o = A$ for all regions o), and that the three regions have symmetric trade costs with respect to each other.[7] I consider the comparative statics of a local change around this symmetric equilibrium, where it is straightforward to show that the model makes the following predictions:

[7] An alternative means of obtaining analytical predictions would be to invoke the commonly-used assumption that one commodity can be traded at zero trade cost, and is important enough to be produced in positive quantities everywhere and always. (This equates r_o to the nominal productivity in this zero trade cost sector). Unfortunately, it is difficult to imagine a commodity that satisfies these conditions in colonial India.

Prediction 1: Railroads Reduce the Responsiveness of Prices to Local Productivity Shocks

This first prediction concerns the extent to which local prices respond to local productivity shocks, and how this responsiveness changes as trade costs fall. For a local change around the symmetric equilibrium considered in this chapter:

(a) $\frac{d}{dT_{YX}}\left(\frac{dp_X}{dA_X}\right) < 0$: The responsiveness of prices in a region (say, X) to productivity shocks *in the same region* (ie $\frac{dp_X}{dA_X} < 0$) is weaker (ie less negative) when the region has low trade costs to another region (say, Y).

(b) $\frac{d}{dT_{YX}}\left(\frac{dp_X}{dA_Y}\right) > 0$: The responsiveness of prices in a region (say, X) to productivity shocks *in other any other region* (say region Y, so the price responsiveness of interest here is $\frac{dp_X}{dA_Y} < 0$) is stronger (ie more negative) when the cost of trading between these two regions (ie T_{YX}) is low.

Prediction 2: Railroads Increase the Responsiveness of Nominal Incomes to Local Productivity Shocks

This prediction concerns the effect of an exogenous change in productivity on a region's nominal income. If the exogenous productivity terms are stochastic (as in my empirical setting) then an increase in the responsiveness of nominal incomes to this stochastic production technology will increase nominal income volatility. Let r_o be the nominal income (the nominal land rental rate, per unit land area) of district o (relative to the numeraire good, taken in Chapter 1 to be the land rental rate of some reference region). Then around the three-region symmetric equilibrium:

(a) $\frac{d}{dT_{XY}}\left(\frac{dr_X}{dA_X}\right) < 0$: The effect of productivity (A_X) on nominal income in a region (say, X) rises when the cost of trading between this region and any other region (say, Y) falls.

Prediction 3: Railroads Reduce the Responsiveness of Real Incomes to Local Productivity Shocks

This prediction is analogous to Prediction 2, but concerns a region's real income rather than its nominal income. Let W_o be the real income of district o (per unit land area), which is equal to the nominal land rental rate divided by an appropriate consumer price index (ie $W_o = \frac{r_o}{\widetilde{P_o}}$ where $\widetilde{P_o}$ is the CPI). Then around the three-region symmetric equilibrium:

(a) $\frac{d}{dT_{XY}}\left(\frac{dW_X}{dA_X}\right) > 0$: The effect of productivity (A_X) on real income in a region (say, X) falls when the cost of trading between this region and any other region (say, Y) falls.

Prediction 4: Railroads Reduce the Responsiveness of the Mortality Rate to Local Productivity Shocks

In this static model, consumption equals income, and hence consumption responsiveness is equal to income responsiveness. Absent data on consumption, my empirical analysis below exploits the mortality rate as a proxy for per-capita consumption (and a measure of economic welfare that is also of its own interest). This model does not specify how consumption affects the probability of dying. But it is trivial to introduce a consumption-mortality relationship that is monotonically increasing, as in, for example, Ravallion (1997). This relationship is typically assumed to be concave. And it is plausible that, in the current setting of colonial Inda, the level of average consumption lies in a region of the consumption-mortality relation in which the slope of the relationship is particularly high. No distinction is made in Ravallion (1997) between individual and aggregate consumption-mortality relationships, though, as argued earlier, it is plausible that the aggregate relationship is also monotonic and likely to be even stronger due to the scope for externalities (such as disease). With this sort of monotonic consumption-mortality relationship added to the model we would expect the mortality rate to respond to productivity shocks and trade costs in the same manner as real income above—but with opposite sign. That is, a positive productivity shock would reduce the mortality rate, but a reduction in trade costs would reduce the extent to which the mortality rate is sensitive to this positive productivity shock.

2.4 Empirical Results

I now present a series of empirical results that are motivated by the theoretical analysis in the previous section, and that aims to shed new light on the relationship between trade openness and real income volatility. These results proceed in four steps, which trade through the four predictions in Section 2.3 sequentially. I illustrate how openness to trade, brought about by railroads, (i) dampens the responsiveness of local prices, (ii) raises the responsiveness of local nominal incomes, and reduces the responsiveness of (iii) local real incomes and (iv) mortality rates, to local productivity shocks.

2.4.1 Railroads and Price Responsiveness

Following prediction 1 of the model, I test the hypothesis that railroads reduced the responsiveness of local agricultural prices to local rainfall (an exogenous determinant of local productivity). In a small open economy (SOE), price responsiveness is zero since local prices are equal to the (exogenous) 'world' price level. However, as trade costs rise and an economy departs from the SOE limit, price responsiveness in that economy should rise (as in prediction 1). The extent of price responsiveness in a district is therefore a novel and powerful test of its openness to trade, which motivates the empirical exercise in this section.[8]

Empirical Strategy

Prediction 1 has two parts: (a) when a district is connected to the railroad network, agricultural goods prices in that district will be less responsive to productivity shocks in that district; and (b), when a railroad line connects two districts, agricultural goods prices in a district will be more responsive to productivity shocks in the other district.

I test this prediction by estimating the following linear specification (which can be interpreted as a linear approximation to the model to the model around a symmetric point):

$$\begin{aligned}
\ln p^k_{dt} = \beta^k_d &+ \beta^k_t + \beta_{dt} + \chi_1 RAIN^k_{dt} + \chi_2 RAIL_{dt} \times RAIN^k_{dt} \\
&+ \chi_3 (\tfrac{1}{N_d}) \sum_{o \in \mathbf{N}_d} RAIN^k_{ot} \\
&+ \chi_4 (\tfrac{1}{N_d}) \sum_{o \in \mathbf{N}_d} RAIN^k_{ot} \times RAIL_{odt} + \varepsilon^k_{dt}. \qquad (2.1)
\end{aligned}$$

Here, p^k_{dt} represents the retail price of agricultural crop k in district d and year t. $RAIN^k_{dt}$ is the amount of crop-specific rainfall that fell in district

[8]To my knowledge, this is a novel test for assessing a change in market integration. However, two papers are closely related. First, Shiue (2002) examines how the price correlation (over many years) between pairs of markets in 19th Century China is related to the weather correlation (over the same years) in these pairs, comparing this correlation in inland locations to that along rivers or the coast. Second, Keller and Shiue (2007a) estimate formally, in the same setting, how the spatial dependence of weather shocks (on prices) varies between inland and water-accessible regions. Neither of these papers focuses on the responsiveness of local prices to local rainfall, nor on whether the spatial transmission of weather shocks is different along some transportation links (such as railroads) than along others (such as roads), in the manner I do here.

d in year t (this variable is described in full in section 1.5 where it was first used). The variable $RAIL_{dt}$ is a dummy variable equal to one when the railroad network enters the boundary of district d, while the variable $RAIL_{odt}$ is a dummy variable equal to one when it is possible to travel from district o to district d using only the railroad network. Finally, the variable $RAIN^k_{ot}$ represents the amount of crop-specific rainfall in district o, where district o is a neighbor of district d—one of the the N_d districts ($o \neq d$) in district d's neighborhood $\mathbf{N}_d$ (taken to be all districts that lie even partially inside a 250 km radius of district d's centroid).[9] The summation terms in equation (2.1) are divided by the number of districts N_d in the neighborhood $\mathbf{N}_d$ to reflect an average effect.

I estimate equation (2.1) using fixed effects for each district-year (β_{dt}), which control for any unobservable variables affecting prices that are constant across crops within a district and year. This means that I identify price responsiveness through variation in how a given amount of annual rainfall in a district affects each of that district's crops differently. I also include fixed effects for each district-crop (β^k_d) to control for unobservables that permanently affect a district's productivity of a given crop (such as the district's soil type), and fixed effects for each crop-year (β^k_t) to control for country-wide shocks to the price of each crop.

To the extent that rainfall is a significant determinant of productivity (as I found to be the case revealed in trade flows, in Chapter 1), the coefficients χ_1 and χ_3 will be negative. Prediction 1 (a) states that the

[9]While in principle the rainfall in any district o could affect prices in district d, the model outlined in Chapter 1 suggests that these effects are likely to die out quickly over distance. In a partial equilibrium sense (that is, without allowing for the land rental rate r_o to adjust), this can be seen easily in equation (1.4). Here, each distant district's productivity term A^k_o affects local prices p^k_d in a manner proportional to $(T^k_{od})^{-\theta_k}$, where T^k_{od} is the trade cost (proportional to distance) and θ_k was estimated in Step 2 as 3.8. I therefore restrict the effect of non-local rainfall on district d's prices to that in a short (250 km) range, though my results are insensitive to using smaller (eg 100 km) or larger (eg 500 km) ranges.

coefficient χ_2 is positive (prices in district d are less responsive to rainfall in district d if district d is on the railroad network). And prediction 1 (b) states that the coefficient χ_4 is negative (lower transport costs should make prices in district d more responsive to rainfall shocks in neighboring districts to d). A positive coefficient χ_2 is consistent with railroads increasing the extent of market integration in India.

Data

I estimate equation (2.1) using annual data on the retail price of 17 agricultural commodities, in 239 districts, from 1861-1930. These prices were collected by district officers who visited the 10-15 largest retail markets in each district once every two weeks. India-wide instructions were issued to each province to ensure that prices of each commodity were recorded in a consistent manner across the provinces. The other variables used to estimate equation (2.1), concerning railroads and rainfall, were described in sections 1.4 and 1.5, respectively.

Results

Table 2.2 presents results from OLS estimates of equation (2.1). Column 1 begins by regressing (log) agricultural prices in a district on the district's crop-specific local rainfall. The coefficient on local rainfall is negative and statistically significant, suggesting that rainfall has a positive impact on crop output, and this increase in supply transmits into local retail prices. This is indicative of imperfect market integration (ie, non-zero trade costs) in these agricultural commodities *on average* over the time period 1861-1930 in India. The coefficient estimate implies a large amount of price responsiveness on average over the period: a one standard deviation (ie 0.604 m) increase in a crop's crop-specific rainfall decreases that crop's prices by approximately 15 percent.

Column 2 of Table 2.2 then tests the first part of prediction 1: that a district's prices will be less responsive to local rainfall after the district is connected to the railroad network. In this specification the coefficient on local rainfall (χ_1 in equation (2.1)) represents price responsiveness before railroads penetrate a district. The estimated coefficient is negative, statistically significant, and demonstrates a great deal of price responsiveness in the pre-railroad era. Further, in line with prediction 1, the coefficient χ_2 on rainfall interacted with a dummy for railroad access ($RAIL_{dt}$) is positive and statistically significant. The sum of the coefficients χ_1 and χ_2 represents the extent of price responsiveness after the district is brought into the railroad network. The estimated coefficients sum to -0.014 which implies that prices are still responsive to local rainfall, but in a dramatically reduced sense when compared to the coefficient of -0.428 that measures price responsiveness in the pre-rail era. However, I cannot reject the null of zero price responsiveness in the post-rail era. These findings suggests that the imperfect market integration from 1861-1930 found in column 1 reflects an average of two extreme regimes separated by the arrival of a railroad line in a district: a first regime of imperfect integration before the railroad arrives (where local supply shocks have large effects on local prices), and a second regime of near-perfect integration after the railroad arrives (where local supply shocks have a negligible effect on local prices).

Column 3 repeats the specification in column 1, but with the inclusion of average rainfall in neighboring districts. The effect of neighboring districts' rainfall on local prices is negative and statistically significant, which implies that, on average over the period from 1861-1930, neighboring districts' supply shocks affected local prices, as is consistent with

Dependent variable: Log agricultural price	(1)	(2)	(3)	(4)
Local rainfall in sowing and growing periods	-0.256 (0.102)**	-0.428 (0.184)***	-0.215 (0.105)**	-0.402 (0.125)***
(Local rainfall) x (Railroad in district)		0.414 (0.195)**		0.375 (0.184)**
Rainfall in neighboring districts			-0.051 (0.023)**	-0.021 (0.018)
Average: (Rainfall in neighboring districts) x (Connected to neighboring districts by railroad)				-0.082 (0.036)***
Crop x Year fixed effects	YES	YES	YES	YES
Crop x District fixed effects	YES	YES	YES	YES
District x Year fixed effects	YES	YES	YES	YES
Observations	73,000	73,000	73,000	73,000
R-squared	0.891	0.892	0.894	0.899

Table 2.2: The Effect of Railroads on Agricultural Price Responsiveness: Notes: OLS Regressions estimating equation (2.1) using data on 17 agricultural crops (listed in Appendix A), from 239 districts in India, annually from 1861 to 1930. 'Local rainfall in sowing and growing period' (abrev. 'local rainfall') refers to the amount of rainfall (measured in meters) in the district in question that fell during crop- and district- specific sowing and harvesting dates. 'Railroad in district' is a dummy variable whose value is one if any part of the district in question is penetrated by a railroad line. 'Rainfall in neighboring districts' is the variable 'local rainfall' averaged over all districts within a 250 km radius of the district in question. 'Connected to neighboring district' is a dummy variable that is equal to one if the district in question is connected by a railroad line to each neighboring district within 250 km. Data sources and construction are described in full in Appendix A. Heteroskedasticity-robust standard errors corrected for clustering at the district level are reported in parentheses. *** indicates statistically significantly different from zero at the 1 % level; ** indicates 5 % level; and * indicates 10 % level.

some degree of market integration.

However, the estimates in column 4 demonstrate that, as was the case in column 2, the average effect in column 3 is masking the behavior of two different regimes. Column 4 estimates equation (2.1) in its entirety by including an interaction term between each neighboring district's rainfall and a dummy variable for whether that district is connected to the 'local' district by railroad (ie $RAIL_{odt}$). As is consistent with the second part of prediction 3, the coefficient on this interaction term is negative and statistically significant. Furthermore, the coefficient on neighboring districts' rainfall (which is now the effect of rainfall in districts not connected by railroad to the local district) is not significantly different from zero. Column 4 therefore suggests that local prices do respond to neighboring districts' supply shocks when those neighbors are connected to the local district by railroads; however, neighboring districts' supply shocks are irrelevant to local prices when there is no railroad connection.

To summarize the results from this section, I find that railroads played a dramatic role in facilitating market integration, as revealed by price responsiveness, among the 17 agricultural goods in my sample. This is consistent with both parts of prediction 1 of the model, and suggests that railroads significantly aided the trade of agricultural items across districts in colonial India, so much so that local scarcities brought about by local rainfall shortages were rapidly filled by supply from surrounding regions.

2.4.2 Railroads and Nominal Income Responsiveness

The results in the preceding section suggest that prior to the arrival of railroads in a district, agricultural prices were considerably responsive to adverse rainfall shocks in that district. However, after railroads arrive in a district, the responsivenes of local prices to local rainfall shocks virtually disappears. As discussed above, this is to be expected if railroads reduced trade costs and allowed a district's local prices to be determined by supply conditions in many regions, rather than just locally.

This finding implies that when a locality receives a negative productivity shock in agriculture, local quantities of agricultural goods produced will fall, but there will be no offsetting price effect. As a result, the effect of these shocks on the value of output—that is, on nominal agricultural incomes—should become *more* responsive to these shocks. In this section I explore the extent to which this logic—the logic behind Prediction 2—holds empirically in the case of colonial India.

Empirical Strategy

Prediction 2 states that railroads will increase the responsiveness of nominal agricultural income in a district to its own rainfall shocks. Since rainfall was a stochastic input to production, an increase in the responsiveness of nominal income to this input will increase nominal income volatility. To test prediction 2 I estimate the following specification:

$$\begin{aligned} \ln r_{ot} = \beta_o + \beta_t \quad &+ \quad \gamma RAIL_{ot} + \psi_1(\tfrac{1}{N_o}) \sum_{d \in \mathbf{N}_o} RAIL_{dt} \\ &+ \quad \psi_2 RAIN_{ot} \\ &+ \quad \psi_3 RAIL_{ot} \times RAIN_{ot} + \varepsilon_{ot}. \end{aligned} \tag{2.2}$$

The dependent variable here is the (log) land rental rate, r_{ot}, which is the measure of nominal income in the model, as discussed above. The results in Chapter 1 suggest that the presence of railroads in a district can raise incomes and that the presence of railroads in neighboring districts can harm incomes. These findings motivate the inclusion of the regressors $RAIL_{ot}$ (a dummy variable for railroad access) and $\frac{1}{N_o}) \sum_{d \in \mathbf{N}_o} RAIL_{dt}$ (the average of $RAIL_{dt}$ within district o's 'neighborhood' of 250 km, $\mathbf{N}_o$). This is in line the specification used in Chapter 1.

Prediction 2 suggests that rainfall will raise nominal incomes, and that railroad access will increase the extent to which nominal incomes depend on rainfall. This motivates the inclusion of the regressors rainfall, $RAIN_{ot} \equiv \left[\sum_k \frac{\hat{\mu_k}}{\hat{\theta_k}} \hat{\kappa} RAIN_{ot}^k\right]$, and the interaction of railroad access $RAIL_{ot}$ and rainfall in equation (2.2). The rainfall variable used here is the weighted sum of crop-specific rainfall measures (introduced in section 2.4.1), where the weights are suggested by equation (1.10) of the model. Because the weights depend on the parameters θ_k, μ_k and κ, I use the values of these parameters estimated in Chapter 1 to calculate the weights. (The weights sum to 0.1, not to one, because of the presence of κ and θ_k). Finally, equation (2.2) is estimated with the inclusion of district fixed effects (β_o) and year fixed effects (β_t).

Earlier results in Chapter 1 suggested that rainfall is an important input to agricultural production; the coefficient ψ_2 is therefore expected to be positive. Prediction 2 states that the coefficient ψ_3 will also be positive, implying more responsiveness of nominal agricultural income to rainfall variation when a district has railroad access. Finally, in line with the results in Chapter 1, the coefficients γ and ψ_1 are expected to be positive and negative, respectively.

Data

The dependent variable in equation (2.2) is the nominal agricultural income in a district, denoted by the land rental rate, r_{ot}. There are no available data on land rental rates in this setting, but in a perfectly competitive, constant-returns, one-factor setting such as the model used here, the land rental rate is equal to nominal agricultural output per unit land area. As discussed in Chapter 1, I have collected data on both nominal agricultural output (the price-weighted sum of physical output quantities over 17 principal agricultural crops) from 1870-1930. I then divide this measure of nominal output by the total land area under cultivation in a district to create a measure of nominal agricultural incomes accruing to the representative owner of one unit of cultivable land.

Results

Table 2.3 presents the results of this test of prediction 2. Column 1 confirms that rainfall is an important determinant of agricultural production, and therefore nominal agricultural income. This is in line with my results from Chapter 1 (where high rainfall was found to promote exporting success) and from the previous section (where high rainfall was found to decrease prices).

However, the results in column 2 demonstrate that rainfall is a stronger determinant of nominal agricultural income once a district gains railroad access than before. That is, the coefficient on rainfall ('rainfall in district') is smaller than in column 1 (and still statistically significant). This coefficient represents the responsiveness of nominal agricultural income to rainfall before the district is connected to the railroad network. By contrast, the effect of rainfall on nominal agricultural income after a district gains railroad access (represented by the sum of the coefficients on the 'rainfall in district' term and the interaction term between railroad access and rainfall) is larger, indicating heightened nominal income responsiveness.. Further, the pattern of all four coefficients in column 2 is

Dependent variable:	Log nominal agricultural income per acre		Log real agricultural income per acre		Log total mortality rate	
	(1)	(2)	(3)	(4)	(5)	(6)
Railroad in district	0.241	0.168	0.186	0.252	0.080	0.143
	(0.114)**	(0.082)**	(0.085)**	(0.132)*	(0.061)	(0.078)*
Rainfall in district	1.410	0.532	1.248	2.434	0.064	0.145
	(0.632)***	(0.249)**	(0.430)***	(0.741)***	(0.032)**	(0.062)***
(Railroad in district) x (Rainfall in district)		0.901		-1.184		-0.123
		(0.444)**		(0.482)**		(0.059)**
Railroad in neighboring district	-0.056	-0.041	-0.031	-0.022	0.012	-0.003
	(0.048)	(0.052)	(0.021)*	(0.027)	(0.015)	(0.017)
District fixed effects	YES	YES	YES	YES	YES	YES
Year fixed effects	YES	YES	YES	YES	YES	YES
Observations	14,340	14,340	14,340	14,340	13,512	13,512
R-squared	0.771	0.775	0.767	0.770	0.642	0.647

Table 2.3: The Effect of Railroads on Nominal and Real Agricultural Income, and the Mortality Rate: Notes: OLS Regressions estimating equation (2.2) using agricultural income constructed from crop-level data on 17 principal agricultural crops (listed in Appendix A), or the total (all-causes) mortality rate, from 239 districts in India, annually from 1870 to 1930. 'Railroad in district' is a dummy variable whose value is one if any part of the district in question is penetrated by a railroad line. 'Rainfall in district' is a weighted sum of a district's crop-specific rainfall amounts (in meters), summed over all 17 crops with weights as suggested by my model, as in equation (2.2). 'Railroad in neighboring districts' was defined in the notes to Table 1.4. Data sources and construction are described in full in Appendix A. Heteroskedasticity-robust standard errors corrected for clustering at the district level are reported in parentheses. *** indicates statistically significantly different from zero at the 1 % level; ** indicates 5 % level; and * indicates 10 % level.

in line with that predicted by prediction 2.[10]

The results in columns 1 and 2 of Table 2.3 are supportive of Prediction 2. They suggest that the move to trade openness due to railroads in India made nominal agricultural incomes more vulnerable to the vagaries of the monsoon rains. This may have been harmful to individuals who consume primarily non-agricultural goods. But for the majority of citizens in this low-income economy the heightened volatility of nominal incomes, while still harmful, will have been offset by the dampened volatility of the prices of the agricultural goods that they consume in large measure. That is, the net effect of railroads on the volatility of these poor consumers' real income streams will reflect a combination of the nominal income responsiveness results seen in this section and the price responsiveness results from the previous section. In the following section I go on to explore the net effect of these two offsetting changes in the responsiveness of rainfall shocks due to India's new railroad network.

2.4.3 Railroads and Real Income Responsiveness

I now turn to the set of empirical results motivated by Prediction 3, that the effect of local productivity shocks on local *real* agricultural income should be lower in an open economy than in a closed economy. This is in contrast to the results of the previous section, where *nominal* agricultural incomes were found to become more responsive once districts opened to

[10]It is possible that while railroad access increased the responsiveness of a district's nominal income to its own rainfall, railroad connections could also have altered the responsiveness to neighboring districts' rainfall (as I found in the case of prices, in Table 2.2 above). I have estimated a specification similar to equation (2.2) but with an extension to include dependence on neighboring districts' rainfall and an interaction term for neighboring districts that are bilaterally connected by rail to the district of observation. However, the coefficients on these two additional terms (neighbors' rainfall and neighbors' rainfall for railroad connected neighbors) are small and not statistically different from zero (jointly or individually) so I have not pursued this further. This result is perhaps unsurprising given the weak price spillover effects estimated in Table 2.2.

rail trade.

As discussed above, the effect of openness on the responsiveness of real incomes to productivity shocks is in principle ambiguous. It is natural to expect openness to reduce the responsiveness of prices to productivity shocks—as expected in Prediction 1 and found empirically in Section 2.4.1 above. Given this, it is then natural to expect trade openness to increase the responsiveness of nominal incomes to productivity shocks—as expected in Prediction 2 and found empirically in Section 2.4.2 above. However, the net effect of openness on the responsiveness of real incomes—the ratio of nominal incomes to a consumer-based price index—to productivity shocks will then depend on the relative strength of the counter-veiling nominal income and price index effects. Prediction 3 was unambiguous about this net effect, predicting that real income responsiveness should be reduced by trade openness. But in a wider class of models it is likely that this result could be overturned.

In this section I therefore explore what the data from colonial India have to say about the relationship between trade openness and real income responsiveness to productivity shocks. That is, I examine the extent to which railroads altered the mapping from a given rainfall shock to the level of real agricultural income. As discussed in section 1.6.5, the 1880 Famine Commission recommended railroad construction in exactly the hope that this mapping would change for the better.

Empirical Strategy

To test prediction 3 I estimate a specification that is identical to that in the previous section (ie equation (2.2)) but use *real* agricultural income as the dependent variable. That is, I estimate the following equation:

$$\begin{aligned}\ln\left(\frac{r_{ot}}{\widetilde{P}_{ot}}\right) = \beta_o + \beta_t &+ \delta_1 RAIL_{ot} + \delta_2(\tfrac{1}{N_o}) \sum_{d \in \mathbf{N}_o} RAIL_{dt} \\ &+ \delta_3 \left[\sum_k \tfrac{\hat{\mu_k}}{\hat{\theta}_k} \hat{\kappa} RAIN^k_{ot}\right] \\ &+ \delta_4 RAIL_{ot} \times \left[\sum_k \tfrac{\hat{\mu_k}}{\hat{\theta}_k} \hat{\kappa} RAIN^k_{ot}\right] + \varepsilon_{ot}. \quad (2.3)\end{aligned}$$

Prediction 3 suggests that the coefficients δ_1, δ_2 and δ_3 will be of the same signs as in previous sections and in Chapter 1 (positive, negative and positive, respectively). However, the coefficient δ_4 is expected to be negative according to Prediction 3, indicating that real agricultural income is now less responsive to rainfall shocks once a district is able to trade by the use of railroads.

Data

The data used here are identical to those used in the previous section, except that the dependent variable is now real rather than nominal agricultural income. I construct a measure of real agricultural income just as in Chapter 1—by dividing nominal agricultural income by a consumption-based price index, defined over the 17 main agricultural crops for which price data are available. To construct consumption weights I follow the procedure used in Chapter 1 (which uses trade data and output data to estimate consumption patterns at the trade block level, and then assigns each trade block's consumption weights to all of the districts in the given trade block).

Results

The results from this test of Prediction 3 are presented in columns 3 and 4 of Table 2.3. These columns are the analogues of columns 1 and 2 respectively, but are based on the dependent variable of real agricultural income rather than nominal agricultural income.

Column 3 of Table 2.3 demonstrates that, on average, rainfall played a considerable role in determining the level of real agricultural income in colonial India. The results in column 4, however, illustrate how railroads mitigate the effect of agricultural productivity shocks on real agricultural income. That is, in line with Prediction 3, the coefficient on the interaction between rainfall and railroads is negative—which implies that an equivalent negative rainfall shock did less damage to real agricultural incomes after the arrival of railroads than before. Of course, a symmetric implication of the results from this symmetric specification is that a given *positive* shock did less to *raise* real income after the arrival of railroads as well. This result suggests that transportation infrastructure projects like India's railroads can bring about significant real income insurance. In the next and final section of this chapter I investigate whether this real income insurance appears to have given rise to consumption insurance too.

2.4.4 Railroads and Mortality Responsiveness

The results in the previous section have demonstrated that, in colonial India, railroads had an economically and statistically significant effect on reducing real income volatility, by reducing the extent to which a given productivity shock affected the real value of output in terms of what a consumer could buy with this output.

A natural follow-up to this investigation into real income volatility would lead to real consumption volatility; unfortunately, however, there exist no district-level consumption series from this time period in colonial India. But it is likely, in this setting, that a reduction in the volatility of real income may have passed through into consumption volatility, as most citizens had no access to formal insurance or banking facilities. For example, Roy (2001) describes how even even the wealthiest members of society in colonial India (outside of major cities) resorted to money-boxes and jewelery as the only means to save. Rosenzweig and Binswanger (1993) and Rosenzweig and Wolpin (1993) document limited access to insurance in post-Independence India. The gains from reduced consumption volatility may have been even more important to poor consumers due to subsistence concerns (or if risk aversion decreases with income more generally).

In this section I present results relating to the mortality rate—both because the mortality rate should provide a window on consumption in this low-income setting, and because the results relating to mortality are, I believe, important in their own right. They are especially relevant in the current setting where consumption volatility was apparently so severe as to give rise to eleven famines in just 46 years. In a sense, the results below relating to the mortality rate put a human face on earlier results which documented how lower trade costs reduced real income volatility.

Empirical Strategy

I estimate a specification that is precisely analogous to that in equation (2.3), but with the log total (all-causes) mortality rate as the dependent variable rather than log real agricultural income. Prediction 4 suggests that the pattern of coefficient signs should be the same as in Prediction 3.

Data

I use data on the total mortality rate for each district and year for which this measure is available post-1870. While the raw data list multiple separate causes of death, I aggregate over this information to create a measure of the number of deaths due to all causes. I then divide this figure by district population (interpolated exponentially between census decadal years) to construct a mortality rate per 1,000 inhabitants.

Results

The results from this test of prediction 4 are reported in Table 2.3, columns 5 and 6. These two columns are analogous to the previous two, but use the log mortality rate as the dependent variable rather than log real income. A similar pattern of coefficient estimates exists in columns 5 and 6 as in columns 3 and 4, suggesting that the death rate tracks real income reasonably closely.

The results in column 6 of Table 3.2 are particularly striking. They suggest that while local rainfall variation played a considerable role in determining the local death rate prior to the arrival of railroads in a district, after the railroad's arrival there is virtually no effect of local rainfall on local death. This suggests an important but novel insurance role played by railroads in this setting, and is, to the best of my knowledge, a first test of the claim that railroads did much to rid India of famine, as was hoped by the 1880 Famine Commission.

2.5 Concluding Remarks

This chapter has been modest in its aims. I have conducted a preliminary investigation into the extent to which railroads reduced the responsiveness of real agricultural income to agricultural productivity shocks, and attempted to trace this finding through from prices, to nominal incomes, to real incomes.

I have shown that railroads reduced the responsiveness of prices, and increased the responsiveness of nominal incomes, to productivity shocks. These findings are sensible and uncontroversial in the existing empirical and theoretical literatures. I have gone on to show, however, that real incomes respond less to a given productivity shocks in an open economy than in a relatively closed one, and this is in contrast to the existing evidence from contemporary cross-country regressions in di Giovanni and Levchenko (2009). My setting is, of course, potentially very different from these authors' and this may explain the discrepancy. But I believe that my results, which allow me to actually identify the source of income volatility, and which draw on the rich variation in trade costs brought about by railroads, allow me to go a step beyond di Giovanni and Levchenko (2009).

In a final set of results I have documented how railroads appear to have dramatically reduced the extent to which the mortality rate in colonial India depended on the randomness of the monsoon rains. This is a novel finding, which sheds light on the remarkable disappearance of peace-time famine from India after 1906.

A central concern with the results of this chapter is in their interpretation. Railroads may have brought many changes to a district (such as the ability to migrate at lower cost, to learn from distant districts, public good provision, or the spread of disease), and cheaper goods transportation is just one of these changes. Teasing apart these competing mechanisms is an important priority for future research.

The next and final chapter of this thesis examines the continued effect of rainfall on mortality in the modern, post-independence period in India. Consistent with the findings in this chapter that pertain to the post-railroad period, I find very little evidence in Chapter 3 of a relationship between rainfall and death. Chapter 3 also extends the work in the present chapter by considering the role that temperature variation played in both mortality and agricultural variables. Temperature extremes, in contrast to rainfall, appear matter a great deal today.

Chapter 3

Weather and Death in India: Mechanisms and Implications of Climate Change[1]

3.1 Introduction

The climate is a key ingredient in the earth's complex system that sustains human life and wellbeing. This is especially so in poor countries located in hot regions of the Earth. In these places, human wellbeing is doubly exposed to the weather—both directly, as weather extremes can place human physiological systems under stress, and indirectly, as weather extremes damage plant life and hence the economic livelihoods of the majority of citizens. Further, the low levels of income in these countries can limit opportunities for adaptation in response to weather shocks. The urgency of the challenges posed by the climate in these countries is underscored by the growing consensus that emissions of greenhouse gases due to human activity are altering the earth's climate, most

[1]This chapter is based on joint work with Robin Burgess, Michael Greenstone and Olivier Deschenes. I am grateful to these co-authors for their input into this work.

notably by causing temperatures to increase (International Panel on Climate Change 2007).

This chapter estimates the impact of inter-annual variation in weather on well-being in India with newly collected data from 1957-2000 in order to quantify the extent to which households in developing countries such as India are exposed to the weather. It is the first such large-scale study for a developing country that I am aware of. My primary outcome variable is the mortality rate as this is the ultimate measure of individuals' abilities to smooth consumption and more generally withstand health and income shocks.

Our main results indicate a striking relationship between daily temperatures and annual mortality rates that is without comparison in the developed world. This is illustrated in Figure 3.1. This figure, which will be described in more detail below, plots coefficient estimates that correspond to the effect of an additional day in each of 10 temperature ranges on the extent of death with in a year, relative to a day in the reference category bin of $50° - 60°$ F (whose coefficient is normalized to zero). The solid line, representing coefficient estimates from India, demonstrates that hot days lead to significantly more death—for example, one additional day with a mean temperature above $90°$ F, relative to a day with a mean temperature in the $50° - 60°$ C range, increases the annual mortality rate by roughly 0.5 %. (The coefficient on cold days between $10° - 20°$ F is high but not statistically significant, as I discuss below.)

By contrast, the coefficient estimates from a similar exercised performed in the United States—shown with a dotted line in Figure 3.1—indicates that temperature extremes have very little effect on death in

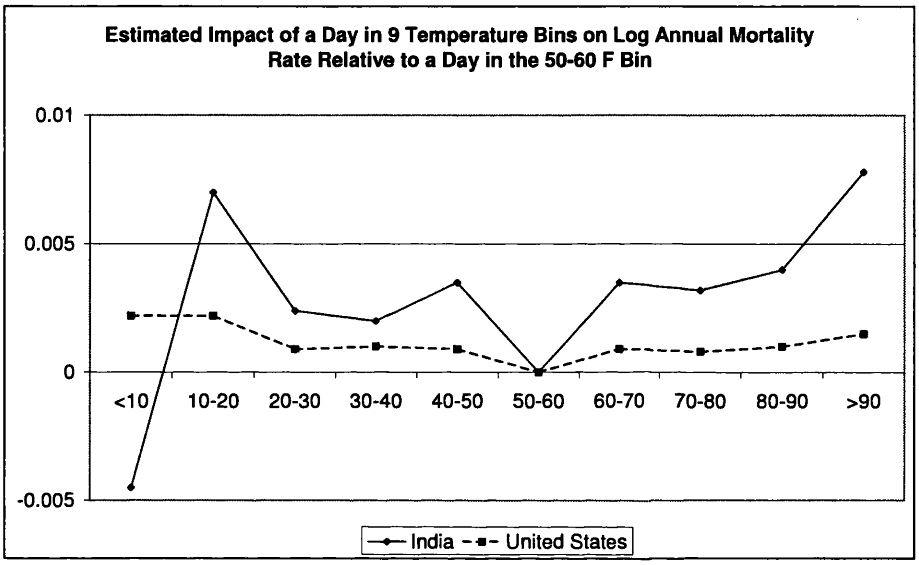

Figure 3.1: Mortality Impact of Temperature in India and United States: Note: The two plotted lines reports 10 coefficient estimates, representing the effect on annual mortality of a single day in each of the corresponding 10 temperature bins. The solid line reports these estimates for India, and the dotted line reports these estimates for the United States. The methodology used to estimate these coefficients is fully explained in Section 3.4.1.

that country.[2] Because these results from India are surprisingly large and have not, to the best of my knowledge, been documented before I repeat my analysis on a completely independent source of mortality data and find very similar coefficient estimates. It thus appears that the ambient air temperature plays an important role in determining the death rate in India.

I then go on to probe two potential mechanisms that might explain this large effect of weather on death in India. One potential explanation is that high temperatures harm human health directly, by placing cardiovascular systems under stress. A second explanation is that temperature and precipitation extremes place plants under stress and reduce agricultural yields–and that these reduced yields harm agricultural workers' incomes and possibly also their consumption levels; in the presence of poor background health, the malnutrition brought about by consumption shortfalls may be so severe as to cause death.

I document a cluster of results that are uniformly consistent with the second of these two potential mechanisms—that working through agricultural incomes–and seems inconsistent with the first. These results are:

1. *There is no effect of weather on death in urban areas of India.* This is true even among infants (those under the age of one) in urban areas, a population that is thought to be particularly vulnerable.

2. *There is no effect of weather during the non-growing months of the year on death.* This is true both in rural and urban areas, and even among rural infants.

3. *There is a similar pattern of effects of weather on agricultural in-*

[2] I obtain the coefficient estimates in Figure 3.1 for the United States from Deschenes and Greenstone (2008).

comes, but not urban incomes. That is, in hot or dry years, agricultural yields fall, agricultural prices rise, agricultural wages fall, and agricultural labor appears not to adjust. Again, this is true of weather shocks that occur during the growing months of the year but not true for equivalent shocks that occur outside of the growing season. However, there is no effect of weather fluctuations (at any time of the year) on wages in the manufacturing sector in India.

In summary, I document a large effect of weather on death among India's rural population and that this effect appears to stem from the fact that this population's economic welfare is almost entirely drawn from agricultural income and is therefore highly exposed to inter-annual variation in weather.

These results suggest that the weather plays an important role in the economic lives and health status of India's rural citizens. Because climatologists predict that the weather in India and elsewhere is likely to change in the coming decades—in particular, for there to be many more hot days during which human and plant physiology may suffer—in a final section of this chapter I use my estimated coefficients of the within-sample temperature-death relationship in India to investigate the mortality predictions implied by two leading climatological models of climate change. It is important to bear in mind that this chapter relies on inter-annual variation in temperature. This will produce an overestimate of the costs of climate change, because individuals can engage in a limited set of adaptation in response to inter-annual variation.

With this caveat in mind, my preferred mortality estimates suggest an increase in the overall Indian annual mortality rate of approximately 11% - 35% by the end of the century. The estimated increase in rural areas ranges between 52% and 69%. As a reference point, a similar exercise suggests that climate change will lead to a roughly 2% increase in the US by the end of the century (Deschenes and Greenstone 2008).

These mortality impacts are large. This is true regardless of whether one views them as the current impact of weather shocks on mortality in India or as informative about the costs of climate change.

The remainder of this chapter proceeds as follows. The next section outlines a conceptual framework that describes the mechanisms through which weather might be expected to lead to death either directly or indirectly, as well as the trail of predictions that would emerge from these two views of the world. Section 3 describes the background features of India in my sample period from 1957-2000, as well as the data on weather, death and economic variables that I have collected in order to conduct my analysis. Section 4 outlines my empirical method and presents baseline results of the weather-death relationship and on the mechanisms behind this relationship. Section 5 discusses what these estimates imply for predicted climate change scenarios in India, and finally Section 6 concludes.

3.2 Conceptual Framework

In this section I discuss the potential mechanisms through which extreme weather—in the form of high temperatures or insufficient rainfall—could lead to excess mortality in developing countries. I draw a distinction between two classes of causal channels relating weather to death: a 'direct' channel, in which human health suffers because of extreme weather conditions that put *human* physiology under stress; and an 'indirect' channel, in which human health suffers because of the stress placed on *agricultural* physiology, on which real incomes may depend.[3] The health of individuals in developing countries is therefore likely to be doubly exposed to the weather relative to those in richer countries. First, they may lack the technology, wealth and public goods necessary to mitigate direct effects. And second, they may earn their incomes in the agricultural sector and consume goods that are produced in this sector, and thereby be especially vulnerable to indirect effects.

3.2.1 Direct Effect of Weather on Death: Human Physiology

A large public health literature discusses the potential direct effects of extreme weather on human physiology. In this literature the effect of high temperatures on human health looms large (see, for example, Basu and Samet (2002) for a comprehensive review),[4] while the effects of excess or surfeit rainfall on human health is thought to be of secondary consideration. Periods of excess temperature place additional stress on cardiovascular and respiratory systems due to the demands of body temperature regulation. This stress is known to impact on the elderly and the very young with particular severity, and can, in extreme cases, lead

[3]A third potential channel relating weather to death works through the physiology of disease vectors, such as insects (or, in the case of plague and typhus, through rats). First and foremost among vector-borne diseases in India is malaria. Malaria, however, is rarely fatal in South Asia (because of the near-eradication of the particularly lethal version of the malarial parasite, *plasmodium falciparum*, which is endemic in much of Sub-Saharan Africa). In ongoing work I have assembled a district-level panel dataset on malarial incidence (rather than mortality) from India's National Malaria Eradication Programme in order to study the role of weather fluctuations in affecting malarial morbidity. A fourth potential weather-to-death mechanism affecting humans via disease could work through the physiology of the bacteria or virus itself.

[4]Extremely cold temperatures can also affect human health adversely through cardiovascular stress due to vasoconstriction and increased blood viscosity. Deschnes and Moretti (2009) find evidence for a moderate effect of extreme cold days on mortality (especially among the elderly) in the United States, though this effect is concentrated among days below 10° F (ie −12° C). Days in this temperature range are extremely rare in India.

to death (Klineberg 2002, Huynen, Martents, Schram, Weijenberg, and Kunst 2001).

The timing of these direct mortality impacts of high temperature days, however, is less well understood. That is, while it is clear that mortality spikes on days of extreme temperature, the significance of these deaths for a reduction in life expectancy is much less clear. If extreme temperatures precipitate the death of individuals whose health systems are already compromised, then the extreme temperature shock may only temporally displace (or "harvest", in the public health terminology) the number of mortality events within a given time window. Deschnes and Moretti (2009) use daily weather and daily mortality data from the United States to shed light on the timing of these effects. They indeed find evidence for the harvesting effect of high temperature days on mortality, to the extent that over a 30-day window there is virtually no lasting impact of heat waves on mortality. For this reason in this chapter I follow Deschenes and Greenstone (2008) and use annual mortality data in order to aggregate over the vast majority of deaths due to a given shock during a year.

The public health and economic literatures that have examined the direct temperature-mortality relationship in developing countries is particularly thin. One study (Hajat, Armstrong, Gouveia, and Wilkinson 2005) however, investigates the effect of heat waves on death in three different cities—Delhi, Sao Paulo and London. Among these, the largest effects (of a roughly equivalent heat wave) were seen in Delhi, but these effects were extremely small.

There are two implications of this direct mechanism that bear noting here because they are at odds with the empirical results I document below. First, a direct effect of temperature on mortality would be expected to work in equal force in both rural and urban areas of India, since these regions face the same weather (if anything, urban areas are hotter). By contrast, I find strong effects of weather on death in rural India and virtually no effect (not even among infants) in urban India. A second implication of this direct effect of temperature on mortality is that it would be expected to work equally (for a given weather shock) at all times of the year. By contrast, I document below that in my data the effect of weather on death is strong during the nine months of the year during which crops are in the soil and virtually non-existent during the three months of the year when they are not.

These two stark empirical failings of the 'direct' channel relating weather to death, coupled with the relatively small heat wave deaths found in Delhi by Hajat, Armstrong, Gouveia, and Wilkinson (2005), underscore the need to look beyond this channel and explore the empirical implications of other channels. I turn to this now and discuss an agricultural income channel relating weather to death.

3.2.2 Indirect Effect of Weather on Death: Plant Physiology and Real Incomes

A second potential channel through which weather extremes can affect human health is through their effect on the productivity of the agricultural economy. The climate is of course an essential input into agricultural production, especially in unirrigated environments. This is likely to be especially important in developing countries, where agricultural production is often unirrigated, where many households' nominal incomes depend directly on agricultural productivity, where households' consumption bundles contain many agricultural food items, and where aggregate income shocks are likely to pass through significantly into consumption because of a difficulty in smoothing consumption from period to period.

Here I discuss this channel sequentially, as it would be expected to unfold from the agricultural production process through to human mortality. I also discuss the trail of predictions that it would leave on observable variables for which data are available in India.

Agricultural yields

The starting point of an agricultural channel linking weather to death would work through agricultural yields. Both water and ambient temperature are essential inputs to agricultural production. Within reasonable ranges, water and heat are conducive to plant growth, but extremes can be damaging. In the Indian context, the two relevant extremes are too much heat (that is, frost is rare), and either too little or too much water (Kumar, Kumar, Ashrit, Deshpande, and Hansen 2004). I therefore expect that relatively hot, dry, or flooded years (measured in the appropriate manner, as I discuss below) to be years with relatively low agricultural yields. Finally, heat and precipitation outside of the growing season, when crops are not in the soil, should have no effect on plant yields. This is helpful for identifying signs of the agricultural channel at work.

Agricultural Commodity Prices

In a closed economy, if agricultural yields fall (and land use is limited, as it plausibly is in the short run) then the prices of agricultural commodities should rise. By contrast, in a small (that is, so small that it acts like a price taker relative to surrounding economies) and perfectly open economy, local agricultural yield shocks should have no bearing on local goods prices—but any departure from this small open economy limit, however, would generate an impact of local productivity shocks on local commodity prices. Our empirical units of observation are districts in India. These are surely not closed economies, but to the extent that each district's agricultural commodity market is imperfectly integrated with the market in neighboring districts (that is, there is a non-zero cost of trading between districts), there may be some effect of reduced local yields in raising local prices.[5] With this in mind, I expect relatively hot or dry weather in a district to cause relatively high prices for agricultural commodities in that district. Again, the effects of non-growing season weather should be irrelevant for agricultural prices. Finally, it is natural to expect prices at urban markets to respond to a weaker extent than do rural market prices, as rural transportation infrastructures are relatively poor and most inter-district (or longer distance) trade will first occur between cities and then from city to 'hinterland'.

Real Rural Wages and Incomes

Rural incomes in India are predominantly determined by the agricultural sector, so the above effects of diminished agricultural yields will act to reduce nominal incomes accruing to those engaged in the agricultural sector. These can be expected to pass through into nominal wages for agricultural laborers, to the extent that agricultural labor is supplied inelastically as one would expect in subsistence settings (Jayachandran 2006).

[5] One theoretical model where this can be seen very clearly (in a partial equilibrium setting) is the Ricardian model of trade due to Eaton and Kortum (2002). Here, the expression for equilibrium prices in district d is given by $p_d = [\sum_o A_o w_o^{-\theta} T_{od}^{-\theta}]^{-1/\theta}$, where the summation is over all potential supplying regions o (including the home district d itself, A_o is a supplying region's technology level, w_o a supplying region's wage, T_{od} is the cost of trading between o and d, and θ is a fixed technological parameter. This suggests that if region d can be reached at low trade costs (that is with low T_{od}) by many large (ie high A_o or low w_o) districts o, then there will be very little effect of local productivity shocks (ie A_d) on local prices. However, if all of district d's potential trading partners are separated by high trade costs then the importance of A_d in determining p_d in this equation will rise.

Further, since rural households in India are known to devote a majority of their expenditure to food items (see, for example, a recent discussion in Deaton and Dreze (2009)), I expect *real* incomes and wages to be even more diminished by adverse weather, to the extent that agricultural prices increase. Again, the effect of high temperatures and surfeit or excessive rainfall on real agricultural incomes and wages should be absent in non-growing seasons.

Real Urban Wages and Incomes

The urban economy in India is primarily non-agricultural. If rural and urban factor markets are perfectly integrated then a productivity shock in the agricultural (ie rural) sector should affect rural and urban incomes equally. However, this is unlikely to be true in India (Foster and Rosenzweig 2008), especially in the short-run as is relevant to my analysis of inter-annual variation. To the extent that rural and urban factor markets are imperfectly integrated, and assuming that high temperatures and deficient rainfall have no effect on worker productivity in non-agricultural (ie urban) sectors, one would expect no effect of weather variation on wages or output in these sectors. Naturally, there should also be no differential effect depending on whether the timing of a weather shock (in terms of the growing and non-growing seasons) within a year.

Infant and Total Mortality

Agricultural households in rural India may be unable to smooth consumption perfectly in the face of income shocks due to the weather. The resulting consumption shortfalls in hot and dry years may be so severe, or occur against a background level of such poverty, that they affect household members' health levels through malnutrition. These shocks may even be so severe, or occur against a background of such ill-health in average years, that the resulting health shocks lead to death. It is likely that especially vulnerable members of a household, such as infants, may be the first to die in the face of a household-wide income shock. A result of the above logic, from tracing through the agricultural channel relating weather to death, is that the effect of hot and dry weather on infant and total mortality in rural areas should be absent as a result of weather shocks outside of the agricultural growing season. Further, there should be no effect of weather on either infant or total mortality in urban areas; and naturally, the timing of these weather shocks (in terms of the agricultural growing season) should be irrelevant in urban India.

Summary of Predictions from Indirect Effect of Weather on Death

For ease of reference I summarize here the predictions that emerge from an agricultural income channel relating weather to death in India. This mechanism would predict that weather that is adverse for agriculture in India (weather that is extremely hot, and either too dry or too wet) would lead to:

- **Reduced agricultural yields:** Bad growing season weather will reduce agricultural yields but bad non-growing season weather will be irrelevant for agricultural yields.
- **Increased agricultural prices:** Bad growing season weather will raise agricultural prices but this effect should be moderated by the presence of irrigation. Bad non-growing season weather should be irrelevant for agricultural prices, irrespective of irrigation.
- **Reduced agricultural/rural real wages:** Bad growing season

weather will reduce real agricultural wages but bad non-growing season weather should be irrelevant for agricultural wages.

- **No change in urban real wages:** Bad growing season or non-growing season weather will have no effect on urban real wages.

- **Increased rural mortality:** Bad growing season weather will increase mortality rates in rural areas and these effects will be stronger for infants. But bad non-growing season weather will be irrelevant for both total and infant mortality.

- **No change in urban mortality:** Bad growing season or non-growing season weather will have no effect on either total or infant mortality rates in urban areas.

3.3 Background and Data

To implement the analysis in this chapter, I have collected the most detailed and comprehensive district-level data available from India on the variables that the above conceptual framework in Section 3.2 suggests are important. These variables include demographic variables (population, mortality and births), and variables that capture key features of India's urban and rural economies (output, prices and wages). I then study the relationship between these data and high-frequency daily data on historical weather that I have assembled. In this section I describe these data, their summary statistics, and the essential features of the background economy they describe.

Throughout this chapter I draw heavily on the implications of the differential weather-death relationship in urban and rural areas. I therefore begin with a short discussion of the essential differences between these regions. Despite the dramatic extent to which the world has urbanized in the last sixty years, the extent of urbanization in India has been relatively slow. 72.2 percent of Indians in 2001 lived in rural areas. The overriding distinction between economic life in rural and urban India is the source of residents' incomes. 76 % of rural citizens belong to households that draw their primary incomes from employment in the agricultural sector, while only 7 % of those in urban areas do so. Another distinction between rural and urban areas lies in their consumption of food—that is, in their exposure to fluctuations in the prices of foodstuffs. Deaton and Dreze (2009) draw on consumption surveys to report that in 2001, 58 % of the average rural residents' budget was spent on food, while only 45 % of the average urban budget was devoted to food. Naturally, these consumption differences may represent differences in the level of household per capita incomes between rural and urban areas. Urban households are, on average, richer than rural households: in 2001 urban residents were 69 % richer on average than rural residents, according to Deaton and Dreze (2009).

3.3.1 Mortality and Population Data

The cornerstone of my analysis is mortality data taken from the *Vital Statistics of India* (VSI) publications for 1957-2000, which were digitized for this project. The VSI data represent the universe of registered deaths in each year and registration was compulsory in India throughout my sample period. This source contains the most detailed possible panel of district-level mortality for all Indian citizens (though because this variable is so important for my analysis, I confirm the robustness of my analysis by performing it separately on an independent source of mortality data, which is discussed below).

Death tallies in the VSI are presented for infants (deaths under the age of one) and all others (deaths over the age of one), and by rural and urban areas separately.[6] From these I construct two measures of mortality: an infant mortality rate, defined as the number of deaths under the age of one per 1000 live births; and an 'age 1+' mortality rate, defined as the number of deaths over the age of one normalized by the population in 1000s.

Table 3.1 (which contains all of the summary statistics for data used in this chapter) summarizes the VSI data from the 1957-2000 period that I use in this chapter, which comprise 315 districts spanning 15 of India's largest states (and account for over 85 % of India's population).[7] The table reveals that measured mortality rates are high throughout this period. For example, the average infant mortality rate is 40.5 per 1,000 live births. Geographically, infant mortality rate ranges from 17.7 per 1,000 in Kerala to 71.3 per 1,000 in Orissa, revealing the substantial heterogeneity. As a basis of comparison, the mean US infant mortality rate over these years was roughly 12 per 1,000. The Indian overall mortality rate was 6.6 per 1,000. It is important to stress that these mortality rates are almost surely underestimates of the extent of mortality in India. Despite compulsory registration of births and deaths, many areas of the country suffer from significant under-reporting.[8]

[6]The rural/urban assignment is based on the following criteria, used throughout official Indian statistics (and the DHS data I use in this chapter): "(a) all places with a Municipality, Corporation or Cantonment or Notified Town Area; and (b) all other places which satisfied the following criteria: (i) a minimum population of 5,000, (ii) at least 75% of the male working population was non-agricultural, and (iii) a density of population of at least 400 per sq. Km. (i.e. 1000 per sq. Mile)."

[7]These states are (in 1961 borders and names): Andhra Pradesh, Bihar, Gujarat, Himachal Pradesh, Jammu and Kashmir, Kerala, Madhya Pradesh, Madras, Maharashtra, Mysore, Orissa, Punjab, Rajasthan, Uttar Pradesh, and West Bengal. These are the states with a consistent time series of observations in the VSI data. The results in this chapter are largely insensitive to the inclusion of all observations in the VSI data.

[8]According to the National Commission on Population of India, only 55 % of the births and 46 % of the deaths were being registered in 2000.

	Rural Areas				Urban Areas			
	1957-1969	1970-1979	1980-1989	1990-2000	1957-1969	1970-1979	1980-1989	1990-2000
Total Death Rate Per 1,000 Population	10.74	8.35	4.26	4.30	11.44	8.78	6.21	6.02
	(6.70)	(4.67)	(2.73)	(2.57)	(5.21)	(3.96)	(2.96)	(2.30)
Infant (<1) Death Rate Per 1,000 Live Births (VSI)	69.2	49.6	29.1	17.8	61.2	51.0	33.1	18.4
	(55.6)	(42.8)	(16.2)	(11.1)	(41.9)	(38.5)	(17.2)	(10.5)
Infant (<1) Death Rate Per 1,000 Live Births (DHS)	167.2	147.7	104.7	84.9	68.2	93.4	73.2	43.9
	(252.8)	(106.9)	(67.7)	(69.8)	(205.6)	(139.7)	(90.2)	(99.5)
Agricultural Yield Index (kg/hectare)	24.5	30.9	39.2	--	--	--	--	--
	(11.6)	(16.3)	(22.7)					
Agricultural Price Index (Rs/kg)	8.0	7.8	7.1	--	--	--	--	--
	(1.5)	(1.5)	(0.9)					
Agricultural Real Wages (Rs/day)	24.80	27.22	(33.96)	--	--	--	--	--
	(9.85)	(11.00)	(14.05)					
Agricultural Labor Supply (Million man-days)	90.48	106.08	116.42	--	--	--	--	--
	(52.43)	(64.60)	(63.71)					
Manufacturing Earnings Per Worker (Rs/annum)	--	--	--	--	28,330	29,982	32,595	24493
					(4741)	(7597)	(6378)	(6149)
Annual Degree-Days (over 32 C)	63.32	61.22	69.45	67.78	55.83	55.29	66.23	61.18
	(57.21)	(58.66)	(61.97)	(62.98)	(57.87)	(58.25)	(62.63)	(61.73)
Annual Total Precipitation (cm)	107.73	110.90	105.11	104.52	103.67	107.24	102.01	104.37
	(37.94)	(40.94)	(41.68)	(42.28)	(37.18)	(40.62)	(41.40)	(42.41)

Table 3.1: Summary Statistics: Note: All statistics are weighted by total district-area (ie rural/urban) population, with the exception of the Crop Productivity Index, which is weighted total crop area. Standard deviations in parentheses. Monetary values are in year 2000 Rs, deflated by urban/rural-specific deflators.

Table 3.1 also documents the time variation in the two mortality rates. There is a remarkable decline in both mortality rates in rural and urban regions. For example, the overall mortality rate declines from roughly 12 in 1957 to about 4 in rural areas and 6 in urban areas by 2000. The decline in the infant mortality rate is also impressive, going from about 100 per 1,000 in 1957 to roughly 13.5 per 1,000 in 2000. In Section 3.4.1, I describe my strategy to avoid confounding these trends in mortality rates with any time trends in temperatures.

Because the mortality data form the centerpiece of my analysis, and because vital statistics are notoriously difficult to collect accurately in developing countries, I assess the robustness of my results to using an independently collected measure of mortality. This measure comes from the Demographic and Health Surveys (DHS) conducted in India in 1992-93 and 1998-99.[9] The DHS is a nationally-representative survey of all mothers in India between the age of 15 and 49 at the age of the survey. Mothers are asked to report on the year and month of birth of each of the children to which they have given birth, as well as—if applicable—the year and month of death of any of these children. Following Young (2005) I exploit the birth history recall feature of this survey to construct a death rate for each district and year (separately by urban and rural areas). These death rates based on recall data will be unbiased estimates of their population analogues if mothers recall without systematic bias, mothers' migration is orthogonal to death rates, and mothers alive at the time of the survey (1992-93 or 1998-99) are an unbiased sample of

[9]This series of surveys are also referred to as the National Family Health Surveys (NFHS). The microdata from these surveys in India and many other countries are available at www.measuredhs.com. I are grateful to Vikram Pathania for providing us with a district code-name correspondence relevant to the DHS microdata. An additional round of surveys is available from India in 2004-05 but this survey does not report district identifiers for confidentiality concerns associated with new questions on HIV status added to the 2004-05 round.

mothers alive at points in the past. I aggregate the two DHS surveys together in order to create a panel of district death rates running from 1957 (the first birth reported in the 1992-93 survey) to 1999. Clearly, the death rates constructed from this source are subject to a very different type of sampling error from that in the VSI data, so I find it encouraging that these very independent and different types of sources of mortality data imply similar results in my analysis.

Table 1 also reports on the equivalent infant (under age one) death rates inferred from the DHS survey. The summary statistics from the DHS data concur reasonably well with those from the VSI data, which is reassuring but in no way preordained given the very different manners in which these two data sources are collected.

3.3.2 Data on India's Rural and Urban Economies

Agricultural Yields

It is natural to expect that the weather plays an important role in the agricultural economy in India. In turn, the agricultural economy may play an important role in the health of rural citizens who draw their incomes from agriculture. To shed light on these relationships I draw on the best available district-level agricultural data in India. The data on agricultural outputs, prices, wages, and employment come from the 'India Agriculture and Climate Data Set', which was prepared by the World Bank.[10] This file contains detailed district-level data from the Indian Ministry of Agriculture and other official sources from 1956 to 1987.[11] From this source I utilize four distinct variables on the agricul-

[10]The lead authors are Apurva Sanghi, K.S. Kavi Kumar, and James W. McKinsey, Jr.

[11]In ongoing work I am working to assemble the data required to extend this resource from 1988 to the present.

tural economy: yields, prices, wages, and employment.

I construct a measure of annual, district-level yields by aggregating over the output of each of the 27 crops covered in the World Bank dataset (these crops accounted for over 95 percent of agricultural output in 1986). To do this I first create a measure of real agricultural output for each year (using the price index discussed below) and then divide this by the total amount of sown area in the district-year. Table 3.1 reports on the resulting yield measure for the 271 districts contained in the World Bank dataset, over the period from 1956 to 1987. All of the major agricultural states are included in the database, with the exceptions of Kerala and Assam. It is worth noting that there is significant geographical heterogeneity in yields across India. The average value of agricultural output (in 2000 Rs per hectare) is 5,218, but this measure of agricultural productivity varies dramatically across states, ranging from 10,764 in Punjab to 2,398 in Rajasthan. This variation reflects may reflect differences in cultivation practice, irrigation use, and climate.

Agricultural Prices

Because rural households spend so much of their budgets on food, food prices are an important determinant of rural welfare in India. I construct an agricultural price index for each district and year which attempts to provide a simple proxy for the real cost of purchasing food in each district-year relative to a base year. Our simple price index weights each crop's price (across the 26 crops in the World Bank sample) by the average value of district output of that crop over the period.[12] Table 3.1 reports on the level of this price index in rural India. (The price data used in the World Bank source are 'farm harvest prices', so I prefer to interpret

[12]Annual, district-level consumption data, which would be required to construct a more appropriate *consumption*-based price index, are not available in India.

these as rural prices rather than urban prices.) These figures and their accompanying standard deviations show that prices are not as variable over space and time as the yield figures in Table 3.1, potentially reflecting a degree of market integration across India's districts (so that a market's price is determined by supply conditions both locally and further afield).

Real Agricultural Wages

A second important metric of rural incomes (in addition to agricultural productivity, discussed above) is the daily wage rate of agricultural laborers. The World Bank dataset contains information on the daily wage rate, as collected by government surveys of randomly chosen villages in each district and year. All figures are given in nominal wages per day, and are then converted into equivalent daily rate to reflect the small degree of variation in the number of hours worked per day across the sample villages.[13] I divide the reported, nominal wage rate by the agricultural price index described above to construct an estimate of the rural agricultural wage in each district-year.[14] The resulting real agricultural wage measure is reported in Table 3.1. The level of real wages is low throughout the period—never rising above 33.96 (in 2000 Rupees).

Agricultural Labor Supply

A potentially important margin along which agricultural households may be able to adjust in the face of shocks to their real wages is in the amount of labor they choose to supply. Data on the number of hours of labor

[13]This variation in hours worked per day is too small to exploit as a measure of labor supply adjustment that might vary across space and time. In the data the vast majority of villages and years report an 8 hour day as standard.

[14]A better real wage measure would of course also incorporate price information on non-agricultural items in the rural consumption basket. Unfortunately, the price and quantity information that would be required to do this are unavailable annually at the district level in India.

supplied in the agricultural sector (in each district and year) are not available. But the census, which reports the number of people who work in the agricultural sector (as their primary occupation), offers a series of decadal snapshots of agricultural labor supply as measured by the number of people who choose to work in that sector. The World Bank data then converts this measure of persons working in agricultural into a measure of annual man-days worked in the agricultural sector using a standard (and time-invariant) conversion factor of the number of working days per year. Table 3.1 documents the resulting measure, the number of reported number of man-days worked by agriculturally-engaged employees at three decadal intervals.

Real Urban Wages

While exploring the possibility of an effect of weather on the rural economy, which works through the effect of weather on plant physiology, I also investigate the extent to which urban wages appear to be affected by the weather. To do this I draw on the best available measure of urban incomes, state-level reports of the nominal manufacturing wage, deflated by an urban-specific CPI. This has some justification as a measure of an urban income, as India's manufacturing sector is almost entirely concentrated in urban regions throughout my sample period. Our measure of real manufacturing wages is taken from Besley and Burgess (2004), who collected the data from the annual *Indian Labour Yearbook* publication. Table 3.1 reports the average manufacturing earnings per worker throughout this period

3.3.3 Weather Data

A key finding from Deschenes and Greenstone (2008) is that a careful analysis of the relationship between mortality and temperature requires *daily* temperature data. This is because the relationship between mortality and temperature is highly nonlinear and the nonlinearities would be missed with annual or even monthly temperature averages. This message is echoed in the agronomic and agricultural economics literatures (as emphasized, for example, by Deschenes and Greenstone (2007) and Schlenker and Roberts (2008)).

Although India has a system of thousands of weather stations with daily readings dating back to the 19th century, the geographic coverage of stations that report publicly available temperature readings is poor (and interestingly there are more of these stations prior to 1970). Further, there are many missing values in the publicly available series so the application of a selection rule that requires observations from 365 days out of the year would yield a database with very few observations.

As a solution, I follow Guiteras (2008) and use data from a gridded daily dataset that uses non-public data and sophisticated climate models to construct daily temperature and precipitation records for 1° (latitude) × 1° (longitude) grid points (excluding ocean sites). This data set, called NCC (NCEP/NCAR Corrected by CRU), is produced by the Climactic Research Unit, the National Center for Environmental Prediction / National Center for Atmospheric Research and the Laboratoire de Mtorologie Dynamique, CNRS. These data provide a complete record for daily average temperatures and total precipitation for the period 1950-2000. I match these gridpoints to each of the districts in my sample by taking weighted averages of the daily mean temperature and total precipitation variables for all grid points within 100 KM of each district's geographic center. The weights are the inverse of the squared distance from the district center.[15]

To capture the distribution of daily temperature variation within a year, I use two different variables. The first of these temperature variables assigns each district's daily mean temperature realization to one of fifteen temperature categories. These categories are defined to include daily mean temperature less than 10° C (50° F), greater than 36° C (96.8° F), and the thirteen 2° C wide bins in between. The 365 daily weather realizations within a year are then distributed over these fifteen bins. This binning of the data preserves the daily variation in temperatures, which is an improvement over previous research on the relationship between weather and death that obscures much of the variation in temperature. Figure 3.2 illustrates the average variation in daily temperature readings across the fifteen temperature categories or bins over the 1957-2000 period. The height of each bar corresponds to the mean number of days that the average person in the vital statistics data experiences in each bin; this is calculated as the weighted average across district-by-year realizations, where the district-by-year's total population is the weight. The average number of days in the modal bin of 26°-28° C is 72.9. The mean number of days at the endpoints is 3.7 for the less than 10° C bin and 3.4 for the greater than 36° C bin.

As a second approach to capturing the influence of temperature, I draw on a stark non-linearity in the relationship between daily temperatures and both human and plant physiology that is well known in the

[15]On average, there are 1.9 grid points within the 100 KM radii circles. The subsequent results are insensitive to taking weighted averages across grid points across distances longer than 100 KM and using alternative weights (e.g., the distance, rather than the squared distance). After the inverse distance weighting procedure, 339 out of a possible 342 districts have a complete weather data record. The three districts that are dropped in this procedure are Alleppey (Kerala), Laccadive, Minicoy, and Amindivi Islands, and the Nicobar and Andaman Islands.

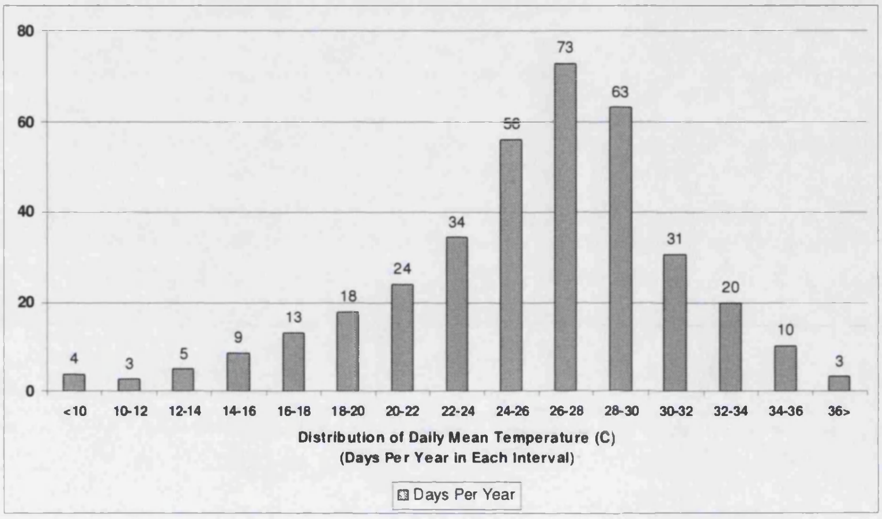

Figure 3.2: Distribution of Daily Mean Temperature, 1957-2000: Note: Mean daily temperature for each district and year, averaged while weighting by population. Each of the 365 daily realizations of the mean temperature in a year are placed into one of 15 bins. This plot illustrates the number of days in a year on (weighted) average over my sample period.

public health and agronomy literatures: temperatures above 32° C are particularly severe. I therefore construct a measure of the cumulative number of degrees-times-days that exceed 32° C in a district and year. This 'degree-days' measure has the advantage of collapsing a year's 365 daily temperature readings down to one single index, while still doing some justice to what is known about the non-linear effects of temperature. Table 3.1 reports on summary statistics of this measure. The national average is approximately 62 degree-days per year over 32° C, which implies an average of two months during which the daily mean temperature is at 34°.

Our approach to precipitation is more parsimonious: I simply use the total amount of annual precipitation (and the square of this variable). This approach aggregates over the daily precipitation data in the simplest possible manner. Further, this approach is in line with the agronomic literature which stresses the fact that water, unlike heat, is storable in plants and soil (and in groundwater and reservoirs/tanks, through irrigation). This implies that plants are significantly sheltered from the day-to-day variation in precipitation (but not temperature). Table 3.1 reports average total annual precipitation (in centimeters). The national average is roughly 1 meter, and this measure ranges from 60.1 centimeters per year on average in Rajasthan to 171.5 centimeters in West Bengal.

3.4 Method and Results

This section is divided into several subsections. I first describe the econometric method that I use to analyze the weather-death relationship and accompanying relationships in this chapter. I then go on to present my results of the effects of weather on death and to dig into how these effects differ across age groups, across rural and urban areas, across different times of the year, and when using alternative sources of mortality data.

The set of results that emerges is more support of the 'indirect' mechanism relating weather to death (that is, a mechanism working through agricultural incomes) that the 'direct' mechanism (where human health suffers through heat stress or disease). I therefore in a final set of results explore the extent to which auxiliary predictions of the indirect, agricultural income channel (outlined in Section 3.2.2) are consistent with the data.

3.4.1 Econometric Method

This section describes the econometric models used to estimate the impact of weather on mortality and other outcomes in India. I pursue two different approaches to modeling the temperature-death relationship, but my approach to the precipitation-death relationship is held constant throughout.

Our first estimating equation uses a flexible specification to model the temperature-death relationship:

$$\begin{aligned} Y_{dt} &= \sum_j \theta_j TMEAN_{dtj} + \delta_1 PREC_{dt} + \delta_2 PREC_{dt}^2 \\ &+ \alpha_d + \gamma_t + \lambda_r^1 t + \lambda_r^2 t^2 + \lambda_r^3 t^3 + \varepsilon_{dt}, \end{aligned} \tag{3.1}$$

where Y_{dt} is the log mortality rate (or an alternative outcome variable) in district d in year t. I use the log of the death rate (or of alternative outcome variables) in order to draw straightforward comparisons across different outcome variables, but my results are largely unchanged if I instead use the level of the death rate (or alternative outcome) rather than its log as my dependent variable. The r subscript refers to a 'climatic region' (explained below). The last term in the equation is a stochastic error term, ε_{dt}.

The variables of interest here are the measures of temperature and precipitation. The variables $TMEAN_{dtj}$ denote the number of days in district d and year t where the daily mean temperature falls in the jth of the fifteen bins used in Figure 3.2. I estimate separate coefficients θ_j for each of these temperature bins. Thus, the only functional form restriction on a day's mortality impact is that it is confined to the impact of the daily mean temperature, and this impact is on the annual mortality rate is constant within 2° C degree intervals. The choice of fifteen temperature bins represents an effort to allow the data, rather than parametric assumptions, to determine the mortality-temperature relationship, while also obtaining estimates that are precise enough that they have empirical content. This degree of flexibility and freedom from parametric assumptions is only feasible because of the use of district-level data from 44 years. The variable $PREC_{dt}$ denotes the total amount of annual precipitation in district d and year t, and $PREC^2_{dt}$ is the square of this variable. As discussed earlier, I use a daily approach to modeling the effects of temperature, but an annual approach to modeling the effect of precipitation, because water is largely storable across days while heat is not.

The estimating equation includes a full set of district fixed effects, α_d, which absorb all unobserved district-specific time invariant determinants of the mortality rate. So, for example, permanent differences in the supply of medical facilities will not confound the weather variables. The equation also includes unrestricted year effects, γ_t. These fixed effects control for time-varying differences in the dependent variable that are common across districts (for example, changes in health related to the

1991 economic reforms). The assumption that shocks or time-varying factors that affect health are common across districts is unlikely to be valid. Consequently, equation (3.1) includes separate cubic time trends for each of the five climatic regions r of India (groupings of states with similar climates according to India's Meteorological Department.) Since the underlying weather data only varies for 1° (latitude) × 1° (longitude) squares, it isn't possible to control for time-varying local determinants of health as flexibly as would be ideal.

Our second approach to modeling the temperature-death relationship estimates fewer parameters while still doing some justice to the non-linear nature of this relationship. This second approach, which I refer to as the 'single-index' approach, estimates the parameters in:

$$Y_{dt} = \beta CDD32_{dt} + \delta_1 PREC_{dt} + \delta_2 PREC^2_{dt} + \alpha_d + \gamma_t + \lambda^1_r t + \lambda^2_r t^2 + \lambda^3_r t^3 + \varepsilon_{dt}, \tag{3.2}$$

where the variable $CDD32_{dt}$ is the number of cumulative degree-days in district d and year t that exceeded 32° C.[16] This is similar to the approach pursued in equation (3.1) except that it collapses the 15 temperature bin parameters θ_j in equation (3.1) into one parameter, β. Our assumptions in doing so are that: (i) on days during which the mean temperature is below 32° C, temperature is irrelevant for determining the outcome variable (eg mortality) Y_{dt}; and (ii), the effect of days whose mean temperatures exceed 32° C is linearly increasing (at the rate β) in the mean daily temperature. This is broadly in line with a large public health and agronomy literature that uses the cumulative degree-day approach, and as we shall see from plots of the 15 coefficient estimates, θ_j, this does not appear to do a great injustice to reality in India. The advantage

[16]For example, if a given district-year had only two days over 32° C, one at 34° C and the other at 36° C, its value of $CDD32_{dt}$ would be 6.

of this single-index approach is that by estimating one coefficient rather than 15 I have more statistical power for teasing out the heterogeneous effects of temperature in order to learn more about the weather-death relationship.

The validity of this chapter's empirical exercise rests crucially on the assumption that the estimation of equations (3.1) and (3.2) will produce unbiased estimates of the θ_j, β and δ parameters. By conditioning on the district fixed effects, year fixed effects, and cubic polynomial time trends specific to each climatic region, these parameters are identified from district-specific deviations in weather about the district averages after controlling for the portion of shocks that remains after adjustment for the year effects and cubic time polynomials. Due to the randomness and unpredictability of weather fluctuations, it seems reasonable to presume that this variation is orthogonal to unobserved determinants of mortality rates.

There are two further points about estimating equations (3.1) and (3.2) that bear noting. First, it is likely that the error terms are correlated within districts over time. Consequently, the chapter reports standard errors that allow for heteroskedasticity of an unspecified form and that are clustered at the district level. Second, I fit weighted versions of equations (3.1) and (3.2), where the weight is the square root of the population in the district for two complementary reasons. The estimates of mortality rates from large population counties are more precise, so it corrects for heteroskedasticity associated with these differences in precision. Further, the results reveal the impact on the average person, rather than on the average district, which I believe is more meaningful.

3.4.2 Basic Results: The Weather-Death Relationship

Results for All Ages and All-India

I begin by discussing my baseline estimates of the temperature-death relationship in India from 1957-2000. The 15 temperature bin regression coefficient estimates (ie $\widehat{\theta}_j$) obtained by fitting equation (3.1) are best presented graphically—as they were in Figure 3.1. To this end, Figure 3.3 plots these regression coefficients from the estimation of a regression of a district's total (log) mortality rate (regardless of age or rural/urban designation) on my weather variables. These coefficient estimates were obtained while controlling for the effects of precipitation (as in equation (3.1), but I defer discussion of the precipitation coefficient estimates to below). Since the number of days per year is always 365 (I dropped the 366 day in leaping years), I must normalize the coefficient for one of the bins. The bin associated with 22°-24° C was normalized to zero, so each measures the estimated impact of an additional day in bin j on the log annual mortality rate, relative to the impact of a day in the 22°-24° C range. The figure also plots the estimated $\widehat{\theta}_j$ coefficients and their associated 95 percent confidence intervals, so their precision is evident.

The coefficient estimates in Figure 3.3 demonstrate that mortality risk is highest at the hottest temperatures. Indeed, the response function shows a significant and increasing relationship between log mortality rates and temperature beginning with days that exceed 30° C. The largest coefficient is for the highest temperature bin ($> 36°$ C). The magnitude is nearly 0.01, so exchanging a *single* day in the 22°-24° C range for one in the $> 36°$ C range would lead to an increase in annual mortality rates of 1%. It is noteworthy that the null of equality with the base category

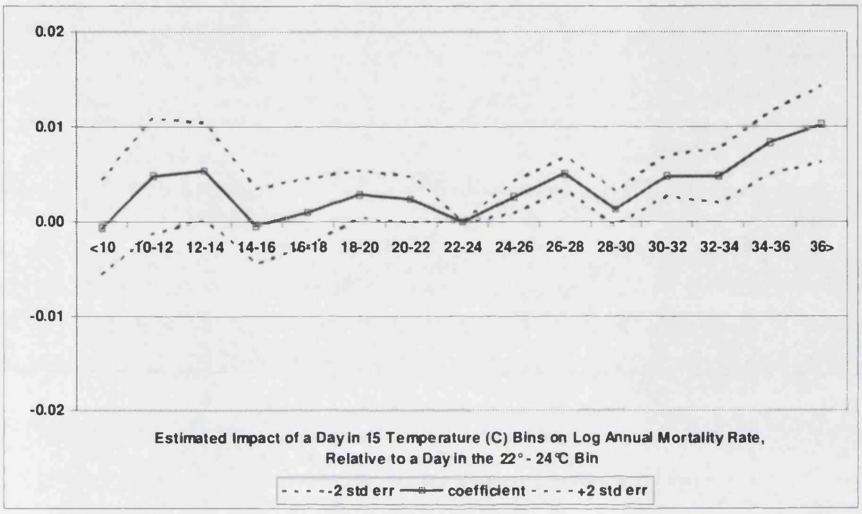

Figure 3.3: The Effect of Daily Temperature Exposure on the Log Annual Mortality Rate (All-India, All Ages): Notes: Exposure window defined based on calendar year. The dependent variable is the log annual all-age mortality rate calculated from VSI data. The model also includes controls for total precipitation and its square, and controls for unrestricted year effects, quadratic region-times-year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

can be rejected at the conventional significance levels for all bins above the reference category, with the exception of the 28°-30° C bin. Finally, the coefficients associated with the temperatures bins below the reference category are all smaller in magnitude and are estimated with less precision. That is, while an application of the present methodology to the United States by Deschenes and Greenstone (2008) has found that cold days lead to significantly more death, these effects do not appear to be salient in India (presumably because it is rarely old enough in India for these effects to take hold.)

Column 1 of Table 3.2 reports the coefficient estimates of $\widehat{\delta}_1$ and $\widehat{\delta}_2$ from equation (3.1), which describes the effect of precipitation on mortality (however, unlike in equation (3.1), I first estimate δ_1 and δ_2 without controlling for the effect of temperature). It is clear that the precipitation-death relationship is more nuanced than that between temperature and death. Column 2 repeats this exercise while controlling for the effect of temperature via the inclusion of the 15 temperature bin regressors, as in equation (3.1). There is no dramatic change in the coefficient on rainfall, implying that there is very little correlation between temperature and rainfall. There is no statistically-significant effect of annual precipitation or its square on total annual mortality. It thus appears that while the effect of temperature on death in India can be severe, the effect of rainfall is weak or non-existent. However, it is also possible that this specification is not capturing rainfall in the appropriate manner. In future work I am exploring alternative functional forms that do better justice to the timing and intra-annual distribution of rainfall, which may be important.

Dependent variable:	Log Total Mortality Rate		Log Infant Mortality Rate		Total Infant Mortality Rate	
	(1)	(2)	(3)	(4)	(5)	(6)
Total annual precipitation	0.00147	0.00134	0.00183	0.00164	-0.0451	-0.0867
	(0.0009)	(0.00101)	(0.00131)	(0.00142)	(0.0582)	(0.122)
Total annual precipitation, squared [x 1000]	-0.00429	-0.00418	-0.00689	-0.00616		
	(0.00481)	(0.00524)	(0.00751)	(0.00818)		
Total annual precipitation, squared					0.000412	0.000310
					(0.000288)	(0.000317)
Controlling for 15 temperature bins	NO	YES	NO	YES	NO	YES
Mortality data source	VSI	VSI	VSI	VSI	DHS	DHS
Number of observations	21,810	21,810	21,810	21,810	12,241	12,241
R-squared	0.63	0.64	0.58	0.59	0.18	0.20

Table 3.2: The Effect of Precipitation on Mortality in India: Notes: Exposure window defined based on calendar year. The mortality data source 'VSI' refers to data in the *Vital Statistics of India*, while that referred to as 'DHS' refers to the data in India's Demographic and Health Surveys. The model also controls for unrestricted year effects, cubic polynomial region*year trends and unrestricted district-times-area effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

Results for Infants Only

I now present results that repeat the above analysis but for infant deaths (those before the age of one) only. Figure 3.4 plots the 15 temperature bin coefficients. A similar pattern to that for adult deaths (in Figure 3.3) emerges, but the coefficients on high temperature bins for infant deaths all lie above their counterparts for adult deaths. That is, as expected, infants appear to be especially vulnerable to the pernicious effects of extreme temperatures when compared to adults, as more of the high temperature coefficients (those over 26° C) are large and statistically significant than was the case for adults. And they are extremely vulnerable: my results suggest that for every day that moves from the 22°-24° C range to anywhere in the *¿26°* C range, approximately 0.6 % more of these infants under the age of one will die. Columns 3 and 4 of Table 3.2 demonstrate the effects of precipitation in a year—whether entered on its own (in column 3) or alongside the 15 temperature bin regressors (in column 4)—has little statistically significant bearing on infant mortality. This mirrors my results for adults.

Robustness of Results to an Independent Mortality Data Source

The results reported above regarding the effect of hot days on both adult and infant mortality are striking and have not, to the best of my knowledge, been documented before in either the public health or economics literatures. It is important, therefore, to be sure that they reflect something true about reality in India rather than, say, spurious effects due to the way that the VSI mortality data were collected. For this reason, Figure 3.5 presents a similar plot to those in Figures 3.3 and 3.4—that is, a plot of the 15 temperature bin coefficients from estimating equation 3.1, but using the death rate constructed from the DHS data as the de-

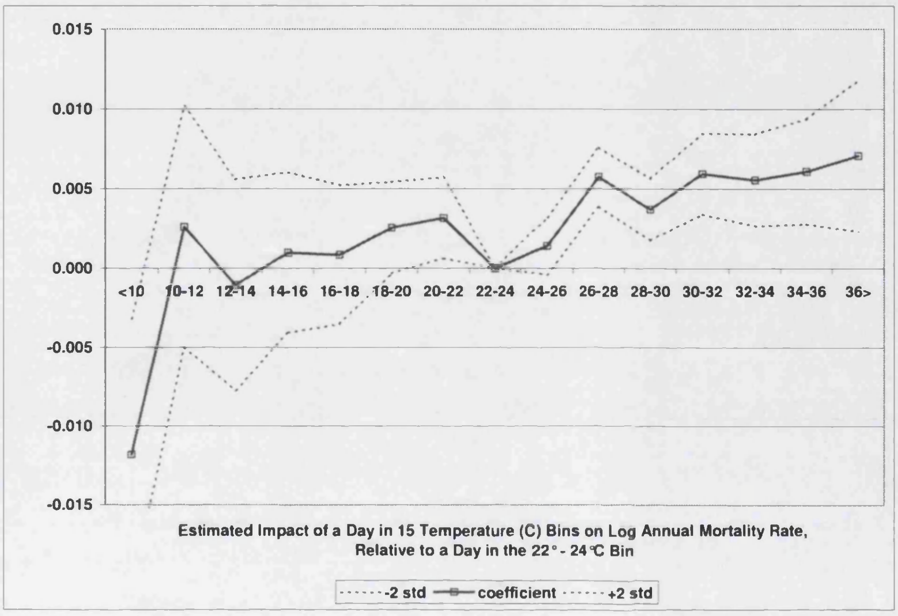

Figure 3.4: The Effect of Daily Temperature Exposure on the Log Annual Infant Mortality Rate (All-India): Notes: Exposure window defined based on calendar year. The dependent variable is the log annual infant (under the age of one) mortality rate as calculated from VSI data. The model also includes controls for total precipitation and its square, and controls for unrestricted year effects, quadratic region-times-year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

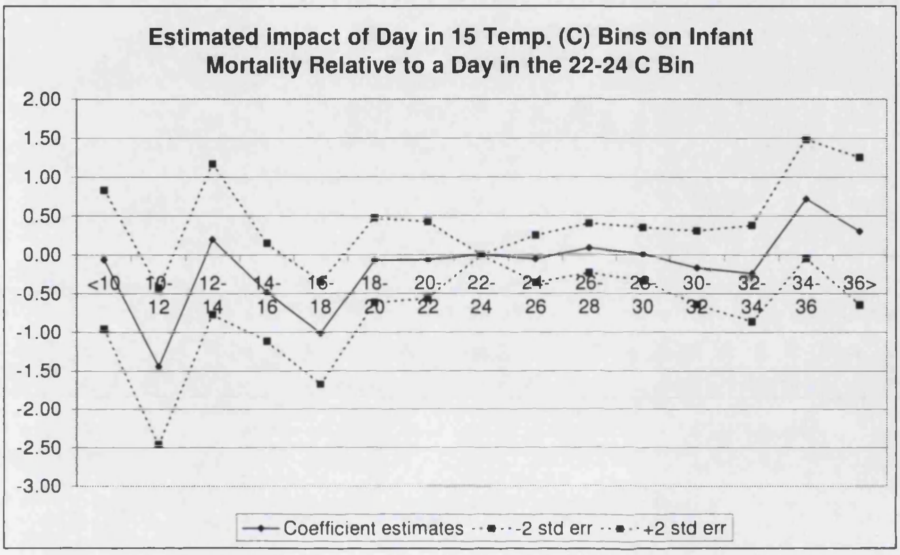

Figure 3.5: The Effect of Daily Temperature Exposure on the Annual Infant Mortality Rate (All-India) using DHS mortality data: Notes: Exposure window defined based on calendar year. The dependent variable is the annual infant (under the age of one) mortality rate as calculated using DHS data. The model also includes controls for total precipitation and its square, and controls for unrestricted year effects, quadratic region-times-year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

pendent variable instead of that constructed from the VSI data. Because there are many zeroes in the DHS data (district-year cells during which no deaths were reported by mothers surveyed in 1992 or 1998) I use the death rate rather than the log death rate as the dependent variable here.

As discussed in Section 3.3, the VSI and DHS sources are both designed to report representative district-level death rates, but both are likely to fall short of this goal; however, these two sources were collected in very different ways (universal registration and mother recall surveys, respectively). It is therefore reassuring that the pattern of coefficients in Figure 3.5 (which uses DHS data) is similar to that in Figures 3.3 and 3.4 (which use VSI data). The coefficient estimates are also quantitatively similar once adjustment is made for the fact that the DHS data regressions use the death rate, rather than the log of the death rate, as the dependent variable. For example, the coefficient on the $34-36°$ C bin is estimated at approximately 0.7; this implies that each additional day in this bin, relative to the reference bin, will lead to 0.7 more deaths (before the age of one) per 1,000 live births. The average infant death rate in the DHS data is 107.6, so this implies a proportional change in the death rate of approximately 0.7 %; by comparison, the infant death coefficients on high temperatures estimated in Figure 3.4 were in the region of 0.07 % impacts.

Finally, columns 5 and 6 of Table 3.2 report the equivalent precipitation effects obtained when using the DHS data to construct infant death rates rather than the VSI data. These are not similar, and of opposing signs, from their VSI data counterparts; but like their VSI counterparts these coefficients are imprecisely estimates to it is difficult to draw lessons from the coefficient estimates themselves. This is again suggestive of an inappropriately modeled precipitation regression.

In summary, the results in this and previous sections—summarized in Figures 3.3 through 3.5 and Table 3.2—demonstrate that temperature extremes play an important role in determining the annual mortality rate in India, but that precipitation extremes, captured in the way I have measured them here, do not. Hot days appear to matter especially among infants relative to among adults. And reassuringly, the pattern of temperature effects that I estimate is similar across the VSI and DHS data sources. The next section examines how these effects differ across rural and urban India and finds robust evidence that the effect of weather on death is essentially a purely rural phenomenon. Ensuing sections attempt to understand why.

Differential Impacts Across Rural and Urban Sectors

Figures 3.6 and 3.7 present estimated response functions between log annual mortality rate and temperature exposure, estimated separately for rural and urban sectors. Again, these results are estimated while controlling for the effects of precipitation. These models pool across age groups and pertain to the total population of an area (ie the urban or rural segment of a district).

Figure 3.6 shows the rural response function. This plot shows a significant and increasing relationship between log mortality rates and temperature beginning with days that exceed 24C. The largest coefficient is for the highest temperature bin ($> 36°$ C), and the magnitude is 0.012, so exchanging a single day in the 22°-24° C range for one in the $> 36°$ C range would lead to an increase in annual mortality rates of 1.2% in the rural sector. The statistical precision of the coefficients above the reference category is evidence as shown by the 95% confidence interval that is bounded away from zero. However, the coefficients associated

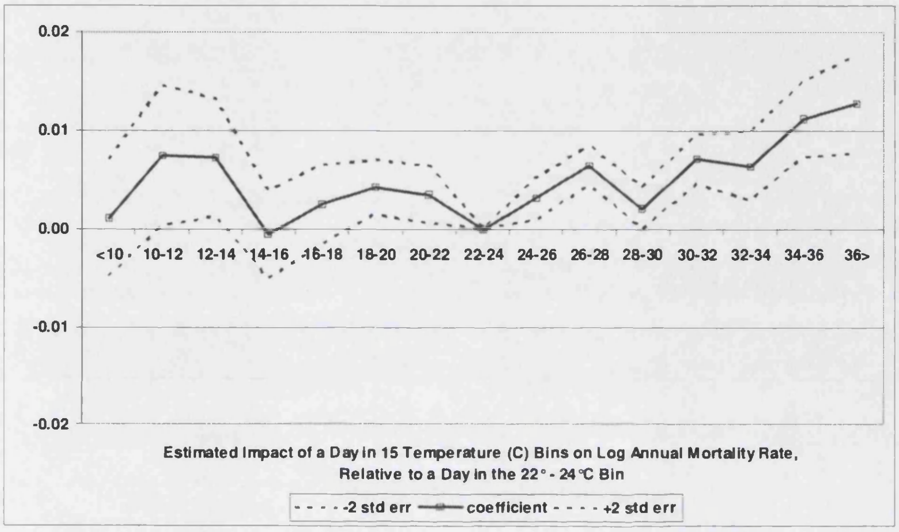

Figure 3.6: The Effect of Daily Temperature Exposure on the Annual Total Mortality Rate (Rural Areas Only): Notes: Exposure window defined based on calendar year. The dependent variable is the total mortality rate in rural areas as calculated using VSI data. The model also includes controls for total precipitation and its square, and controls for unrestricted year effects, quadratic region-times-year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

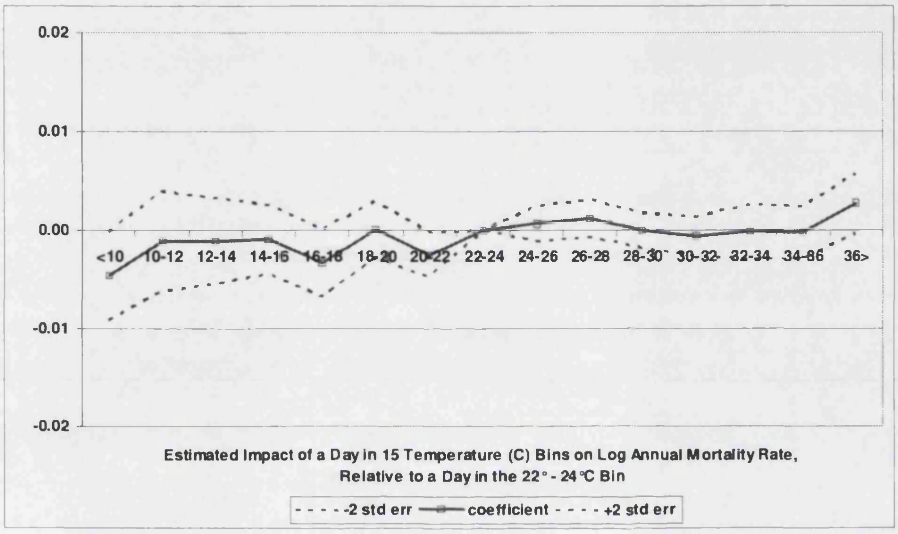

Figure 3.7: The Effect of Daily Temperature Exposure on the Annual Total Mortality Rate (Urban Areas Only): Notes: Exposure window defined based on calendar year. The dependent variable is the total mortality rate in urban areas as calculated using VSI data. The model also includes controls for total precipitation and its square, and controls for unrestricted year effects, quadratic region-times-year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

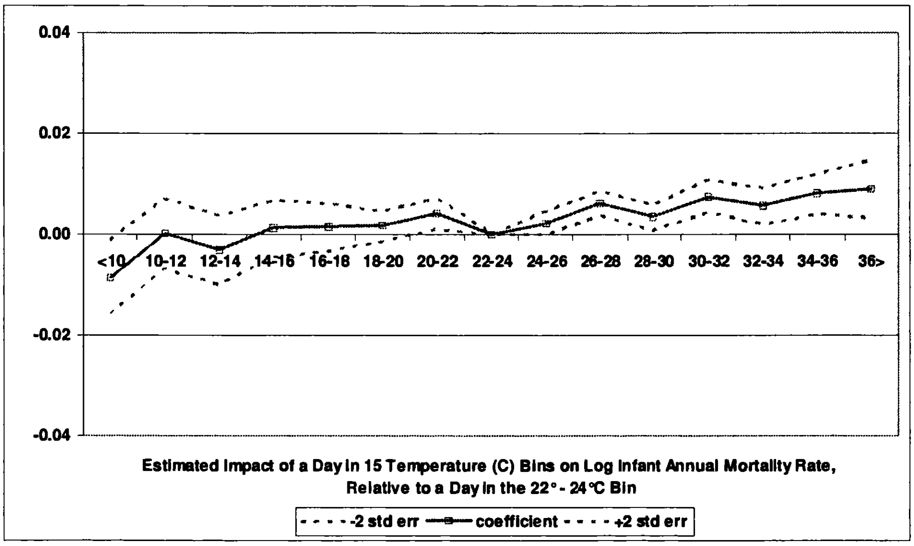

Figure 3.8: The Effect of Daily Temperature Exposure on the Annual Infant Mortality Rate (Rural Areas Only): Notes: Exposure window defined based on calendar year. The dependent variable is the infant (under age one) mortality rate in rural areas as calculated using VSI data. The model also includes controls for total precipitation and its square, and controls for unrestricted year effects, quadratic region-times-year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

with the temperatures bins below the reference category are all smaller in magnitude and not statistically different from zero.

Figure 3.7 shows the response function estimated from the urban population. The results are remarkably different from those in the rural areas. The largest coefficient is for the highest temperature bin (> 36° C), and the magnitude is only 0.003—an order of magnitude smaller than its rural equivalent of 0.012. Further, it is notable that to none of the other temperature effects are statistically significant, and all are relatively small in magnitude.

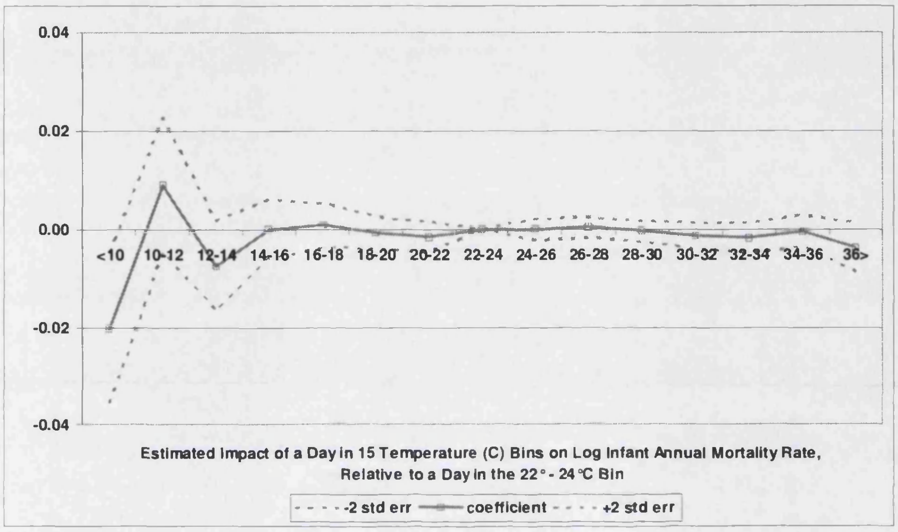

Figure 3.9: The Effect of Daily Temperature Exposure on the Annual Infant Mortality Rate (Urban Areas Only): Notes: Exposure window defined based on calendar year. The dependent variable is the infant (under age one) mortality rate in urban areas as calculated using VSI data. The model also includes controls for total precipitation and its square, and controls for unrestricted year effects, quadratic region-times-year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

Differential Impacts Across Rural and Urban Sectors

Figures 3.8 and 3.9 illustrate that the same is true for infant death rates—and that, as for the all-India results, rural infants appear to be more vulnerable than adults. It is remarkable that even India's urban infants, a group that is widely though to be a fragile population and that is the concern of an enormous public health literature, are seemingly immune to temperature extremes. As such the estimates of the response function in urban areas, both for adults and infants, suggest either that urban citizens are better positioned to adapt to temperature shocks, or perhaps also to a weaker connection between extreme temperatures and economic welfare and consumption levels in these areas.

As a corroboration of these findings I again check their robustness to using an alternative source of mortality data, that from the DHS. Figures 3.10 and 3.11 plot the 15 temperature death coefficients (obtained using DHS data) separately for rural and urban infants, in the same manner as in Figures 3.8 and 3.9 (which used VSI data to construct infant death rates). It is notable that even in the DHS data source, there is no significant (in magnitude or statistically) effect of extreme temperatures on death in urban areas, but a large and statistically significant effect in rural areas. That is, the distinct results in rural and urban areas does not appear to be driven by features of the manner in which the VSI data were collected.

Finally, Table 3.3 collects results relating to the effect of precipitation on death (while controlling for temperature) in rural and urban India separately. Like the all-India counterparts in Table 3.2, there is rarely evidence of statistically significant effects of rainfall on mortality in either rural or urban areas. One exception is among urban adults where both precipitation and its square are individually statistically distinguishable

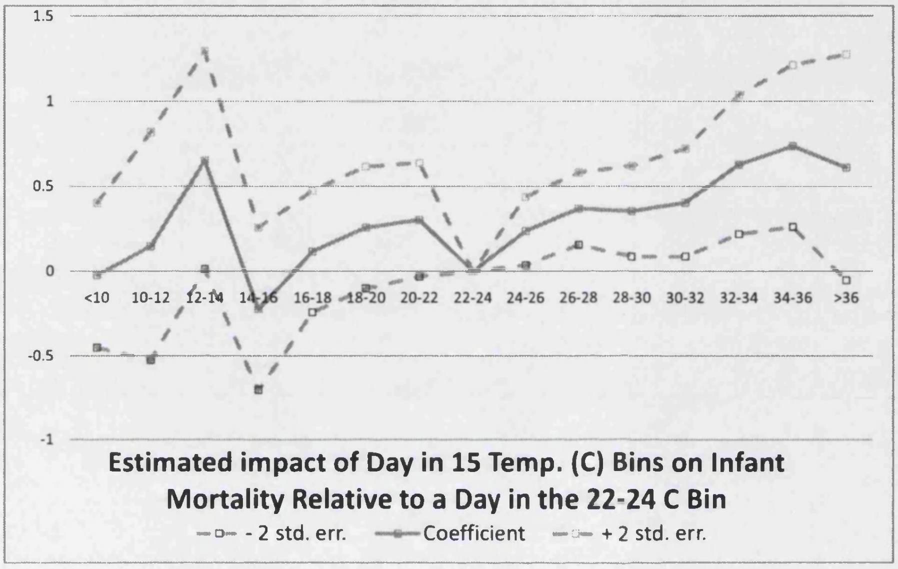

Figure 3.10: The Effect of Daily Temperature Exposure on the Annual Infant Mortality Rate (Rural Areas Only) using DHS data: Notes: Exposure window defined based on calendar year. The dependent variable is the infant (under age one) mortality rate in rural areas as calculated using DHS data. The model also includes controls for total precipitation and its square, and controls for unrestricted year effects, quadratic region-times-year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

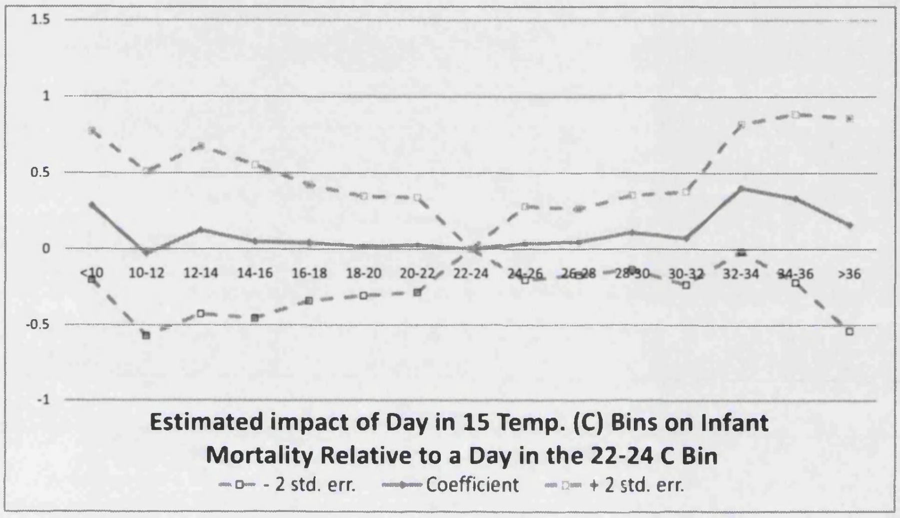

Figure 3.11: The Effect of Daily Temperature Exposure on the Annual Infant Mortality Rate (Urban Areas Only) Using DHS data: Notes: Exposure window defined based on calendar year. The dependent variable is the infant (under age one) mortality rate in urban areas as calculated using DHS data. The model also includes controls for total precipitation and its square, and controls for unrestricted year effects, quadratic region-times-year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

Dependent variable:	Log Total Mortality Rate		Log Infant Mortality Rate		Total Infant Mortality Rate	
	Rural	Urban	Rural	Urban	Rural	Urban
	(1)	(2)	(3)	(4)	(5)	(6)
Total annual precipitation	0.00101	0.00206	0.00156	0.00254	-0.0245	0.0418
	(0.00091)	(0.00078)	(0.00121)	(0.00110)	(0.0782)	(0.0923)
Total annual precipitation, squared [x 1000]	-0.0026	-0.00831	-0.0038	-0.00980		
	(0.0037)	(0.00309)	(0.0039)	(0.00497)		
Total annual precipitation, squared					0.00002	-0.00043
					(0.00008)	(0.0014)
Mortality data source	VSI	VSI	VSI	VSI	DHS	DHS
Number of observations	10.869	10,941	10.869	10,941	12.769	10,854
R-squared	0.65	0.66	0.55	0.55	0.17	0.17

Table 3.3: The Effect of Precipitation on Mortality in India, by Rural and Urban Areas: Notes: Exposure window defined based on calendar year. The mortality data source 'VSI' refers to data in the *Vital Statistics of India*, while that referred to as 'DHS' refers to the data in India's Demographic and Health Surveys. The model also controls for unrestricted year effects, cubic polynomial region*year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

from zero at conventional levels. But the non-robustness of this result among urban infants (whether measured using the VSI or DHS data) suggest that these results be interpreted with caution.

To summarize, the results in this section demonstrate that those in the previous section—which referred to all-India averages—masked a striking heterogeneity between rural and urban India. In rural areas, ambient temperatures play an important role in determining the starkest aspect of health, the probability of dying. But in urban areas of India, this effect is largely absent, even among presumably vulnerable children under the age of one. That is, even though rural and urban residents experience the same weather extremes, these extremes have a dramatically different effect on these two populations.

The Timing of Weather and Death: Growing vs Non-Growing Seasons

Our analysis so far has documented a strong effect of a given years's weather on a given year's death rate. But it is natural to expect the effect of weather on mortality to differ according to the seasons. In particular, as argued in Section 3.2, if the weather causes mortality in rural areas because it harms rural residents' agricultural incomes, then it is weather during the agricultural growing seasons that should matter for death in rural areas while weather during non-growing seasons should be irrelevant for determining rural mortality. However, under this hypothesis, a simple but important falsification exercise is to compare the effects of growing and non-growing season weather in urban areas; any difference would suggest that something other an an agricultural income mechanism may be at work.

To test this hypothesis I take a parsimonious approach to determining the 'growing' and 'non-growing' seasons of Indian agriculture. The agricultural calendar in India is driven by the arrival of the southwest monsoon rains, on roughly June 1st of every year. Since over 90 % of annual rainfall arrives in June, July and August, June 1st defines the traditional start of the agricultural calendar. Prior to the monsoon's arrival, there is little water with which seeds can germinate. Two cropping seasons are possible in most of India, the first (the *kharif*) running from approximately June to October, and the second (the *rabi*) running from approximately November to February. By March 1st, therefore, most crops in most regions have been harvested. Based on this simple characterization of the agricultural calendar, I therefore define June 1st - February 28th as the growing season and March 1st-May 31st as the non-growing season.

Using this definition, Table 3.4 presents results that demonstrate the differential effects of weather on death at these two distinct times of the year. Because this split (and further splits used below) of the data entail a loss of precision, I move to the single-index specification introduced in equation (3.2); as discussed above, this has the great advantage of estimating only one temperature coefficient rather than 15 coefficients while still capturing the essential features of non-linearity evident from the 15 coefficient estimates in Figure 3.2. Column 1 of Table 3.4 reports the simple single-index temperature coefficient estimate relating to the entire calendar year. This specification captures effects of temperature on death through the variable 'cumulative degree-days over 32° C', and continues to incorporate the effects of precipitation through a simple measure of total annual precipitation and its square. Unsurprisingly, given the pattern of temperature coefficients in Figure 3.3, the number of degree-days exceeding the 32° C threshold in a year has a strong effect

Dependent variable: log total mortality rate	All India		Rural	Urban	Rural	Urban
	(1)	(2)	(3)	(4)	(5)	(6)
Cumulative degree-days over 32 C	0.011		0.013	0.0004		
	(0.0042)		(0.0030)	(0.0002)		
Total annual precipitation	0.0013		0.0010	0.002		
	(0.0009)		(0.0009)	(0.0007)		
Total annual precipitation, squared [x 1000]	-0.0034		-0.0026	-0.0083		
	(0.0029)		(0.0037)	(0.0030)		
Cumulative degree-days over 32 C, in growing season		0.0022			0.0026	0.0008
		(0.004)			(0.0004)	(0.0004)
Cumulative degree-days over 32 C, in non-growing season		-0.00003			-0.00009	0.00007
		(0.0001)			(0.0004)	(0.0003)
Total annual precipitation, in growing season		0.0009			0.0005	0.0022
		(0.0007)			(0.0008)	(0.0008)
Total annual precipitation, in non-growing season		0.0023			0.0011	0.0070
		(0.0029)			(0.0034)	(0.0035)
Total annual precipitation squared, in growing season [x 1000]		0.0001			0.0008	-0.0114
		(0.0025)			(0.0038)	(0.0039)
Total annual precipitation squared, in non-growing season [x 1000]		-0.184			-0.171	-0.271
		(0.105)			(0.0789)	(0.110)
Number of observations	10,869	10,869	10,869	10,941	10,869	10,941
R-squared	0.63	0.63	0.66	0.66	0.66	0.66

Table 3.4: The Effects of Temperature and Precipitation on Log Mortality in India, by Agricultural Season: Notes: Mortality data are drawn from the *Vital Statistics of India*. The model also controls for unrestricted year effects, cubic polynomial region*year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

on the death rate in that year. This is consistent with the existing public health literature, as well as with the findings in Deschnes and Moretti (2009) and Deschenes and Greenstone (2008), that the ill effects of high temperatures on human health are particularly concentrated in the range over 32° C. Finally, turning to the results in column 1 relating to the effects of annual rainfall on death, it is reassuring that the coefficients estimated here are largely invariant to the manner in which temperature is included. That is, the coefficient estimated here is similar in magnitude to that in column 1 of Table 3.2.

I now go on to investigate the differential effects of weather on death throughout the agricultural year. Column 2 of Table 3.4 estimates the effect of weather on death in the growing and non-growing seasons separately, for all of India. It is clear that an equivalent weather shock (whether in the form of high temperatures or insufficient rainfall) has very different effects at different times of the agricultural cycle. Temperature shocks during the growing season have large and statistically significant effects on the death rate, but in contrast, equivalent shocks during the non-growing months of the agricultural cycle have small and statistically insignificant effects on death. It is important to note that, if anything, the most extreme temperatures (both daily average temperatures and the number of degree-days over 32° C) in an average year in India occur late in the *non*-growing season. So the particularly weak (and statistically insignificant) temperature effects seen in the non-growing season are not being driven by an absence of extremely hot days in the non-growing season.

To probe this result further, columns 3 through 6 of Table 3.4 repeat the specifications of columns 1 and 2 respectively, but by urban and rural areas separately. The resulting temperature coefficient estimates convey a stark contrast between rural and urban areas: in urban areas, neither growing nor non-growing season temperature extremes have any significant bearing on mortality; but in rural areas, growing season temperature is an important determinant of the mortality rates while non-growing season temperature is largely irrelevant for determining the rural death rate.

In summary, the results in Table 3.4 paint a coherent picture about the temperature-death relationship in India. This relationship is strong on adults, stronger on infants, confined (even among infants) to the rural population, and within rural areas (even among rural infants), to weather shocks that occur when crops are in the soil. This is precisely the pattern of results that would emerge naturally from an agricultural income mechanism relating weather to death, as discussed in section 3.2. Further, it is difficult to interpret these results in the context of 'direct' human physiological channels, since these would be expected to operate both in urban and rural areas that face the same shocks, on urban infants (who a large public health literature has identified as particularly vulnerable on average), and among rural adults and infants both during the growing and non-growing seasons of the calendar year. Yet none of these predictions of the 'direct' channel appear to be borne out in the data. The cluster of results presented thus far therefore provide supportive reduced-form evidence for an agricultural income channel relating weather shocks to death. I now go on to explore this potential channel further by looking directly at the extent to which agricultural incomes are indeed affected by weather fluctuations in India.

3.4.3 Mechanisms: Weather and Agricultural Incomes

I now investigate the plausibility of a mechanism relating weather to death that works through the income and consumption stress of those who work in the agricultural sector. Section 3.2 outlined a set of predictions on agricutural variables (yields, prices and wages) that would be expected if weather extremes do harm agricultural incomes; I look at these variables in turn here. I also look at additional predictions from Section 3.2 that would be expected if income shocks lead to death in rural areas but not in urban areas, such as the lack of a weather-income relationship in urban areas.

Agricultural Yields

An agricultural income channel relating weather to death in India would begin, as discussed in section 3.2, with an effect of weather shocks on agricultural productivity. Using the data on agricultural yields introduced in Section 3.3, I estimate this weather-yields relationship in this section. I model temperature and precipitation in precisely the same manner as when estimating the weather-death relationship, as laid out in equations (3.1) and (3.2); that is, temperature is modeled using either 15 temperature bins or the single-index approach (ie cumulative degree-days over 32° C), and precipitation is modeled using the total annual amount and its square.[17]

[17] One small difference here, when compared to the death regressions, is my adjustment of the timing of the weather data when relating it to agricultural outcomes (that is, to yields in this section as well as to prices and wages in following sections). Our agricultural yield data are based on measures of the total amount of output produced during the agricultural year (defined as running from June 1st to May 31st). If the weather in 'year *t*' is to matter for agricultural output in 'year *t*', it is important to define 'year t' in the same way across both the weather and agricultural output data. In the agricultural regressions in this and following sections, I therefore re-label the years in the weather data so that weather on dates from January 1st to May 31st are

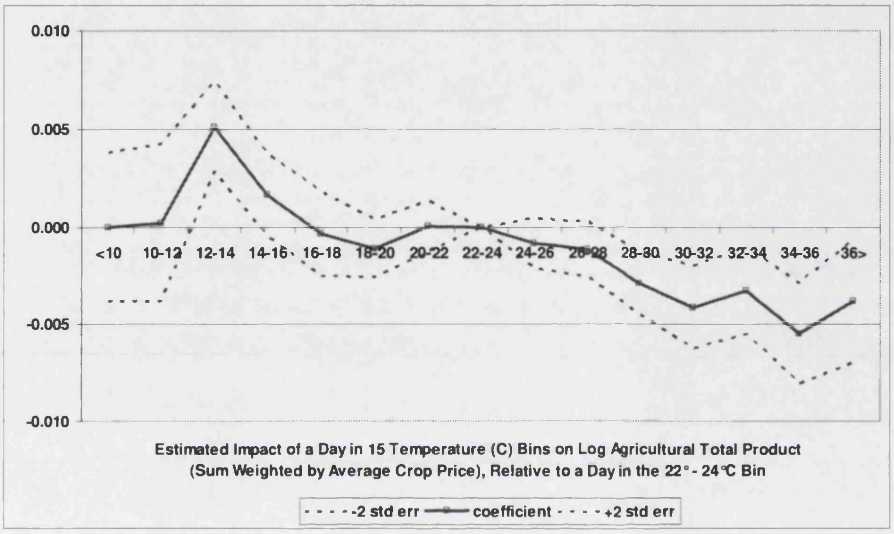

Figure 3.12: The Effect of Daily Temperature Exposure on Log Agricultural Yields: Exposure window defined on agricultural year, June 1st through May 31st. The model also includes controls for monthly total precipitation and controls for unrestricted year effects, cubic polynomial region-year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

Figure 3.12 plots the 15 temperature bin coefficients when these variables (as well as total annual precipitation) are regressed on agricultural yields. The pattern of coefficients that emerges is strikingly similar to that between temperature and death presented in Figure 3.3, only it is inverted because high temperatures damage plants and therefore reduce yields. High temperature days reduce agricultural yields significantly—for example, the coefficient estimate for days with mean temperatures exceeding 36° C implies that every *single* day in this category (relative to a day in the 22° - 24° C reference category) reduces agricultural yields by 0.04 %. And the point estimates on each of these temperature bin coefficients are individually statistically different from zero above the 28° C mark.

Table 3.5 reports (in column 1) the corresponding rainfall effect that was established in the same regression as the 15 coefficients reported in Figure 3.12. As expected, this coefficient is positive—more rainfall produces more output per acre—and statistically significant, but the effect of rainfall is diminishing, as indicated by the negative and statistically significant rainfall squared term.

Column 2 of Table 3.5 repeats the baseline analysis between weather and agricultural yields but this time using the single-index measure of temperature (the cumulative number of degree-days in the year that exceed 32° C). As implied by the less parametric results in Figure 3.12, this single-index measure captures the significant negative impact of high temperature days on agricultural yields. And notably, the coefficients on rainfall are quite stable between columns 1 and 2.

In column 3 I break up the calendar year, as in the previous section, into the growing and non-growing seasons within an agricultural cycle.

lagged by a year. Put another way, when estimating equations (3.1) and (3.2) here, the year t is defined as the 365 days beginning on June 1st of any given calendar year.

Dependent variable:	Log agricultural output per acre			Log agricultural price index		
	(1)	(2)	(3)	(4)	(5)	(6)
Cumulative degree-days over 32 C		-0.00068			0.00014	
		(0.00033)			(0.00007)	
Total annual precipitation	0.00984	0.00929		-0.00142	-0.00161	
	(0.00211)	(0.00160)		(0.0006)	(0.0003)	
Total annual precipitation, squared [x 1000]	-0.0418	-0.0347		0.00548	0.00617	
	(0.00692)	(0.00686)		(0.00211)	(0.00110)	
Cumulative degree-days over 32 C, in growing season			-0.00073			0.00016
			(0.00033)			(0.00008)
Cumulative degree-days over 32 C, in non-growing season			0.00049			-0.00014
			(0.00022)			(0.00006)
Total annual precipitation, in growing season			0.00948			-0.00166
			(0.00163)			(0.00027)
Total annual precipitation, in non-growing season			0.00115			0.00085
			(0.00129)			(0.00039)
Total annual precipitation squared, in growing season [x 1000]			-0.0353			0.00069
			(0.0066)			(0.00011)
Total annual precipitation squared, in non-growing season [x 1000]			-0.0488			-0.0224
			(0.0026)			(0.0094)
Control for 15 temperature bin effects	YES	NO	NO	YES	NO	NO
Number of observations	7,729	7,729	7,729	7,773	7,773	7,773
R-squared	0.87	0.86	0.87	0.77	0.76	0.76

Table 3.5: The Effects of Temperature and Precipitation on Log Agricultural Productivity and Prices in India, by Agricultural Season: Notes: The model also controls for unrestricted year effects, cubic polynomial region*year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by rural census population. See the text for more details.

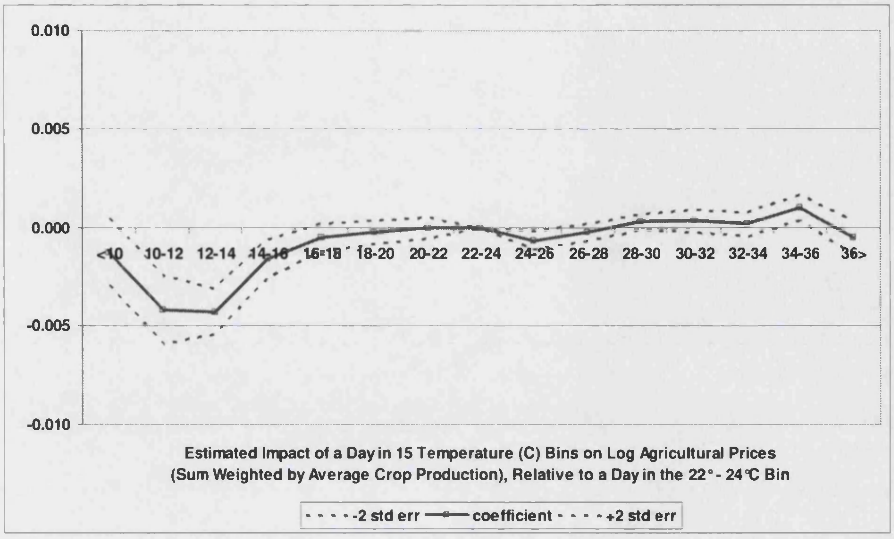

Figure 3.13: The Effect of Daily Temperature Exposure on Log Agricultural Prices: Exposure window defined on agricultural year, June 1st through May 31st. The model also includes controls for monthly total precipitation and controls for unrestricted year effects, cubic polynomial region-year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

As was the case for my mortality estimates in Table 3.4, there is a strong effect of temperature on agricultural yields during the agricultural growing season but, reassuringly, no effect during the non-growing season (the three months of the year during which virtually no agricultural activity takes place).

Agricultural Prices

A natural implication of reduced agricultural yields in a locality—in any type of trading regime shy of the small open economy limit—is a rise in that locality's agricultural prices. In this section I investigate the magnitude of this agricultural price response to weather shocks using the price data discussed in Section 3.3. I present these results in an analogous fashion to the previous results relating to agricultural yields. That is, Figure 3.13 begins by reporting the coefficients on 15 temperature bins. A familiar pattern emerges, which is broadly the inverse of that in Figure 3.12 relating to agricultural yields, with hot days driving up agricultural prices. Notably, the point estimates I obtain here, from the hottest days, are smaller than those on agricultural yields or mortality. This is a natural result of either price- and income-inelastic demand for these essential agricultural food items, or of the presence of trade flows between regions that act to smooth the effects of a local productivity shock on local prices. Both are plausible explanations, but it is beyond the scope of this chapter to disentangle them. Columns 4 through 6 of Table 3.5 repeat the specification in columns 1 through 3 of the same table, but using agricultural prices as the dependent variable rather than agricultural yields. Again, a familiar pattern is on display (but with coefficient estimates of opposite signs to columns 1-3): rainfall shortages raise agricultural prices in a diminishing fashion (column 4); the effect of temperature on price is robust to the way in which it is modeled (column 5); and the effect of weather (both temperature and precipitation) on agricultural prices is largely confined to the agricultural growing season with no strong effects on prices from shocks during the non-growing season.

Agricultural Wages

I documented in Section 3.4.3 that weather fluctuations had strong and statistically significant effects on agricultural yields. It would would be natural to expect that when agricultural productivity falls, so too do real wages in the agricultural sector. I therefore pursue an analogous empirical specification on agricultural real wages to that on agricultural yields estimated above (in Figure 3.12 and columns 1-3 of Table 3.5). Figure

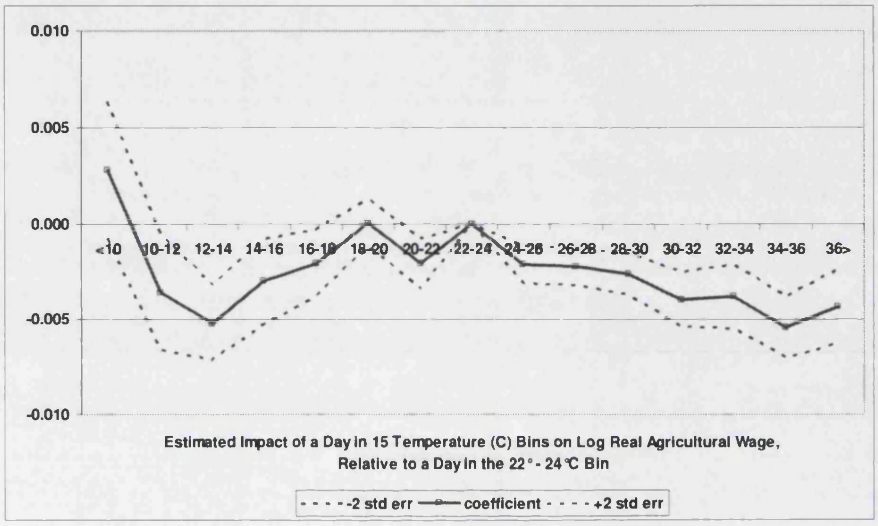

Figure 3.14: The Effect of Daily Temperature Exposure on Log Real Agricultural Wages: Exposure window defined on agricultural year, June 1st through May 31st. The model also includes controls for monthly total precipitation and controls for unrestricted year effects, cubic polynomial region-year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

3.14 plots the 15 temperature bin coefficients that result from doing this, and a now-familiar pattern continues to hold. Once again, there is a strong effect of weather on the agricultural economy, measured this time with real agricultural wages. Finally, Table 3.6 (columns 1-3) repeats the empirical specifications of Table 3.5 (columns 1-3), using agricultural wages as the dependent variable rather than agricultural. Reassuringly, the effect of weather on the measure of agricultural productivity used here, real wages, lines up with a second measure of agricultural productivity used earlier, yields (ie real output per acre). That is, hot years drive down agricultural wages, and rainfall drives up wages but in a diminishing manner—but these effects are only true for weather shocks during the growing season.

Agricultural Labor Supply

Agricultural workers—who count themselves among the poorest segments of India's rural society—care of course about both their hourly wage (studied in the previous section) and their total *wage bill.* In this section I estimate the extent to which agricultural labor supply adjusts in the face of an agricultural productivity shock, in order to say something about the effect of these weather shocks on the wage bill. I are circumscribed in my ability to do so, however, by the lack of data on total hours of supplied labor in each district and year. As discussed in Section 3.3 above, the best available data in this regard is a measure of the number of people who report their primary occupation as 'agricultural worker' (defined as anyone whose primary job is to earn an agricultural wage, rather than an owner-cultivator or sharecropper) in the year in question. This of course misses labor supply adjustment on the margins of days supplied per year and hours supplied per day. Further, this measure is

Dependent variable: log real agricultural wage	(1)	(2)	(3)
Cumulative degree-days over 32 C		-0.00032 (0.00015)	
Total annual precipitation	0.00152 (0.00055)	0.00160 (0.00047)	
Total annual precipitation, squared [x 1000]	-0.00581 (0.00247)	-0.00643 (0.00191)	
Cumulative degree-days over 32 C, in growing season			0.00036 (0.00015)
Cumulative degree-days over 32 C, in non-growing season			0.00014 (0.00014)
Total annual precipitation, in growing season			0.00169 (0.00049)
Total annual precipitation, in non-growing season			-0.00174 (0.00108)
Total annual precipitation squared, in growing season [x 1000]			-0.00701 (0.00190)
Total annual precipitation squared, in non-growing season [x 1000]			0.00658 (0.03150)
Control for 15 temperature bin effects	YES	NO	NO
Number of observations	7,733	7,733	7,733
R-squared	0.78	0.77	0.77

Table 3.6: The Effects of Temperature and Precipitation on Log Real Agricultural Wages in India, by Agricultural Season: Notes: The model also controls for unrestricted year effects, cubic polynomial region*year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by rural census population. See the text for more details.

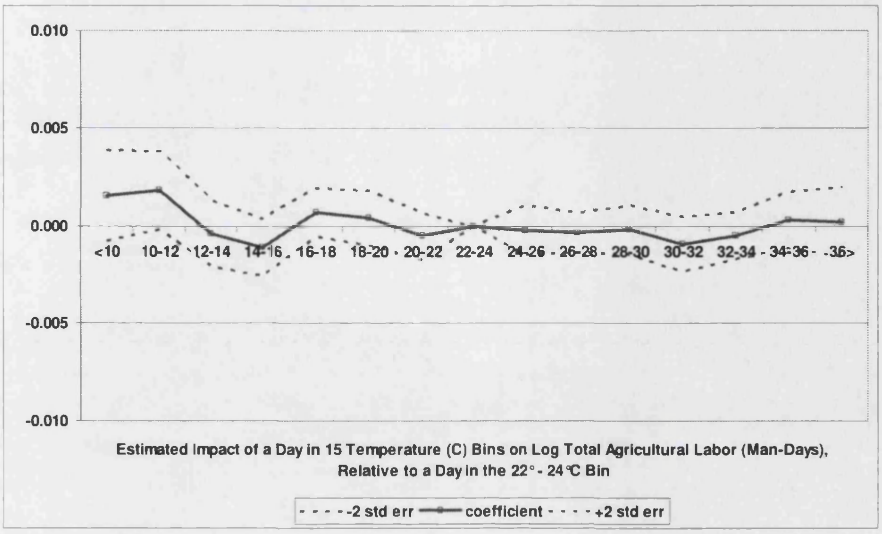

Figure 3.15: The Effect of Daily Temperature Exposure on Agricultural Labor Supply: Labor supply is measured as the number of workers who report their primary occupation to be an agricultural laborer; this measure if only available decadally, from the Census of India. Exposure window defined on agricultural year, June 1st through May 31st. The model also includes controls for monthly total precipitation and controls for unrestricted year effects, cubic polynomial region-year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

only available decadally, and it would be inappropriate to interpolate this measure between decadal census year. This is a shortcoming that will affect the precision of estimates based on this variable but which should leave my coefficient point estimates unbiased and appropriately comparable (bearing in mind the above caveats about missing margins of adjustment) with my estimates in earlier sections that used annual data.

Figure 3.15 illustrates my estimated temperature bin coefficients when using this decadal agricultural employment variable as a measure of labor supply. As expected, the standard error bounds on the 15 coefficient estimates are considerably wider than those in the preceding figures, which drew on annual outcome variables. However, the key finding in Figure

3.15 is that the 15 point estimates here are all close to zero, implying very little effect of high temperatures on the labor supply adjustment margin of a worker's occupational choice between agriculture and non-agriculture (or not working at all). It is important to stress, however, that I can say nothing about the extent of the labor supply response to weather shocks along other margins of adjustment.[18]

One interpretation of the findings in Figures 3.12, 3.14 and 3.15 is that the labor supply curve in the agricultural sector is essentially vertical and so a change in labor demand induced by temperature and precipitation shocks leads to an adjustment only on the wage margin and not on the labor supply. As such increased exposure to extreme weather conditions is likely to cause agricultural income to decline substantially. This is in line with the findings of Jayachandran (2006) who argues that structural features of developing countries like India (such as subsistence living standards, credit constraints and high migration costs) give rise to low elasticities of labor supply in equilibrium.

Urban Wages and Incomes

I have documented in Sections 3.4.3 and 3.4.3 that adverse weather shocks in the agricultural growing seasons (but not those in non-growing seasons) play a significant role in driving down agricultural productivity (as measured by yields or wages). These results suggest that the incomes of rural residents, those who draw their livelihoods from agricultural activities, are likely to suffer in periods of adverse weather. I now investigate

[18]I have also investigated the effects of the single-index temperature index, as well as precipitation (and its square), on agricultural labor supply. In regressions that parallel columns 1-3 of Table 3.6, but which use agricultural labor supply rather than real agricultural wages as the dependent variable, there is no evidence for any economically significant coefficient estimates relating temperature or precipitation to my measure of labor supply. (Naturally, the standard errors in these regressions are large relative to those in Table 3.6; but the key point is that point estimates are all uniformly small relative to those in Table 3.6.)

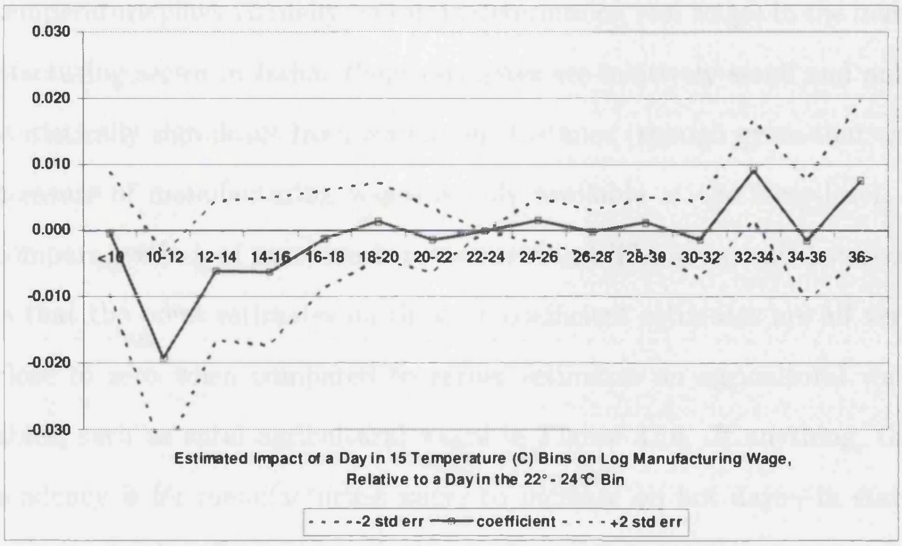

Figure 3.16: The Effect of Daily Temperature Exposure on Real Manufacturing Wages: The real manufacturing wage is measured as the total annual earnings of laborers in the registered manufacturing sector in India, deflated by an urban-specific CPI. This measure is only available at the state-level. The model also includes controls for monthly total precipitation and controls for unrestricted year effects, cubic polynomial region-year trends and unrestricted district effects. Standard errors are clustered by district. Regressions weighted by census population. See the text for more details.

whether there exist similar effects of weather on productivity or income among *urban* residents. As discussed in Section 3.3, there are no systematic district-level data on a representative 'urban wage'. I therefore focus on wages in manufacturing, an activity that is almost entirely conducted in urban areas in India. In particular, I use the best available annual data on real manufacturing wages, which is presented at the state level.

Figure 3.16 plots the 15 temperature bin coefficient estimates that I obtain when regression state-level real manufacturing wages on state-level weather variables.[19] The resulting coefficient estimates suggest that

[19] Our measures of state-level weather are population-weighted analogues of their district-level counterparts.

temperature plays virtually no role in determining real wages in the manufacturing sector in India. Point estimates are relatively small and only statistically significant from zero in one instance (though given that the measure of manufacturing wages is only available at the state-level, a comparative lack of precision is not surprising.) The key point, however, is that the point estimates on these 15 coefficient estimates are all very close to zero when compared to earlier estimates on agricultural variables, such as rural agricultural wages in Figure 3.14. If anything, the tendency is for manufacturing wages to *increase* on hot days—in stark contrast to the effects shown earlier in Figure 3.14 relating to wages in the agricultural sector.

This finding is important for several reasons. First, it increases my confidence in the fact that my earlier agricultural wage results were not spuriously driven by some feature of the weather data. For example, a natural concern might be that my region-specific cubic polynomial time trends are not controlling for enough of the background trends in temperature that might be spuriously correlated with trending economic variables such as rural or urban wage rates; to the extent that these same spurious effects would be expected in urban areas, this does not seem to be the case. Second, this finding of a significant rural wage-weather relationship but no urban wage-weather relationship suggests that these are two separate labor markets, at least in the short run (though it is also possible that the agricultural labor market employs entirely different types of workers than the manufacturing labor market). And third, this finding supports the hypothesis that rural residents—but not urban residents—die in the face of weather shocks that harm agricultural output because their real incomes (and consumption levels) suffer, unlike urban residents whose incomes are not determined by the weather.

3.5 Implications of Climate Change

Our results in Section 3.4.2 above suggest that weather extremes, in the form of hot or dry years, have strong effects on mortality in rural areas. Likewise, my results in Section 3.4.3 above suggest that these same weather extremes leave a remarkably similar pattern of results on markers of economic welfare among the rural population, such as agricultural yields, real agricultural wages, and agricultural food prices. I believe that both of these sets of results are important in their own right as they suggest that weather fluctuations may matter a great deal for poor citizens in developing countries. However, in an era when climatologists are increasingly confident that the world's climate is changing and will continue to change, my estimates of the weather-death relationship obtained above can also be used to provide—with considerable caution, as I stress below—upper-bound estimates of some of the health costs of this predicted climate change.

To shed light on this I have obtained data on the climate change in India's climate that emerges from two leading global circulation models (GCMs), the models that climatologists use to make predictions about likely climate change scenarios. I refer to these models as 'Hadley 3' (the preferred model in use by the Hadley Centre, which provided all climate change predictions in the influential Stern Review), and 'CCSM 3' (the preferred model in use by the National Center for Atmospheric Research). These models were used in the most recent Intergovernmental Panel on Climate Change (IPCC) report. In the Data Appendix I describe the construction of these models in more detail.

These models make predictions about the evolution of daily weather at finely spaced gridpoints all over the world, on every day for the next 100 years. I use these predictions (averaged over hundreds of simulations of the models) to construct a set of temperature predictions, one for each of the two GCM models, for each of India's districts using a procedure detailed in the Data Appendix. In particular, in order to align the predicted climatic variation with the inter-annual climatic variation I used in Section 3.4.2 to estimate the within-sample weather-death relationship, I use the GCM models to predict the average number of days in which the mean temperature will fall into each of the 15 temperature bins between the years 2070 and 2099. (I choose to average over 30 years of predicted values in order to smooth out prediction error in these climate models). This generates a variable that I denote $TMEAN_{dj}^{2070-2099}$, the climate change model's prediction for the number of days on which the mean temperature in district d will fall into temperature bin j on average over the 2070-2099 period.

These models also make predictions about changes to the distribution of rainfall in India. Because there is less agreement in the climatological literature about these changes (particularly in India, where the complex dynamics of the monsoon are not well understood), however, I focus my analysis on the increase in temperatures.

Before proceeding, it is important to underscore that the validity of the chapter's estimates of the impacts of climate change depend on the validity of the climate change predictions. The state of climate modeling has advanced dramatically over the last several years, but there is still much to learn, especially about the role of greenhouse gases on climate (Karl and Trenberth 2003). Thus, the Hadley 3 A1FI and CCSM 3 A2 predictions should be conceived of as two realizations from a superpopulation of models and scenarios. The sources of uncertainty in these models and scenarios are unclear, so it cannot readily be incorporated into the below estimates of the impacts of climate change. Nevertheless,

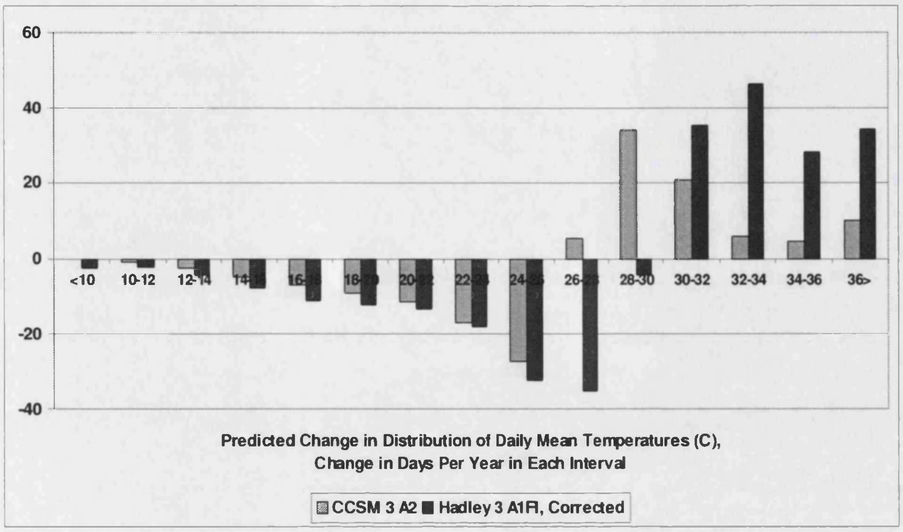

Figure 3.17: Predicted Change in Distribution of Daily Temperatures in India to 2070-2099: Temperature distributions (daily averages) for two leading global circulation models under an increase in greenhouse gas emissions and a 'business as usual' scenario. Mean temperatures are weighted by average district population between 1957 and 2000. See the text for more details.

the use of two sets of daily business as usual climate change predictions provides some sense of the variation.

Figure 3.17 provides an opportunity to understand how climate change is expected to change the full distributions of daily mean temperatures in India. In this figure I compare the predicted distribution of daily mean temperatures (across the 15 temperature bins), ie $TMEAN_{dj}^{2070-2099}$ averaged over districts d, with the actual historical average equivalent over the observed period used in this chapter (ie 1957-2000). I denote the historical average number of days in which the mean temperature in district d fell into temperature bin j between 1957 and 2000 by $TMEAN_{dj}^{1957-2000}$.

Figure 3.2, discussed earlier, plotted the distribution of daily temperatures into 15 temperature bins $TMEAN_{dj}^{1957-2000}$ averaged over all districts d. In Figure 3.17 I therefore plot the predicted *change* in the average daily temperature distribution going out to 2070-2099, ie $\Delta TMEAN_{dj} \equiv TMEAN_{dj}^{2070-2099} - TMEAN_{dj}^{1957-2000}$, averaged over districts d. Since the two GCMs make different predictions about this distributional change I plot both of their predictions in Figure XX. The resulting plot reveals that there will be large reductions in the number of days in the 14° to 28° C range. These reductions are predicted to be offset by increases in days with temperatures exceeding 28° C. Thus, the mortality impacts of climate change rest on the impact of the days in the 14° to 28° C range, relative to days at higher temperatures. Due to India's already warm climate, it is unlikely to get much benefit from reductions in the number of days in its left tail of the temperature distribution, which stands in stark contrast to Russia and other relatively cold countries. That is, under both predicted climate change scenarios, India will exchange days that I estimated (in Figure 3.3) to be relatively low mortality days for days that I have estimated (in Figure 3.3) to be high mortality ones.

I now turn to a more precise calculation of the predicted mortality impacts of climate change on India. Table 3.7 presents estimates based on the estimation of equation (3.1) for the various subsamples. The predictions are based on the Hadley 3 A1FI and CCSM 3 A2 models, and pertain to the years 2070-2099. The predicted impact in district d is based on district-level predictions calculated as:

$$\Delta \widehat{Y_{dt}} = \sum_j \widehat{\theta}_j \Delta TMEAN_{dj}, \tag{3.3}$$

where $\Delta\widehat{Y_{dt}}$ is the predicted change in the log mortality rate, and $\widehat{\theta}_j$ is the estimated coefficient on temperature bin j obtained in Section 3.4.2, $\Delta TMEAN_{dj}$ is the predicted (according to the Hadley 3 or CCSM 3 A2 model) change in the number of days on which the mean temperature will fall into temperature bin j by 2070-2099. This is the predicted impact of climate change, according to these models and an extrapolation of my estimated weather-death relationship, in district d. In order to construct a meaningful total impact for all of India, I report the population-weighted average of each district d's predicted impact. The standard error of this prediction is calculated accordingly.

Columns 1-3 of Table 3.7 summarize this calculation for three daily mean temperature categories, those for $< 16°$ C, $16° - 32°$ C, and $> 32°$ C respectively. Column 4 then reports the total temperature impact obtained by summing the impacts in columns 1-3. Finally, the rows of Table 3.7 correspond to different statistical models and different climate change models.

For each climate change model, I calculate the predicted percentage change in annual mortality for rural areas, urban areas, and India as a whole. All models are based on the pooled age specification. The top panel reports the Hadley 3 A1FI results and suggests that climate change would lead to a 40.8% increase in the annual mortality rate in India. These estimates are precise and importantly the null hypothesis of a zero effect is rejected at conventional significance levels. Examination of column 3 shows that the increased mortality is entirely attributable to the increase in the number of very hot days (where the mean temperature exceeds 32° C).

The next columns break down the analysis by rural/urban area. As before, the results are sharply different for urban and rural areas. For rural areas, annual mortality rates are predicted to increase by 55.1% and this estimate is precise, with robust t-statistics in excess of 3. Again, the increased mortality is almost entirely attributable to the increase in

	Impact of Change in Days with Temperature:			Total Temperature Impact
	< 16 C	16C - 32 C	>32 C	
	(1)	(2)	(3)	(4)
A. Based on Hadley 3, A1F1				
Pooled	-0.011	-0.144	0.671	0.516
	(0.032)	(0.047)	(0.126)	(0.127)
Rural Areas	-0.029	-0.168	0.846	0.648
	(0.039)	(0.056)	(0.154)	(0.154)
Urban Areas	0.046	0.006	0.143	0.195
	(0.032)	(0.058)	(0.111)	(0.097)
B. Based on CCSM3, A2				
Pooled	-0.008	0.033	0.155	0.180
	(0.014)	(0.043)	(0.029)	(0.062)
Rural Areas	-0.014	0.064	0.195	0.245
	(0.017)	(0.050)	(0.035)	(0.073)
Urban Areas	0.012	0.036	0.041	0.089
	(0.014)	(0.039)	(0.023)	(0.050)

Table 3.7: The Predicted Impacts of Climate Change on Log Annual Mortality Rates in India, 2070-2099: Based on regression models that control for unrestricted year effects, region-specific cubic polynomials in time, and unrestricted district-area effects, and use coefficient estimates that are weighted by census population. Projections compare historical period average temperatures (averaged over 1957-2000) with the end of the century (averaged over 2070-2099). Reported numbers correspond to elements of equation (3.3), averaged over all districts (weighted by average population between 1957 and 2000). Standard deviations are based on regression standard errors that were are clustered by district. See text for more details.

the number of very hot days (where the mean temperature exceeds 32° C). Taken alone, the change in precipitation predicts a significant 13% decline in annual mortality rates in rural areas. Column 3, which focuses on urban areas tells a completely different story. The predicted change in annual mortality is 17.9%, and is not statistically distinguishable from zero at the conventional level.

The lower panel shows the results derived from the CCSM 3 A2 model as opposed to the Hadley model. The predicted increases in annual mortality are smaller than those from the Hadley model in Panel A, but still large and concentrated in the rural areas, ranging from 7.3% to 17.7%. The discrepancy between the Hadley and CCSM predictions reflects in part the fact that the A1FI scenario is associated with larger increases in temperature than the A2 scenario. The overall CCSM impacts are marginally significant, but like in Panel A, it is clear that the increase in annual mortality is caused by the predicted increase in exposure to extreme temperatures. It is noteworthy that the segment of the temperature distribution that is predicted to increase the most (days above 16° C) is associated with large and significant increase in annual mortality rates.

These results suggest that the health costs of predicted climate change in India could be severe—when standard models of climate change are used in combination with my estimates of the weather-death relationship, these models predict strikingly large rises in the death rate in India by 2080. Because my focus has been on mortality rather than on morbidity, the effects of weather on wider health indicators in India are likely to be understated by my estimates. Further, an analogous exercise based on my estimates of the weather-agricultural income relationship would suggest that predicted climate change in India could be economically harmful as well as damaging to health.

However, it is important to bear two caveats in mind when interpreting these findings. First, I have estimated the effect of weather on death using inter-annual variation, so my estimates are best thought of as short-run impacts. As such they are likely to provide only an upper-bound to the impact of long-run climate change of the sort predicted by standard climatological models. This is because individuals are likely to be better able to adapt to long-run change, for example through migration, technology adoption, or occupational change away from climate-exposed industries such as agriculture. Second, the climatological models whose climate change predictions I have used here do not incorporate any possibility of catastrophic change in India's climate as a result of a rise in greenhouse gas emissions. That is, while some climatological models predict that modest rises in temperatures may have catastrophic knock-on effects (eg rises in ocean temperature, widespread melting of Himalayan glaciers, reversal of trade winds, or cessation of the Southwest monsoon), I have deliberately obtained my climate predictions from climatological models that ignore these catastrophic, but highly uncertain and controversial effects.

3.6 Conclusion

This study has produced the first large-scale study of the impact of weather shocks on mortality and adaptations for a developing country of which I am aware. It is based on the finest geographical data available on mortality for India over the period 1957-2000, augmented with rich high-frequency data on historical daily weather realizations.

The results are striking and indicate a large and nonlinear relationship between daily temperatures and annual mortality rates. This is in stark contrast to a similar exercise performed on data from the United States, as shown in Figure 3.1. For example, I find that a *single* additional day with a mean temperature above 32° C, relative to a day with a mean temperature in the 22° - 24° C range, increases the annual mortality rate by roughly 0.5%. This effect is even larger in the rural regions of India where even now more than two thirds of the population lives—but the effect of weather on death is largely absent, even among infants, in urban India. Further, these differential results across rural and urban areas continue to hold when controlling for the wealth differences that distinguish rural and urban areas, and are entirely concentrated in the nine months of the year during which crops are in the soil in India.

One plausible explanation for this finding, which I pursue in detail, is that rural incomes are dependent on agricultural income—and hence the weather—in a way that urban incomes are not. I find a pattern of results that is uniformly consistent with this explanation, in the form of significant effects of growing season (but never non-growing season) weather on agricultural yields, prices and wages, but no effects of weather on manufacturing incomes.

Finally, the chapter uses the estimated response functions between temperature and mortality to provide some predictions on the impacts of climate change on mortality in India. The resulting estimates, while likely to be uppers bounds on the impact of a long-run change in India's climate, suggest that standard climatological models of climate change predict an increase in the overall Indian annual mortality rate of approximately 12% - 41% by the end of the century.

These mortality impacts are large and they are likely to be accompanied by income impacts (some of which I have estimated) and morbidity impacts (which I have been unable to investigate) too. These estimated impacts therefore imply a largely neglected role for health policy in developing countries such as India to assist citizens in mitigating the effects of weather variation on their well-being.

Bibliography

ACEMOGLU, D., S. JOHNSON, AND J. ROBINSON (2005): "The Rise of Europe: Atlantic Trade, Institutional Change, and Economic Growth," *American Economic Review*, 95(3), 546–579.

ADES, A. F., AND E. L. GLAESER (1999): "Evidence On Growth, Increasing Returns, and The Extent Of The Market," *Quarterly Journal of Economics*, 114(3), 1025–1045.

AGGARWAL, S. (1937): *The Salt Industry in India.* Delhi: Government of India Press.

AHUJA, R. K., T. L. MAGNANTI, AND J. B. ORLIN (1993): *Network Flows: Theory, Algorithms, and Applications.* New Jersey: Prentice Hall.

ALCALA, F., AND A. CICCONE (2004): "Trade and Productivity," *Quarterly Journal of Economics*, 119(2), 613–646.

ANDERSON, J. E. (1979): "A Theoretical Foundation for the Gravity Equation," *American Economic Review*, 69(1), 106–116.

ANDERSON, J. E., AND E. VAN WINCOOP (2003): "Gravity with Gravitas: a Solution to the Border Puzzle," *American Economic Review*, 93(1), 170–196.

——— (2004): "Trade Costs," *Journal of Economic Literature*, 42(3), 691–751.

ANDRABI, T., AND M. KUEHLWEIN (2005): "Railways and Market Integration in British India," mimeo Pomona.

ANDREW, W. P. (1883): *Indian Railways.* London: W. H. Allen & Co.

ARKOLAKIS, K., P. KLENOW, S. DEMIDOVA, AND A. RODRIGUEZ-CLARE (2008): "The Gains from Trade with Endogenous Variety," *American Economic Review Papers and Proceedings*, 98(4), 444–450.

ARNOLD, D. (1993): *Colonizing the Body: State Medicine and Epidemic Disease in 19th-century India.* Berkeley: University of California Press.

ASCHAUER, D. A. (1989): "Is Public Expenditure Productive?," *Journal of Monetery Economics*, 23(2), 177–200.

ATACK, J., M. R. HAINES, AND R. A. MARGO (2008): "Railroads and the Rise of the Factory: Evidence for the United States, 1850-1870," NBER working paper 14410.

BANERJEE, T. (1966): *Internal Market of India (1834-1900).* Calcutta: Academic Publishers.

BARTHOLOMEW, J. G. (1893): *Constable's Hand Atlas of India.* London: A. Constable.

BASU, R., AND J. M. SAMET (2002): "Relation Between Elevated Ambient Temperature and Mortality: A Review of the Epidemiologic Evidence," *Epidemiologic Review*, 24 (December), 190–202.

BAYLY, C. A. (1983): *Rulers, Townsmen and Bazaars: North Indian Society in the Age of British Expansion, 1770-1870.* Cambridge: Cambridge University Press.

BERNARD, A., J. EATON, J. B. JENSEN, AND S. KORTUM (2003): "Plants and Productivity in International Trade," *American Economic Review*, 93(4), 1268–1290.

BESLEY, T., AND R. BURGESS (2004): "Can Labor Regulation Hinder Economic Performance? Evidence From India," *Quarterly Journal of Economics*, 119(1), 91–134.

BHATIA, B. M. (1967): *Famines in India.* Delhi: Konark Publishers.

BLACK, W. R. (2003): *Transportation: A Geographical Analysis.* Guilford: The Guilford Press.

BLYN, G. (1966): *Agricultural Trends in India, 1891-1947: Output, Availability, and Productivity.* Philadelphia: University of Pennsylvania Press.

BOURNE, J. (1849): *Indian River Navigation.* London: W. H. Allen.

BRANDER, J., AND P. R. KRUGMAN (1983): "A "Reciprocal Dumping Model of International Trade"," *Journal of International Economics*, 15, 313–321.

BRODA, C., AND D. E. WEINSTEIN (2007): "Understanding International Price Differences Using Barcode Data," mimeo Chicago GSB.

CHANDLER, A. D. (1977): *The Visible Hand: The Managerial Revolution in American Business.* Cambridge: Harvard University Press.

CHANEY, T. (2008): "Distorted Gravity: the Intensive and Extensive Margins of International Trade," *American Economic Review*, 98(4), 1707–1721.

CHAUDHURI, K. N. (1983): "Chapter X: Foreign Trade and Balance of Payments (1757-1947)," in *The Cambridge Economic History of India, Volume 2*, ed. by D. Kumar. Cambridge: Cambridge University Press.

CHETTY, R. (2008): "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," NBER Working Paper 14399.

COSTINOT, A., AND I. KOMUNJER (2008): "What Goods Do Countries Trade? A Structural Ricardian Model," mimeo MIT.

DAVID, P. A. (1969): "Transport Innovation and Economic Growth: Professor Fogel on and off the Rails," *The Economic History Review*, 22(3), 506–524.

DAVIDSON, E. (1868): *The railways of India : with an account of their rise, progress and construction; written with the aid of the records of the India Office.* London: E. and F. Spon.

DAVIS, K. (1951): *The Population of India and Pakistan.* Princeton: Princeton University Press.

DAVIS, M. (2001): *Late Victorian Holocausts: El Nino Famines and the Making of the Third World.* New York: Verso.

DEARDORFF, A. (1998): "Determinants of Bilateral Trade: Does Gravity Work in a Neoclassical World?," in *The Regionalization of the World Economy*, ed. by J. A. Frankel. Chicago: University of Chicago Press (for NBER).

DEATON, A., AND J. DREZE (2009): "Food and Nutrition in India: Facts and Interpretations," *Economic and Political Weekly*, XLIV (7), 42–65.

DELOCHE, J. (1994): *Transport and Communications in India Prior to Steam Locomotion: Volume I: Land Transport (French Studies on South Asian Culture and Society, No VIII)*. Oxford: Oxford University Press.

——— (1995): *Transport and Communication in India Prior to Steam Locomotion: Volume 2: Water Transport (French Studies in South Asian Culture and Society VII)*. Oxford: Oxford University Press.

DERBYSHIRE, I. (1985): "Opening Up the Interior: The Impact of Railways upon the North Indian Economy and Society, 1860-1914," Ph.D. thesis, University of Cambridge, Department of History.

DESCHENES, O., AND M. GREENSTONE (2007): "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather," *American Economic Review*, 97(1), 354–385.

——— (2008): "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US," *Mimeo, MIT.*

DESCHNES, O., AND E. MORETTI (2009): "Extreme Weather Events, Mortality and Migration," *Review of Economics and Statistics*, Forthcoming.

DI GIOVANNI, J., AND A. LEVCHENKO (2009): "Trade Openness and Volatility," *Review of Economics and Statistics.*

DIEWERT, W. E., AND C. J. MORRISON (1986): "Adjusting Output and Productivity Indexes for Changes in the Terms of Trade," *The Economic Journal*, 96(3), 659–679.

DINKELMAN, T. (2007): "The Effects of Rural Electrification on Employment: New Evidence from South Africa," mimeo Princeton.

DIRECTORATE OF ECONOMICS AND STATISTICS (1967): *Indian Crop Calendar.* Delhi: Government of India Press.

DJANKOV, S., C. L. FREUND, AND C. S. PHAM (2006): "Trading on Time," World Bank Policy Research Paper No. 3909.

DORNBUSCH, R., S. FISCHER, AND P. A. SAMUELSON (1977): "Comparative Advantage, Trade, and Payments in a Ricardian Model with a Continuum of Goods," *American Economic Review*, 67(5), 823–839.

DUFLO, E., AND R. PANDE (2007): "Dams," *Quarterly Journal of Economics*, 122(2), 601–646.

DUTT, R. (1904): *The Economic History of India in the Victorian Age.* London: Kegan Paul.

DYSON, T. (1991): "On the Demography of Famines: Part I," *Population Studies*, 45(2), 279–297.

EATON, J., AND S. KORTUM (2002): "Technology, Geography and Trade," *Econometrica*, 70(5), 1741–79.

EATON, J., S. KORTUM, AND F. KRAMARZ (2005): "An Anatomy of International Trade: Evidence from French Firms," mimeo Chicago.

EVANS, C. J., AND J. HARRIGAN (2005): "Distance, Time, and Specialization: Lean Retailing in General Equilibrium," *American Economic Review*, 95(1), 292–313.

FEENSTRA, R. C. (2003): *Advanced International Trade: Theory and Evidence.* Princeton: Princeton University Press.

FISHLOW, A. (1965): *Railroads and the Transformation of the Ante-Bellum Economy.* Cambridge: Harvard University Press.

——— (2000): "Internal Transportation in the Nineteenth and Early Twentieth Centuries," in *Cambridge Economic History of the United States (vol II)*, ed. by S. L. Engerman, and R. E. Gallman. Cambridge University Press.

FOGEL, R. (1964): *Railways and American Economic Growth.* Baltimore: Johns Hopkins Press.

——— (1979): "Notes on the Social Saving Controversy," *The Journal of Economic History*, 39(1), 1–54.

FOSTER, A. D., AND M. R. ROSENZWEIG (2008): *Handbook of Development Economics, volume 4* chap. Economic development and the decline of agricultural employment. Elsevier.

FRANKEL, J. A., AND D. ROMER (1999): "Does Trade Cause Growth?," *American Economic Review*, 89(3), 379–399.

GADGIL, S., M. RAJEEVAN, AND P. A. FRANCIS (2007): "Monsoon variability: Links to major oscillations over the equatorial Pacific and Indian oceans," *Current Science*, 93 (2), 182–194.

GANDHI, M. (1938): *Indian Home Rule.* Ahmedabad: Navajivan Publishing House.

GREENSTONE, M., AND E. MORETTI (2004): "Bidding for Industrial Plants: Does Winning a 'Million Dollar Plant' Increase Welfare?," MIT Department of Economics Working Paper.

GUITERAS, R. (2008): "The Impact of Climate Change on Indian Agriculture," mimeo University of Maryland.

HAJAT, S., B. G. ARMSTRONG, N. GOUVEIA, AND P. WILKINSON (2005): "Mortality Displacement of Heat-Related Deaths: A Comparison of Delhi, Sao Paulo, and London," *Epidemiology*, 16 (5), 613–620.

HASAN, R., D. MITRA, AND K. RAMASWAMY (2007): "Trade Reforms, Labor Regulations, and Labor-Demand Elasticities: Empirical Evidence from India," *Review of Economics and Statistics*, 89(3), 466–481.

HEAD, K., AND A.-C. DISDIER (2008): "The Puzzling Persistence of the Distance Effect on Bilateral Trade," *Review of Economics and Statistics*, 90(1), 37–48.

HEADRICK, D. R. (1988): *The Tentacles of Progress: Technology Transfer in the Age of Imperialism, 1850-1940.* New York: Oxford University Press US.

HECKMAN, J. J., AND J. H. ABBRING (2007): "Econometric Evaluation of Social Programs, Part III: Distributional Treatment Effects, Dynamic Treatment Effects, Dynamic Discrete Choice, and General Equilibrium Policy Evaluation," in *Handbook of Econometrics, vol 6B*, ed. by J. J. Heckman, and E. E. Leamer. Elsevier.

HELPMAN, E., M. MELITZ, AND Y. RUBINSTEIN (2008): "Estimating Trade Flows: Trading Partners and Trading Volumes," *Quarterly Journal of Economics*, 123(2), 441–487.

HESTON, A. (1973): "Official Yields per Acre in India, 1886-1947: Some Questions of Interpretation," *Indian Economic and Social History Review*, 10(40), 303–334.

HESTON, A. (1983): "National Income," in *The Cambridge Economic History of India, Volume 2*, ed. by D. Kumar, and M. Desai. Cambridge: Cambridge University Press.

HOUDE, J.-F. (2008): "Spatial Differentiation in Retail Markets for Gasoline," mimeo Wisconsin.

HOUSE OF COMMONS PAPERS (1853): "Correspondence from Governor General in India, relative to Railway Undertakings," *British Parliamentary Papers*, 1852-53 (787), LXXVI.481.

——— (1884): "Select Committee on East India Railway Communication Report, Proceedings, etc.," *British Parliamentary Papers*, 1884(284)., XI.1.

HUMMELS, D. (2007): "Transportation Costs and International Trade in the Second Era of Globalization," *Journal of Economic Perspectives*, Summer 2007, 131–154.

HUMMELS, D., AND A. SKIBA (2004): "Shipping the Good Apples Out? An Empirical Confirmation of the Alchian-Allen Conjecture," *Journal of Political Economy*, 112(6).

HURD, J. (1975): "Railways and the expansion of markets in India, 1861-1921," *Explorations in Economic History*, 12(30, 263–288.

——— (1983): "Chapter VIII.2: Railways," in *Cambridge Economic History of India, volume 2*, ed. by D. Kumar. Cambridge: Cambridge University Press.

HUYNEN, M., P. MARTENTS, D. SCHRAM, M. P. WEIJENBERG, AND A. E. KUNST (2001): "The Impact of Heat Waves and Cold Spells on Mortality Rates in the Dutch Population," *Environmental Health Perspectives*, 109 (May), 463–70.

INTERNATIONAL PANEL ON CLIMATE CHANGE (2007): *Climate Change 2007: the Physical Science Basis.* Published for the International Panel on Climate Change.

IRWIN, D. A., AND M. TERVIO (2002): "Does Trade Raise Income? Evidence from the Twentieth Century," *Journal of International Economics*, 58(1), 1–18.

JAYACHANDRAN, S. (2006): "Selling Labor Low: Wage Responses to Productivity Shocks in Developing Countries," *Journal of Political Economy*, 114(3).

JENSEN, R. T. (2007): "The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector," *Quarterly Journal of Economics*, 122(3), 879–924.

JOHNSON, J. (1963): *The Economics of Indian Rail Transport.* Bombay: Allied Publishers.

KARL, T. R., AND K. E. TRENBERTH (2003): "Modern Global Climate Change," *Science*, 302 (December 5), 1719–1723.

KEHOE, T. J., AND K. J. RUHL (2008): "Are Shocks to the Terms of Trade Shocks to Productivity?," *Review of Economic Dynamics*, 11, 804–819.

KELLER, W., AND C. H. SHIUE (2007a): "The Origins of Spatial Interaction," *Journal of Econometrics*, 140(1), 304–332.

——— (2007b): "Tariffs, Trains, and Trade: The Role of Institutions versus Technology in the Expansion of Markets," mimeo Colorado.

KLINEBERG, E. (2002): *Heat Wave: A Social Autopsy of Disaster in Chicago*. Chicago: Chicago University Press.

KORTUM, S. (1997): "Research, Patenting, and Technological Change," *Econometrica*, 65(6), 1389–1419.

KRISHNA, P., AND M. Z. SENSES (2009): "International Trade and Labor Income Risk in the United States," NBER working paper 14992.

KRUGMAN, P. (1980): "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review*, 70(5), 950–970.

——— (1991): "Increasing Returns and Economic Geography," *Journal of Political Economy*, 99(3), 483–499.

KRUGMAN, P., AND A. J. VENABLES (1995): "Globalization and the Inequality of Nations," *The Quarterly Journal of Economics*, 110(4), 857–883.

KUMAR, D. (1983): "The Fiscal System," in *The Cambridge Economic History of India, Volume 2*, ed. by D. Kumar. Cambridge: Cambridge University Press.

KUMAR, K. K., K. R. KUMAR, R. ASHRIT, N. DESHPANDE, AND J. HANSEN (2004): "Climate impacts on Indian agriculture," *International Journal of Climatology*, 24, 1375–1393.

LEAKE, H. M. (1923): *The Foundations of Indian Agriculture*. Cambridge: W. Heffer and Sons.

LIN, J. Y., AND D. T. YANG (2000): "Food Availability, Entitlements and the Chinese Famine of 1959-61," *Economic Journal*, 110 (460), 136–158.

MACPHERSON, W. J. (1955): "Investment in Indian Railways, 1845-1875," *The Economic History Review*, 8(2), 177–186.

MAHARATNA, A. (1996): *The Demography of Famines: An Indian Historical Perspective.* Delhi: Oxford University Press.

MCALPIN, M. B. (1979): "Dearth, Famine, and Risk: The Changing Impact of Crop Failures in Western India, 1870-1920," *Journal of Economic History*, 39(1), 143–157.

——— (1983): *Subject to Famine: Food Crises and Economic Change in Western India, 1860-1920.* Princeton: Princeton University Press.

MCCALLUM, J. (1995): "National Borders Matter: Canada-U.S. Regional Trade Patterns," *American Economic Review*, 85(3), 615–623.

MELITZ, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71(6), 1695–1720.

MELITZ, M. J., AND G. I. P. OTTAVIANO (2007): "Market Size, Trade, and Productivity," *Review of Economic Studies*, 75(1), 295–316.

MICHAELS, G. (2008): "The Effect of Trade on the Demand for Skill — Evidence from the Interstate Highway System," *Review of Economics and Statistics*, 90(4).

MOLLISON, J. (1901): *A Text-Book on Indian Agriculture.* Bombay: "Advocate of India" Steam Press.

MUKERJI, N. (1915): *Handbook Of Indian Agriculture.* Calcutta: Thacker, Spink and Co.

MUNSHI, K., AND M. ROSENZWEIG (2007): "Why is Mobility in India so Low? Social Insurance, Inequality, and Growth," mimeo Brown.

NAIDU, N. (1936): *Coastal Shipping in India.* Madras.

NAMJOSHI, M., AND B. SABADE (1967): *Chambers of Commerce in India.* Poona: Gokhale Institute of Politics and Economics.

NEWBERY, D. M., AND J. E. STIGLITZ (1981): *The Theory of Commodity Price Stabilization : a Study in the Economics of Risk.* Oxford: Clarendon Press.

NGO-DUC, T., J. POLCHER, AND K. LAVAL (2005): "A 53-year Forcing Data Set for Land Surface Models," *Journal of Geophysical Research*, 110, 1–13.

O'BRIEN, P. (1977): *The New Economic History of the Railways.* Routledge.

O'GRADA, C. (2007): "Making Famine History," *Journal of Economic Literature.*

——— (2008): "Sufficiency, sufficiency and sufficiency: revisiting the Bengal Famine of 1943-44," mimeo University College Dublin.

PADMANABHAN, S. Y. (1973): "The Great Bengal Famine," *Annual Review of Phytopathology*, 11, 11–24.

PAVCNIK, N. (2002): "The Long and Short of the Canada-U.S. Free Trade Agreement," *Review of Economic Studies*, 69(1), 245–276.

PONCET, S., AND S.-J. WEI (2005): "Imagine China as a Collection of Economies: Estimating the Effect of Trade Costs on Growth," mimeo Paris School of Economics.

RABALLAND, G., AND P. MACCHI (2008): "Transport Prices and Costs: The Need to Revisit Donors Policies in Transport in Africa," mimeo World Bank.

RAMONDO, N., AND A. RODRIGUEZ-CLARE (2008): "The Gains from Openness: Trade, Multinational Production and Diffusion," mimeo Penn State.

RAUCH, J. (1999): "Networks Versus Markets in International Trade," *Journal of Economic Literature*, 48(1), 7–35.

RAVALLION, M. (1997): "Famines and Economics," *Journal of Economic Literature*, 35(3), 1205–1242.

RISLEY, H., AND E. GAIT (1903): *Census of India 1901, Volume 15, Madras.* Calcutta: Superintendent of Government Printing.

RODRIGUEZ, F., AND D. RODRIK (2001): "Trade Policy and Economic Growth: A Skeptic's Guide to the Cross-National Evidence," in *NBER Macroeconomics Annual 2000.* MIT Press.

RODRIK, D. (1997): *Has Globalization Gone Too Far?* Institute for International Economics.

ROSENZWEIG, M. R., AND H. BINSWANGER (1993): "Wealth, Weather Risk and the Composition and Profitability of Agricultural Investments," *Economic Journal*, 103, 56–78.

ROSENZWEIG, M. R., AND O. STARK (1989): "Consumption Smoothing, Migration, and Marriage: Evidence from Rural India," *The Journal of Political Economy*, 97(4).

ROSENZWEIG, M. R., AND K. I. WOLPIN (1993): "Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investments in Bullocks in India," *Journal of Political Economy*, 101, 223–244.

ROY, T. (2001): *The Economic History of India, 1857-1947.* Oxford: Oxford University Press.

RUBIN, D. B. (1978): "Bayesian Inference for Causal Effects: The Role of Randomization," *Annals of Statistics*, 6, 34–58.

SAMUELSON, P. (1954): "The Transfer Problem and Transport Costs, II: Analysis of Effects of Trade Impediments," *The Economic Journal*, 64(2), 264–289.

SANYAL, N. (1930): *Development of Indian Railways.* Calcutta: University of Calcutta Press.

SCHLENKER, W., AND M. ROBERTS (2008): "Estimating the Impact of Climate Change on Crop Yields: The Importance of Nonlinear Temperature Effects," NBER Working Paper 13799.

SCHWARTZBERG, J. E. (1978): *A Historical Atlas of South Asia.* Chicago: University of Chicago Press.

SEN, A. (1981): *Poverty and Famines: An Essay on Entitlement and Deprivation.* Oxford: Oxford University Press.

SETTAR, S. (ed.) (1999): *Railway Construction in India: Selected Documents (3 vols).* New Delhi: Northern Book Centre.

SHEPARD, D. (1968): "A Two-Dimentional Interpolation Function for Irregularly-Spaced Data," *Proceedings of the National Conference, ACM*, 23, 517–524.

SHIUE, C. H. (2002): "Transport Costs and the Geography of Arbitrage in Eighteenth-Century China," *American Economic Review*, 92(5), 1406–1419.

SINGH, R., AND J. BANTHIA (2004): *India administrative atlas, 1872-2001.* New Delhi: Conroller of Publications.

SOKOLOFF, K. L. (1988): "Inventive Activity in Early Industrial America: Evidence from Patent Records, 1790-1846," *Journal of Economic History*, 48, 813–850.

SRIVASTAVA, H. (1968): *The History of Indian Famines and Development of Famine Policy, 1858-1918.* Sri Ram Mehra.

STOCK, J. H., J. H. WRIGHT, AND M. YOGO (2002): "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments," *Journal of Business & Economic Statistics*, 20(4), 518–529.

STOKES, E. (1959): *The English Utilitarians and India.* Oxford University Press.

STONE, I. (1984): *Canal Irrigation in British India.* Cambridge: Cambridge University Press.

STRACHEY, J., AND R. STRACHEY (1882): *The Finances and Public Works of India from 1869 to 1881.* London: K. Paul Trench.

TAUGER, M. (2003): "Entitlement, Shortage and the 1943 Bengal Famine: Another Look," *The Journal of Peasant Studies*, 31, 45–72.

THORNER, D. (1950): *Investment in Empire; British Railway and Steam Shipping Enterprise in India, 1825-1849.* Philadelphia: University of Pennsylvania Press.

TIMES OF INDIA (1938): *Upper India Chamber of Commerce: 1888-1938.* Bombay: Times of India Press.

TIRUMALAI, R. (1986): *The Voice of Enterprise: 150 Years of the Madras Chamber of Commerce & Industry.* Bangalore: Macmillan India.

TOPALOVA, P. B. (2005): "Trade Liberalization, Poverty, and Inequality: Evidence from Indian Districts," NBER Working Paper 11614.

TREFLER, D. (2004): "The Long and Short of the Canada-U.S. Free Trade Agreement," *American Economic Review*, 94(4), 870–895.

TYSON, G. (1953): *The Bengal Chamber of Commerce & Industry: 1853-1953, a Centenary Survey.* Calcutta: D.A. Lakin.

UNITED PROVINCES OF AGRA AND OUDH (1868): *Report on the Administration of the Northern India Salt Revenue Department.* Allahabad: Salt Revenue Department.

UNITED STATES CENSUS OFFICE (1902): *United States Census; Volume VI (Agriculture), Part II.* Washington: United States Census Office.

VISARIA, L., AND P. VISARIA (1983): "Population (1757-1947)," in *The Cambridge Economic History of India, Volume 2*, ed. by D. Kumar. Cambridge: Cambridge University Press.

WALLACE, R. (1892): *Indian Agriculture.* Calcutta: Thacker, Spink and Co.

WATT, G. (1889): *A dictionary of the products of India.* London: J. Murray.

WAUGH, M. E. (2007): "International Trade and Income Differences," mimeo Minneapolis Federal Reserve.

WELLINGTON, A. M. (1877): *The Economic Thory of the Location of Railways.* New York: John Wiley & Sons.

WHITCOMBE, E. (1983): "Chapter VIII.1: Irrigation," in *Cambridge Economic History of India, volume 2*, ed. by D. Kumar. Cambridge: Cambridge University Press.

WILLIAMSON, J. G. (1974): *Late Nineteenth-Century American Development: A General Equilibrium History.* Cambridge: Cambridge University Press.

WORLD BANK (2007): *Evaluation of World Bank Support to Transportation Infrastructure.* Washington DC: World Bank Publications.

YANG, A. (1999): *Bazaar India: Markets, Society, and the Colonial State in Gangetic Bihar.* University of California Press.

YOUNG, A. (2005): "The Gift of the Dying: The Tragedy of AIDS and the Welfare of Future African Generations," *Quarterly Journal of Economics*, 120 (2), 423–466.

Appendix A

Data Appendix

A.1 Data From Colonial India for Chapters 1 and 2

This section provides further information on the data used in Chapters 1 and 2 of this thesis (which relate to the period from 1861 to 1930).

A.1.1 Sample of Districts

The data I use in Chapters 1 and 2 of this thesis cover the areas of modern-day India, Pakistan and Bangladesh, most of the area known as British India. I work with a panel of 239 geographic units of analysis that I refer to as districts, for as much of the period 1861 to 1930 as possible. The majority of British India was under direct British control, and was divided into nine large, administrative units known as provinces. Each province was further sub-divided into a total of 223 districts, which are the units of analysis that I track from 1861 to 1930. Areas not under direct British control were known as 'Princely States'. For administrative purposes these were grouped into divisions similar to the provinces and districts described above, so in princely state areas I use the lower administrative units as my units of analysis and refer to them as districts, following the *Indian Administrative Atlas* (Singh and Banthia 2004). There were 251 of these districts in my sample area, but data collection in the princely states was extremely incomplete and I include only 16 districts from the princely state regions in my final sample.

A.1.2 Trade Cost Proxy Variables

I construct trade cost proxy variables using a newly constructed GIS database on the Indian transportation network, from 1851 to 1930. The database covers four modes of transportation: railroads, roads, rivers and coastal shipping. To construct this database, I begin with a GIS database that contains the locations of contemporary railroad, river and coast lines from the *Digital Chart of the World.* Each segment (approximately 20 km long) of the railroad network is coded according to the year in which it was opened. To do this I use the publication *History of Indian Railways, Constructed and in Progress* (1918 and 1966 volumes), the 1966 volume of which refers to railway lines in modern-day India only. To obtain years of opening for line segments in modern-day Pakistan and Bangladesh from 1919 to 1930 I use the annual *Railway Reports* published by the Railways Department, which list all line section openings in each year. For river transport I keep only those rivers that are reported in Schwartzberg (1978) or Bourne (1849) as navigable in 1857. The final component of the colonial India GIS database that I construct is the location of each district and salt source. To calculate district locations I digitize a map of the district borders in India (as they existed in 1891). I use the maps in the *Indian Administrative Atlas* and *Constable's Hand Atlas of India* (Bartholomew 1893) to create this digital map. I use this to calculate district centroids, which I take to be the 'location' of each district. Finally, I obtain the location of each salt source from contemporary maps.

I then convert the GIS database of transportation lines and district/salt source locations into a graph of nodes and arcs, as is common in the transportation literature (Black 2003). I work with a simplified graph representation of the Indian transportation network, where the number of nodes and the sparsity of arcs is low enough for network algorithms to be feasibly operated on it using a desktop computer (the resulting network has 7651 nodes). To do this, I use the 'simplify' command in ArcGIS. A line in ArcGIS is a series of vertices connected by straight lines. The 'simplify' command removes vertices in such a way as to minimize the sum of squared distances between the original line and a the simplified line. The original *Digital Chart of the World* railway layer, for example, consists of approximately 33,000 vertices; I simplify the railway layer to one of only 5616 vertices. Because the density of informal roads was extremely high (Deloche 1994), I allow road transport to occur along the straight line between any two nodes on the network, but only if the two nodes either represent districts or salt sources, or the two nodes are within 1000 km of each other. Allowing straight-line road travel between any two nodes would yield a network with over 58 million arcs. The shortest path between each of the nodes on such a dense network cannot be calculated using a desktop computer, so I restrict many of these arcs to be non-existent; the result is that the 7651-by-7651 matrix representing the network can be stored as a sparse matrix, and analyzed using sparse matrix routines (that increase computation speed dramatically) in Matlab. The result is a network with 7651 nodes, 5616 of which represent the railroad network, 660 of which represent the navigable river network, 890 of which represent coastal shipping routes, 477 of which represent the centroids of the 477 districts in India (in 1891 borders), and 8 of which represent the locations of the sources of 8 different types of salt. Because the railroad arcs are coded with a year of opening indicator, this network can be restricted to represent the transportation network for any year from 1851 to 1930 by simply turning these arcs on or off.

Finally, I use this network representation of the Indian transportation system to calculate the two trade cost proxy variables described in section 1.4. One such proxy is a measure of the cost of traveling between any two points (where a point is either a district or a salt source) in a year using the lowest-cost route along the network (available in that year). The lowest-cost route depends on the value of the relative freight rates, $\boldsymbol{\alpha}$, and the available network, $\mathbf{R}_t$. Conditional on values of $\boldsymbol{\alpha}$, I use a standard algorithm from graph theory and transportation science (Dijkstra's algorithm) to calculate the shortest path between every pair of points, along the transportation network available in each year from 1861 to 1930. The resulting measure, $LCR_{odt}(\boldsymbol{\alpha}, \mathbf{R}_t)$, is in units of railroad-equivalent kilometers. The second proxy variable for trade costs is a dummy variable that indicates when it is possible to travel between two districts without leaving the railroad network. This is easily constructed using the transportation network representation and digital map of Indian districts described above.

A.1.3 Bilateral Trade Flows

The data I use on bilateral trade flows was collected from a variety of different sources, one for each mode of transportation. I describe each of these modes in turn, and then how they were combined into aggregate data on trade flows.

Data on *railroad trade* within India were published separately for each province. The geographic unit of analysis in these records is the 'trade block', which spans between four and five districts. Trade blocks split into smaller blocks over time, but I aggregate over these splits to maintain constant geographic units. The trade blocks were always drawn so as to include whole numbers of districts. The railroad trade flow data, like that on all modes of transportation described below, represents final shipments between two regions (even if a shipment changed railroad companies). All bilateral block-to-block intra-provincial trade flows were published, except that from a block to itself (which was always unreported). Inter-provincial trade flows were published from each internal block to each external province (and vice versa), but not by trade block within the external province. I therefore create a full set of inter-provincial block-to-block flows by following an analogous procedure to that used to prepare bilateral trade data on provincial-state trade between Canada and the United States in McCallum (1995) and Anderson and van Wincoop (2003). This method assigns a province's trade block's imports from each of another province's trade blocks in proportion to the exporting blocks' stated exports to the entire importing province (and vice versa for exports). In order to match the internal block-level railroad trade data to international trade data (leaving via specified ports, as described below), I apply a similar proportionality method. This is possible because the railroad trade data differentiate railroad trade to/from principal ports (in each province) from trade bound for non-port consumption. Only if a shipment was taken off the railroad system and re-shipped onwards would it be counted as two separate shipments. I collect this data from various annual, provincial publications from 1880 onwards. The titles of these publications changed over time, from *Returns of the Rail [and River-borne] Trade of [Province]* to *Report on* *the trade carried by rail [and river] in [Province]* to *Report on Inland Trade of [Province].* In the province of Madras, these statistics were only published from 1909 onwards. Railroad trade statistics were not published by the princely states themselves, but each province's external trade to/from each of the large princely states were published. I therefore treat each large princely state (Central India Agency, Hyderabad, Mysore, Rajputana and Travancore) as a single trade block. Data on *river-borne trade* within India was published in a similar manner to the railroad trade data, for the Brahmaputra, Ganges and Indus river systems. I collecte the river-borne trade data from the railroad trade statistics publications for the provinces of Assam, Bengal, Northwestern Provinces, and Sind. Data on trade within India that occurred via *coastal shipping* was published by each of the coastal provinces (Bengal, Bombay, Madras and Sind) in a similar manner to the railroad trade data. I collected the coastal trade data from various annual, provincial publications. The coastal trade data were published in publications whose titles changed from *Annual Statement of the Sea-borne Trade and Navigation of [Province]* to *Report on the Maritime Trade of [Province].*

Data on *international trade* leaving India was published separately for trade by maritime shipping and by roads. Each province published its own maritime international trade statistics, with each reporting the trade to and from its major and minor ports. The maritime international trade data was published in the same publications as those containing the coastal trade data, described above. This international maritime trade data was presented disaggregated into over 30 foreign countries, but to maintain consistent geographic units over time I aggregate these 30 countries into 24 foreign regions. Foreign trade by land occurred (in extremely small volumes) between Bengal, Northwest Provinces and

Punjab provinces and neighboring foreign countries (modern-day Nepal, China, Afghanistan and Bhutan). This trade data was published by each of these provinces, disaggregated by the border post through which trade left or arrived. I collect this data from various annual, provincial publications. The overland international trade data was published in: *Annual Report on the Trans-frontier Trade of Bihar and Orissa with Nepal*, *Bengal Frontier Trade: Trade of Bengal with Nepal, Tibet, Sikkim and Bhutan*, *Accounts Relating to the Trade by Land of British India with Foreign Countries*, *Annual Report on the Foreign Trade of the United Provinces*, and *Report on the External Land Trade of the Punjab.* I assign each of these border posts to the internal trade block in which it is located, and assume that all of the foreign land trade came to/from these blocks only.

Trade data by all modes of transport discussed above was published disaggregated by commodities. The railroad and river-borne trade data reported 85-100 commodities (depending on the year and province), the coastal shipping data 200-400 commodities, and the international maritime shipping data over 400 commodities. In order to compare commodities across these different levels of aggregation, I aggregate all data to the 85-commodity level. I use the commodity classification in the international maritime shipping publications (used to organize the over 400 commodities in these publications) to do this. Finally, I aggregate the trade data on each of the modes (for each commodity separately) into one trade dataset. Wherever relevant, I treat the regions of modern-day Afghanistan, Myanmar and Sri Lanka as foreign countries, since they are outside of the region on which I have other data from India. All of the above trade data are available from (at least) 1861 to 1930 (and beyond), except for the railroad trade data. The railroad trade data only starts in a coherent manner in 1880, and was discontinued in 1920. I therefore use bilateral trade data from 1880 to 1920 only.

A.1.4 Rainfall Data

A thick network of 3614 rain gauges at metorological stations (illustrated in Figure 1.1) recorded daily rainfall amounts from 1891-1930. From 1901 onwards, these records have been digitized by the Global Historical Climatology Network (Daily) project; the GHCN dataset also provides the latitude and longitude of each station. For the years 1891-1900, I collect the data from the publication *Daily Rainfall for India in the year....* In the years 1865 to 1890, very little daily rainfall data was published in colonial India, but monthly data from 365 stations (spread throughout India) was published by each province. These publications included the *Administration Reports* for each province, described in the agricultural price data section below. I use additional data (to increase the number of stations) that was published in selected provinces' *Sanitary Reports.* I convert monthly station-level data to daily station-level data using a procedure that is common in the meteorological statistics literature (eg, Ngo-Duc, Polcher, and Laval (2005)). Using daily data from 1891 to 1930, I estimate the district-specific relationship between the pattern of monthly rainfall in a year and the rainfall on any day of that year; I then use these estimated relationships to predict the rainfall on any day in a given district and year from 1865 to 1890, conditional on the pattern of monthly rainfall actually observed in that district and year. While these daily rainfall predictions are likely to be imprecise, much of the imprecision is averaged over when I construct crop-specific rainfall shocks, which are measures of the total rainfall in a given period (a length ranging from 55 to 123 days.) I convert station-level data to district-level data by simply averaging over the many stations in each district. If a given district-day has no reported rainfall observations I impute this missing observation by using an inverse distance-weighted average of that day's rainfall in the 5 closest reporting stations (know as "Shepard's method" in the meteorological literature (Shepard 1968).

A.1.5 Prices of Salt and Agricultural Commodities

I use data on eight different types of salt for each of the six provinces in Northern India. These eight salt types are those from: the Bombay sea salt sources near the city of Bombay, salt from the UK distributed via Calcutta, the Didwana salt source in Punjab, the Kohat mines in Punjab (principally the Jatta mine, according to Watt (1889)), the Mandi mine in Punjab, the Salt Range mines in Punjab (principally the Mayo mine, according to Watt (1889)), the Sambhar Salt Lake in Rajputana, and the Sultanpur source in the Central India Agency. And I use data on 17 agricultural commodities, namely: bajra, barley, bengal gram, cotton, indigo, jowar, kangni, linseed, maize, opium, ragi, rape and mustard seed, rice, sesamum, sugarcane, tur and wheat. The agricultural commodity price data, unlike the salt price data, are available throughout India. I collect this price data from various annual, provincial publications. These publications were: *Prices and Wages in India*; *Administration Reports* from all provinces; the *Salt Report of Northern India*; the *Statistical Atlas of Andhra State* with agricultural price data (for the Madras Presidency); the *Season and Crop Reports* from various provinces with agricultural price data; and the *Sanitary Reports* from various provinces with data on prices of food grains. Prices reported in these publication were an average of observations taken by district officers once per fortnight at each of 10-15 leading retail markets per district.

A.1.6 Real Agricultural Income

I use data that present the area under each of 17 crops (the 17 for which price data are available), and the yield per acre for each of these crops, in each district and year. These data were published in *Agricultural Statistics of India* from 1884 to 1930. For the years 1870-1883 I use data on crop areas and yields in the provincial *Administration Reports*, as described in the agricultural prices data section above. Data on agricultural output was published in each province's *Administration Report* except for Punjab and Bengal. For supplementary data I use each province's *Season and Crops Report* between 1904 and 1930. While Blyn (1966) and Heston (1973) have discussed the potential for measurement error in these sources, these authors have not been concerned with mechanisms through which measurement error might be correlated with the regressors I use in Chapters 1 and 2 of this thesis. I take the product of each area and yield pair to create a measure of real output for each crop, district and year. I then evaluate this bundle of 17 real output measures at the prices prevailing for these crops (from the agricultural price data described above), in each district and year, to create a measure of total nominal agricultural output for each district and year. Finally, I divide nominal output by a consumer price index (the Fisher ideal index) to create a measure of real income. In order to compute this consumer price index I use district and year specific consumption weights from the internal trade data, computing consumption as output minus net exports.

A.2 Climate Change Prediction Data for Chapter 3

To obtain predictions on the manner in which India's climate is predicted to change by the end of the century I use the output of two leading general circulation models. The first is the Hadley Centre's 3rd Coupled Ocean-Atmosphere General Circulation Model, which I refer to as Hadley 3. This is the most complex and recent model in use by the Hadley Centre. I also use predictions from the National Center for Atmospheric Research's Community Climate System Model (CCSM) 3, which is another coupled atmospheric-ocean general circulation model (NCAR 2007). The results from both models were used in the 4th IPCC report (IPCC 2007).

Predictions of climate change from both of these models are available for several emission scenarios, corresponding to 'storylines' describing the way the world (population, economies, etc.) may develop over the next 100 years. I focus on the A1FI and A2 scenarios. These are "business-as-usual" scenarios, which are the appropriate scenarios to consider when judging policies to restrict greenhouse gas emissions.

I obtain daily temperature predictions for grid points throughout India from the application of A1FI scenario to the Hadley 3 model for the years 1990-2099 and the A2 scenario to the CCSM 3 for the years 2000-2099. The Hadley model gives daily minimum and maximum temperatures, while the CCSM model reports the average of the minimum and maximum. Each set of predictions is based on a single run of the relevant model and available for an equidistant set of grid points over land in India.

I calculate future temperature realizations by assigning each district a daily weather realization directly from the Hadley and CCSM predictions. Specifically, this is calculated as the inverse-distance weighted average among all grid points within a given distance from the county's centroid. These daily predicted temperature realizations are used to develop estimates of predicted end of century climate. The Hadley 3 model has predictions for the years 1990 through 2099. I utilize the historical predictions to account for the possibility of model error. In particular, I undertake the following multiple step process:

1. For each Hadley 3 grid point, I calculate the daily mean temperature for each of the year's 365 days during the periods 1990-2000 and 2070-2099. These are denoted as $T^H_{gt,2070-2099}$ and $T^H_{gt,1990-2000}$, respectively, where the 'H' superscript refers to Hadley 3, g indicates grid point and t references one of the 365 days in a year.

2. I calculate the grid point-specific predicted change in temperature for each of the 365 days in a year as the difference in the mean from the 2070-2099 and 1990-2000 periods. This is represented as $\Delta T^H_{gt} = (T^H_{gt,2070-2099} - T^H_{gt,1990-2000})$.

3. I then take these grid-point specific predicted changes for all 365 days and assign district-specific predicted changes by taking weighted averages within 250 KM of the district centers. Again, the weight is the inverse of the square of distance. This procedure yields a predicted change in the daily mean temperature for all 365 days for each district or ΔT^H_{dt}, where d denotes district.

4. Using the NCC weather data, I calculate the grid-point specific daily mean temperature for each of the 365 days over the 1957-2000 period. I then take weighted averages of these daily mean temperatures for all grid points within 100 KM of each district's

geographic center, with the same weights as above. This yields $T^{NCC}_{dt,1957-2000}$.

5. The predicted end of century climate for each day of the year is equal to $T^{NCC}_{dt,1957-2000} + \Delta T^{H}_{dt}$. To preserve the daily variation in temperature, I apply the fifteen temperature bins from above to these 365 daily means. The resulting distribution of temperatures is the Hadley 3 predicted end of century distribution of temperatures that is utilized in the subsequent analysis.

In the case of the CCSM 3 predictions, I am unable to account for model error because these predictions are only available for the years 2000 through 2099, so there are no historical years available to remove model error.