

ESSAYS ON INDIA'S ECONOMIC DEVELOPMENT.

by

SAURABH GUPTA

A DISSERTATION

Presented to the Department of Economics
and the Division of Graduate Studies of the University of Oregon
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

June 2022

DISSERTATION APPROVAL PAGE

Student: Saurabh Gupta

Title: Essays on India's Economic Development.

This dissertation has been accepted and approved in partial fulfillment of the requirements for the Doctor of Philosophy degree in the Department of Economics by:

Dr. Shankha Chakraborty	Chair
Dr. Anca D. Cristea	Core Member
Dr. Alfredo Burlando	Core Member
Dr. Nagesh Murthy	Institutional Representative

and

Krista M. Chronister	Vice Provost for Graduate Studies
----------------------	-----------------------------------

Original approval signatures are on file with the University of Oregon Division of Graduate Studies.

Degree awarded June 2022

© 2022 Saurabh Gupta
All rights reserved.

DISSERTATION ABSTRACT

Saurabh Gupta

Doctor of Philosophy

Department of Economics

June 2022

Title: Essays on India's Economic Development.

This dissertation is on the economic development of India during the past three decades with a focus on its changing industrial and household structure. Chapter 1 provides a brief introduction of the Indian economy and motivates the theme of the dissertation.

In Chapter 2, I study the effects of transportation infrastructure on regional manufacturing activity. I exploit geographical and temporal variation in project implementation to argue for causal effects on the regional industrial outcomes. I investigate how highways can improve market competition between firms situated in geographically distant locations, and as a result, create incentives to invest in activities that improve productivity. The results show that highways had no direct effect on India's manufacturing output growth and led to a decline in average manufacturing productivity. I argue that these results can be attributed to lack of improvements in allocative efficiency within regions, and slow movement of skilled labor into the manufacturing sector. These results are contrary to some recent work on India but in line with evidence presented in the wider literature on low-income countries

In Chapter 3, I investigate the relationship between highways and female labor force participation (FLFP). Using census level data from India, I estimate

how the construction of highways may have opened up market opportunities for households and consequently affected FLFP. I find that the effects are heterogeneous across districts with some districts experiencing an increase while others experiencing a decline in FLFP. The decline was driven mostly by married and educated women withdrawing from the manufacturing and services sectors. I also find suggestive evidence that highways led to an increase in labor force participation of low skilled women.

In Chapter 4, jointly with Dr. Shankha Chakraborty, we examine how household decision making can explain India's declining FLFP over the last three decades. We propose a tractable analytical model in which married women respond to opportunity costs of their labor hours when dividing their time between household and market production. The model incorporates cultural costs attached to female work and its negative effect on female labor supply. We highlight competing mechanisms at play that suggest a U-shaped pattern of FLFP in response to economic growth.

Finally, Chapter 5 summarizes the results of the dissertation and presents a concluding remarks.

CURRICULUM VITAE

NAME OF AUTHOR: Saurabh Gupta

GRADUATE AND UNDERGRADUATE SCHOOLS ATTENDED:

University of Oregon, Eugene, Oregon, USA
Central University of Tamil Nadu, Chennai, Tamil Nadu, India
Panjab University, Chandigarh, Punjab, India

DEGREES AWARDED:

Master of Science, Economics, 2015, Central University of Tamil Nadu
Bachelor of Engineering, Mechanical Engineering, 2011, Panjab University

AREAS OF SPECIAL INTEREST:

Applied Econometrics
Development Economics
Industrial Economics
International Trade

PROFESSIONAL EXPERIENCE:

Graduate Employee, Department of Economics, University of Oregon, 2017-2022
Research Associate, Indian Institute of Management, Bangalore, 2017
Research Assistant, Madras School of Economics, 2016
Data Scientist, Scienaptic Systems, 2015

GRANTS, AWARDS AND HONORS:

Graduate Teaching Fellowship, University of Oregon, 2017-2022
Homan Fellowship Award, University of Oregon, 2021

PUBLICATIONS:

Gupta, S., & Bhaduri, S.N. (2019). Skin in the Game – Investor Behavior in Asset Pricing, the Indian Context. *Review of Behavioral Finance*, Vol. 11 No. 4, pp. 373-392

ACKNOWLEDGEMENTS

I thank the members of my committee for their guidance in conceiving and executing my research agenda. In particular, I would like to thank Shankha Chakraborty for his invaluable mentorship, encouragement, and patience. I thank Anca Cristea, Alfredo Burlando, and Nagesh Murthy for many helpful comments and advice. I extend my gratitude to Carol and family for providing support during my stay in Eugene. A special thanks to Sarthak Bhatia for his efforts in helping me access relevant data. Finally, I thank my friends and fellow graduate students for always being there for me.

TABLE OF CONTENTS

Chapter	Page
I INTRODUCTION	1
II DO HIGHWAYS AFFECT MANUFACTURING GROWTH? A CAUTIONARY TALE FROM INDIA	5
2.1 Introduction	5
2.2 Relevant Literature	7
2.3 The Golden Quadrilateral Project	10
2.4 Data Preparation	11
2.5 Empirical Strategy	13
2.5.1 Basic Setup	13
2.5.2 Main Specification	14
2.5.3 Identification Issues	15
2.6 Empirical Results	18
2.6.1 Exploring Selection Bias	18
2.6.2 Main Estimates: Manufacturing Output	23
2.6.3 Manufacturing Productivity	26
2.7 Robustness	33
2.7.1 Local Average Treatment Effect	33
2.7.2 Sampling Errors: Sub-sample Analysis	36
2.8 Conclusion	38
III NATIONAL HIGHWAYS AND FEMALE LABOR FORCE PARTICIPATION	40

Chapter	Page
3.1 Introduction	40
3.2 Relevant Literature	42
3.3 The Highway Project	45
3.4 Empirical Strategy	45
3.5 Data	49
3.6 Results	50
3.6.1 Baseline Comparison	51
3.6.2 Main Results	52
3.7 Discussion	64
3.8 Conclusion	66

IV FEMALE LABOR FORCE PARTICIPATION IN

INDIA	68
4.1 Introduction	68
4.2 Background	73
4.2.1 Additional Facts	73
4.2.2 Cultural values	76
4.2.3 Theoretical works	77
4.3 Household Decisions	79
4.3.1 Case I: $I = 0$	81
4.3.2 Case II: $I = 1$	82
4.3.3 Extensive margin decision: to work or not	83
4.4 Aggregate Female Labor Force Participation	89
4.5 Numerical Illustration	92
4.6 Conclusion and future plan	95

Chapter	Page
V CONCLUSION	97
 APPENDICES	
A CHAPTER II APPENDIX	100
A.1 Baseline Levels of Outcome	100
A.2 Tables and Figures	102
B CHAPTER III APPENDIX	106
B.1 Tables and Figures	106
REFERENCES CITED	106

LIST OF FIGURES

Figure		Page
1	Share of India’s GDP by industry	2
2	Straight line and Least Cost Path Instrument Variables	16
3	Comparing manufacturing sector outcomes for different district groups.	20
4	Comparing annual changes in outcomes after controlling for district fixed effects.	21
5	Women’s Employment Share by Industry Sectors	55
6	Women’s Employment Shares by Industry and Skill level	59
7	Simulation of Literacy Interaction Effects	60
8	Female Labor Force Participation Rate, India	69
9	Fertility rate 1990-2018	70
10	FLFP 1990-2019	70
11	FLFP and per capita income, Afridi, Bishnu and Mahajan, 2019	72
12	Secondary enrollment 1990-2019	74
13	LFPR by education (urban, married and not married, age 20-45) Afridi <i>et al.</i> (2020)	74
14	Changes in female educational attainment in India	75
15	Changes in values if men should have more right to jobs	77
16	Changes in values if university is important for a boy or a girl	78
17	LFP decision at the household level	87

Figure	Page
18 Aggregate FLFP, $\chi \in \{\chi^L, \chi^H\}$	90
19 L_f with continuum of types, effect of culture	91
20 Numerical Simulation for LFP decision at the household level	94
21 Numerical Simulation for Aggregate FLFP, $\chi \in \{\chi^L, \chi^H\}$	94

LIST OF TABLES

Table		Page
1	Golden Quadrilateral - Progress (in kms.)	11
2	GQ construction main specification results	23
3	GQ treatment effect on productivity changes within districts	30
4	GQ treatment effect on changes in labor composition within districts	33
5	Testing for Heterogeneous Effects	36
6	Main results with Census firms: Manufacturing Output	37
7	Baseline - Placebo Treatment Effects	52
8	Average Treatment Effects	54
9	Parameter Values for Simulation	93
A.1	Summary Statistics	102
A.2	First Stage Results - Straight Line IV	103
A.3	Probability of Treatment	104
A.4	GQ construction reduced-form results	105
B.1	Summary Statistics	107
B.2	Treatment Effects - FLFP with Literacy Interactions by Industry . . .	108
B.3	Treatment Effects - Female Employment Levels by Industry	109
B.4	Baseline - Placebo Treatment Effects with Literacy Interactions . . .	110
B.5	Average Treatment Effects - FLFP with Road Connectivity Interactions by Industry	111
B.6	Average Treatment Effects - FLFP with Caste Diversity Interactions by Industry	112

Table	Page
B.7 Treatment Effects - Education Wise Female Employment Levels . . .	113
B.8 Treatment Effects - Female Employment Levels by Industry and Married Status	114

CHAPTER I

INTRODUCTION


India has experienced rapid growth in the past four decades starting with a structural shift in the late 1980s. During this time, India liberalized its economy and introduced gradual trade reforms that paved the way for growth in international trade and investment. A dismal GDP growth of 3.4% during the decade of 1970-1980 transformed into an average growth rate of 7.5% after the late 1980s, peaking at 8.9% during 2000-2010. This economic transformation was accompanied by secular improvements in the country's development indicators and large-scale structural shift away from traditional agriculture.

India's growth and transformation has been unique in one respect. It has largely been driven by the services sector instead of growth in the manufacturing sector, a pattern at odds with the experiences of other developing countries such as Bangladesh, China, or South Korea. Decomposing this growth further, Figure 1 plots the average sector-wise decomposition of value added to the GDP between 1950s and 2010s. These estimates show that between the decades of 1980s and 2010s the contribution of manufacturing sector towards the GDP increased from 14.5% to 17.7% whereas the contribution by the services sector increased by a much larger margin from 46% to 61%. Not surprisingly, contribution by the agriculture and allied activities decreased from 35% to 16.5% during the same time period.¹

Several reasons have been proposed in the literature for this slow growth in manufacturing. For example, the literature has identified a "missing middle" in India's manufacturing sector (Hsieh and Klenow, 2009): manufacturing firms are unable to scale up due to which markets are dominated by a few very large

¹These estimates have been compiled from Reserve Bank of India's annual publication on the Indian economy, Handbook of Economics and Statistics

Figure 1. Share of India's GDP by industry



firms and numerous small firms. These conditions are further exacerbated by a large number of very small firms operating in the informal sector. The stunting of manufacturing growth has also been attributed to India's strict labor laws making it difficult for firms to hire cheap labor, and bureaucratic hurdles such as complicated tax laws that make it difficult to run businesses (see Mitra and Ural, 2008, Rao, 2005).

In addition to India's unique sectoral growth, transformation in the labor markets and demographics is also worth noting. The economic growth in the past four decades has been associated with high educational attainment among younger people and declining fertility rates for women. As per World Bank estimates, the

average literacy rate in India has grown from 48% in the year 1991 to 74% in the year 2018. Simultaneously, the average fertility rate went down from 4.83 child births per woman in 1980 to 2.22 child births per woman in 2018. We may expect this fall in fertility and increase in education to encourage more women to join the market work; however, the female labor force participation (FLFP) rate has declined from 30% in 1990 to 19% in 2019.²

In this dissertation, my goal is to understand these patterns of slow manufacturing growth and female labor force participation in conjunction with investments in transportation infrastructure. The decades when India experienced significant economic growth coincided with the construction of large-scale transportation infrastructure projects. I hypothesize that construction of highways may have helped transform regions closer to it in a manner different than what we see through the national trends. For example, highways could have improved market competition and efficiency, and hence, comparatively higher positive gains for the manufacturing sector in regions closer to the highways. Similarly, if highways improved access to schools, jobs, business opportunities for households, they could have encouraged female labor force participation.

Exploration of these ideas situates this dissertation in the broader policy framework that tries to understand the role played by transport infrastructure in economic development. Significant efforts at the global and the country level are devoted towards improving public infrastructure that can increase overall productivity of the economy. Transport infrastructure is believed to be one of the most significant factors that influences the economy in multiple ways. For example, Woetzel et al. (2016) predict that global infrastructure demand will be \$2.5 trillion

²These empirical facts have been compiled from World Bank's DataBank that can be accessed at this website. <https://data.worldbank.org/country/IN>

annually between 2016-2030. In 2021 alone, the World Bank approved transport projects worth \$3.6 billion. These facts emphasize the growing importance of transport infrastructure in the global economy.

In Chapter 2 and 3, I investigate the role played by the construction of highways on regional manufacturing industry and female labor force participation in India. The results show that highways did not significantly impact the manufacturing output in regions that were located closer to the highways; instead, I notice a decline in average manufacturing productivity. In the labor markets, the highways seem to have affected districts based on their initial market conditions. Districts with lower literacy rate, higher road connectivity, and lower caste fractionalization experienced an increase in FLFP.

Chapter 4, jointly with Dr. Chakraborty, investigates the economic reasons behind the falling FLFP in India and proposes an analytical model that emphasizes the role of gender norms and mother's time devoted towards children's human capital production.

CHAPTER II

DO HIGHWAYS AFFECT MANUFACTURING GROWTH? A CAUTIONARY TALE FROM INDIA

2.1 Introduction

Transport infrastructure is a crucial element of public investment. Understanding the potential of transportation to facilitate economic gains is of growing policy relevance. Intuitively, we can assume that railroads, highways, and rivers provide easy access to markets across regions within a country, improve market competition and contribute to welfare gains. They may also help alleviate cost impediments to trade between firms located in different regions, thereby providing an integrated market economy to firms as well as consumers. A lot remains to be understood about this relationship between transport infrastructure and the market economy.

The current literature offers mixed results on this relationship in the context of low-income countries. For example, in the case of China, highways do not seem to significantly stimulate any economic growth. Some studies find that locations closer to the highways experienced negative GDP growth in manufacturing and all other sectors combined (see Faber, 2014; Banerjee, Duflo and Qian, 2020). In contrast, evidence from sub-Saharan Africa and India show robust positive growth for locations situated near the highways (see Storeygard, 2016, Ghani, Goswami and Kerr, 2016). The purpose of this paper is to revisit the case of India to estimate the relationship between transport infrastructure and economic growth using robust methodological tools. In addition, it also attempts to shed light on barriers that may hold back economic gains from the highways.

I first examine the link between highway construction and subsequent changes in the manufacturing sector. I use a large highway project, the Golden Quadrilateral (GQ) Project, constructed in India during the period 2000-2009, to examine the effects of national highways on local district-level manufacturing output growth. To counter the problem of endogenous placement of the highway, a combination of the difference-in-differences and instrument variable strategies is used. Two hypothetical highway routes are constructed using detailed geographical information to instrument the original GQ project, and bring quasi-random variation in the placement of highways among districts. Finally, the estimates are conditioned on a variety of district-level covariates including socio-economic and geographical characteristics. The results suggest that contrary to previous work on the effect of highways in India, districts closer to the GQ project did not experience significant manufacturing growth compared to districts further away.

I then explore potential mechanisms that may have limited the effects of the GQ project on manufacturing growth. Two notable roadblocks to India's manufacturing growth have been identified in the literature. First is the higher degree of resource misallocation in India (Hsieh and Klenow, 2009). One implication of resource misallocation is that market shares are not efficiently allocated across firms; specifically, low productivity firms may command a higher market share. I examine how the GQ project may have alleviated this problem. The results suggest that for districts closer to the GQ project, market shares adjusted towards less productive firms as opposed to more productive firms. Additionally, these districts experienced a significant decline in their overall productivity.

Another challenge to India's manufacturing sector is labor market rigidity (T. Besley and R. Burgess, 2004). Given India's labor abundance, I hypothesize that market integration through the GQ project would provide better access to factors resources, incentivizing firms to adopt more labor to be more competitive. Following the stylized results established in trade theory, this part of the chapter investigates whether the removal of barriers between locations increases the proportion of skilled labor in skill-abundant regions and non-skilled labor in skill-scarce regions. We learn that labor composition did not change significantly due to the construction of new highways, at least through the channels considered in the chapter.

The chapter is divided into the following sections. Section 2 gives a brief account of the relevant literature. Section 3 provides background information about the GQ project I use as a quasi-natural experiment. Section 4 details the data description and preparation. Section 4 lays out the empirical strategy. Section 6 presents the main results followed by some robustness checks in Section 7. Finally, Section 8 presents the conclusion of the chapter.

2.2 Relevant Literature

Low-income countries provide a unique setup for analysis because of the higher potential for returns on infrastructure investment. On the contrary, they also suffer from market distortions and low quality of construction technology that makes the efficacy of highway networks non-comparable to those in developed countries. Therefore, returns on investment in infrastructure may vary along a wide margin across countries.¹

¹The seminal paper by Chandra and Thompson (2000) leads the way for many studies in recent decades, they study the interstate highway system in the US and their results show that counties connected with the national highway experience significant growth in economic activity.

Two influential studies addressing the economic impact of highways in China are Faber (2014) and Banerjee, Duflo and Qian (2020). Faber (2014) analyze a more comprehensive highway network than the one I study in this chapter. He finds that China's National Trunk Highway System (NTHS), a highway project connecting all provinces with an urban population of more than 500,000, decreased economic activity in non-targeted small peripheral counties. He argues that the reason for this decreased activity is lower trade costs which shift economic output towards larger metropolitan centers. To address the issue of endogenous placement of highways he constructs a hypothetical highway network between counties based on a minimum cost path network, and euclidean straight lines to instrument the original highway.

Banerjee, Duflo and Qian (2020) focus on long term growth effects. Confirming Faber's results, they show that highways did not have any significant impact on economic growth in the long run. While they did affect other economic variables such as GDP per capita levels, number of firms, and firm profits, the effects were too small to be considered of policy relevance. Their methodology also involves constructing straight lines to proxy for highway and rail networks, while controlling for a variety of other geographical controls.

In contrast, for the case of India we have have three studies who find robust economic growth in locations close to the GQ project. The conclusion of this research is quite different from these earlier studies. Ghani, Goswami and Kerr (2016) use similar data and methods to find robust manufacturing growth in districts closer to the GQ project whereas I find lack of evidence to support this claim. I argue that difference arises due to different period of analysis and methods

On the contrary, the adjacent counties experience a decrease in economic activity. They establish identification by assuming that non-metropolitan cities accidentally receive the highway network.

used to identify the causal parameters of interest. I restrict the study period to two years before the GQ project construction started, allowing for better geographical granularity of districts, and consistent definitions across years. In contrast, Ghani, Goswami and Kerr use the year 1994, six years before construction, to validate their results. Therefore, it becomes essential for them to aggregate a few districts to keep the sample consistent across years. The second difference is that I do not use the baseline levels of the outcome variable as an additional control for the main specifications. I argue that this additional control leads to an alternate set of identifying assumptions consequently changing the results with a wide margin.

A recent study by Chatterjee, Lebesmuehlbacher and Narayanan (2021) uses firm-level manufacturing data similar to this chapter, to understand the returns on public investment in the manufacturing sector. They proxy public investment with the GQ project. A major deviation from this chapter is their identification strategy, which uses variation in project completion times and geographical proximity to identify the effects of the GQ project. Their results corroborate the evidence provided by Ghani, Goswami and Kerr, and finds that firms located closer to the GQ project had higher productivity. Khanna (2016) uses luminosity data to argue for increased urbanization in more narrowly-defined locations near the GQ project. He even finds spillover effects of the GQ project over time. The luminosity data captures overall economic activity which may include effects due to other sectors, public amenities, or agglomeration. I attempt to provide a more complete picture of the manufacturing sector which may not be captured by the luminosity data.

The chapter is also related to the literature on resource reallocation after a reduction in trade barriers. Harrison, Martin and Nataraj (2013) explore a related concept, the reallocation of market shares among Indian firms after trade

liberalization in 1991. My approach in understanding the misallocation of resources is very similar to their paper. The main difference is that their treatment variable is the trade liberalization episode whereas I consider the construction of highways as an event that reduces trade barriers. Their results emphasize within-firm productivity improvements as opposed to the reallocation of market shares as a channel through which aggregate productivity has improved in India in the 1990s and early 2000s.

None of the studies in the literature directly investigate the relationship between labor composition and highways for low-income countries. Michaels (2008) investigates this relationship in the US using an elaborate Heckscher-Olin model. My objective is tangential to his approach and is limited to providing insights on labor market rigidity and the subsequent impact on skill composition through trade-related channels.

2.3 The Golden Quadrilateral Project

The Golden Quadrilateral (GQ) project was announced by the Government of India in 1998. The objective of the project was to connect four major metro cities in India, Delhi, Mumbai, Chennai, and Kolkata, situated at opposite geographical ends of the country. Figure 2 presents the GQ project on the Indian map. The GQ project was a significant part of a bigger initiative by the government called the National Highway Development Project (NHDP). The objective of this broader project was to introduce and upgrade the vast network of national highways in India. The other major project of this scheme was the North South East West Corridor (NS-EW), which was completed almost 9 years after the GQ project but was started at a similar time.²

²Interestingly, the GQ project also overlapped with an ancient Indian road system called the Grand Trunk Road (GTR) which supported majority trade and labor mobility in ancient and

The GQ project took almost 11 years to complete. A total stretch of 5846 kilometers was introduced/upgraded with 4/6 lane high-speed highways. The planning for the project started in the year 1998. However, the actual construction started only between late 2000 and early 2001. A substantial portion of project execution was done during the years 2002-2005. Finally, it was around 95% complete by 2007. A full account of the project progression is given in Table 1.³

Table 1. Golden Quadrilateral - Progress (in kms.)

	2000-01	2001-02	2002-03	2003-04	2004-05	2005-06	2006-07
Completed	643	1063	1327	2612	4697	5278	5556
Under Constr.	2093	4594	4383	3234	1149	568	290
% Complete	10.9%	18.1%	22.7%	44.7%	80.3%	90.2%	95%
Total Length	5846	5846	5846	5846	5846	5846	5846

Note: The measures are calculated from the official annual reports of the National Highway Authority of India (NHAI). The year 2000-01 represents the financial year which runs from 1 April, 2000 to 31 March, 2001. All values are in kilometers

2.4 Data Preparation

I use firm-level survey data on formal Indian manufacturing published every year by the Ministry of Statistics and Program Implementation (MOSPI), India. The survey is called the Annual Survey of Industries (ASI) and has been used in numerous influential firm-level studies related to India (for example, Harrison, Martin and Nataraj, 2013, Martin, Nataraj and Harrison, 2017). It is a nationwide survey of all formal manufacturing firms employing 10 or more workers. The data classify firms into standardized industries defined as the National Industrial

medieval India. The British upgraded this road system by introducing metaled roads and at least one-fourth of the GQ project was engaged in further upgrading this ancient transport system to 4/6 lane highways.

³I exclude the analysis of the NS-EW corridor in this chapter, to better understand the results from the GQ project in isolation, and without the contamination of treatment effects due to the NS-EW corridor.

Classification (NIC) which is regularly revised to meet the international industry classification standards.

MOSPI collects the ASI data as a representative sample. The sampling procedure varies and can be classified into two categories, “Census” or “Sample”. A standard procedure is followed each year, a complete list of registered firms is collected from the state-level authorities, and firms are broadly categorized into “Census” or “Sample” sectors based on their size. The Sample firms constitute a representative sample of the population of firms in a given state and NIC-4 digit level industry. Only a certain percentage (20% for recent years) of total firms are recorded in the data and are assigned appropriate survey weights to match the population at the industry-state level. The Census firms include all firms with more than 200 workers.⁴

In addition, the ASI data is disseminated in two formats, a panel set of firms and a cross-section repeated over years. The panel data is protected behind a paywall whereas the cross-sectional format is publicly available since 2019. For this chapter, I use the publicly available cross-sectional data. A panel data analysis would have been ideal since it gives information about the evolution of firm-level parameters over time. It can also accurately predict the entry and exit rates of firms in the market, and precisely estimate firm-level productivity.

The cross-sectional data set comes with an elaborate description of firm-level characteristics. For example, product level input-output information, kinds of labor employed, capital assets, and other important firm-level information. It allows the construction of several firm-related parameters, I particularly make use of total output, total input, fixed capital net of depreciation, total emoluments

⁴For a more detailed exposition of the ASI data please refer to Martin, Nataraj and Harrison (2017).

paid to labor including bonus, other benefits, and total labor used. I clean the data to eliminate any outliers and errors due to data entry. Observations having zero or negative values for the variables of interest are dropped from the sample. I only keep firm observations that are marked as open. Therefore, the results in this chapter will apply only to the incumbent firms.

I gather other geographical data from different sources. The actual shapefiles of the GQ and NSEW project are from Ghani, Goswami and Kerr (2016). The shapefile data on the railway network and the Indian coastline is collected from MIT GeoWeb where it has been published by the International Steering Committee for Global Mapping. While this data was recorded in the year 2007, it is an accurate description of railways in the 1990s since the rail network in India has not changed significantly over the last three decades.

The descriptive statistics of the data are provided in the Table A.1. The last three rows of the table show that the average distance from railways, coastline, and the GQ project are positively correlated with each other. Therefore, adding these geographical controls are necessary to properly identify the effects of the GQ project. Second, the average output is the highest for the nodal districts, which is expected because they represent the biggest metropolitan districts of India. In addition, the districts closer to the GQ project have a substantially higher output as compared to districts 40 km or 120 km away.

2.5 Empirical Strategy

2.5.1 Basic Setup. The basic strategy is to compare the changes in outcome variables from the pre-construction period to the post-construction period when the GQ project was completely built in difference-in-differences setup. The unit of analysis is the sub-national regions in India called districts. I assume

that the districts closer to the highways would experience a higher impact of the highway as compared to districts further away. This assumption motivates the categorization of districts into four district groups based on their proximity to the GQ project. The first group consists of nodal districts which were originally intended to be connected with the GQ project. The second group is defined as the districts between (0-40) km away, the third group has districts (40-120) km away, and the fourth group has districts (>120) km away from the GQ project.

The main focus of the analysis is the (0-40) distance group which I consider as the treated group because it is closest to the highway. I compare (0-40) distance group with the (>120) distance group which acts as control. The second group (40-120) km away acts as a buffer between the two groups to capture spillover effects due to the highway, and also helps provide an overall picture of shifts in the manufacturing activity. For example, one possible scenario is that the introduction of the highway has heterogeneous effects across industries, for some industries the effects may be captured within the (0-40) range whereas for others the effects may extend beyond 40 km. To account for such differences I take into account an alternate treatment group (40-120). In any case, I assume that the highway effects do not extend beyond 120 km. Therefore, I consider (>120) as the control group.

2.5.2 Main Specification. The main specification follows the long difference approach. I compare changes in the years 2007 and 2009 with the baseline year 2000, when the GQ project construction had started. The baseline specification is of the form:

$$\ln\left(Y_i^{2009/2007}\right) - \ln\left(Y_i^{2000}\right) = \alpha + \sum_{g \in G} \beta_g GQdistgroup_g + \eta X_i + \epsilon_i \quad (2.1)$$

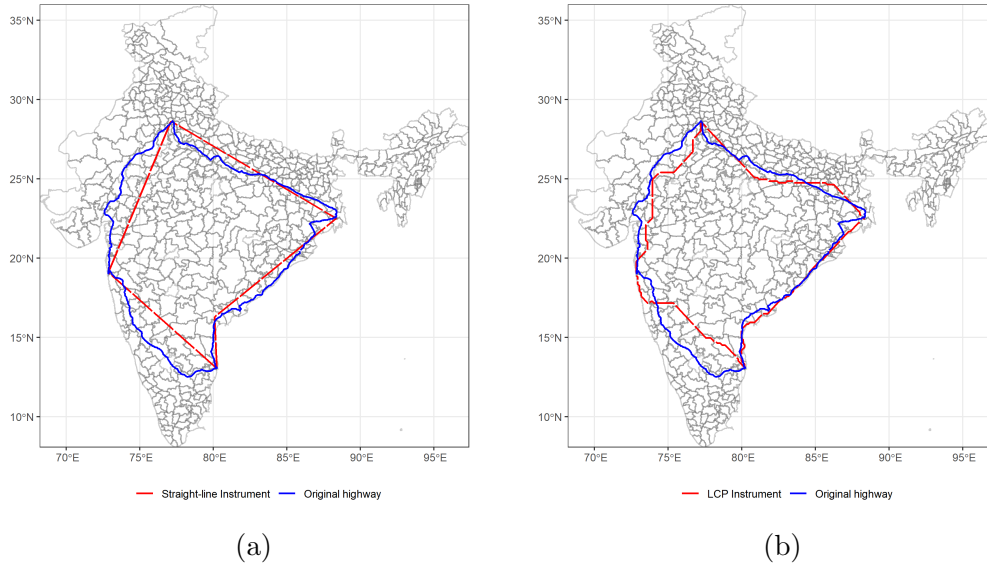
here, Y_i is the outcome of interest, subscript i represents the district, and g represents the four distance groups. $GQdistgroup_g$ is a dummy which captures which distance group does the district belongs to. X_i represents district-level characteristics such as literacy rate, rural population, population density, and SC/ST population. These district characteristics are recorded in the year 2001. It also includes geographical controls such as distance from the nearest rail network and distance from the Indian coastline. The coefficient of interest is β_g which captures the decade-long effects of the GQ project on the distance group g when compared with (>120) distance group. I consider the average of the variables in the years 2007 and 2009 as the post-construction outcomes since most of the highway was complete by the year 2006 (see Table 1). The average of two post-construction years helps overcome issues arising due to serial correlation.⁵ The year 2000 is chosen as the baseline since the GQ project construction started at the end of that year. The error term ϵ_i could be correlated across years for districts. Therefore, I cluster standard errors for all specifications at the district level.

2.5.3 Identification Issues. The most challenging issue while investigating the relationship between highways and economic activity is the problem of reverse causality. In most cases, highways are planned with the growth potential of locations in mind. For the GQ project, the policymakers likely selected districts based on certain economic characteristics. Annual reports and geographical maps show that the highway passes through some important cities in India. Descriptive statistics in the previous section show that (0-40) district group had significantly higher manufacturing output than districts further removed. This makes it hard to identify the true effects of the GQ project from other

⁵This is after assuming that the highways impacted all treated districts similarly. Another reason is to compare the results with Ghani, Goswami and Kerr (2016)

potential mechanisms or confounders which may have existed before the highway construction started.

Figure 2. Straight line and Least Cost Path Instrument Variables



Note: Hypothetical networks used as instruments for the original highway are presented in red and the original GQ project is presented in blue. Figure 2a presents the straight line instrument and Figure 2b presents the least cost path instrument. The linear cost function for the least cost path network is computed using factors such as elevation, land cover and wetlands. The background map shows district boundaries in India. First stage estimation results are reported in the appendix.

A related problem is selection bias, where the selection of districts into the treatment group could be driven by potential confounders. I control for a variety of factors under X_{ig} to take care of socio-economic characteristics and particulars related to geographical location. Additionally, the long differences help control for any time-invariant district fundamentals confounding the estimates. However, there may be other unobservable factors that may remain uncaptured. A simple way to check if these unobservables are leading to any kind of bias is by examining the growth trajectory of outcome variables before the construction of the GQ project conditional on the observed covariates. The data restrictions do not allow for such

an extension before 1998, without discounting the quality of the data.⁶ However, I do provide some empirical evidence in the next section that endogeneity is a crucial concern.

To resolve the problem of endogeneity due to the non-random selection of districts into districts groups, I use an instrument variable strategy commonly used in the highways literature. The first instrument is a connection of four straight lines between the nodal districts which were initially supposed to be connected with the GQ project. An additional kink is added to the eastern side of the instrument in order to keep it on land. Figure 2a presents this instrument on the Indian map alongside the original GQ project. This instrument is similar to the Euclidean network constructed by Faber (2014). The second instrument is also a hypothetical construction based on the least cost path between the four nodal districts. The least-cost path instrument (LCPIV) is constructed by assigning a cost structure to a high-resolution location grids representing the map of the whole country. For every grid on the map, the total cost is assumed to be a function of surface elevation, land cover, and coverage due to wetlands.⁷ Using this cost structure a path is constructed which minimizes the overall costs of travel between the nodal districts. Figure 2b depicts this instrument on the map alongside the original GQ network.

The identification assumption is that the two hypothetical instruments affect the outcome variables of interest only through the original GQ network. To

⁶The district definitions change very frequently in the 1990s because of the creation of three new states. The reshuffling of districts makes it hard to keep the district definitions consistent over a long period. As a result, districts need to be aggregated to keep them consistent throughout which leads to a loss in geographical granularity and sample size.

⁷More precisely, the cost function follows Faber (2014). For every grid n , I compute a cost function such that $c_n = Elevation + 25 * (LandCover - Dummy) + 25 * (Wetland - Dummy)$

control for factors that may potentially violate the exclusion restriction, I include the two sets of control variables discussed before, district-level characteristics and geographical characteristics in all specifications estimated in the chapter. Table A.2 presents the first stage results with and without the additional controls. The first stage is estimated using the two-stage least squares method. The first stage results show that both instruments are highly correlated with the original highway with F-statistic for the straight line IV of 38.8 and for the least cost path IV as 65.27.

In addition to the socio-economic and geographical characteristics, I also estimate the main specification after including the baseline levels of district output as an additional control. I do this only for the case of manufacturing output. This additional control serves two purposes. First, it helps compare the results in this chapter with Ghani, Goswami and Kerr (2016), who use similar data and methodology. Second, it helps shed light on the sensitivity of results under two cases, that is, with and without the baseline levels as an additional control. The identification assumptions under these two cases are very different, discussed more elaborately in the appendix. In a nutshell, including this additional control amounts to conditioning the estimates on lagged dependent variables whereas excluding it conditions on district fixed effects.

2.6 Empirical Results

2.6.1 Exploring Selection Bias. To provide a complete picture of the long-term effects, I plot the dynamic trajectory of the raw outcome variables (output, labor, capital and number of firms) ranging from the pre-construction period in 1998 till the year 2009 when the GQ project was complete. Figure 3 plots these graphs. A visual inspection of these plots shows that there is no significant increase in any of these variables. The growth trends are similar for every distance

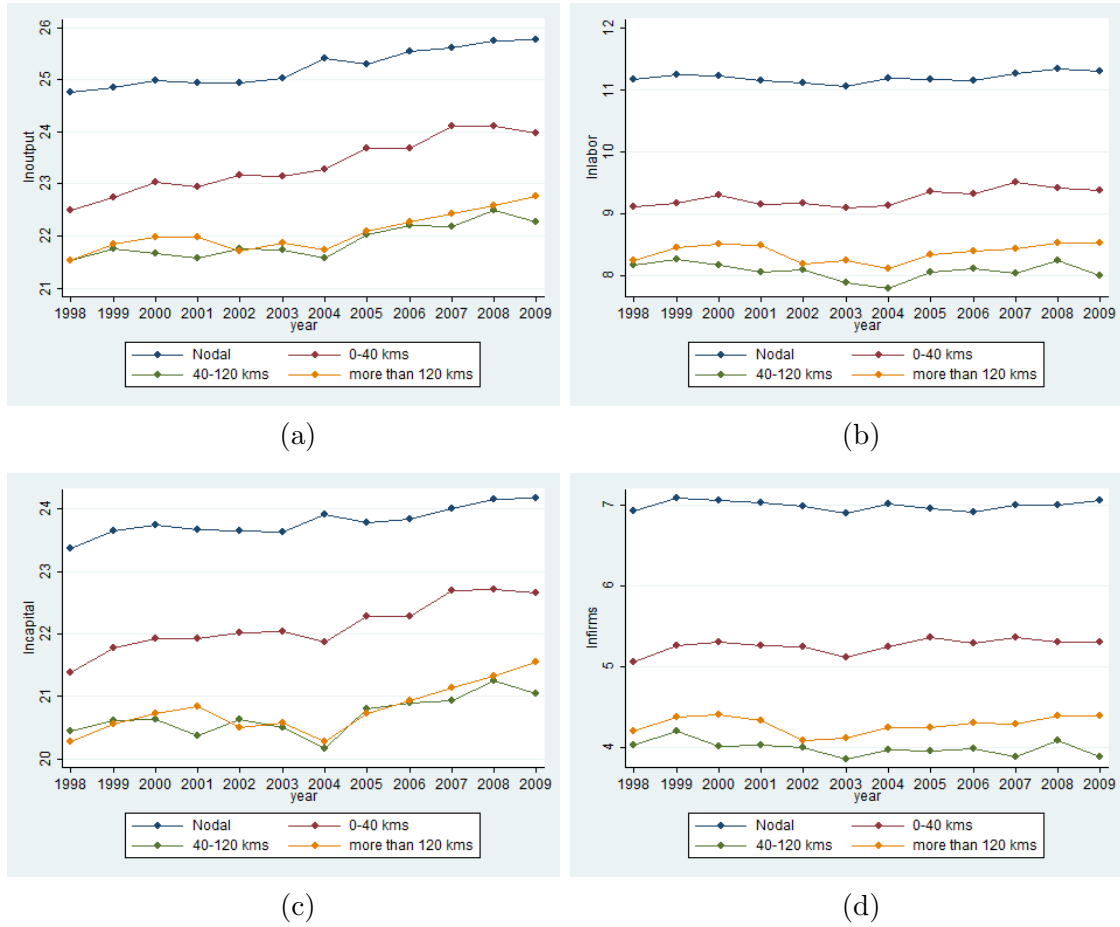
group. The level differences between distance groups confirm evidence from the descriptive statistics that the distance groups are characteristically different from each other. The raw plots show that the magnitude of effects from the GQ project, if any, were not large enough to decisively change the growth trajectory of districts. However, the (0-40) district group shows a marginally higher increase in total output (Figure 3a) and fixed capital (Figure 3c) after the year 2000 when compared with other distance groups. To put this into context, the (0-40) distance group experienced grew at an average rate of 10% per annum between 2000 and 2009 highest when compared with other distance groups. This exploratory analysis shows that the GQ project did not disproportionately affect the overall growth trajectory of manufacturing activity in districts close to it. However, we do see comparatively higher growth in output and capital in the (0-40) distance group.

In the next step, I investigate the more promising growth in manufacturing output. It is important to understand which factors contributed to this growth and how much of it is driven by the construction of the GQ project. I identify the effects of the GQ project by controlling for as many district-level factors as possible. The first set of specification assumes a dynamic panel framework and estimates a fixed effects model with districts as the unit of analysis. I include district and year fixed effects and re-estimate the growth trajectory of manufacturing output. I estimate the following equation,

$$\ln(Y_{igt}) = \alpha + \sum_{t=1998}^{2009} \beta_{gt} GQdistgroup_{ig} \times Year_t + \lambda_t + \mu_i + \epsilon_{igt} \quad (2.2)$$

here, subscript i represents the district, subscript g represents the distance groups, and subscript t represents the year. $GQdistgroup_g$ represents a dummy for the distance category of either (0-40) km or (40-120) km away, two separate regressions

Figure 3. Comparing manufacturing sector outcomes for different district groups.




Note: This figure plots raw manufacturing sector variables for districts categorized into different distance groups based on their proximity to the GQ project. Nodal represents the group of four main districts connected by the GQ project, that is, Delhi, Mumbai, Chennai, and Kolkata. It also includes the sub-urban districts located near the main ones. All variables are log transformations.

are run for both distance groups with (>120) as the base category. $Year_t$ is a dummy for all years between 1998 and 2009 with the year 2000 as the base year. Finally, μ_i represents district fixed effects and λ_t represents year fixed effects. The district fixed effects help capture district fundamentals such as resource endowments, demographic and cultural aspects whereas year fixed effects capture overall changes to the Indian economy. The coefficients β_{gt} capture the effect of the

GQ project on the distance groups for every year t . These coefficients are plotted separately for (0-40) and (40-120) distance groups.

Figure 4. Comparing annual changes in outcomes after controlling for district fixed effects.



Note: This figure plots coefficients β_{gt} from a fixed effects model defined by Equation 2.2 for two separate distance groups. Figure 4a compares treated districts with the control districts and Figure 4b compares intermediate districts with control districts. The black dotted lines represent confidence intervals and the significance levels are marked alongside the nodes following conventional norms (* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$).

The results show wide differences in growth between different distance groups with districts (0-40) km away showing increased improvements in manufacturing output. Figure 4a and Figure 4b show these contrasts. The point estimates β_{gt} show a growing trend for the (0-40) district groups whereas the trends are almost flat throughout the sample period for (40-120) district groups. These coefficient estimates are not significant except for a few at the end of the sample period but have a positive magnitude suggesting a positive effect.

The major cause of concern in these estimates is the lack parallel trends before the construction of the GQ project. Figure 4a shows that the growth trend starts in the year 2000 as opposed to the years after 2000 when the GQ construction started. This makes it hard to attribute the differences after 2000

to the GQ project specifically. The unobservable factors driving growth before the construction of the GQ project could be driving the future trajectory of growth in (0-40) district groups. More pre-construction periods could help give more information on how these trajectories evolved, what underlying factors could be driving them, and if we can specifically attribute this growth to the GQ project or not.

I check the extent of bias due to the self-selection of districts into the treatment group by estimating a logit model which predicts the construction of the GQ project using economic characteristics recorded before the highway was constructed. I use the dummy for the distance group (0-40) km away as the dependent variable and regress it on a battery of district-level characteristics recorded in the year 2001, and manufacturing industry characteristics recorded in the year 1999. Ideally, the district characteristics must also be recorded before 2000, but due to lack of data, I assume that the demographic variation across districts does not change significantly over these two years. The equation I estimate is as follows,

$$Pr(GQdistgroup_{i(0-40)} = 1) = \Lambda(\psi_0 + \psi_1 X_{i1}^{1999} + \psi_2 X_{i2}) \quad (2.3)$$

Here, X_{i1}^{1999} represents district characteristics such as literacy rate, rural population, population density, distance to the rail network, and distance to the coastline. X_{i2} represents manufacturing industry characteristics recorded in the year 1999, these include the outcome variables studied in the following sections, total output, share weighted average productivity, skill intensity, and urban share.

Table A.3 presents results from the above model. The coefficient estimates on total output are statistically significant even after controlling for other district characteristics. This means that districts with higher output had a higher

Table 2. GQ construction main specification results

	OLS			SLIV			LCPIV		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	<i>No Controls</i>	<i>Controls</i>	<i>Controls</i>	<i>No Controls</i>	<i>Controls</i>	<i>Controls</i>	<i>No Controls</i>	<i>Controls</i>	<i>Controls</i>
Nodal = 1	-0.414*** (0.132)	-0.083 (0.208)	0.282 (0.251)	-0.354** (0.168)	0.110 (0.311)	0.604 (0.381)	-0.281 (0.177)	-0.078 (0.256)	0.308 (0.319)
Treated = 1	0.068 (0.123)	0.139 (0.130)	0.374*** (0.138)	0.281 (0.592)	0.459 (0.677)	0.760 (0.678)	-0.237 (0.339)	-0.117 (0.353)	0.139 (0.377)
Intermediate = 1	0.177 (0.181)	0.213 (0.180)	0.241 (0.161)	0.199 (0.949)	0.174 (1.057)	0.560 (0.905)	1.033 (0.635)	1.022 (0.655)	1.219** (0.602)
Baseline			-0.249*** (0.043)			-0.266*** (0.060)			-0.219*** (0.053)
N	284	284	284	284	284	284	284	284	284
R sqr.	0.012	0.041	0.19

Note: This table presents the main results for the effects of the GQ project on manufacturing output growth. The coefficient estimates are produced by running Equation 2.1 with change in logged levels of output between 2000 and 2007/2009 as the dependent variable. Columns (1)-(3) present the OLS estimates and (4)-(9) present the IV estimates. Standard errors are clustered at the district level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

probability of getting treated with the GQ project. The negative sign in column (3) on the variables representing the distance from the nearest rail network and the coastline show that the GQ project was constructed close to these geographical entities. The results show that manufacturing industry characteristics of districts recorded in the year 1999 played a crucial role in determining the treatment of districts with the GQ project. Therefore, the estimates from the long differences presented in the next section will be biased. Therefore, the results presented by the instrument variables will be crucial for causal interpretation of results.

2.6.2 Main Estimates: Manufacturing Output. The results from the main specification defined in Equation 2.1 are presented in Table 2. The OLS estimates presented in columns (1)-(2) show that the district groups (0-40) and (40-120) experienced higher output growth than the (>120) district group. The coefficient estimates are positive but statistically insignificant. Additionally, the magnitude of estimates is higher for the (40-120) district group when compared with the district group (0-40). This is contrary to the expectation that the treated

district group (0-40) will experience the highest growth due to the construction of the GQ project.

The qualitative aspects of the OLS estimates are repeated in columns (4)-(5) where the original highway is instrumented with the hypothetical highway built along a straight line. The point estimates remain positive but the magnitude increases, pointing at the fact that the GQ project targeted more prosperous regions biasing the OLS estimates downward whereas the actual estimates may be larger. In contrast, the least-cost path instrument produces different results. The coefficient estimates on the (0-40) district group are almost zero with controls and negative without controls whereas the (40-120) district group shows statistically significant coefficients with higher magnitudes. These results predict that both district groups (0-40) and (40-120) experienced positive output growth when compared with (>120), but higher growth is seen in intermediate districts. However, noise and uncertainty makes these point estimates non-significant.

The IV results for the straight line and the least-cost path instruments are very different. This difference could be due to two possible reasons. Either these two instruments are targeting fundamentally different types of districts and these differences are not controlled for in the control variables, or, the implicit assumption stating that the effects of the GQ project are homogeneous across districts may not be true. The GQ project may have had heterogeneous effects across districts along dimensions that may be difficult to identify. The two instruments may pick up some of these heterogeneous effects separately. Later in the chapter, I check how far we can extend these local average treatment effects to the average treatment by showing differences in characteristics of complying districts under these two instruments.

The differences in magnitudes of the coefficient estimates for (0-40) and (40-120) district groups warrant further discussion. One potential reason is diminishing returns, due to which economic growth could be negatively correlated with the baseline output levels. In other words, districts with higher baseline output levels may experience lower economic growth than a district with lower baseline output levels. Therefore, we may need to estimate the main specification while conditioning on the baseline levels of district output. Columns (3), (6), and (9) show results after controlling for the baseline levels of output for each district. The results in column (3) are statistically significant, and the magnitude is higher for the district group closer to the highways.

An important point to note is that adding baseline output levels changes the conditional assumptions underlying the main specification. As discussed in subsection 2.5.3, columns (2) and (3) are estimated under a different set of assumptions. Column (2) conditions the regression equations on district-level controls and time-invariant confounders. Column (3) on the other hand, conditions based on district-level controls and the baseline level of output. As is evident in columns (2) and (3) the point estimates are sensitive to these two alternative specifications. It is difficult to know which of the controls correctly identify the effects of the highways since both specifications are independent and not nested into each other.

Ghani, Goswami and Kerr (2016) estimate long differences while controlling for baseline outcome levels and find positive significant effects of the GQ project on districts as shown in column (3). Based on that they conclude that districts near highways gained significantly from the GQ project. This conclusion may be

misleading for the alternate specification for which, the intermediate district group seems to gain more from the highways than the treated group.

2.6.3 Manufacturing Productivity. The previous section presents suggestive evidence that the positive effects of the GQ project may not be as large as we have seen in the literature. The effects are more pronounced in the intermediate districts than the treated districts. In this section, I explore other channels through which highways may have affected manufacturing output in the districts. I explore two channels more relevant for the Indian economy. These channels can be summarized under two basic questions I try to answer. First, do highways shift economic activity towards more productive firms? Second, do highways shift factors of production towards more favorable spatial locations?

These questions have been investigated in two different strands of the literature. Hsieh and Klenow (2009) highlight the importance of misallocation of factor resources in India's manufacturing sector. They propose a theoretical model which captures market distortions due to misallocation of resources among firms and use it to compare the magnitude aggregate productivity between different countries. They argue that misallocation helps explain the vast differences in aggregate productivity between India and other developed countries. Resource misallocation could emerge due to market distortions created by inefficient labor markets, trade barriers, industrial policy, and other policy-related factors.⁸ As a consequence, less productive firms can hold a vast proportion of labor and capital resources. The introduction of the GQ project has the potential to reduce this misallocation by optimally allocating factors of production across space. As a

⁸The literature has always struggled to find a dominant source of misallocation. The factors at play are more context-dependent. A set of potential factors are explored in a survey by Restuccia and Rogerson (2017)

result, we may see efficient firm selection into the market due to the increased competitiveness of the manufacturing industry.

The second potential reason for negligible growth comes from frictions in India's large labor markets. Factor mobility in India has been low and to what extent do highways alleviate this issue is a question of crucial policy importance. Some policy reports provide evidence that labor migration across states is very restricted, whereas a large portion of internal migration in India happens between locations within states. In a different regulatory context, Banerjee, Duflo and Qian (2020) point to a similar issue in the case of China, where factor mobility is restricted due to peculiar government regulations which they argue could be one of the critical factors why we see no effects of highways on economic growth.

In the backdrop of low factor mobility, I test the main results from the two-factor Heckscher-Olin model which holds only under complete factor mobility. More specifically, I test a version of Rybczynski theorem which states that regions with higher skill endowment will specialize in skill-intensive goods, and vice versa. I test whether the introduction of the GQ project changed the skilled labor composition of districts along these lines.

District Productivity and Firm Selection. To capture shifts in the firm selection, I construct two measures of district-level productivity. The first measure is a simple average of productivity for all firms in a given district. The second measure weights each firm's productivity with its output share. The prime focus of this section is the second measure, which helps capture the reallocation of output shares among firms based on their productivity. If highways bring trade-related competitive effects to regions, we will see more productive firms selecting into the market and gaining higher market shares (Melitz, 2003). This measure

could shift due to two reasons. First, changes in average levels of firm productivity, which represent within-firm enhancements in production efficiency. Second, due to the reallocation of market shares between firms.

To construct this measure, I first compute firm-level productivity using a method proposed by Levinsohn and Petrin (2003) (LP). The residuals of the log-linearized Cobb-Douglas production function are used as a measure of total factor productivity and can be estimated according to the following equation.

$$\ln(GVA_{njt}) = \alpha_j + \beta_{1j}\ln(L_{njt}) + \beta_{2j}\ln(K_{njt}) + \phi_{njt} \quad (2.4)$$

Here, the subscript n represents the firm, j represents NIC 3-digit level industry, and t represents the year.⁹ GVA_{njt} is the gross value added by the firm every year, L_{njt} is the total labor used, K_{njt} is the fixed capital assets of the firm, and ϕ_{njt} is the residual which captures the total factor productivity at the firm level. In the above specification, the capital and labor coefficients estimated by ordinary least squares may be biased. A common issue of endogeneity arises due to simultaneity bias, where the level of output may determine firms' level of fixed capital adoption. The LP method estimates consistent labor β_{1j} and capital β_{2j} coefficients for each industry while accounting for endogeneity in capital usage. They use control inputs like materials, fuel, to inform about the productivity dynamics. More specifically, firm productivity is assumed to follow a first-order Markov process and is proxied by a polynomial expressed in capital and control inputs. This helps extract a consistent labor coefficient which I then use to compute the capital coefficient

⁹NIC stands for national industrial classification, a standard for industry classification in India which is comparable to the international industry classification standards. A 3-digit NIC classification refers to the broader industry, for example, textile industry, paper industry, or, food products.

assuming constant returns to scale.¹⁰ Due to the cross-sectional nature of the data I cannot directly implement this approach. Therefore, I estimate one period lagged values of relevant variables for each firm by taking the average of closest industry-size-location matches from one period lagged years, a method similar to Sivadasan (2009). After estimating labor and capital coefficient for each 3-digit level NIC industry. I take the residuals as the measure of a firm’s total factor productivity. Consequently, I estimate district-level share weighted average productivity as follows,

$$\Phi_{it} = \sum_n s_{nit} \phi_{nit} \quad (2.5)$$

here, s_{nit} is the market share of the firm n in a district i . There are two things to note about this measure. First, it implicitly captures district-level changes in productivity due to the entry/exit of firms and doesn’t parse these effects separately. Second, it captures both within the industry and across industry market share shifts. The literature argues that trade liberalization generally shifts market shares between firms within well-defined industries (Melitz, 2003). Focusing on within industry-district market share shifts would be more plausible. However, as discussed before the data sample is not designed to represent the population at the industry-district level unit. The industries become very sparse at the district level and many of them are not recorded for the majority of the years, making it difficult to track manufacturing activity within industry-district units. As an alternative, I look at the overall changes within the district. Therefore, shifts in output shares could come from within or across industry reallocations of market shares.

¹⁰Structural methods like these are regularly used in the industrial organization literature to overcome endogeneity issues while estimating production functions.

Table 3. GQ treatment effect on productivity changes within districts

	OLS		SLIV		LCPIV	
	(1) <i>No Controls</i>	(2) <i>Controls</i>	(3) <i>No Controls</i>	(4) <i>Controls</i>	(5) <i>No Controls</i>	(6) <i>Controls</i>
<i>Panel A: Δ Productivity (share-weighted mean)</i>						
Nodal	0.041 (0.233)	-0.374 (0.230)	-0.152 (0.295)	-0.728** (0.350)	0.146 (0.256)	-0.309 (0.259)
Treated	-0.059 (0.127)	-0.203 (0.145)	-0.272 (0.394)	-0.676 (0.517)	-0.164 (0.209)	-0.276 (0.252)
Intermediate	-0.339** (0.149)	-0.371** (0.157)	-0.748 (0.679)	-0.575 (0.605)	0.058 (0.478)	0.070 (0.429)
<i>Panel B: Δ Productivity (mean)</i>						
Nodal	0.096 (0.177)	-0.316* (0.178)	0.006 (0.222)	-0.489** (0.249)	0.271 (0.207)	-0.261 (0.216)
Treated	-0.243*** (0.088)	-0.316*** (0.099)	-0.192 (0.286)	-0.504 (0.361)	-0.326** (0.162)	-0.437** (0.188)
Intermediate	-0.203** (0.102)	-0.195* (0.110)	-0.513 (0.459)	-0.402 (0.411)	0.385 (0.367)	0.331 (0.319)

Note: This table follows a similar setup to Table 2 but with share weighted productivity as the dependent variable. Standard errors are clustered at the district level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

The results are presented in Table 3. Panel (A) presents results from the second measure representing changes in market share weighted district productivity. The OLS specification results in columns (1) and (2) show negative and statistically significant estimates for treated as well as intermediate districts, meaning that market shares were distributed to less productive firms in these district groups. Instrumenting the original highway with the straight line and least-cost path switches the sign for intermediate districts but the treated districts retain their negative sign and significance. These results show that the GQ project forced the market shares to be reallocated towards less productive firms. One possible explanation is that more productive firms are gaining ground in districts away from the highways.

Panel (B) further solidifies these results, it shows the effects on average district productivity without weighting on market shares. This measure captures changes within-firm productivity and provides more insights into the overall productivity of the districts. The results show that average productivity declined decisively for the treated districts with negative point estimates significant under OLS and IV specifications. Three possible scenarios explain these average productivity results. First, incumbent firms in the treated group are leaving for more favorable locations. Second, incumbent firms experience a decline in their productivity which seems highly unlikely. Third, highways attract more economic activity but due to market distortions, low productivity firms can enter the market.

Two plausible reasons could potentially bias these estimates. First, there are strong assumptions on the way productivity is estimated. For example, firm-level characteristics other than location, size, and industry could influence the firms' adoption of resources, making the estimated lagged variables for firms inaccurate. Second, the nature of the production function could be different. Following the literature, I assume that the production function exhibits constant returns to scale, this may or may not be appropriate in the Indian context.

District Labor Composition. In the previous exercise, I looked at the firm selection and changes in district productivity. In this section, I explore India's low labor mobility and rigid labor laws which add to the challenge of efficient allocation of resources across firms. The purpose of this section is not to directly assess the labor mobility across India as that would require an elaborate analysis beyond the scope of this chapter. Instead, I check for evidence of labor mobility through trade-related channels more often associated with the highways. I use results from the Hecksher-Olin (HO) model of trade to study the changing nature

of labor composition in districts. One of the important results from the HO model is the Rybczynski theorem, which states that under an open economy with trade, locations favor the production of goods for which they have a higher abundance of relevant factor endowments. For example, districts with higher skill endowments will start specializing in more skilled labor-intensive goods. These results only hold under the assumption that there is non-frictional labor mobility across space and the labor laws favor easy adoption of skilled labor.

I investigate whether districts close to the GQ project attract more skilled labor into manufacturing.¹¹ I construct a simple skill intensity measure by taking the ratio of skilled labor to total labor engaged in formal manufacturing for a given district year. Here, the treatment variable is the interaction between the dummy for the distance groups and pre-construction endowment of skilled labor, measured by districts' literacy rate recorded in the year 2001. Specifically, I estimate the following equation,

$$\ln\left(S_{ig}^{2009/2007}\right) - \ln\left(S_{ig}^{2000}\right) = \alpha + \sum_{g \in G} \beta_g GQ distgroup_{ig} \times Literacy_{ig}^{2001} + X_{ig} + \epsilon_{ig} \quad (2.6)$$

here $Literacy_{ig}^{2001}$ represents district literacy rate for districts in the year 2001.

The results are presented in Table 4. The results are too noisy to draw any useful conclusions. The signs of point estimates change depending on which instrument is used. The extreme volatility of results points at two possible explanations. First, the effects of the GQ project on labor markets in the districts are heterogeneous across regions. Second, the specification defined in Equation 2.6 may not be correct

¹¹Michaels (2008) empirically tests the predictions of the Hecksher-Ohlin (HO) model and concludes that, contrary to the predictions of the HO models, highways in the US didn't affect demand for high skilled workers relative to low skilled workers. However, he finds that it did increase the wage bill of high-skilled workers in counties with relatively high skilled intensity before the highway construction.

Table 4. GQ treatment effect on changes in labor composition within districts

	OLS		SLIV		LCPIV	
	(1)	(2)	(3)	(4)	(5)	(6)
	<i>NoControls</i>	<i>Controls</i>	<i>NoControls</i>	<i>Controls</i>	<i>NoControls</i>	<i>Controls</i>
Nodal	0.080 (0.049)	0.080 (0.050)	0.122* (0.067)	0.222 (0.313)	0.098** (0.049)	0.094 (0.061)
Nodal X Literacy	-0.115 (0.078)	-0.105 (0.086)	-0.186* (0.113)	-0.355 (0.572)	-0.153* (0.079)	-0.137 (0.109)
Treated	-0.016 (0.035)	-0.007 (0.033)	0.963 (2.186)	1.116 (2.377)	-0.002 (0.081)	0.006 (0.087)
Treated X Literacy	0.058 (0.064)	0.051 (0.060)	-1.492 (3.452)	-1.721 (3.782)	0.023 (0.132)	0.025 (0.147)
Intermediate	0.009 (0.030)	0.019 (0.031)	-0.273 (0.933)	-0.290 (1.036)	0.089 (0.058)	0.111 (0.068)
Intermediate X Literacy	0.028 (0.054)	0.017 (0.057)	0.396 (1.367)	0.382 (1.490)	-0.138 (0.114)	-0.182 (0.134)
Literacy	0.081*** (0.028)	0.076*** (0.029)	0.152* (0.087)	0.116 (0.089)	0.119*** (0.035)	0.101*** (0.029)
N	284	284	284	284	284	284
R sqr.	0.12	0.19

Note: This table estimates Equation 2.6 with the ratio of skilled labor to total labor as the dependent variable. Standard errors are clustered at the district level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

and other functional forms may define the relationship between skill composition and highway construction.

Overall, the results from this exercise confirm the findings in Michaels (2008) who finds similar insignificant results in the case of US. These results also suggest that the rigid labor markets in India may have restricted the incentives for firms to adopt skilled labor as a resource.

2.7 Robustness

2.7.1 Local Average Treatment Effect. A potential concern is that the IV estimates presented before only address the treatment effects for a subset

of “complier” districts.¹² The complier districts are those which are randomly assigned to the GQ project and whose nature of treatment was decided by their allocation into the treated group based on the distance from straight line or least cost path instruments. In a counterfactual scenario, they would not have received the highway had they been assigned to the control group by these instruments.¹³ The characteristics of the compiler group define the local average treatment effects (LATE) estimated in the previous sections. If this group is significantly different from the overall population of districts then this is not a a valid instrument. The LATE might not inform us about the average treatment effects (ATE). An implicit assumption throughout the IV specifications has been that the effects are homogeneous across all districts. However, there may exist heterogeneous treatment effects with a potential to limit the results presented in the chapter. To investigate this problem, I check whether the complier districts are different from the rest of the population.

It is not possible to directly identify the complier districts in the data. However, we can estimate the proportion of compliers in the population using the following expression.

$$Pr(G_{1i} > G_{0i} | G_i = 1) = \frac{Pr(Z_i = 1)(E(G_i | Z_i = 1) - E(G_i | Z_i = 0))}{Pr(G_i = 1)} \quad (2.7)$$

Where G_i is a dummy for whether the original GQ project categorizes a district into the treated group or not, G_{1i} and G_{0i} represent the potential for treatment

¹²Here the terminology “compliers” is similar to what is used in the program evaluation literature, see Angrist and Pischke (2009).

¹³Contrast this with “never takers” who do not receive the GQ project irrespective of their assignment to treatment or control groups, and “always takers” who receive the project under all conditions. An implicit assumption is that there are no “defiers”, meaning a district does not have the power to go against its initial assignment.

and only one of these observed in the real data. Pr and E are probability and expectation operators, respectively. Z_i is the dummy for whether the instrument predicts treatment into the treated or control group. Using the above expression, the proportion of compliers for the straight line instrument is 31.3% and for the least cost path network is 36.2% of the districts.

The next crucial question is how different are these complier districts from the rest of the population? Although it is not possible to identify the complier districts directly, Angrist and Pischke (2009) show that we can still know about the characteristics of the compliers. To implement their strategy, I first create dummies for district characteristics by categorizing districts into above or below median value of a characteristic. Following this, we can measure the relative likelihood of a complier district belonging to the above median category of a district characteristic by taking the ratio of the first stage for the districts with above median characteristics to the overall first stage. The results of this exercise are reported in Table 5.

For the straight line instrument, column (2) and (7) show that the complier districts had comparatively higher literacy rates and were generally closer to the coast line. For the least cost path instrument, column (2), (3), (6), and (7) show that districts close to this network had higher literacy rate, more rural population, and were closer to both, existing rail network and the coastline. This robustness check shows that there may be significant heterogeneity in the way highways affected the districts. It also helps understand the main results where both instruments show different estimated coefficient. One potential reason for these differences could be characteristically different complier districts being picked up by the two instruments.

Table 5. Testing for Heterogeneous Effects

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	<i>FullSample</i>	<i>literacy</i>	<i>rural</i>	<i>pop_dens</i>	<i>scst_pop</i>	<i>H_rail</i>	<i>H_coast</i>
<i>Panel A: First Stage: Straight Line Instrument</i>							
GQ_40kms_iv	0.626*** (0.068)	0.720*** (0.069)	0.807*** (0.044)	0.662*** (0.056)	0.631*** (0.085)	0.658*** (0.123)	0.823*** (0.039)
N	208	107	97	97	110	111	104
R sqr.	0.25	0.33	0.38	0.32	0.24	0.22	0.26
<i>Panel B: First Stage: Least Cost Path Instrument</i>							
GQ_40kms_iv2	0.702*** (0.056)	0.784*** (0.044)	0.838*** (0.042)	0.681*** (0.056)	0.738*** (0.062)	0.832*** (0.069)	0.849*** (0.037)
N	208	107	97	97	110	111	104
R sqr.	0.32	0.39	0.47	0.35	0.39	0.49	0.37

Note: The first-stage results are presented for the treated and control district groups only. The first column presents full sample first-stage regression. The rest of the columns run separate first-stage regressions after dividing the data into sub-samples based on district with above median characteristics. Standard errors are clustered at the district level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

2.7.2 Sampling Errors: Sub-sample Analysis. The problem with the manufacturing data I use is that due to unknown sampling errors it is not representative of the population of firms at the district level. For the main specifications, I assume that the sampling errors are small such that they do not impact the main results presented in the chapter. However, if the sampling errors are significant then the estimates could be biased in two potential ways. First, if the proportion of firms sampled in the observed data is lower than the real population of firms in districts closer to the highways, the results presented could be downward biased. Second, if this proportion is higher the results could be upward biased.

I make use of sampling details provided along with the ASI data to check the extent of the bias generated by the sampling errors. As discussed before, the ASI data is disseminated in two different categories. The Census category of firms

Table 6. Main results with Census firms: Manufacturing Output

	OLS		SLIV		LCPIV	
	(1)	(2)	(3)	(4)	(5)	(6)
	$\Delta \ln Output$	$\Delta \ln Output$	$\Delta \ln Output$	$\Delta \ln Output$	$\Delta \ln Output$	$\Delta \ln Output$
Nodal	-0.360** (0.145)	-0.152 (0.227)	-0.308* (0.182)	0.174 (0.335)	-0.252 (0.187)	-0.121 (0.269)
Treated	0.044 (0.117)	0.088 (0.128)	0.508 (0.633)	0.714 (0.732)	-0.177 (0.302)	-0.070 (0.320)
Intermediate	0.211 (0.192)	0.250 (0.194)	-0.065 (1.031)	-0.076 (1.170)	0.877 (0.615)	0.890 (0.635)
N	284	284	284	284	284	284
R sqr.	0.012	0.037

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

Note: This table follows a setup similar to Table 2 with manufacturing output as the dependent variable. The data considered is a sub-sample of firms with more than 200 workers. These firms represent the whole population of bigger sized firms, therefore, sampling errors are very minimal. Standard errors are clustered at the district level.

represents the population of manufacturing firms with more than 200 workers, no sampling is done for this category and the whole population of firms is recorded.¹⁴ I drop the firms belonging to the Sample category which contributes the most to the sampling errors, as only 10-20% of these firms are sampled every year. I then re-estimate the main specifications from Equation 2.1 on manufacturing output growth only using the Census firms. This helps confirm the robustness of results presented in subsection 2.6.2, but only for firms with more than 200 workers.

If the results are driven by sampling errors, then this exercise should output different results than those presented before. The results are presented in Table 6. The magnitude and sign for the OLS and IV coefficients in all columns is qualitatively similar to the main results presented in Table 2. This provides some reassurance that the results presented before are not significantly affected by the sampling errors.

¹⁴This may still have other measurement errors commonly encountered in any random sample of a data.

2.8 Conclusion

This chapter finds no direct evidence related to the positive effects of highways on manufacturing output growth. On the contrary, it finds that districts located closer to the highways experienced a decline in their manufacturing productivity and market shares shift towards less productive firms. Additionally, no significant changes in labor composition are found in response to the highways.

The results in this chapter help consolidate a wider picture of growth in low-income countries. Services and the informal sector has played an important role in India's GDP growth over the past years. The double-digit growth during the decade of 2000s was propelled by large productivity gains in the services sector. In addition, the informal sector employs a significantly large proportion (80% of non-agricultural labor) of India's labor markets.¹⁵ The effects of the highways may not be as large in the manufacturing sector as has been found in other studies. However, the GQ project may have led to developments in the service or informal sector.

There are three potential ways to forward our understanding of the effect of highway effects on economic activity in low-income countries. First, for the manufacturing sector, other potential channels should be explored. For example, labor market rigidity and resource misallocation are likely to be important margins along which highways affect growth in low-income countries. Second, for an overall picture, it is important to know how highways may affect the labor markets, and the kind of opportunities they create for individual households by providing access to markets. A tentative step in this direction is taken in Chapter 3. Third, since highways bring large scale urbanization, it is worth exploring its impact on

¹⁵Rough estimates as per the International Labor Organization found here

the informal sector, a major component of employment in low-income countries.

Exploring these additional channels can provide a more clear picture of the returns to infrastructure investment in developing countries.

CHAPTER III

NATIONAL HIGHWAYS AND FEMALE LABOR FORCE PARTICIPATION

3.1 Introduction

Women may face asymmetrically higher socioeconomic costs to enter the labor force compared to men due to unfavorable gender norms, and lack of suitable job opportunities. These can negatively affect their economic prospects despite significant economic growth. When it comes to female labor force participation (FLFP), several mechanisms have been proposed regarding the effect of economic growth. For example, a common argument is that growth has a U-shaped non-monotonic effect on FLFP. Besides the culture-based explanation pursued in Chapter 4, one explanation is that during the initial stages of economic growth, women transition from necessary subsistence work to household production because of an income effect. As the economy grows and real wages for women continue to rise, the opportunity costs for household production increase and women substitute household work for market work.

Due to regional disparities in opportunities for market work, the role of transport infrastructure in this type of transition is critical. Roads can provide access to schools, colleges, conducive work environment and better paying jobs. Therefore, roads may not only help accelerate the movement along the U-shape but also help with structural transformation in regions where growth is contingent upon geographical connectivity. In this chapter, I provide suggestive evidence that in rural areas of districts located closer to national highways, skilled women working in manufacturing and services sector withdrew from market production. In contrast, the agricultural sector shows signs of an increase in labor supply of low skilled female workers. This evidence aligns with U-shaped relationship between

FLFP and economic growth, additionally suggesting that highways do open up market opportunities for female workers.

India has seen a secular decline in FLFP during the past three decades. This is surprising because at the same time fertility levels have fallen, and for married women, that should have made more time available for market work. Moreover, increases in labor productivity due to higher female educational attainment would have opened up job opportunities. The rural regions have contributed more to India's FLFP decline whereas the urban centers have experienced stagnation (Afridi, Dinkelman and Mahajan, 2018, Klasen and Pieters, 2015).

In this chapter, I explore how the construction of the Golden Quadrilateral (GQ) Project may have led to FLFP decline in rural areas. The analysis shows that the overall treatment effect of the GQ project is non-significant on FLFP in districts located closer to it. Interestingly, pre-construction district characteristics play an important role in determining the future impact of the project on FLFP in non-agricultural sectors. In particular, I find that districts located closer to the project that had higher literacy rates, lower road connectivity, and higher caste fractionalization before the construction of the highway project experienced a sharper decline in FLFP. These changes are primarily due to decreased female employment rather than increase in working age population from migration towards districts with higher literacy. The results are robust to several district level controls and can be interpreted causally.

I establish these results using a difference-in-differences methodology. I use the GQ project constructed in India between 2000-2009 as a quasi-natural experiment and study its effects on districts located closer to it (treated group), in comparison to districts located further away (control group). Consequently, the

universe of districts within India are divided into groups based on their proximity to the GQ project. Treated and control groups are defined, and an incidental group is considered to account for any spillover effects. FLFP is then compared between these groups starting in the year 1999 (before construction) and 2009 (after construction). To estimate district level measures, I make use of data from a wide variety of sources. I use the National Sample Survey (NSS) dataset to capture employment information on households in districts, augmenting this with Census data to add district level economic characteristics such as literacy rate, population characteristics, and finally, extracting infrastructure related variables from the publicly available SHRUG data set.

This chapter contributes to the literature in significant ways. First, I provide evidence that large scale highway projects may have asymmetric effects across genders. Second, I highlight the importance of prevailing market conditions in determining how highways may impact a given location. Third, I add to the U-shaped FLFP literature by providing evidence at a more dis-aggregated level and show how the transformation may take place in the regional economies. Fourth, I emphasize the role played by increased market access brought upon by the construction of the highways.

3.2 Relevant Literature

Several literatures are relevant to this chapter. First, my work contributes to the limited evidence on asymmetric effects of national highways across genders in developing economies. Lei, Desai and Vanneman, 2019 address this issue on the intensive margin, that is, hours worked. They use a household level panel dataset, the India Human Development Survey (IHDS) collected between 2004-2005 and 2011-2012, to predict the probability of individual working in the non-

farm as compared to traditional farm sector. They find that access provided by paved or unpaved roads and frequent bus service increases the likelihood of employment in the non-agriculture sector for both men and women. These effects seem to be stronger for women. Interestingly, they also find that women's non-farm employment increases in communities with more egalitarian gender norms, measured by the practice of *purdah*.¹ The results in this chapter are estimated at the district level and capture overall changes in local labor market rather than looking at changes at the household-level. The results confirm the broader argument of their paper. Similar to their conclusions, I show that even at the district level rural areas with better road connectivity saw an increase in female labor force participation due to the construction of the GQ project.

Melecky, Sharma and Subhash (2018) use a methodology similar to this chapter to estimate the wider economic benefits of the GQ project. They also argue that initial local market conditions play an important role in determining how certain locations benefit from the highways.² They emphasize that districts with higher levels of secondary education attainment prior to the project construction experienced a significant transition from farm to non-farm employment. In addition, they look at a wider set of outcomes dealing with poverty, household consumption, and environmental impacts. This study extends their analysis by thoroughly examining female labor force participation in particular and investigating the role played by the GQ project in its decline in India for the past

¹According to Lei, Desai and Vanneman (2019), *purdah* is the practice of covering one's face with a piece of cloth like a scarf or shawl. This practice is followed in the presence of men and while stepping out of the household.

²They also cover another national highway project built alongside the GQ project but that was significantly delayed, the North South East West (NSEW) Corridor. I remove districts treated with the NSEW corridor from my sample, and focus solely on the GQ project.

three decades. To improve upon their methodology, I incorporate a decade long pre-construction time series to test for parallel trends between treated and non-treated district groups.

In addition to gender specific effects, this chapter contributes to the literature on the effects of transport infrastructure on rural economic development. One of the few studies relevant in this context, specifically for India, is by Asher and Novosad (2020). They argue that local rural roads facilitate structural transformation by incentivizing workers to move out of agriculture. However, they don't find any major changes in agricultural outcomes, such as income, or assets. Similarly, Shamdasani (2021) finds that road connectivity for remote villages leads to crop portfolio diversification, and farmers start to invest in more labor and capital intensive production methods. Overall, she augments the structural transformation results presented by Asher and Novosad (2020) but adds that agriculture sector grew as opposed to the non-farm sectors for some villages despite road connectivity. I show that the gender specific effects are important, highways may impact certain demographics through radically different channels. I provide evidence that structural shifts from agriculture to non-agriculture sectors may not be happening for women, even when the regional economies are experiencing significant economic growth. I find that the effects of highways is heterogeneous across districts with some districts experiencing a decline whereas others experiencing an increase in FLFP. I propose potential channels that may explain these shifts.

Finally, I provide evidence in support of the U-shaped theory of female labor force participation proposed by Goldin (1990). More specifically, results in this chapter support the predictions of the theory that during the initial period

of economic growth women's labor supply might decline. In the context of India, Afridi, Dinkelman and Mahajan (2018) demonstrate that increased education levels among rural population of men and women leads to a fall in FLFP. They argue that relatively higher returns to home production compared with market production may be driving these results. The results presented in the study suggest that these effects might be more localized and may as well be applicable to the regional economies.

3.3 The Highway Project

To understand the relationship between highways and the local labor markets, I use the introduction of a large infrastructure project in India, the Golden Quadrilateral (GQ) project, as a quasi-experiment to parse out the effects of highways on Indian districts. The GQ project was introduced with the objective of connecting four major metro cities in India located at the four corners of the country. I will use this project to study how it may have affected household incentives through local market connectivity. Briefly, the project was proposed in the year 1999 and the construction was started only in the second half of the year 2000. The project was more than 90% complete by the year 2006. Therefore, I consider the year 1999 as the pre-construction period and the year 2009 as the post-construction period. A more detailed description of the project is given in Chapter 1 of this dissertation.

3.4 Empirical Strategy

This section lays out the empirical strategy to analyse the impact of the GQ project on the regional labor markets. I rely on a standard difference in difference approach to estimate the causal effects of the GQ project. The publicly available

data makes it easier to construct a long time-series of outcomes variables before the construction of the project that can be used to check for pre-construction trends.

I conduct the analysis at the district level. According to the census of 2001, India had 593 districts in total, a few of which were carved out of bigger districts since the country's independence. To keep a consistent panel of time series at the district level, the districts are collapsed to their parent regions and 423 individual districts remain in the data set for the analysis. I allocate districts into four distinct groups based on their distance from the GQ project. The nodal districts lie at the four points that form the corners of the GQ project. The first group contains districts that are 0-40 kms away from the highway, the second group contains districts that are 40-120 kms away from the highway and finally, the last group contains districts 120-500 kms away from the highway. The focus throughout the analysis is on the first group that is assumed to benefit the most from the GQ project. I also consider the second group to take into account any spillover effects from the highway construction. The third group is the control group.³

I follow a simple identification strategy. I compare the labor market outcomes in treated and control district groups before and after the construction of highways. The hypothesis is that district groups closer to the GQ project (0-40kms) will show significant effects in the outcome variables as compared to district groups (120-500kms) further away from the highways. To check for these temporal changes in labor market outcomes at the district level, I estimate the following equation .

³This classification of districts is similar to Chapter 1, to keep a consistent and comparable sample I keep the same classification for this Chapter as well.

$$\Delta Y_i = \sum_{g=1}^3 \beta_g D_{ig} + X_i + \epsilon_i \quad (3.1)$$

Here Y_i is the dependent variable such as district level FLFP. The subscript i represents district, g represents distance group to which the district is allocated based on its distance from the highway. For example, $g = 1$ represents a dummy for districts falling on the nodal points of the highway, $g = 2$ represents a dummy for districts 0-40 kms away from the GQ project. X_i represents district level controls recorded before the highway construction was started.

As discussed later in this chapter, treatment heterogeneity will play a crucial role in determining the causal effects of the GQ project. For example, as per the U-shaped hypothesis, rural areas of the districts along the project with higher presence of the manufacturing industry may experience a decline in FLFP. I capture this treatment heterogeneity by interacting district characteristics with the treatment dummies. Specifically, I interact the D_{ig} dummy with a set of district level characteristics recorded before the GQ construction had started. Consequently, I will also estimate the following equation.

$$\Delta Y_i = \sum_{g=1}^3 \beta_g^1 D_{ig} + \sum_{g=1}^3 \beta_g^2 D_{ig} \times Int_i + Int_i + X_i + \epsilon_i \quad (3.2)$$

Here, Int_i represents the interaction term for each district that captures the pre-construction economic characteristics of the districts such as literacy rate, road connectivity, and caste fractionalization.

Endogeneity Issues: A common issue while estimating the causal effects of highways on economic activity is that highways are usually built where there is already a potential for economic activity to grow. In the context of this project,

female labor participation may not directly influence the placement of highways but it may influence the incidence of highways through other channels. For example, if we assume that gender norms based on cultural factors such as societal outlook towards women participating in market work is positively correlated with the economic output of the region, then regions with higher economic output may see more women participating in the labor force. It is likely that highways will be placed in regions with high economic output or with gender norms favorable for women's participation in the labor force. Therefore, highway placement may not be completely random and may be governed by the district characteristics correlated with the outcome variables of interest.

The general approach to circumvent this identification problem in a difference-in-differences framework is to test for parallel trends between treated and control regions in the pre-treatment period. A precise time series analysis of this kind is not possible because the survey data I use is conducted every 5 years. However, I run placebo tests that mimic the GQ project construction before the actual construction and check if any significant trends are prevalent before the GQ construction started. I discuss this more in the next sections. In addition to this, I take several other precautions to make sure that the estimated coefficients are not biased. I consider regions which were "accidentally" treated with the highway project. In other words, I exclude all nodal districts (the four metro cities) and adjacent regions which were supposed to be connected with the highway project from the treated group. Additionally, I add a set of control variables accounting for district level characteristics that may potentially affect determine how districts respond to the construction of the highway.

3.5 Data

The data I primarily use for this chapter is the household level survey conducted after every 5 years in India, the national sample survey (NSS). The survey is divided into four parts covering various aspects of the Indian household. These parts are named schedules, and each schedule is roughly surveyed every 5 years. I make use of the schedule that covers the employment and unemployment information of each member of the sampled household. The survey is conducted separately for rural as well as urban areas of the country. The rural sample is representative at the district level whereas the urban sample is representative at the sub-state level but not at the district level.⁴ The survey helps construct labor force participation rates and other labor market characteristics for the rural areas at the district level for almost two decades between 1987-2009. I focus on the surveys that cover periods during and around the time when the GQ project was constructed. Surveys relevant for this study are National Sample Survey Office (NSSO)'s 43rd round (July 1995-June 1996), 52nd round (July 1995-June 1996), 55th round (July 1999-June 2000), and 64th round (July 2007-June 2008), 66th round (July 2009-June 2010), 68th round (July 2011-June 2012). In this version of the chapter, I only make use of the 43rd, 55th, and 66th rounds to assess the impact of the GQ project.

To construct the measures for labor force participation at the district level, I make use of relevant survey responses from the household members. For each member, the survey records “current activity status” which has more than 10 different categories that signify the status of an individual. For example, “worked as a casual wage earner”, “sought work”, “did not seek work”, etc. I

⁴The sampling for the urban areas is done from the NSS regions, that are generally defined as an aggregate of 4-5 districts within a state.

use a specific set of responses to categorize each sampled individual into whether she participated in the labor force or not.⁵ The data also allows for capturing the industry in which the worker is employed. For everyone employed, the data provides industry classifications according to the National Industry Classification (NIC). I use these NIC codes to identify the broader sectoral categories such as agriculture, manufacturing, and services where an individual works. Additionally, I use the education and marriage information of the household members to further decompose the effects of the GQ project on different categories. In addition to the NSS data, I also use census information from the 2001 survey to construct district characteristics such as literacy rates, rural population, and caste fractionalization. I extract the census information from Martin, Nataraj and Harrison (2017) publicly provided along with their study. To check the baseline results before the construction of the GQ project I also use census information from the 1991 survey made available in the SHRUG data set.⁶ The SHRUG data set also helps provide information about the number of colleges, schools, and other public amenities (local roads, etc.) at the district level that I use to check heterogeneous impacts of the GQ project. Summary statistics of selected variables are provided in the Table B.1

3.6 Results

This section presents the main results of the study. I first check the trends of the outcome variables in treated versus control district groups before the construction of the GQ project. Next, I present the main results. I start with

⁵For example, I consider an individual to be in the labor force if he/she belongs to either of these categories defined by NSS 55, “self-employed”, “employer”, “unpaid family worker”, “wage employee”, “casual wage laborer”, “other types of work”, and “did not work but was seeking”.

⁶SHRUG stands for Socioeconomic High-resolution Rural-Urban Geographic Platform for India. It is an initiative by economists to consolidate data on the Indian economy from a variety of sources for easy access. More details about the data can be found at its website here.

average treatment effects that simply compare the treated districts with the control districts. This is followed by an analysis of heterogeneous treatment effects where I interact district economic characteristics and the treatment dummies.

3.6.1 Baseline Comparison. To causally interpret the results it is critical to assess the pre-construction behavior of the districts. As mentioned earlier, the difference-in-differences methodology stipulates that the treated and control groups exhibit parallel trends in outcome variables before the treatment takes effect. For this project, data restrictions make it impractical to perform such an analysis. The primary data used for this study, specific schedules of the NSS data, are surveyed only after every 5 years making it harder to visually inspect the trajectory of outcome variables across time. Alternatively, I can run placebo regressions. Placebo regressions repeat the main analysis of comparing the treated and control districts but using a different time period, where the treatment variable remain the same. Any significant changes from the main results will shed light on pre-trends that can be a possible cause of worry for causal inference. I run these placebo regressions for the period between 1987 and 1999 with the same treatment dummies that are used in the main regressions. I use Equation 3.1 to estimate these regressions.

The results of this exercise are presented in the Table 7. The regressions are run separately for men and women. I also add a set of controls that account for distance from the railways and the coastline.⁷ The coefficient estimates are non-significant, except for women in the treated districts. This implies that the districts closer to the GQ project were experiencing a decline in female labor

⁷I add more controls in the main equations that account for several other district characteristics. Since data on district characteristics is not available for the year 1987, I have not included those controls here.

Table 7. Baseline - Placebo Treatment Effects

	Male		Female	
	(1)	(2)	(3)	(4)
	$\Delta FLFP$	$\Delta FLFP$	$\Delta MLFP$	$\Delta MLFP$
Nodal	0.046** (0.020)	0.020 (0.029)	0.015 (0.012)	0.007 (0.012)
Treated	-0.032 (0.024)	-0.054** (0.024)	-0.003 (0.011)	-0.010 (0.009)
Intermediate	-0.010 (0.029)	-0.016 (0.028)	0.013 (0.011)	0.012 (0.008)
District Controls	No	Yes	No	Yes
N	294	294	294	294
R sqr.	0.0057	0.021	0.0036	0.013

Note: This table presents results from placebo regressions run for the period 1987-1999 with the GQ project as treatment and female labor force participation rate as the outcome variable. The coefficient estimates are produced by running Equation 3.1 and standard errors are clustered at the district level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

force participation. These trends are in line with the overall changes that were happening in the female labor markets during that time (see Afridi, Dinkelman and Mahajan, 2018). As I will show in the later sections, this falling trend on average is reversed for the treated districts but hastened after the construction of the GQ project for districts with higher literacy rates .

3.6.2 Main Results. I estimate the average treatment effects by comparing treated and control districts. The results are estimated using Equation 3.1 with labor force participation as the outcomes variable, estimated separately for both men and women in the sample. The results are presented in Table 8. The coefficient estimates of average treatment effects for both men and women in treated districts is positive after controlling for the set of controls. The magnitude of these positive effects is very low. For example, in the case of women, the highway construction only led to 2 percentage points increase in the labor force

participation over the decade during which the GQ project was constructed. It is hard to make any causal claims from these estimates since they are not statistically significant, and consequently indistinguishable from zero. Taking into account the placebo effects estimated earlier, the magnitude of effects in treated districts is now opposite. This shows that there was a significant shift in the trends of labor supply due to the construction of the GQ project. The secular decline for all districts between the years 1987 and 1999 was replaced by a positive or zero net effect between 1999 and 2009.

It is not surprising that the average treatment effects from the GQ project are not significant. The high standard deviations of district economic characteristics presented in Table B.1 suggests that they (treated, intermediate, etc.) were different in many respects. These differences could have played an important role in determining the future impact of the GQ project. For example, the marginal benefits from the highway construction could have been much higher for districts that had relatively low market opportunities available for women before the construction. In other words, there may be diminishing returns to the highway construction with low developed districts gaining the most. In addition, there could be a variety of other channels where initial market conditions interact with the GQ project construction to bring about specific changes that might not get captured in the average treatment effects. I next explore these district-level heterogeneities in terms of economic characteristics of districts recorded prior to the GQ project construction and investigate how they may have determined the trajectory of changes in the local labor markets.

The most important margins along which the highways affect districts are highlighted when we look at the different production sectors separately. Much

Table 8. Average Treatment Effects

	Women		Men	
	(1) $\Delta FLFP$	(2) $\Delta FLFP$	(3) $\Delta MLFP$	(4) $\Delta MLFP$
Nodal	0.062* (0.034)	0.092** (0.040)	0.003 (0.034)	0.026 (0.040)
Treated	-0.019 (0.031)	0.020 (0.033)	-0.012 (0.014)	0.008 (0.013)
Intermediate	-0.013 (0.028)	0.023 (0.028)	-0.026** (0.012)	-0.011 (0.011)
District Controls	No	Yes	No	Yes
N	322	322	322	322
R sqr.	0.0038	0.079	0.0092	0.041


Note: This table presents the main results for the effects of the GQ project on female labor force participation (FLFP). The coefficient estimates are produced by running Equation 3.1 with change in FLFP between 1999 and 2009 as the dependent variable. Standard errors are clustered at the district level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

empirical research has emphasized the structural change in local economies due to road construction (see Asher and Novosad, 2020, Shamdasani, 2021, Aggarwal, 2018). This change shifts surplus labor from agriculture to manufacturing or services due to increases in agricultural productivity facilitated by economic channels opened up by road infrastructure. The decade through which the GQ project was built seems to have gone through similar economic shifts. This is shown in Figure 5 where I plot the overall change in the share of women’s employment for the three production sectors before and after the construction of the GQ project. The shift in the labor share from agriculture to manufacturing and services is very evident.

Taking this into account, I dissect the district-level employment data into three primary production sectors and calculate female labor force participation for each separately. To take into account the role played by initial economic

Figure 5. Women's Employment Share by Industry Sectors



characteristics I then interact the variables representing these characteristics with the district group dummies and finally estimate the regression coefficients defined in Equation 3.2. I estimate these equations separately for the agriculture, manufacturing, and services sector. I consider a variety of initial district level characteristics as interaction terms but only present results for literacy rate, paved roads, and caste diversity. I briefly talk about other parameters at the end but do not emphasize those results because of the lack of statistical significance.

Literacy Rate: I construct the rural literacy rate of districts using the Indian census data of 2001. I take the ratio of the total “literate” population in the rural areas of a district to the total rural population. The definition of literacy in India at an individual level that one must be able to read and write with understanding. The expectations to accomplish this are low, and, on average, anyone with primary education is considered literate. I interact these literacy rates with treatment dummies and run regressions as per Equation 3.2 for each production sector separately. The results are presented in Table B.2. The coefficient estimates are significant (at 1%) for the treated districts. Broadly, the results suggest that FLFP declined more for districts that exhibit higher literacy rates before the construction of the GQ project. This holds for manufacturing as well as the services sector. On the contrary, the agricultural sector saw positive shifts but the coefficient estimates are non-significant.

To interpret these results statistically and compute the overall effects, we can take the coefficient on the interaction terms, multiply it with the literacy rate in a district, and add the overall coefficient representing the effects on district groups. For example, the services sector results outlined in column (6), there is a 3.12 percentage points decrease in FLFP due to the GQ project for districts with

a literacy rate of 70% ($0.057 - (0.126) \times 0.7$). Similarly, one point increase in the literacy rate of a district leads to a decline of 0.126 percentage points in the FLFP. For a clearer exposition, I simulate these interaction effects in Figure 7 for all sectors. These simulations clearly show that FLFP in the manufacturing and services sectors is declining for high-literacy treated districts. The threshold literacy rate lies between 0.4 and 0.6. This implies that, on average, districts with literacy rates higher than the threshold on average will experience a decline in FLFP, whereas, districts with lower literacy rates on average will experience an increase in FLFP.

Since the FLFP is the ratio of labor force and total population in a district, changes in FLFP could be due to shifts in population across districts. A potential concern could be that labor is migrating towards higher literacy districts to access better education and other related amenities. Therefore, the decline in FLFP presented in Table B.2 could be driven by population shifts rather than changes in employment levels or contraction of labor force. To test what is driving these results, I look at changes in the employment levels of districts. The results presented in Table B.3 show that employment levels also decreased in the manufacturing and services sector where rural literacy was higher.

To test whether these results are due to the GQ project instead of unobservable factors already prevalent before the construction, I run the placebo regressions with the treatment dummies interacted with the literacy rate of districts in the year 1991. The outcome variable is similar to before, that is, the change in FLFP between 1987 and 1999. For this analysis, I run the regression Equation 3.2. The results are presented in Table B.4. Here, none of the coefficient estimates are statistically significant. This provides confidence in the fact that there were no


alternate unobservable factors in the pre-treatment period that affected the treated districts.

A possible explanation for these results is the U-shaped behavior of the female labor force participation. The theory first proposed by Goldin (1990), posits that in the initial stages of economic growth the female labor supply declines as production moves from individual households to businesses in the market. This happens because low skilled women generally have strong social stigma attached to manual work that is available outside the household. Decline in labor participation may also be seen for educated women as they may choose to focus more on investing in children's education, an argument that we pursue in Chapter 4.⁸ The U-shaped hypothesis has been tested in the case of India by Afridi, Bishnu and Mahajan (2019) where they use the same data as used in this chapter but perform a country-wide study and argue that the U-shaped behavior is at play in India as well. Their decomposition exercise shows that FLFP in rural India declined between 1999 and 2011. They argue that this decline can be explained by the fact that more number of young women are pursuing education and married women are finding it more beneficial to invest their time in household production.

The literacy interaction results confirm this idea at the district level and show that the infrastructure investments like the GQ project accelerate the movement along the U-shaped curve. Significantly different results for agriculture and non-agricultural sectors help magnify the movement along the U-shape curve. Figure 6 shows that the average skill level of female workers in these sectors is quite different. The services sector has the highest share of labor with graduate education and the agriculture sector has the higher share of people with primary or

⁸It is assumed that women disproportionately bear the responsibility of childbearing in a household and draw away from market production.

Figure 6. Women's Employment Shares by Industry and Skill level

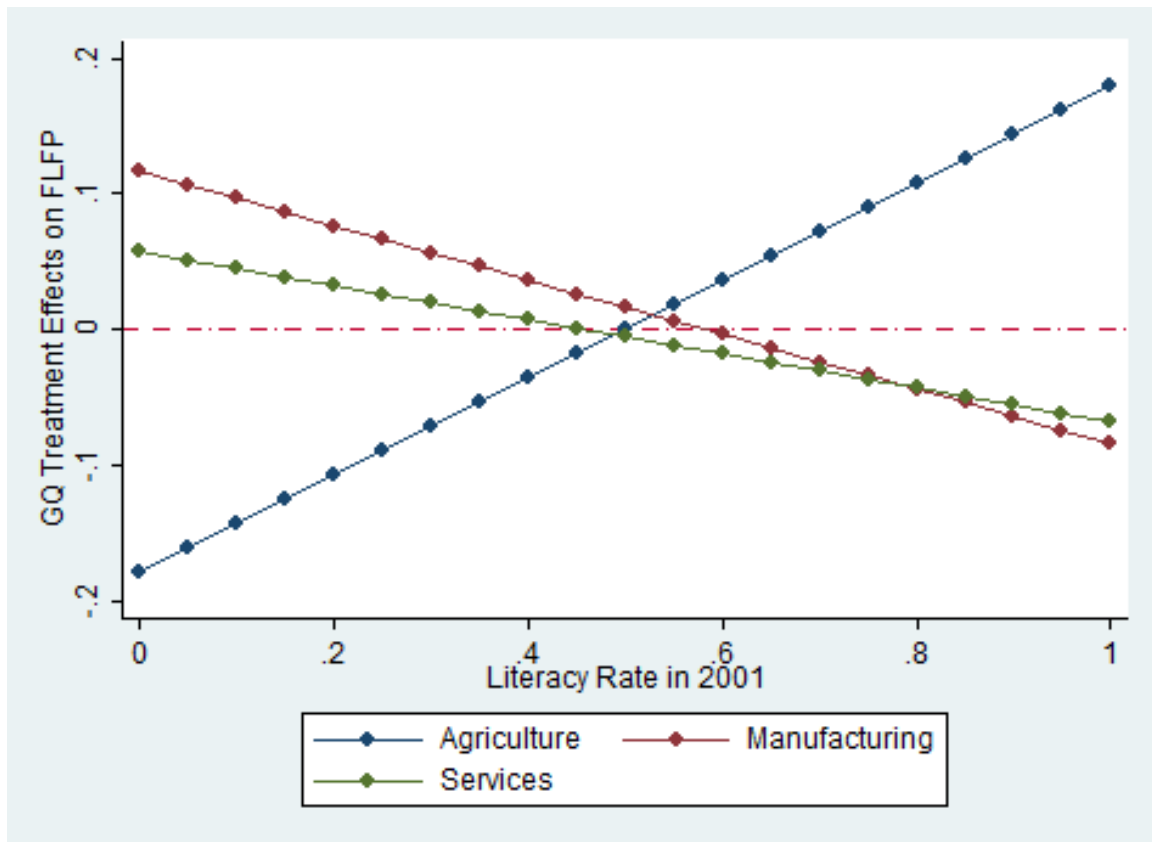


no education. In light of this distribution, we can assume that the declining female labor supply in the services and the manufacturing sector is primarily driven by more skilled women. In contrast, women with low skill levels seem to be engaging in agricultural labor in larger numbers.

In summary, the GQ project led to a movement along the U-shaped curve with FLFP in the manufacturing and services sector declining significantly, and the agriculture sector drawing in more unskilled women into the markets.

Rural Infrastructure: Next, I check whether the existing road infrastructure in the rural areas had a significant role to play in the way districts benefited from the GQ project. I construct district aggregates that capture paved road

Figure 7. Simulation of Literacy Interaction Effects



connectivity to villages in 2001. I use the SHRUG data set to construct these measures. The data provides binary response variables for every village describing whether a village has a paved road connectivity or not. I aggregate these binary responses to the district level and use it to proxy for rural road connectivity of the district. In other words, the measure captures the number of villages that have paved road connectivity in a district. I take log of this variable to interpret the results in percentages.

The results are presented in the Table B.5. Similar to the literacy interaction effects, the effects are more evident in the non-agricultural sectors. However, two differences stand out. First, the coefficients on the interaction terms

are positive for all the sectors, suggesting that female labor force participation increased in districts with better rural connectivity before the construction of the GQ project. Second, the coefficients for the manufacturing and agricultural sectors are non-significant, suggesting that the services sector was the most affected through this channel. Examining the coefficient estimates in column (6) we can see that the magnitude of positive change is high. For example, a one percent increase in rural road connectivity increases the impact of the GQ project on services FLFP by 0.5%. These results could mean that village road connectivity was augmented by the GQ project, opening up wider range of market opportunities for women. Since a major share of high-skilled workers is in the services sector, the results suggest that the nature of these opportunities was specific to the skilled workers.

These results are opposite to the previous exercise with literacy interactions. Ideally, economic growth is assumed to be positively correlated with the literacy rate, rural road infrastructure, and development related characteristics of districts. Therefore, the results in the above analysis should be similar to the literacy interaction effects. In other words, if the FLFP is declining for districts with higher literacy rates then it must decline for districts with better rural road infrastructure.

However, the best possible explanation for the contradictory results is that two different economic forces are counter-acting each other. To understand this, we must note that the cross-sectional variation in district-level literacy rates and rural road infrastructure may not be positively correlated with each other. These economic characteristics might be a product of district and state-level budget allocations where there are trade-offs between education and road infrastructure investments. A simple correlation exercise shows that these two characteristics are negatively correlated in the case of India in the year 2001. Therefore, the GQ

project might be affecting FLFP through different channels that are actuated based on the pre-construction district characteristics. I previously discussed the U-shaped phenomena in the case of literacy interaction effects. The other channel, different than the U-shaped phenomena discussed earlier, could be the opening up of market opportunities. The GQ project could have augmented the already existing rural road infrastructure and opened access to opportunities that may not be accessible through local roads, as a consequence increasing FLFP. This argument is aligned with the literature on intra-country trade where highways provide access to remote markets.

Caste Fractionalization: The literature has actively explored the role played by identity in economic decision-making. In the context of labor supply decisions in India, Oh (2021) argues that these decisions are influenced by the caste association of an individual. She finds that individuals favor jobs that are associated with their caste/community and do not take up even a higher-paying jobs if the job under consideration is traditionally performed by a lower caste. Similarly, theoretical research has pointed out caste networks as a valuable resource that opens up opportunities in urban areas (see Munshi, 2019). These social factors could significantly influence the female labor force participation decisions in the treated districts. It is interesting to see how these social factors interact with the construction of the GQ project and affect the local labor markets. For example, at the district level, a more homogeneous caste population will be able to benefit more from caste networks and increase FLFP. Additionally, this would also mean that individuals have a higher probability of finding jobs that are associated with their own caste group, and hence have lower social costs to bear while taking up a job offer.

To test this hypothesis, I construct a caste fractionalization index using the Indian census data from the year 2001. The index represents caste diversity within a district. Specifically, I use a Herfindahl concentration method to construct caste concentration based on the population share of the scheduled castes, scheduled tribes, and the general category. I subtract this concentration index from 1 and use it as a measure for caste fractionalization. Higher levels of the index indicates more caste diversity. More diversity restricts the set of caste specific job choices female labor can choose from when they participate in the labor force. Therefore, having a negative effect on average district-level FLFP. By the same logic, lower levels of the index may have positive effects on FLFP.

The results presented in Table B.6 suggest that the caste fractionalization did not play a significant role in districts located next to the GQ project. Nevertheless, regressions for the services sector show some support for the evidence underlining the mechanisms discussed by Oh (2021). The magnitude of coefficient estimates is negative for all the sectors, meaning that higher caste fractionalization (more diversity) leads to lower levels of FLFP. The interaction terms are statistically significant only for the services sector, meaning that one index point increase in the diversity of the district leads to 6.6% drop in female labor force participation.

These results can be further interpreted in light of the evidence provided by Oh (2021). She runs a field experiment with 640 male participants to study the effect of identity on labor supply decisions. The study focuses on male subjects, whereas the results presented in this chapter study similar effects for female labor supply. Caste fractionalization is generally interpreted as the probability that two randomly selected people come from different caste backgrounds. Higher

probabilities in this context would lead to more labor market conflicts while taking up a job offer, negatively affecting the overall supply of female labor force in the region.

3.7 Discussion

In the previous section, we saw that the GQ project on average did not have any significant effects on the treated districts located 0-40 kms away. However, we learn that significant heterogeneities across districts determined the future gains from the GQ project. Higher literacy rates and caste fractionalization within districts before the construction of the GQ project led to a decline in FLFP in the treated districts. On the contrary, better road connectivity in the rural areas of the districts led to positive changes in FLFP from the project. In this section, I discuss some of these results in more detail and provide further support for the arguments presented in the previous section.

I argued that the reason for decline in FLFP in districts with higher literacy rates is that the GQ project could be accelerating movement along the U-shaped curve. The main idea behind this theory is that relative returns to home production increase (compared with market production) with increased economic growth and activity due to the highway construction. To further strengthen this argument, it may help to highlight changes in how women choose to allocate time for different activities. The NSS data provides information on whether a household member is pursuing education or engaged in domestic activities, among other activities. However, this data becomes very thin at the district level and makes the analysis infeasible. However, I can test whether this decline is primarily driven by educated versus less educated and married versus unmarried women.

I first categorize district-level female employment according to women's education characteristics. I recognize three categories of education, primary schooling and no education, secondary schooling, and graduate-level education. For each of these education categories, I compute the share of their employment in year 1999 and 2009. I then run similar regressions as before, defined in Equation 3.2. The results of this exercise are presented in Table B.7. As expected, there is a significant decline in the share of women's employment with graduate degrees in districts that had higher literacy rates. Moreover, districts with higher literacy rates also show an increased share of women employed with secondary education. The coefficients on these interaction terms are statistically significant, providing robust evidence in favor of the U-shaped hypothesis.

The U-shaped path of female labor force participation could explain the significant effects we see on educated women. The theory underlines two competing forces that govern women's labor market decisions. First, women are motivated to join market production since it increases household utility by bringing in more income. Second, the opportunity cost of market production is time spent in the household production raising children and performing other domestic duties. The construction of the GQ project led to economic growth in the treated districts (see Khanna, 2016). Households with lower income levels and human capital will respond to this growth by joining the labor force. On the contrary, women in households with higher income levels and human capital will withdraw from market production as they devote more time in children's education that becomes more important with economic growth.

I also check whether this decline is concentrated among married women. Since the choice of the market versus home production would more often arise in

the case of married women, married women would contribute more towards the declining FLFP. Similar to the exercise above, I categorize women employed into married vs non-married categories and compute the FLFP rates. For the married category I also include women who are either divorced or widowed. Finally, I test the change in employment levels of married vs unmarried women caused by the GQ project. The results are presented in Table B.8. Here too, the married women show a significantly larger decline in the female labor force participation in the services and manufacturing sector, augmenting the previously presented results and strengthening the U-shaped hypothesis.

3.8 Conclusion

This chapter presents evidence that the construction of a national highway project led to heterogeneous effects on female labor force participation (FLFP) in districts located closer to the highways. Districts with higher literacy rates, lower levels of road connectivity, and higher caste fractionalization experienced a significant decline in FLFP. This decline is predominantly concentrated in the manufacturing and services sector.

The results are consolidated by proposing two broader theories of structural transformation. First, suggestive evidence is presented in favor of acceleration downward along a U-shaped transformation of FLFP. The highway-led economic growth may have led to appreciation in the returns to children's education, making investment in home production more beneficial. Consequently, skilled women may be withdrawing from the labor force to engage in child rearing duties, lowering the district level FLFP. Second, the highway construction shows signs of opening up market opportunities for households. Districts with lower literacy rates seem to have gained the most from these new opportunities. In addition, the national

highways potentially augmented rural road connectivity, because districts with better road infrastructure before highway construction also saw an increase in FLFP.

In future work, I plan to extend the analysis in this chapter in three important dimensions. First, the key implicit assumption that holds together the explanation for the decline in FLFP is that highways create economic growth. I plan to check this association by capturing economic growth through the luminosity data from the SHRUG dataset and providing robust evidence on how the labor market may have changed on the demand side. Along with this, I also plan to strengthen the asymmetric effects on women by incorporating an analysis of change in income levels and household consumption patterns data collected from the NSS surveys. This will help shed light on the household level decision making may have influenced changes at the extensive margin, at the district level. Second, I wish to explore the kind of market opportunities that were created due to the GQ project and how they might be different from the market access provided by the local rural roads. Third, robustness checks are needed to provide confidence in the causal claims made in the chapter. The first set of robustness checks will incorporate more data from the NSS surveys that are not covered in this chapter. This will provide an up-to-date complete trajectory of FLFP in time and help with additional placebo regressions. Taking into account significant heterogeneity in the treatment effects, I also plan to estimate the causal effects using alternative techniques such as a combination of matching and difference-in-difference methods.

CHAPTER IV

FEMALE LABOR FORCE PARTICIPATION IN INDIA

This chapter is co-authored with Dr. Shankha Chakraborty. The empirical analysis and numerical simulations were done by me under the supervision of Dr. Chakraborty. The analytical model was developed by Dr. Chakraborty, and I provided proof-checking assistance. The writing of sections 4.2.1, 4.2.2, and 4.5 was done by me, and the rest of the chapter was written by Dr. Chakraborty.

4.1 Introduction

A central goal of economic development is to improve the living standards and opportunities of a population. In practice, those improvements are unevenly spread across socio-economic groups. Of interest is how they affect disadvantaged groups, particularly women.

This chapter studies how economic development facilitates the participation of women in the market economy and dissolves gender norms. It is reasonable to think that economic development will raise wages for women and, through reduced fertility, shift their time allocation from demanding childcare to compensated market work. One measure of such work is the female labor force participation (FLFP) rate, defined as the proportion of women 15 years and older who are working. During 1880-1998, income (GDP) per person in the U.S. increased by a factor of 11, full-time earnings of American women rose from 46% to 67% of men's earnings and fertility rates for white women fell from about 3.5 children to 1-2 (Bailey and Hershbein, 2018; Goldin, 1990; Maddison, 2001). During that same period, the FLFP rate steadily rose from 2% to more than 70% in the U.S. (Fernández, 2013). This general increase in FLFP is observed in many other advanced economies.

Contrast this to the case of India, an emerging economy. Between 1990 and 2016, India's GDP per person grew by a factor of 3.5 and fertility fell from 4.5 to 2.33 per woman. Over the same period, India's FLFP rate dropped from 30.38% to 23.18% (Figures 8 and 9). Why?

Recent work has been documenting how this drop in India's FLFP is occurring in several sectors, education groups and religious communities. Some researchers also provide explanations relying on within-family resource allocation and weak job opportunities for women (for example, Afridi, Bishnu and Mahajan, 2019, Klasen and Pieters, 2015, Neff, Sen and Kling, 2012). What is lacking, though, is a comprehensive understanding of women's labor supply decisions in an economy that is undergoing rapid structural change. In particular, some commentators have equated the falling FLFP with worsening gender equality (Pande and Moore, 2015) and to treat the Indian experience as unique. Neither is necessarily true.




Figure 8. Female Labor Force Participation Rate, India

To understand why not, compare the Indian case with contemporary and historical ones. During 1990-2016, as China's GDP per person grew by a factor of 11.8, its FLFP rate fell from 73.20% to 60.64%, even more than India's. Morocco's




Figure 9. Fertility rate 1990-2018

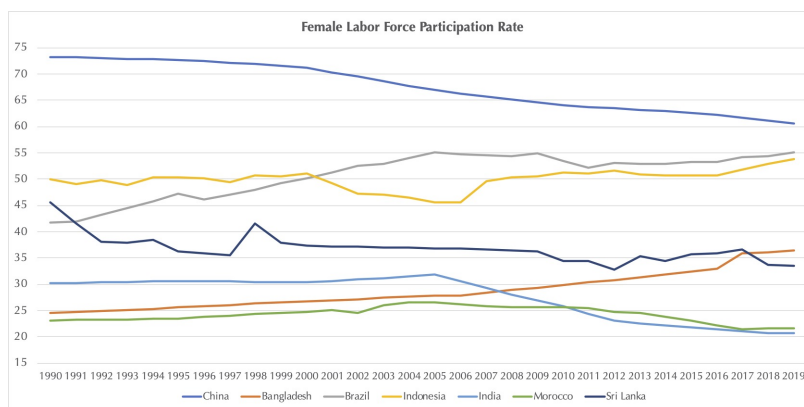


Figure 10. FLFP 1990-2019

FLFP rate, on the other hand, first rose from 23.18% (1990) to 26.56% (2004) before falling to 22.27% (2016) during a time when its GDP per capita increased by a factor of 1.86.¹ The pattern across countries, as Figure 10 shows, is not uniform and falling FLFP is hardly unique to India. In fact, falling FLFP is also observed in historical data. Goldin (1990) finds, for example, that the labor force participation rate of married – as opposed to all – women in the U.S. declined from the late-nineteenth to early-twentieth centuries before increasing. A similar decrease was observed in late-nineteenth century England when more and more educated women were staying home to raise families (Clark, 2015). Viewed in this light, decreasing FLFP may be a normal pattern of economic development as long as it is temporary, that is, soon followed by increasing participation. To the extent that one can extrapolate time series implications from cross-sectional data, Figure 11 shows that the association between FLFP and economic development has been non-monotonic.

In this chapter, we propose a tractable analytical framework of household decision-making that identifies factors behind a potentially U-shaped FLFP. The model has several ingredients. Decisions are made in a unitary household on consumption, fertility, child investment and labor supply by the wife. The latter is subject to social costs: specifically, the very decision to work outside the house involves a social stigma that the household internalizes.² The wife's decision to work therefore depends on balancing the obvious benefits – higher household income – with the costs – time away from childbearing and care giving, and the

¹Contemporary FLFP data from the International Labor Organization, income and fertility data from World Development Indicators (World Bank).

²Alternatively, the wife internalizes her husband's resistance to the idea of her working. This may not be as far fetched as it sounds; see Lowe and McKelway (2019) for experimental evidence in a different context.




Figure 11. FLFP and per capita income, Afridi, Bishnu and Mahajan, 2019

psychic cost of breaking a social taboo. The tradeoff regarding time away from children is particularly salient: working outside means not just less time in raising young children, but also time away from monitoring them, looking after them and other time investments that augment their human capital. Anecdotal evidence and commentaries, for example, point to these as important margins for urban Indian women (M. B. Das and Zumbyte, 2017).

The tractable household model outlined below shows three things. First, unchanging culture in the sense of unchanging attitudes towards women’s work can interact with economic forces in interesting ways to produce a non-monotonic response of the female labor force participation rate to economic development. Secondly, child quality investment within the household is an important margin – particularly where fertility rates are falling – affecting women’s decision to work. When households are primarily investing mother’s time in quality investment, labor supply can be particularly costly unless wages were to be sufficiently high. Third, FLFP is not necessarily non-monotonic. This depends on the importance of culture

and level of development: societies less resistant to the idea of working wives and mothers and/or at a level of development where the gender wage gap is relatively low would see FLFP rates rise with further development.

Section 4.2 below presents some empirical evidence on India's FLFP and related contributions to the literature. The household model is presented in section 3 which is used to construct the aggregate FLFP rate in section 4. Section 5 outlines the next steps we plan to take in the future and ideas for an extension.

4.2 Background

4.2.1 Additional Facts. Much has been written about the secular increase in FLFP in the advanced economies since industrialization. Goldin's (1990) seminal work on the US labor market since the late nineteenth century has informed modern commentaries about the challenges and nuances of labor force participation of married women.

The quirkiness of India's FLFP during last three decades of robust growth has, not surprisingly, attracted a lot of attention. India's overall FLFPR has declined from an already low of 22% in 1987. If we also consider the status of subsidiary activities then, Sarkar, Sahoo and Klasen (2019) show that the overall FLFP has declined from 51.4% in 1983 to 38.7% in 2011 for the age group of 25-55. These changes when decomposed at the rural and urban level, show contrasting pictures.

In the rural areas, the decline has been more significant among the age group of 25-65 years old. Afridi, Dinkelman and Mahajan, 2018 perform a decomposition analysis of the National Sample Survey (NSS) data and conclude that the higher decline in rural areas is because of an increased participation in domestic work and higher levels of education attainment. They conjecture that

only having primary education is not sufficient, as it provides skills complementary to household domestic work. Therefore, the opportunity cost of participating in the labor force has increased, driving down rural FLFP.

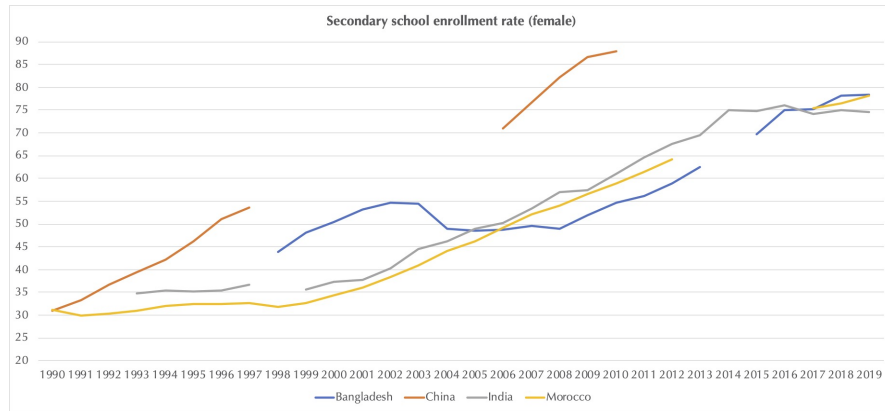


Figure 12. Secondary enrollment 1990-2019




Figure 13. LFPR by education (urban, married and not married, age 20-45) Afridi *et al.* (2020)

Another quirk of India’s FLFP has been the weak association between education and work. It stands to reason that rapid economic growth creates both demand for skilled labor and incentives to acquire education. As in several other Asian economies, female education has been steadily rising in India during 1990-2019 (Figure 12). Figure 14 plots changes in school level educational attainment




Figure 14. Changes in female educational attainment in India

for women in India above the age of 25. There has been a consistent increase in educational attainment at all levels of schooling, primary, secondary and tertiary. Also, a simultaneous decrease in the number of females who never went to school has declined significantly between 1985 and 2010. Yet FLFP of educated women has stagnated.

This is especially clear when we look at urban areas. The urban population has more graduates and fewer illiterates, while the opposite is true in rural areas. Therefore, we would expect an increase in labor supply and as a consequence higher FLFP in the urban areas. FLFP for married women between the age group of 25-55 there fell slightly from 18.5 % in 1987 to 17.9% in 2011, even as it declined significantly in rural regions (Klasen and Pieters, 2015).

Basu and Desai (2012) and M. B. Das and Zumbyte (2017) find that children in one-child families to be advantaged: they are more likely to be sent to private schools and English medium schools, and more likely to receive private

tutoring in addition to schooling, than children from larger families. Yet women in one-child families are less likely to be employed than those in larger families. Following on this theme, M. Das and Desai (2003) argue that standard economic factors such as job opportunities may have more to do with low FLFP in India, than cultural factors.

4.2.2 Cultural values. To get a sense of the role of culture, we look at the changing values and perceptions of Indian households. The *World Values Survey* (WVS) provides comprehensive data on multi-dimensional values and belief systems people hold true in different countries. We look at two specific variables from this dataset to gauge how values may have shifted in past two decades.

The WVS asks relevant questions to the participants and gives them choices to select their answer from. Figure 15 shows responses to the question whether or not men should have more right to a job than a woman. In 2010 more people have shifted from disagreeing with the statement to being neutral, that is, more people became undecided whether or not men have more right to a job than women. In contrast, in 1990 a significant number of people disagreed to the statement. While this change may be due to shifting cultural attitudes against FLFP, it is also possible that it reflects more competition in the labor market and limited job availability. In such a scenario, holding cultural norm constant, men may prefer to be chosen for the job rather than an equally qualified woman.

Another variable of interest in the WVS is how much importance households place on their children's education and specifically, on their daughter's. Figure 16 report survey data on the question whether university education is more important for a boy than a girl. Responses do not conclusively point in one direction. We see that between 1995-2005 there is a decline in population which strongly disagrees



Figure 15. Changes in values if men should have more right to jobs

with the fact that boys deserve more education than girls and simultaneously an increase in population which agrees with the statement.

4.2.3 Theoretical works. Two sets of theoretical works bear upon the analysis in this chapter. One set studies household decision-making under alternative assumptions of the decision-making process. These range from collective models of the household (a special case of this, the unitary model, is widely used in the macro literature and adopted below), to cooperative models where the bargaining strength of the spouses determine allocations (for example, Heath and Tan (2020)), to non-cooperative models where the conflict arises from preference mis-match and the sequence in which decisions are made.³

³A popular framework that has been increasingly adopted is Lundberg and Pollak's (1993) "separate spheres" model where each spouse's outside option involves them spending on their private good and the household public good they care more about.

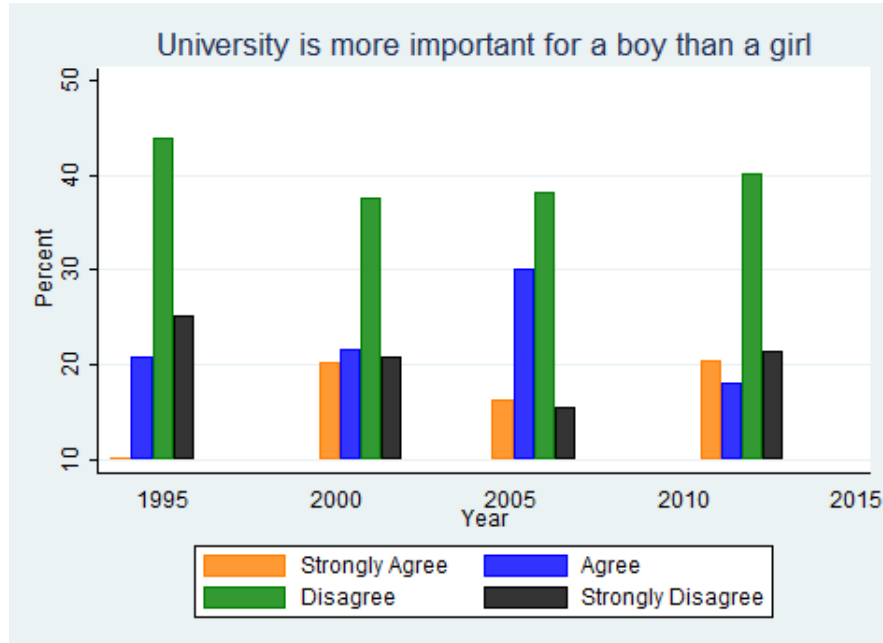


Figure 16. Changes in values if university is important for a boy or a girl

Our work has less to do with this body of theoretical work on household decision-making and more to do with analyses on gender and development itself. This literature is relatively sparse but a few papers are closely related to ours. Hiller, 2014 proposes an OLG model where inegalitarian cultural norms affect parental investment in children’s human capital: “strong norm” parents invest less in daughters. During the early stages of development, two opposing forces are at work. There is a positive income effect that raises educational spending on all children, while a negative cultural effect delays parental investment in daughters’ human capital. The author shows that the interaction between endogenous inegalitarian norm and the gender gap in education can generate multiple equilibria. Initial conditions determine convergence to either an inegalitarian norms development trap, or to an egalitarian norms affluent society. More relevant for our purposes, along convergence path, FLFP follows an U-pattern.

In Chichilniski and Frederiksen's (2009) model of the gender wage gap, women's productivity in market work decreases with an increase in time devoted towards child-bearing activities. The model produces multiple equilibria: expectation of women's lower quality work is self-fulfilling as the lower offered wage prevents women from engaging in market work, thereby keeping their productivity low. The evolution of the gender pay gap is the focus of Galor and Weil's (1996) model. They assume that women specialize in mental labor, men in physical labor. As capital accumulation rises with economic development, returns to women's labor rise faster because mental labor complements capital in the production process. Rising women's relative wages, in turn, reduces fertility by raising the cost of children more than household income. Lower fertility raises capital per worker even more. These effects are complemented by increasing education among women that augments their productivity in mental labor; see also Kimura and Yasui (2010).

4.3 Household Decisions

Consider a static framework where an economy is populated by a continuum of unitary households. Households are heterogeneous along two dimensions, human capital and culture.

Each household consists of a husband (m) and wife (f) who pool their resources and jointly maximize the household's welfare. This involves choosing household consumption, fertility, child investment and wife's labor force participation. Culture enters the model through resistance to the idea of the wife working outside the household. This is parameterized as χ , the subjective cost or social stigma experienced should the wife decide to work. Besides paid market work, women have the option staying home to raise and care for children and/or engaging in unpaid home production, distinct from childcare. Home production of

the *non-market* good is valued differently (imperfect substitutes) from consumption of the market good c . The wife's opportunity cost of raising children depends on her immediate alternative: home production or market production.

Let the human capital of the wife be represented by h_f and of husband by h_m . A household is characterized by the vector (h_f, h_m, χ) . Denote by w_f and w_m the wage rates per unit of human capital for female and male workers, respectively; we will refer to these as efficiency wage rates. The joint decision-problem for the household is

$$\max \ln c + b l_h + \gamma [\theta \ln n + (1 - \theta) \ln q] - \chi I$$

s.t.

$$c = w_m h_m + w_f h_f l_f - \phi e n$$

$$l_h + l_f = 1 - (\tau^n + \tau^q) n$$

$$l_h \geq 0, l_f \geq 0$$

$$q = \max \{q_0, (e + a\tau^q h_f)^\alpha\}, \alpha \in [0, 1]$$

$$I = \begin{cases} 0, & \text{if } l_f = 0 \\ 1, & \text{if } l_f > 0 \end{cases}$$

It is assumed that $\theta > \alpha(1 - \theta)$ and $b > 1$. l_f and l_h are the wife's time allocations to market and home production. Home production requires time away from child rearing, besides withholding labor supply from the market.⁴ We assume that only the wife contributes to child rearing and education time costs. It can be shown that the household will optimally choose to do so as long as $w_m > w_f$.

⁴This serves an important objective: if $l_f = 0$, a mother's opportunity cost of time is no longer the market wage rate and we need her time to be used for something else to get the opportunity cost effect.

The production of child quality q takes as input mother's education time (attention to children's welfare) τ^q , and market-provided educational goods and services (school quality for example) e . These two inputs are perfectly substitutable. This is a simplification: as long as the two inputs are substitutable to some degree, child quality will be time-intensive at low levels of income (specifically w_f), resource-intensive at higher levels of income.⁵

We analyze household decisions sequentially. First we derive their behavior based on whether or not the wife chooses to work, then we look at the female labor supply decision on the extensive margin.

4.3.1 Case I: $I = 0$. For $l_f = 0$, we can simplify the decision problem to be

$$\max \quad \ln c + bl_h + \gamma [\theta \ln n + \alpha(1 - \theta) \ln (e + a\tau^q h_f)]$$

s.t.

$$c = w_m h_m - \phi en$$

$$l_h = 1 - (\tau^n + \tau^q) n \geq 0.$$

from which follows the optimality conditions

$$\frac{\phi e}{c} + b(\tau^n + \tau^q) = \frac{\gamma \theta}{n} \quad (4.1)$$

$$-bn + \alpha\gamma(1 - \theta) \frac{ah_f}{e + a\tau^q h_f} \leq 0, \tau^q \geq 0 \quad (4.2)$$

$$-\frac{\phi n}{c} + \alpha\gamma(1 - \theta) \frac{1}{e + a\tau^q h_f} \leq 0, e \geq 0 \quad (4.3)$$

for n , τ^q and e respectively. This leads to one of two outcomes depending on wages and human capital of the spouses.

⁵Introducing q_0 in the quality production functions nests the possibility of no investment in child quality. This assumption is often used in fertility models to produce a larger fertility gap between parents who do not invest and parents who do. In what follows, we will assume this is not binding. It may become necessary for computational work later.

For $w_m h_m \leq a\phi h_f/b$, the household invests only time in child quality production ($e = 0$) and

$$\begin{aligned}
n &= \frac{\gamma[\theta - \alpha(1 - \theta)]}{b\tau^n} \\
\tau^q &= \frac{\alpha(1 - \theta)\tau^n}{\theta - \alpha(1 - \theta)} \\
c &= w_m h_m \\
q &= \left[\frac{a\alpha(1 - \theta)\tau^n h_f}{\theta - \alpha(1 - \theta)} \right]^\alpha
\end{aligned} \tag{4.4}$$

For $w_m h_m > a\phi h_f/b$ on the other hand, only resources go into child quality production:

$$\begin{aligned}
n &= \frac{\gamma[\theta - \alpha(1 - \theta)]}{b\tau^n} \\
e &= \frac{\alpha(1 - \theta)}{1 + \alpha\gamma(1 - \theta)} \frac{b\tau^n}{\phi[\theta - \alpha(1 - \theta)]} w_m h_m \\
c &= \frac{w_m h_m}{1 + \alpha\gamma(1 - \theta)} \\
q &= \left[\frac{\alpha(1 - \theta)}{\theta - \alpha(1 - \theta)} \frac{b\tau^n}{\phi[\theta - \alpha(1 - \theta)]} w_m h_m \right]^\alpha
\end{aligned} \tag{4.5}$$

Since the wife does not work outside, only the husband's labor income matters for whether or not the household can afford resource investment in children.

4.3.2 Case II: $I = 1$. Under $l_h = 0$, the decision problem

$$\max \ln c + \gamma [\theta \ln n + \alpha(1 - \theta) \ln (e + a\tau^q h_f)] - \chi$$

s.t.

$$c = w_m h_m + w_f h_f l_f - \phi e n$$

$$l_f = 1 - (\tau^n + \tau^q) n \geq 0$$

leads to the FONCs

$$\frac{\phi e + (\tau^q + \tau^n)w_f h_f}{c} = \frac{\gamma\theta}{n} \quad (4.6)$$

$$-\frac{w_f h_f n}{c} + \alpha\gamma(1 - \theta)\frac{a h_f}{e + a\tau^q h_f} \leq 0, \tau^q \geq 0 \quad (4.7)$$

$$-\frac{\phi n}{c} + \alpha\gamma(1 - \theta)\frac{1}{e + a\tau^q h_f} \leq 0, e \geq 0 \quad (4.8)$$

for n , τ^q and e respectively. Similar to above, when $w_f \leq a\phi$, only time is invested in child quality and we have

$$\begin{aligned} n &= \frac{\gamma[\theta - \alpha(1 - \theta)]}{(1 + \gamma\theta)\tau^n} \left(1 + \frac{w_m h_m}{w_f h_f}\right) \\ \tau^q &= \frac{\alpha(1 - \theta)\tau^n}{\theta - \alpha(1 - \theta)} \\ c &= \frac{w_m h_m + w_f h_f}{1 + \gamma\theta} \\ q &= \left[\frac{a\alpha(1 - \theta)\tau^n h_f}{\theta - \alpha(1 - \theta)}\right]^\alpha. \end{aligned} \quad (4.9)$$

For $w_f > a\phi$, $\tau^q = 0$ and we have

$$\begin{aligned} n &= \frac{\gamma[\theta - \alpha(1 - \theta)]}{(1 + \gamma\theta)\tau^n} \left(1 + \frac{w_m h_m}{w_f h_f}\right) \\ e &= \frac{\alpha(1 - \theta)}{1 + \alpha\gamma(1 - \theta)} \frac{\tau^n w_f h_f}{\phi} \\ c &= \frac{w_m h_m + w_f h_f}{1 + \gamma\theta} \\ q &= \left[\frac{\alpha(1 - \theta)}{\theta - \alpha(1 - \theta)} \tau^n w_f h_f\right]^\alpha \end{aligned} \quad (4.10)$$

Here, the husband's income does not matter for whether or not the household invests resources in child quality. Because time investment in children is undertaken only by mothers in this setup, the relevant tradeoff is whether or not the female wage rate is high enough: for a low wage rate, the opportunity cost of mother's time is low and working wives are also in charge of child quality production.

4.3.3 Extensive margin decision: to work or not. Armed with these optimality conditions, the decision on whether or not the wife works outside

the household boils down to comparing the indirect utility functions under $I = 0$ and $I = 1$. Observe, however, that this comparison will depend on which of the parametric conditions for child quality production hold hold. We analyze two scenarios below.

The model predicts gender equality in education, that is, same educational outcomes for boys and girls in a household. We assume on top of that positive assortative matching on the marriage market so that $h_f = h_m = h$. This reduces the dimensionality of the household's characteristics vector to (h, χ) . Also assume that $\mu \equiv w_f/w_m \in (0, 1)$ which we call the *gender wage ratio* or wage gap.

Scenario 1: Case I.A versus Case II.A. Recall that Case I.A applies when $w_m h_m \leq a\phi h_f/b$ and Case II.A when $w_f \leq a\phi$. We are, therefore, considering the scenario where both apply, that is,

$$w_f < w_m < a\phi$$

assuming $b > 1$ and $h_m = h_f$. The wife decides to work as long as

$$\ln\left(\frac{c^1}{c^0}\right) + \gamma \left[\theta \ln\left(\frac{n^1}{n^0}\right) + (1 - \theta) \ln\left(\frac{q^1}{q^0}\right) \right] \geq \chi + b l_h^0$$

where variables with superscript 0 refer to decisions (4.4) and those with superscript 1 to decisions (4.9) under $h_m = h_f = h$. On the RHS of the condition above are the social stigma faced and the utility from home production given up when the wife chooses to work. The first term on the LHS is the utility gain from consumption when the wife works: clearly working outside brings in additional income, and that gain is higher for a higher gender wage ratio:

$$\frac{c^1}{c^0} = \frac{1}{1 + \gamma\theta} \frac{w_m + w_f}{w_m} = \frac{1 + \mu}{1 + \gamma\theta}.$$

The second and third terms on the LHS represent the tradeoffs regarding quantity and quality of children from the wife working. On the one hand, more time to

the labor market detracts from time available to raising children which lowers the household's utility. Here the loss is higher, higher is the wage ratio:

$$\frac{n^1}{n^0} = \frac{b}{1 + \gamma\theta} \frac{1 + \mu}{\mu}.$$

On the other hand, the household may want to compensate for the utility loss from fewer children by investing more intensively in them when the wife works. As it turns out, because in both situations it is the mother who is investing her time in children's human capital formation, the net effect is a wash:

$$\frac{q^1}{q^0} = 1.$$

Simplifying these expressions, we arrive at the decision that $l_f > 0$ as long as

$$G(\mu) \equiv \ln(1 + \mu) + \gamma\theta \ln \left[\frac{1 + \mu}{\mu} \right] > \Gamma \quad (4.11)$$

where $\Gamma \equiv \chi + b + \ln(1 + \gamma\theta) - \gamma\theta$. The associated female labor supply function for the household

$$l_f = 1 - \frac{\theta}{1 + \theta} \left[1 + \frac{1}{\mu} \right] \quad (4.12)$$

is increasing in the female wage rate, decreasing in the male wage rate; the latter is due to the usual opportunity cost margin.

Note from (4.11) that χ and b enter additively: cultural stigma is indistinguishable from the cultural value placed on the enjoyment of home production. Both discourage women's labor force participation similarly. Two points are worth making about this. First, this implies cultural resistance to FLFP can be broader than just social stigma against wives and mothers working outside. Secondly, in some societies the cultural value of household production can be as strong as, if not stronger than, stigma. Take this example from Dasgupta (1997):

... the Pathan tribeswomen of a particular village in the Northwest Frontier Province (Pakistan) have traditionally walked some 8 km each way to a small

aqueduct to wash clothes once a week. In order to ease the plight of these women, a foreign aid project built a traditional washing facility nearby. The project was a failure. While the women admired it, they refused to use it. A female anthropologist, employed to discover why the facility, remained unused, discovered that the tribeswomen are rarely permitted to leave their homes. They spend most of their lives within the confines of the mud castles owned by each Pathan family. These weekly trips were their only chance to talk to the other women in the community, when they could laugh, gossip, and play, The event lasted the whole day, and was much looked forward to. The women also reasoned that if they were to complete the laundry in less time their husbands would merely find more things for them to do.

Rising male wages can allow wives to more intensively specialize in household production of children and goods. Even if the female wage rate is rising at the same time, this margin may be unaffected for households that strongly value home life.

Returning to (4.11), observe that it does not depend on the household's human capital h , an outcome of the positive-assortative-matching assumption. As such the household's potential income, more precisely whether or not the household is (potentially) rich or poor in terms of human capital (h_m, h_f) , has no bearing on the female labor supply decision. What does produce heterogeneity of outcomes is heterogeneity in χ : all else equal, higher χ (Γ) households withhold female labor supply.

Conversely, given a χ , higher wage ratio (lower gender wage gap) has a non-monotonic effect on the labor market participation decision as illustrated in Figure 17. If $\mu < \mu_L$ or $\mu > \mu_H$, then $I = 1$. In this range, l_f increases in w_f from (4.12). The tradeoff between staying at home and working for a wage depends on opportunity cost and productivity in home production, and working mothers' involvement in non-market production depends on whether or not they invest time in children's education. To see this clearly, first note that the non-monotonicity




Figure 17. LFP decision at the household level

of G depends on the increasing function $\ln(1 + \mu)$ and on the decreasing function $\ln(1 + 1/\mu)$. The former comes from the household's relative utility gain from higher household consumption when the wife works (clearly higher when wage ratio is higher), the latter is the relative utility loss from restricting family size were the wife to work. Starting from low values of μ , the latter dominates: because μ is low, the household's consumption gain is not particularly high, while the need to restrict family size owing to the mother's time commitments on the labor market is very costly.

Scenario 2: Case I.B versus Case II.A. Suppose now

$$w_m > a\phi/b \text{ and } w_f < a\phi$$

which requires a lower wage ratio than the previous case. A more complicated tradeoff arises here since the wife would not be investing time into child quality

production if she were not to work. As before, participation is positive as long as

$$\ln\left(\frac{c^1}{c^0}\right) + \gamma \left[\theta \ln\left(\frac{n^1}{n^0}\right) + (1 - \theta) \ln\left(\frac{q^1}{q^0}\right) \right] \geq \chi + bl_h^0$$

or, substituting in optimal decisions, when

$$G(\mu) \equiv \ln(1 + \mu) + \gamma\theta \ln\left(\frac{1 + \mu}{\mu}\right) \geq \Gamma + \kappa(w_m) \equiv \Omega(w_m) \quad (4.13)$$

where

$$\begin{aligned} \kappa(w_m) \equiv & \alpha\gamma(1 - \theta) [\ln w_m - \ln(a\phi)] + (1 + \gamma\theta) \ln(1 + \gamma\theta) \\ & - [1 + \alpha\gamma(1 - \theta)] \ln [1 + \alpha\gamma(1 - \theta)] + \alpha\gamma(1 - \theta) \ln b, \end{aligned}$$

and the associated female labor supply function for the household is the same as (4.12). By participating in the labor market, women face w_f as the time cost of child bearing: as w_f goes up, it lowers fertility and lowers utility. In the previous case the lower fertility was not compensated by an increase in q , now it is. But the increase is less than the decrease because $\theta > \alpha(1 - \theta)$. Hence the gain in consumption has to be high enough to offset the loss of utility from children. That, in turn, requires $\mu = w_f/w_m$ to be high enough.

This illustrates another channel that affects FLFP besides the conventional income channel and culture: the ability to substitute child quality for quantity. Moreover, here this ability depends on ϕ : a higher cost of education lowers q^0 , which *encourages* labor force participation because in doing so, the household switches to mother's time-intensive quality production which can raise child quality significantly if w_f is not high enough. In (4.13), Ω is decreasing in ϕ : all else equal, the constraint is therefore more likely to be satisfied for higher ϕ .

What is especially interesting in Scenario 2 is the Ω is now endogenous through its dependence on the male wage rate w_m . This is because when the wife does not work, it is the husband's labor income that is relevant for financing child

quality investment. In choosing to work, if the female wage rate is not high enough, the wife has to balance her labor time with child quality time: higher is w_m , the more does the wife need to devote time to produce a commensurate level of child quality and, therefore, less time she has to work. Given a w_m , changes in w_f affect the decision to supply labor similar to Figure 17 above: labor supply is positive for $\mu < \mu'_L$ and $\mu > \mu'_H$ where (μ'_L, μ'_H) are suitably defined. Were w_m to change too, it would, *all else equal*, disincentivize female labor supply at the margin.

We leave it to future work to analyze this particular scenario. In light of the evidence discussed in section 2, it may be especially relevant to understanding how urban, relatively educated women respond to changing labor market conditions.

4.4 Aggregate Female Labor Force Participation

Returning to the case above, we next construct the aggregate female labor supply decision. Let us assume that $\chi \sim U[\chi_L, \chi_H]$ with a mean value $\bar{\chi} \equiv (\chi_H - \chi_L)/2$. Given a wage ratio μ , as we saw in Figure 17, those households for whom $G(\mu) \geq \Gamma$.

To understand how this works, consider Figure 18, that shows only the low- and high-types as described by χ_L and χ_H respectively, corresponding to Γ_L and Γ_H . For a wage ratio $\mu < \mu_L$ or a wage ratio $\mu > \mu_H$, both high-cost and low-cost households are willing to supply female labor. For intermediate values of the wage ratio, only the low-cost household is. As a result, the extensive margin at the aggregate level looks like the lower panel of the Figure.⁶

⁶There is an intensive margin response: below μ_L and above μ_H , and increase in w_f holding constant w_m increases the hours supplied by a household.




Figure 18. Aggregate FLFP, $\chi \in \{\chi^L, \chi^H\}$

It follows that with a continuum of household types on $[\chi_L, \chi_H]$, the aggregate labor force participation rate

$$\ell_f = \int_i I_i d\chi_i$$

will be U-shaped with respect to the gender wage ratio as seen in the lower panel of Figure 19. Of course, the aggregate labor supply is a function of both extensive and




Figure 19. L_f with continuum of types, effect of culture

intensive margins, the latter being an increasing function of μ as long as $I = 0$. The figure is drawn under the assumption the extensive margin dominates.

It is important to note that the theory does *not* predict the aggregate FLFPR is always non-monotonic. For example, were Γ^H is below $G(\mu)$ in Fig 17, then all households choose $I = 1$ and therefore, via (4.12), $\ell_f = 1$. Or, if an economy were to start from relatively high gender wage ratio, say higher than μ_l , then an increase in the relative wage can only increase ℓ_f .

Moreover, while cultural resistance by itself does not change in response to μ , it can explain a low level of FLFP rate. An increase in $\bar{\chi}$, for example, shifts the distribution of Γ_i up, making previously indifferent households opt for no participation and lowering the aggregate FLFP rate.

Based on the discussion from the previous section, we summarize the key features of ℓ_f : Comparative statics with respect to b, ϕ, α

- An increase in χ and/or b lowers the net benefit of FLFP, increases μ_L , decreases μ_H and lowers ℓ_f ;
- An increase in ϕ affects FLFP only if the household were spending resources in child quality investment in the absence of FLFP. In this case, an increase in ϕ , can increase ℓ_f ;
- An increase in the return to education, α , has ambiguous effects on FLFP;
- An increase in the gender wage ratio μ from below μ_L (low) to above will produce a non-monotonic effect on ℓ_f ; conversely, ℓ_f will increase if the economy starts with relatively high gender wage ratio.

4.5 Numerical Illustration

This section numerically illustrates the household model presented in the previous sections. The illustration simulates a U-shaped FLFP with respect to

Parameter	Parameter Explanation	Parameter Values
b	Utility from Home Production	0.6
θ	Elasticity of Substitution (Child Quantity and Quality)	0.6
γ	Utility from Children	1.4
χ^h	Negative Utility from Female Market Production (Max.)	0.905
χ^l	Negative Utility from Female Market Production (Min.)	0.897
ϕ	Costs for child-rearing	0.5
α	Child Quality Production Function	0.5
a	Child Quality Production Function	10000
τ_n	Time spent child-bearing	2

Table 9. Parameter Values for Simulation

economic development, where the latter is proxied by the female/wage ratio. An increase in the wage ratio is indicative of better economic development. This exercise does not calibrate the model instead presents the mechanisms at work using arbitrary parameter values. Calibration, together with dynamic simulation, would be the next step of the project. Relevant parameters are considered as per the constraints defined by the model. For example, because $l_f \geq 0$ requires $\mu \geq \theta$ from (4.12) we restrict the gender wage ratio to the interval $[\theta, 1]$. Parameter values are presented in Table 9.

Household Level We first present simulations at the intensive margin. In other words, how households make female labor force participation decisions as cultural costs change. The simulation results are presented in Table 20. The figure plots the threshold condition arrived at in equation 4.11. Here, the relationship between $G(\mu)$ the relative wage of women is plotted. The variation in thresholds for labor force participation at the household level are driven by cultural costs χ_L and χ_H depicted by the horizontal curves. To interpret this, take an example of the highest value of cultural cost χ_H , households facing these costs will not participate in market production for the relative wage rate between $[0.77, 0.98]$.

Aggregate FLFP Extending the intensive margin decision to the aggregate level, we consider that the cultural costs are heterogeneous across the population.

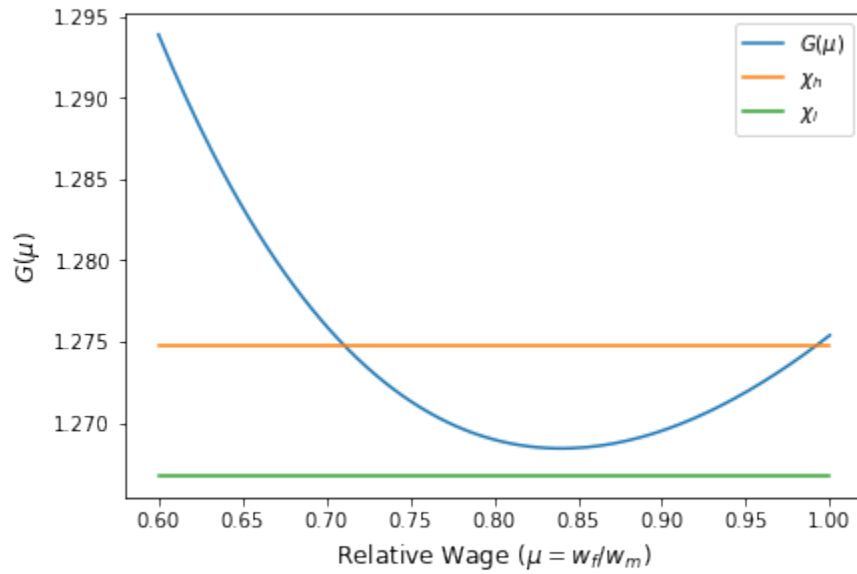


Figure 20. Numerical Simulation for LFP decision at the household level

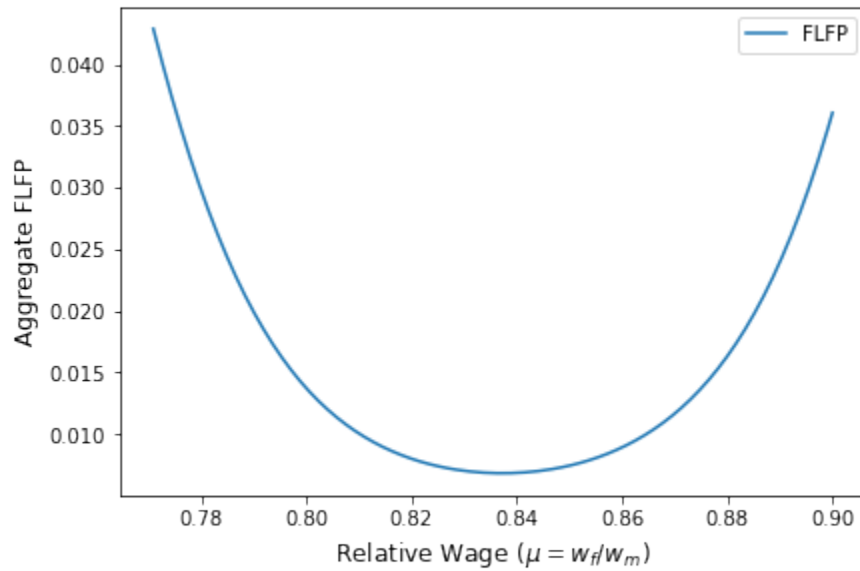


Figure 21. Numerical Simulation for Aggregate FLFP, $\chi \in \{\chi^L, \chi^H\}$

We assume that the population is distributed normally over the space of low and high cultural costs. More specifically, we assume that the population is distributed normally with the mean of $\bar{\chi} \equiv (\chi_H - \chi_L)/2$ and standard deviation of $\chi_{sd} \equiv$

$(\bar{\chi} - \chi_H)/3$ with parameters of χ_H and χ_L presented in Table 9. The simulation results are presented in Figure 21. As per the simulations, a relative wage increase from 0.77 to 0.81 decreases aggregate female labor force participation by 75%.

4.6 Conclusion and future plan

The household decision problem outlined so far identifies three margins that contribute to FLFP: consumption gain from both spouses working, women’s ability to balance child care responsibilities with working outside the home, and the cultural cost society imposes on working wives and mothers. If people are heterogeneous in how they perceive this cost, then they will differ in their willingness to supply labor given the same set of economic conditions. This simple idea produces a naturally U-shaped response of the aggregate FLFP rate to female wages (more precisely, to the gender wage ratio).

Future iterations of this work will pursue two related strategies. The first strategy would be to introduce a role for human capital for the FLFP decision. It is unclear, at this stage, how that can be done without breaking the assumption of positive assortative matching. One possible option is to alter the nature of human capital production. We are also interested in making the model dynamic. That would require explicit assumptions about why the gender wage ratio is less than one and why it increases with economic development. Galor and Weil (1997), discussed in section 2, provide a roadmap for this using an aggregate production function that relies on physical and mental labor. A more parsimonious approach is to assume exogenous productivity growth rates of male and female labor that differ in the short- and medium run but converge in the long run. This “reduced-form” specification would be sufficient to understand the dynamics of FLFP, though not the deeper causes of gender inequality.

A second strategy for future research is to study a more realistic labor market choice. In deciding to work outside home in developing countries, women face a choice between traditional market activities (agrarian, low-skill occupations) and non-traditional ones (manufacturing, formal sector service). The former is organized along family and community lines in which home and market activities are unified (Goldin, 1990; Galor and Weil, 1997). This implies two things. First, through tradition and upbringing, women are well-trained in these activities. Secondly, they have an easier time balancing childcare with market production. In comparison, jobs in non-traditional occupations require information about participating in them, at least initially, and a clear separation of time devoted to home versus market production. The former implies women start out in the new economy with an informational disadvantage, the latter makes it harder for them to give up time from childcare.

When an economy is experiencing rapid structural change, resources gradually shift from the agricultural sector and labor-intensive cottage industries to capital-intensive manufacturing and, then, skill-intensive services. These new jobs offer a higher wage but the time-allocation and information problems can dissuade many women from pursuing these jobs. Recent works, on US FLFP since the nineteenth century, have identified information barriers as deciding factors. Models such as Fernández (2013) and Fogli and Veldkamp (2011) offer feasible ways these can be incorporated in a model of structural change.

CHAPTER V

CONCLUSION

This dissertation presented empirical evidence on the role played by highway investments on India's manufacturing industry and female labor force participation (FLFP), and explored the broader economic reasons behind India's FLFP decline using an analytical model. My research shows that highways did not have significant effects on the manufacturing sector, suggesting that there are other structural reasons behind India's slow manufacturing growth. Similarly, highways did not significantly affect average district-level FLFP, but had significant heterogeneous effects across districts conditional on initial market conditions. This led to FLFP decreasing in some districts while increasing in others. Finally, an analytical model of household decision-making shows that India's FLFP decline can be accounted for by gender specific norms that discourage FLFP, and mothers choosing to invest in their children's human capital over market production.

In Chapter 2, no significant evidence supports the idea that highways lead to growth in the manufacturing output. On the contrary, the results show that districts near the highways experienced a decline in their manufacturing productivity. I speculate that India's strict labor laws and complicated tax regime may have restricted the positive gains from the highway project. These results are different from what other studies have found for India but in line with the evidence collected from China. Therefore, this chapter recommends that exploring specific channels that connect transport infrastructure with manufacturing sector may help consolidate contradictory results in the literature.

The third chapter explored the relationship of highways with FLFP in India. The results are heterogeneous across districts, where districts with higher

literacy rates, lower levels of road connectivity and higher caste fractionalization experienced a relative decrease in FLFP. These effects are predominantly concentrated in the manufacturing and services sector. I study the districts that experienced a decline and present suggestive evidence that this decline is driven by married and skilled women. Future work will closely examine the trajectory of FLFP growth in districts with low and high literacy rate districts separately. This will help understand the heterogeneity across districts and propose mechanisms explaining the variation.

The fourth chapter deconstructs the decline in FLFP by examining the household decision-making process through the lens of an analytical model. A U-shaped behavior of aggregate FLFP with respect to gender wage ratio is simulated. Three margins are identified that contribute to FLFP: consumption gain from both spouses working, women's ability to balance child care responsibilities with working outside the home, and the cultural cost society imposes on working wives and mothers. We assume that people are heterogeneous in how they perceive these costs, therefore, they differ in their willingness to supply labor given similar set of economic conditions. This simple idea produces a naturally U-shaped response of the aggregate FLFP rate to female wages (more precisely, to the ratio of female-to-male wages).

The analyses in Chapters 3 and 4 suggest that when the economy is going through the declining phase of the U-shaped relationship, simply increasing opportunities for market production, for example due to better highway infrastructure, may not be enough to increase FLFP. This is particularly true for regions where returns to investing in children's human capital are high. This may also account for the evidence presented in Chapter 3, that high literacy regions

and more educated mothers experienced a bigger drop in FLFP due to the GQ project.

APPENDIX A

CHAPTER II APPENDIX

A.1 Baseline Levels of Outcome

The chapter discusses the sensitivity of results with and without the baseline level of manufacturing output as an additional control in Equation 2.1. To demonstrate this, consider a simple setup without the instrument variables and district-level characteristics as controls. This simple setup can be extended to two versions, one with district fixed effects and the other with baseline level of manufacturing output as a control variable. In a regression framework these two scenarios can be presented as follows.

$$Y_{it} = \beta_1 GQ_i + \alpha_i + \lambda_t + \epsilon_{it} \quad (\text{A.1})$$

$$Y_{it} = \alpha + \beta_2 GQ_i + \gamma Y_{t-h} + \lambda_t + \epsilon_{it} \quad (\text{A.2})$$

here, Y_{it} is the outcome variable with subscript i for district and t for the year. GQ_i is a dummy for the treated districts, intermediate districts are not considered for simplicity. The district fixed effects are represented by α_i and the baseline outcome variable is represented by Y_{t-h} . λ_t represents year fixed effects, and ϵ_{it} is the residual term. Since we only have two periods, the baseline and the post construction period, the the fixed-effects specification can also be represented as the long differences. Therefore, Equation A.1 can also be written as,

$$Y_{it} - Y_{it-h} = \alpha + \beta_1 GQ_i + \epsilon'_{it} \quad (\text{A.3})$$

here, $\epsilon'_{it} = \epsilon_{it} - \epsilon_{it-h}$. The above equation represents the main specification used to estimate the results in this paper. However, adding a baseline outcome levels

to the above equation makes it algebraically equivalent to Equation A.2. This can be seen by first adding the baseline outcome levels as an additional control to Equation A.3,

$$Y_{it} - Y_{it-h} = \alpha' + \beta_1' GQ_i + \gamma' Y_{it-h} + \epsilon'_{it} \quad (\text{A.4})$$

in the above specification β_1' is algebraically equivalent to β_2 in Equation A.2.

In conclusion, adding a baseline outcome level to the long difference equation eliminates fixed effects and conditions the estimates on baseline levels instead.

Angrist and Pischke (2009) argue that these two alternate scenarios are not nested in each other and are independent. Therefore, for empirical estimation we have to choose between the two options, and there is no clear way to figure out which of the specification is correct.

A.2 Tables and Figures

Table A.1. Summary Statistics

	Nodal	0 - 40 Kms	40-120 kms	> 120 kms
Total Output (Billion Rupees)	135.62	79.31	17.61	22.83
Average Productivity	3.74	3.35	3.47	3.55
Literacy Rate	0.65	0.56	0.50	0.54
(%) Rural Population	0.38	0.73	0.79	0.77
(%) SC/ST Population	0.17	0.24	0.27	0.25
Population Density (100 per sq.km)	68.96	8.64	4.75	5.54
Skill Intensity	0.11	0.10	0.10	0.08
Urban Share (Output)	0.77	0.29	0.30	0.29
Distance from Highway	19.27	18.21	80.13	249.53
Distance from Railways	6.67	10.35	16.45	20.71
Distance from Coast	338.88	314.25	413.86	487.87

Note: The summary statistics are calculated primarily from the ASI data. There are a total of 441 districts with an unbalanced panel of 4885 district-years in total. Districts are the sub-national administrative units after the states, geographically bigger than villages. All districts are divided into separate distance bands based on their proximity to the GQ project. Nodal districts represent Delhi, Mumbai, Chennai, Kolkata, and other smaller districts areas closer to them. Geographical distances from railways and coastline are computed from the data provided by MIT Geodata and socio-economic characteristics are compiled using the 2001 Census data.

Table A.2. First Stage Results - Straight Line IV

	SLIV			LCPIV		
	(1) (0 – 40)km	(2) (0 – 40)km	(3) (0 – 40)km	(4) (0 – 40)km	(5) (0 – 40)km	(6) (0 – 40)km
SLIV	0.621*** (0.068)	0.562*** (0.071)	0.484*** (0.078)			
LCPIV				0.694*** (0.057)	0.647*** (0.064)	0.582*** (0.072)
Literacy Rate		0.003 (0.229)	-0.337 (0.256)		0.004 (0.226)	-0.215 (0.249)
Rural Pop.		-0.057 (0.204)	-0.031 (0.203)		-0.177 (0.215)	-0.118 (0.216)
Pop. Density		0.144*** (0.042)	0.109** (0.045)		0.117*** (0.044)	0.088* (0.046)
SC/ST Population		0.281 (0.246)	0.238 (0.229)		0.046 (0.223)	0.031 (0.221)
Distance to Rail Network			-0.058** (0.024)			-0.059*** (0.023)
Distance to Coast Line			-0.068** (0.026)			-0.048** (0.024)
N	208	208	208	208	208	208
R sqr.	.24	.29	.33	.31	.36	.38
Robust F-stat	84.67	62.54	38.8	148.5	101.15	65.27

Note: First stage results are presented. The dependent variables is a dummy which takes the value 1 if a district is within 0-40 kms away from the highway and 0 otherwise. The districts falling only on the NSEW corridor have been dropped and the districts 40-120 kms have also been dropped. Standard errors are clustered at the district level and are in parentheses below point estimates.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

Table A.3. Probability of Treatment

	(1) (0 – 40)km	(2) (0 – 40)km	(3) (0 – 40)km
Output	0.192** (0.085)	0.231** (0.099)	0.194** (0.098)
Average Productivity (share-weighted)	-0.262 (0.191)	-0.323 (0.213)	-0.307 (0.209)
Average Productivity	-0.117 (0.214)	-0.217 (0.242)	-0.211 (0.261)
Skill Intensity	5.233** (2.282)	5.847** (2.365)	6.016** (2.413)
Capital to Labor Ratio	-0.000 (0.000)	-0.000** (0.000)	-0.000* (0.000)
Urban Share	0.376 (0.492)	0.503 (0.610)	0.442 (0.632)
Literacy Rate		0.803 (1.449)	-1.344 (1.647)
Rural Pop.		2.575** (1.298)	2.278 (1.426)
Pop. Density		0.741*** (0.212)	0.530** (0.236)
SC/ST Population		-0.171 (1.362)	-0.722 (1.298)
Distance to Rail Network			-0.328*** (0.113)
Distance to Coast Line			-0.332** (0.140)
N	264	264	264

Note: This table presents the coefficient estimates after predicting the probability of highway treatment using district level manufacturing sector characteristics, socio-economic characteristics and geographical factors. The estimates are produced by running a logit regression as defined in Equation 2.3. Column (1) presents results with manufacturing characteristics as explanatory variables, Column (2) with socio-economic characteristics added, and finally Column (3) with geographical factors added to the specification.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

Table A.4. GQ construction reduced-form results

	OLS			SLIV			LCPIV		
	<i>No Controls</i>	<i>Controls</i>	<i>Controls</i>	<i>No Controls</i>	<i>Controls</i>	<i>Controls</i>	<i>No Controls</i>	<i>Controls</i>	<i>Controls</i>
Nodal	-0.414*** (0.132)	-0.083 (0.208)	0.282 (0.251)	-0.407*** (0.128)	-0.059 (0.206)	0.296 (0.252)	-0.433*** (0.135)	-0.141 (0.202)	0.200 (0.253)
Treated	0.068 (0.123)	0.139 (0.130)	0.374*** (0.138)	0.174 (0.167)	0.234 (0.177)	0.425** (0.172)	-0.100 (0.179)	-0.052 (0.197)	0.144 (0.198)
Intermediate	0.177 (0.181)	0.213 (0.180)	0.241 (0.161)	0.136 (0.161)	0.176 (0.175)	0.357** (0.172)	0.214 (0.138)	0.231* (0.136)	0.401*** (0.129)
Baseline			-0.249*** (0.043)			-0.252*** (0.045)			-0.247*** (0.044)
N	284	284	284	284	284	284	284	284	284
R sqr.	0.012	0.041	0.19	0.012	0.042	0.19	0.019	0.045	0.19

Note: This table presents the main results for the effects of the GQ project on economics growth. The coefficient estimates are produced by running Equation 2.1 with change in logged levels of output between 2000 and 2007/2009 as the dependent variable. Columns (1)-(3) present the OLS estimates and (5)-(6) present the IV estimates. The IV estimates do not include the nodal districts. Standard errors are clustered at the district level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

APPENDIX B

CHAPTER III APPENDIX

B.1 Tables and Figures

Table B.1. Summary Statistics

	Mean	S.D.	N
<i>Population Statistics</i>			
Total Population	1045.87	755.37	841
Male Population	527.43	380.59	841
Female Population	518.44	378.75	841
Population (Age 15-30)	435.34	318.96	841
Population (Age 30-45)	334.84	248.88	841
Population (Age 45-65)	275.69	207.71	841
<i>Employment Statistics</i>			
Total Employment	622.22	464.91	841
Male Employment	448.43	326.97	841
Female Employment	173.79	187.72	841
Agricultural Employment	430.23	352.85	839
Manufacturing Employment	92.20	94.18	829
Servies Employment	89.92	84.18	838
Employment (Age 15-30)	214.72	168.45	841
Employment (Age 30-45)	231.77	176.52	841
Employment (Age 45-65)	175.74	139.82	841
<i>Labor Force Participation Rates</i>			
Male LFP	0.84	0.08	841
FeMale LFP	0.34	0.23	841
FLFP Agriculture	0.28	0.22	821
FLFP Manufacturing	0.04	0.05	689
FLFP Services	0.03	0.03	758
<i>Distance Measures</i>			
Distance to the GQ project	241.06	216.01	841
Distance to the nearest Railway	34.47	91.64	841
Distance to the nearest Coastline	481.91	340.67	841

Note: The summary statistics are calculated primarily from the NSS data. There are a total of 423 districts with an unbalanced panel of 1236 district-years in total. Districts are the sub-national administrative units after the states, geographically bigger than villages. Geographical distances from railways and coastline are computed from the data provided by MIT Geodata, socio-economic characteristics are compiled using the 2001 Census and the SHRUG data set.

Table B.2. Treatment Effects - FLFP with Literacy Interactions by Industry

	Agriculture		Manufacturing		Services	
	(1)	(2)	(3)	(4)	(5)	(6)
	$\Delta FLFP$	$\Delta FLFP$	$\Delta FLFP$	$\Delta FLFP$	$\Delta FLFP$	$\Delta FLFP$
Nodal	0.309 (0.241)	0.039 (0.257)	-0.094*** (0.027)	-0.109** (0.050)	0.086 (0.071)	0.066 (0.077)
Treated	-0.152 (0.166)	-0.179 (0.168)	0.117*** (0.044)	0.116*** (0.043)	0.057*** (0.021)	0.057*** (0.021)
Intermediate	-0.059 (0.136)	-0.030 (0.136)	0.040 (0.043)	0.029 (0.047)	0.020 (0.019)	0.026 (0.020)
Nodal×Rural Literacy(2001)	-0.473 (0.417)	0.050 (0.456)	0.187*** (0.047)	0.237** (0.097)	-0.137 (0.135)	-0.102 (0.149)
Treated×Rural Literacy(2001)	0.252 (0.302)	0.359 (0.310)	-0.211*** (0.079)	-0.201** (0.081)	-0.124*** (0.040)	-0.126*** (0.040)
Intermediate×Rural Literacy(2001)	0.144 (0.279)	0.109 (0.276)	-0.069 (0.076)	-0.044 (0.087)	-0.045 (0.039)	-0.061 (0.040)
District Controls	No	Yes	No	Yes	No	Yes
N	311	311	218	218	265	265
R sqr.	0.031	0.056	0.027	0.044	0.100	0.13

Note: This table presents the main results for the effects of the GQ project on female labor force participation (FLFP) for each industrial sector separately. The treatment effects are interacted with the literacy rate of districts in the year 2001. The coefficient estimates are produced by running Equation 3.2 with change in FLFP between 1999 and 2009 as the dependent variable. Standard errors are clustered at the district level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

Table B.3. Treatment Effects - Female Employment Levels by Industry

	Agriculture		Manufacturing		Services	
	(1) $\Delta FLFP$	(2) $\Delta FLFP$	(3) $\Delta FLFP$	(4) $\Delta FLFP$	(5) $\Delta FLFP$	(6) $\Delta FLFP$
Nodal	1.072 (1.101)	0.537 (1.512)	-3.799*** (0.862)	-3.688*** (1.005)	-1.117 (6.363)	-1.000 (6.595)
Treated	0.012 (1.053)	-0.117 (1.050)	2.087*** (0.776)	2.093*** (0.784)	1.922** (0.975)	1.997** (0.966)
Intermediate	-0.025 (0.667)	0.074 (0.661)	-0.501 (1.213)	-0.631 (1.209)	0.582 (0.726)	0.704 (0.706)
Nodal×Rural Literacy(2001)	-1.753 (1.896)	-0.740 (2.668)	7.972*** (1.590)	7.964*** (1.915)	4.645 (12.526)	4.115 (13.017)
Treated×Rural Literacy(2001)	-0.777 (2.002)	-0.289 (2.012)	-3.448** (1.487)	-3.418** (1.553)	-3.875* (1.994)	-4.279** (1.981)
Intermediate×Rural Literacy(2001)	-0.014 (1.336)	-0.020 (1.300)	0.908 (2.109)	1.181 (2.157)	-1.037 (1.492)	-1.523 (1.479)
District Controls	No	Yes	No	Yes	No	Yes
N	311	311	218	218	265	265
R sqr.	0.034	0.053	0.044	0.049	0.095	0.12

Note: This table presents the main results for the effects of the GQ project on levels of female employment for each industrial sector separately. The treatment effects are interacted with the literacy of districts measured in the year 2001. The coefficient estimates are produced by running Equation 3.2 with change in female employment levels between 1999 and 2009 as the dependent variable. Standard errors are clustered at the district level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

Table B.4. Baseline - Placebo Treatment Effects with Literacy Interactions

	Agriculture		Manufacturing		Services	
	(1)	(2)	(3)	(4)	(5)	(6)
	$\Delta FLFP$	$\Delta FLFP$	$\Delta FLFP$	$\Delta FLFP$	$\Delta FLFP$	$\Delta FLFP$
Nodal	0.020 (0.057)	0.275* (0.164)	0.024 (0.024)	-0.022 (0.032)	0.045 (0.059)	0.039 (0.054)
Treated	-0.043 (0.091)	-0.053 (0.084)	-0.029 (0.031)	-0.035 (0.032)	-0.011 (0.010)	-0.011 (0.010)
Intermediate	-0.058 (0.084)	-0.057 (0.078)	0.010 (0.026)	0.020 (0.027)	0.007 (0.012)	0.006 (0.012)
Nodal×Total Literacy(1991)	0.033 (0.118)	-0.731* (0.399)	-0.037 (0.057)	0.097 (0.095)	-0.122 (0.123)	-0.101 (0.113)
Treated×Total Literacy(1991)	0.069 (0.219)	-0.024 (0.198)	0.052 (0.076)	0.081 (0.080)	0.037 (0.024)	0.040 (0.024)
Intermediate×Total Literacy(1991)	0.179 (0.225)	0.114 (0.217)	-0.032 (0.059)	-0.049 (0.062)	-0.030 (0.030)	-0.025 (0.030)
District Controls	No	Yes	No	Yes	No	Yes
N	287	287	208	208	239	239
R sqr.	0.0043	0.058	0.0082	0.036	0.053	0.064

Note: This table presents results from placebo regressions run for the period 1987-1999 with the GQ project as treatment, literacy rate in 1991 as the interaction term, and female labor force participation rate as the outcome variable. The coefficient estimates are produced as per Equation 3.2 and standard errors are clustered at the district level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

Table B.5. Average Treatment Effects - FLFP with Road Connectivity Interactions by Industry

	Agriculture		Manufacturing		Services	
	(1) $\Delta FLFP$	(2) $\Delta FLFP$	(3) $\Delta FLFP$	(4) $\Delta FLFP$	(5) $\Delta FLFP$	(6) $\Delta FLFP$
Nodal	0.001 (0.216)	-0.144 (0.232)	-0.067* (0.039)	-0.040 (0.078)	-0.039 (0.063)	-0.042 (0.064)
Treated	-0.143 (0.101)	-0.058 (0.109)	-0.018 (0.036)	-0.007 (0.039)	-0.048*** (0.015)	-0.039** (0.017)
Intermediate	0.173 (0.141)	0.128 (0.138)	0.109* (0.058)	0.124** (0.058)	-0.010 (0.021)	-0.025 (0.023)
Nodal \times ln(Rural Paved Roads(2001))	0.009 (0.029)	0.033 (0.034)	0.012** (0.006)	0.011 (0.012)	0.008 (0.010)	0.009 (0.010)
Treated \times ln(Rural Paved Roads(2001))	0.020 (0.015)	0.010 (0.017)	0.004 (0.005)	0.004 (0.006)	0.007*** (0.002)	0.005** (0.002)
Intermediate \times ln(Rural Paved Roads(2001))	-0.025 (0.021)	-0.016 (0.021)	-0.015* (0.009)	-0.017** (0.009)	0.001 (0.003)	0.003 (0.003)
District Controls	No	Yes	No	Yes	No	Yes
N	306	306	215	215	260	260
R sqr.	0.012	0.057	0.027	0.068	0.050	0.11

Note: This table presents the main results for the effects of the GQ project on female labor force participation (FLFP) for each industrial sector separately. The treatment effects are interacted with a measure that captures road connectivity in rural areas of districts in the year 2001. The coefficient estimates are produced by running Equation 3.2 with change in FLFP between 1999 and 2009 as the dependent variable. Standard errors are clustered at the district level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

Table B.6. Average Treatment Effects - FLFP with Caste Diversity Interactions by Industry

	Agriculture		Manufacturing		Services	
	(1) $\Delta FLFP$	(2) $\Delta FLFP$	(3) $\Delta FLFP$	(4) $\Delta FLFP$	(5) $\Delta FLFP$	(6) $\Delta FLFP$
Nodal	-0.276* (0.151)	-0.164 (0.203)	0.073*** (0.022)	0.122*** (0.044)	0.076 (0.049)	0.088* (0.050)
Treated	0.121 (0.226)	0.226 (0.225)	0.028 (0.048)	0.055 (0.051)	0.010 (0.024)	0.034 (0.024)
Intermediate	-0.145 (0.123)	-0.034 (0.124)	0.053 (0.034)	0.059 (0.037)	-0.018 (0.015)	0.006 (0.016)
Nodal×Caste Index(2001)	0.543** (0.212)	0.370 (0.284)	-0.103*** (0.037)	-0.161** (0.063)	-0.104 (0.083)	-0.123 (0.082)
Treated×Caste Index(2001)	-0.225 (0.365)	-0.360 (0.361)	-0.036 (0.078)	-0.069 (0.082)	-0.029 (0.038)	-0.066* (0.039)
Intermediate×Caste Index(2001)	0.246 (0.201)	0.093 (0.199)	-0.083 (0.056)	-0.087 (0.060)	0.022 (0.025)	-0.016 (0.026)
District Controls	No	Yes	No	Yes	No	Yes
N	311	311	218	218	265	265
R sqr.	0.015	0.059	0.011	0.033	0.033	0.11

Note: This table presents the main results for the effects of the GQ project on female labor force participation (FLFP) for each industrial sector separately. The treatment effects are interacted with the caste fractionalization index constructed for the year 2001. The coefficient estimates are produced by running Equation 3.2 with change in FLFP between 1999 and 2009 as the dependent variable. Standard errors are clustered at the district level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

Table B.7. Treatment Effects - Education Wise Female Employment Levels

	Primary		Secondary		Graduate	
	(1) $\Delta FLFP$	(2) $\Delta FLFP$	(3) $\Delta FLFP$	(4) $\Delta FLFP$	(5) $\Delta FLFP$	(6) $\Delta FLFP$
Nodal	0.831*	0.801*	-1.024***	-0.964***	0.205	0.182
	(0.426)	(0.408)	(0.202)	(0.234)	(0.329)	(0.328)
Treated	-0.035	-0.022	-0.344**	-0.358**	0.504**	0.508**
	(0.201)	(0.205)	(0.163)	(0.167)	(0.213)	(0.218)
Intermediate	-0.268*	-0.246	0.185	0.147	0.078	0.101
	(0.161)	(0.160)	(0.152)	(0.147)	(0.160)	(0.156)
Nodal×Rural Literacy(2001)	-1.381*	-1.407**	1.921***	1.891***	-0.561	-0.534
	(0.705)	(0.688)	(0.370)	(0.433)	(0.544)	(0.559)
Treated×Rural Literacy(2001)	0.125	0.067	0.644**	0.704**	-0.984**	-1.003**
	(0.379)	(0.390)	(0.325)	(0.337)	(0.394)	(0.411)
Intermediate×Rural Literacy(2001)	0.594*	0.520	-0.420	-0.312	-0.158	-0.212
	(0.325)	(0.327)	(0.311)	(0.304)	(0.317)	(0.306)
N	317	317	316	316	240	240
R sqr.	0.017	0.034	0.036	0.060	0.068	0.076

Note: This table presents the results for the effects of the GQ project on the employment share of women belonging to different education categories. The treatment effects are interacted with the literacy rate in districts measured for the year 2001. The coefficient estimates are produced by running Equation 3.2 with change in employment share between 1999 and 2009 as the dependent variable. Standard errors are clustered at the district level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

Table B.8. Treatment Effects - Female Employment Levels by Industry and Married Status

	Agriculture		Manufacturing		Services	
	(1) $\Delta FLFP$	(2) $\Delta FLFP$	(3) $\Delta FLFP$	(4) $\Delta FLFP$	(5) $\Delta FLFP$	(6) $\Delta FLFP$
Panel A: Married Women						
Treated	14.967 (93.339)	5.813 (93.010)	63.565*** (22.965)	64.498*** (22.922)	30.568** (13.170)	31.531** (13.264)
Intermediate	-5.001 (54.291)	1.609 (54.112)	-11.132 (36.804)	-15.157 (36.502)	11.120 (8.939)	13.182 (8.929)
Treated×Rural Literacy(2001)	-91.134 (176.010)	-54.714 (176.787)	-108.740** (45.220)	-109.991** (46.224)	-63.872** (25.580)	-68.458*** (25.854)
Intermediate×Rural Literacy(2001)	9.044 (108.242)	11.372 (106.102)	17.320 (64.609)	25.433 (65.537)	-22.912 (18.272)	-29.501 (18.871)
Panel B: Un-Married Women						
Treated	-32.743 (22.399)	-35.088 (22.873)	-6.976 (7.292)	-6.667 (7.693)	-8.365 (17.334)	-9.442 (16.932)
Intermediate	-1.632 (16.886)	-1.720 (16.372)	6.884 (12.040)	1.378 (13.454)	-23.697 (21.863)	-27.609 (21.614)
Treated×Rural Literacy(2001)	53.835 (46.292)	62.069 (47.016)	13.566 (14.859)	12.668 (15.607)	25.317 (35.104)	24.004 (34.513)
Intermediate×Rural Literacy(2001)	-4.744 (35.758)	-1.369 (33.906)	-9.409 (20.157)	1.078 (24.020)	45.602 (43.082)	50.872 (42.518)
District Controls	No	Yes	No	Yes	No	Yes
N	201	201	76	76	92	92
R sqr.	0.046	0.063	0.18	0.21	0.20	0.24

Note: This table presents the results for the effects of the GQ project on female labor force participation (FLFP) for married and unmarried women in different industrial sectors. The treatment effects are interacted with the literacy rate in districts measured for the year 2001. The coefficient estimates are produced by running Equation 3.2 with change in FLFP between 1999 and 2009 as the dependent variable. Standard errors are clustered at the district level.

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$, Standard errors in parentheses.

REFERENCES CITED

- Akerberg, Daniel A., Kevin Caves and Garth Frazer (2015). “Identification Properties of Recent Production Function Estimators”. In: *Econometrica* 83.6, pp. 2411–2451.
- Afridi, Farzana, Monisankar Bishnu and Kanika Mahajan (2019). *What Determines Women’s Labor Supply? The Role of Home Productivity and Social Norms*. SSRN Scholarly Paper 3468613. Rochester, NY: Social Science Research Network.
- Afridi, Farzana, Taryn Dinkelman and Kanika Mahajan (2018). “Why Are Fewer Married Women Joining the Work Force in Rural India? A Decomposition Analysis over Two Decades”. In: *Journal of Population Economics* 31.3, pp. 783–818. ISSN: 1432-1475.
- Aggarwal, Shilpa (2018). “Do Rural Roads Create Pathways out of Poverty? Evidence from India”. In: *Journal of Development Economics* 133, pp. 375–395. ISSN: 03043878.
- Alder, Simon (2016). “Chinese Roads in India: The Effect of Transport Infrastructure on Economic Development”. In: *SSRN Electronic Journal*. ISSN: 1556-5068.
- Amiti, Mary and Jozef Konings (2007). “Trade Liberalization, Intermediate Inputs, and Productivity: Evidence from Indonesia”. In: *The American Economic Review* 97.5, pp. 1611–1638. ISSN: 0002-8282.

- Angrist, Joshua D. and Jörn-Steffen Pischke (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.
- Asher, Sam, Tobias Lunt et al. (2021). “Development Research at High Geographic Resolution: An Analysis of Night Lights, Firms, and Poverty in India Using the SHRUG Open Data Platform”.
- Asher, Sam and Paul Novosad (2020). “Rural Roads and Local Economic Development”. In: *American Economic Review* 110.3, pp. 797–823. ISSN: 00028282.
- Asturias, Jose, Manuel García-Santana and Roberto Ramos (2019). “Competition and the Welfare Gains from Transportation Infrastructure: Evidence from the Golden Quadrilateral of India”. In: *Journal of the European Economic Association* 17.6, pp. 1881–1940. ISSN: 15424766.
- Bailey, Martha J. and Brad J. Hershbein (2018). “U.S. Fertility Rates and Childbearing, 1800 to 2010”. In: Cain, Louis P., Price V. Fishback and Paul W. Rhode. *Oxford Handbook of American Economic History*. Vol. 1. Oxford University Press. ISBN: 978-0-19-093706-5.
- Banerjee, Abhijit, Esther Duflo and Nancy Qian (2020). “On the Road: Access to Transportation Infrastructure and Economic Growth in China”. In: *Journal of Development Economics* 145, p. 102442. ISSN: 0304-3878.
- Basu, Alaka M and Sonalde Desai (2012). *Middle Class Dreams: India's One-Child Families*, p. 40.

- Baum-Snow, Nathaniel et al. (2017). “Roads, Railroads, and Decentralization of Chinese Cities”. In: *The Review of Economics and Statistics* 99.3, pp. 435–448. ISSN: 0034-6535, 1530-9142.
- Besley, T. and R. Burgess (2004). “Can Labor Regulation Hinder Economic Performance? Evidence from India”. In: *The Quarterly Journal of Economics* 119.1, pp. 91–134. ISSN: 0033-5533, 1531-4650.
- Besley, Timothy and Robin Burgess (2000). “Land Reform, Poverty Reduction, and Growth: Evidence from India”. In: *The Quarterly Journal of Economics* 115.2, pp. 389–430. ISSN: 0033-5533.
- Chandra, Amitabh and Eric Thompson (2000). “Does Public Infrastructure Affect Economic Activity?: Evidence from the Rural Interstate Highway System”. In: *Regional Science and Urban Economics* 30.4, pp. 457–490. ISSN: 0166-0462.
- Chatterjee, Santanu, Thomas Lebesmuehlbacher and Abhinav Narayanan (2021). “How Productive Is Public Investment? Evidence from Formal and Informal Production in India”. In: *Journal of Development Economics* 151, p. 102625. ISSN: 0304-3878.
- Chichilnisky, Graciela and Elisabeth Hermann Frederiksen (2008). “An Equilibrium Analysis of the Gender Wage Gap”. In: *International Labour Review* 147.4, pp. 297–320. ISSN: 1564-913X.
- Das, Maitreyi and Sonalde Desai (2003). *Why Are Educated Women Less Likely to Be Employed in India? Testing Competing Hypotheses*. 0313. World Bank.

- Das, Maitreyi Bordia and Ieva Zumbyte (2017). *The Motherhood Penalty and Female Employment in Urban India*. World Bank, Washington, DC.
- Datta, Saugato (2012). “The Impact of Improved Highways on Indian Firms”. In: *Journal of Development Economics* 99.1, pp. 46–57. ISSN: 0304-3878.
- Dhar, Diva, Tarun Jain and Seema Jayachandran (2019). “Intergenerational Transmission of Gender Attitudes: Evidence from India”. In: *The Journal of Development Studies* 55.12, pp. 2572–2592. ISSN: 0022-0388.
- Donaldson, Dave (2015). “The Gains from Market Integration”. In: *Annual Review of Economics* 7.1, pp. 619–647. ISSN: 1941-1383, 1941-1391.
- (2018). “Railroads of the Raj: Estimating the Impact of Transportation Infrastructure”. In: *American Economic Review* 108.4-5, pp. 899–934. ISSN: 0002-8282.
- Donaldson, Dave and Richard Hornbeck (2016). “Railroads and American Economic Growth: A “Market Access” Approach*”. In: *The Quarterly Journal of Economics* 131.2, pp. 799–858. ISSN: 0033-5533, 1531-4650.
- Duflo, Esther (2001). “Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment”. In: *The American Economic Review* 91.4, pp. 795–813. ISSN: 0002-8282.
- Duranton, Gilles, Peter M. Morrow and Matthew A. Turner (2014). “Roads and Trade: Evidence from the US”. In: *The Review of Economic Studies* 81.2, pp. 681–724. ISSN: 0034-6527.

- Faber, Benjamin (2014). “Trade Integration, Market Size, and Industrialization: Evidence from China’s National Trunk Highway System”. In: *The Review of Economic Studies* 81.3, pp. 1046–1070. ISSN: 0034-6527.
- Fernández, Raquel (2013). “Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century”. In: *American Economic Review* 103.1, pp. 472–500. ISSN: 0002-8282.
- Fogli, Alessandra and Laura Veldkamp (2011). “Nature or Nurture? Learning and the Geography of Female Labor Force Participation”. In: *Econometrica* 79.4, pp. 1103–1138. ISSN: 1468-0262.
- Galor, Oded and David N. Weil (1996). “The Gender Gap, Fertility, and Growth”. In: *The American Economic Review* 86.3, pp. 374–387. ISSN: 0002-8282.
- Ghani, Ejaz, Arti Grover Goswami and William R. Kerr (2016). “Highway to Success: The Impact of the Golden Quadrilateral Project for the Location and Performance of Indian Manufacturing”. In: *The Economic Journal* 126.591, pp. 317–357. ISSN: 1468-0297.
- Gibbons, Stephen et al. (2019). “New Road Infrastructure: The Effects on Firms”. In: *Journal of Urban Economics* 110, pp. 35–50. ISSN: 0094-1190.
- Goldin, Claudia (1990). *Understanding the Gender Gap: An Economic History of American Women*. New York: Oxford University Press. ISBN: 9780195050776.

- Harrison, Ann E., Leslie A. Martin and Shanthi Nataraj (2013). “Learning versus Stealing: How Important Are Market-Share Reallocations to India’s Productivity Growth?” In: *The World Bank Economic Review* 27.2, pp. 202–228. ISSN: 1564-698X, 0258-6770.
- Heath, Rachel and Xu Tan (2020). “Intrahousehold Bargaining, Female Autonomy, and Labor Supply: Theory and Evidence from India”. In: *Journal of the European Economic Association* 18.4, pp. 1928–1968. ISSN: 1542-4766, 1542-4774.
- Hiller, Victor (2014). “Gender Inequality, Endogenous Cultural Norms, and Economic Development”. In: *The Scandinavian Journal of Economics* 116.2, pp. 455–481. ISSN: 0347-0520.
- Holl, Adelheid (2016). “Highways and Productivity in Manufacturing Firms”. In: *Journal of Urban Economics* 93, pp. 131–151. ISSN: 0094-1190.
- Hsieh, Chang-Tai and Peter J Klenow (2009). “Misallocation and Manufacturing TFP in China and India”. In: *The Quarterly Journal of Economics* 124.4, pp. 1403–1448.
- Jedwab, Remi and Alexander Moradi (2016). “The Permanent Effects of Transportation Revolutions in Poor Countries: Evidence from Africa”. In: *Review of Economics and Statistics* 98.2, pp. 268–284. ISSN: 0034-6535, 1530-9142.
- Khanna, Gaurav (2016). “Road Oft Taken: The Route to Spatial Development”.

- Kimura, Masako and Daishin Yasui (2010). “The Galor–Weil Gender-Gap Model Revisited: From Home to Market”. In: *Journal of Economic Growth* 15.4, pp. 323–351. ISSN: 1381-4338, 1573-7020.
- Klasen, Stephan and Janneke Pieters (2015). “What Explains the Stagnation of Female Labor Force Participation in Urban India?” In: *The World Bank Economic Review* 29.3, pp. 449–478. ISSN: 0258-6770.
- Lei, Lei, Sonalde Desai and Reeve Vanneman (2019). “The Impact of Transportation Infrastructure on Women’s Employment in India”. In: *Feminist Economics* 25.4, pp. 94–125. ISSN: 1354-5701, 1466-4372.
- Levinsohn, James and Amil Petrin (2003). “Estimating Production Functions Using Inputs to Control for Unobservables”. In: *The Review of Economic Studies* 70.2, pp. 317–341. ISSN: 0034-6527.
- Lowe, Matt and Madeline McKelway (2019). “Bargaining Breakdown: Intra-Household Decision-Making and Female Labor Supply”.
- Lundberg, Shelly and Robert A. Pollak (1993). “Separate Spheres Bargaining and the Marriage Market”. In: *Journal of Political Economy* 101.6, pp. 988–1010. ISSN: 0022-3808.
- Maddison, Angus (2001). *Development Centre Studies The World Economy A Millennial Perspective: A Millennial Perspective*. OECD Publishing. 385 pp. ISBN: 978-92-64-18998-0.

- Martin, Leslie A., Shanthi Nataraj and Ann E. Harrison (2017). “In with the Big, Out with the Small: Removing Small-Scale Reservations in India”. In: *The American Economic Review* 107.2, pp. 354–386. ISSN: 0002-8282.
- Melecky, Martin, Siddharth Sharma and Hari Subhash (2018). *Wider Economic Benefits of Investments in Transport Corridors and the Role of Complementary Policies*. World Bank, Washington, DC.
- Melitz, Marc J. (2003). “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity”. In: *Econometrica* 71.6, pp. 1695–1725. ISSN: 1468-0262.
- Michaels, Guy (2008). “The Effect of Trade on the Demand for Skill: Evidence from the Interstate Highway System”. In: *Review of Economics and Statistics* 90.4, pp. 683–701. ISSN: 0034-6535, 1530-9142.
- Mitra, Devashish and Beyza P. Ural (2008). “Indian Manufacturing: A Slow Sector in a Rapidly Growing Economy”. In: *The Journal of International Trade & Economic Development* 17.4, pp. 525–559. ISSN: 0963-8199.
- Munshi, Kaivan (2019). “Caste and the Indian Economy”. In: *Journal of Economic Literature* 57.4, pp. 781–834. ISSN: 0022-0515.
- Nataraj, Shanthi (2011). “The Impact of Trade Liberalization on Productivity: Evidence from India’s Formal and Informal Manufacturing Sectors”. In: *Journal of International Economics* 85.2, pp. 292–301. ISSN: 0022-1996.

- Neff, Daniel, Kunal Sen and Veronika Kling (2012). *The Puzzling Decline in Rural Women's Labor Force Participation in India: A Reexamination*. GIGA Working Paper 196. German Institute of Global and Area Studies.
- Oh, Suanna (2021). *Does Identity Affect Labor Supply?* SSRN Scholarly Paper 3998025. Rochester, NY: Social Science Research Network.
- Olley, G. Steven and Ariel Pakes (1996). "The Dynamics of Productivity in the Telecommunications Equipment Industry". In: *Econometrica* 64.6, pp. 1263–1297. ISSN: 0012-9682.
- Pande, Rohini and Charity Troyer Moore (2015). "Opinion — Why Aren't India's Women Working?" In: *The New York Times. Opinion*. ISSN: 0362-4331.
- Pavcnik, Nina (2002). "Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants". In: *The Review of Economic Studies* 69.1, pp. 245–276. ISSN: 0034-6527.
- Rao, M. Govinda (2005). "Tax System Reform in India: Achievements and Challenges Ahead". In: *Journal of Asian Economics* 16.6, pp. 993–1011. ISSN: 1049-0078.
- Redding, Stephen J. and Matthew A. Turner (2015). "Chapter 20 - Transportation Costs and the Spatial Organization of Economic Activity". In: *Handbook of Regional and Urban Economics*. Ed. by Gilles Duranton, J. Vernon Henderson and William C. Strange. Vol. 5. Handbook of Regional and Urban Economics. Elsevier, pp. 1339–1398.

- Restuccia, Diego and Richard Rogerson (2017). “The Causes and Costs of Misallocation”. In: *Journal of Economic Perspectives* 31.3, pp. 151–174. ISSN: 0895-3309.
- Sarkar, Sudipa, Soham Sahoo and Stephan Klasen (2019). “Employment Transitions of Women in India: A Panel Analysis”. In: *World Development* 115, pp. 291–309. ISSN: 0305-750X.
- Shamdasani, Yogita (2021). “Rural Road Infrastructure & Agricultural Production: Evidence from India”. In: *Journal of Development Economics* 152, p. 102686. ISSN: 0304-3878.
- Sivadasan, Jagadeesh (2009). “Barriers to Competition and Productivity: Evidence from India”. In: *The B.E. Journal of Economic Analysis & Policy* 9.1. ISSN: 1935-1682.
- Storeygard, Adam (2016). “Farther on down the Road: Transport Costs, Trade and Urban Growth in Sub-Saharan Africa”. In: *The Review of Economic Studies* 83.3, pp. 1263–1295. ISSN: 0034-6527.
- Syverson, Chad (2011). “What Determines Productivity?” In: *Journal of Economic Literature* 49.2, pp. 326–365. ISSN: 0022-0515.
- Topalova, Petia (2010). “Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India”. In: *American Economic Journal: Applied Economics* 2.4, pp. 1–41. ISSN: 1945-7782.
- Woetzel, Jonathan et al. (2016). “Bridging Global Infrastructure Gaps”. In: *McKinsey Global Institute* 14.