From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Economics

2018

Essays in Trade and Development Economics

Chatterjee, Shoumitro

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp0173666721w

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

ESSAYS IN

TRADE AND DEVELOPMENT ECONOMICS

SHOUMITRO CHATTERJEE

A DISSERTATION

PRESENTED TO THE FACULTY

OF PRINCETON UNIVERSITY

IN CANDIDACY FOR THE DEGREE

OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE

BY THE DEPARTMENT OF ECONOMICS

ADVISOR: ESTEBAN ROSSI-HANSBERG

JUNE 2018

© Copyright by Shoumitro Chatterjee, 2018.

All rights reserved.

Abstract

This dissertation is a collection of essays on two areas related to poverty in the developing world. Chapters 1 and 2 seek to understand a reason behind low incomes of farmers in India. Chapter 3 delves into studying the relationship between economic growth and fertility in the developing world.

In the first chapter, I argue that the market power of intermediaries or "middle-men" contributes to lowering farmer incomes in India. In particular, I study the role of spatial competition between intermediaries in determining the prices that farmers receive in India by focusing on a law that restricts farmers to selling their goods to intermediaries in their own state. I show that the discontinuities in market power generated by the law translate into discontinuities in prices. Increasing spatial competition by one standard deviation causes prices received by farmers to increase by 6.4%.

In the second chapter, I shed light on spatial and aggregate implications of the intra-country trade restriction of agriculture goods in India. I propose and estimate a quantitative spatial model of bargaining and trade. Using this structural model, I estimate that the removal of the interstate trade restriction in India would increase competition between intermediaries substantially, thereby increasing the prices farmers receive and their output. My estimates suggest that average farmer prices and output would increase by at least 11% and 7% respectively. The value of the national crop output would, therefore, increase by at least 18%.

In the third chapter, co-authored with Tom Vogl, we analyze how fertility change related to economic growth during 1960–2010. Although there were grave predictions of Malthusian catastrophe, fertility in the developing world more than halved while living standards more than doubled during this period. Using data on 2.3 million women from 255 household surveys, we find different responses to fluctuations and long-run growth, both heterogeneous over the lifecycle. Fertility was procyclical but declined and delayed with long-run growth; fluctuations late (but not early) in the reproductive period affected lifetime fertility. The

results are consistent with models of the escape from the Malthusian trap, extended with a lifecycle and liquidity constraints.

Acknowledgments

I am incredibly fortunate to have been surrounded by wonderful mentors, teachers, and friends all my life. This dissertation is a result of their belief in me and the support they have provided over numerous years.

I want to express my deepest gratitude to my advisor, Esteban Rossi-Hansberg. He watched from a distance while he channeled my thoughts and energies, steadying the ship when needed. His care and dedication toward his students are profound. I have always been in awe of his intellect, intelligence, and knowledge and it has been my pleasure and fortune to learn the craft from him.

Tom Vogl has taught me all the applied microeconometrics that I know. Since my early days in Princeton, he has selflessly nurtured me into the researcher I am. It is only befitting that my work with him, which forms the third essay, resulted in my first academic publication.

Richard Rogerson has taught me how to think about the simpler problem before delving into the complexities and nuances. Eduardo Morales provided critical feedback and support at crucial stages of my research. The sagacious Bo Honore has saved many a development economists at Princeton, including me. The depth of his knowledge is dazzling, infectious, and inspires to learn more every day. Angus Deaton is the conscience of applied economics. He left an indelible impression on me as a teacher and thinker in the classroom. I hold his work as the model for conducting my research. Kala Krishna has supported my research since its inception. She has pushed me to think out-of-the-box, take greater risks in research, and attempt ambitious projects.

Devesh Kapur, the perfect mentor, has always patiently heard my thoughts and ideas. He encouraged me when I needed it the most. Not only did he nudge me to do quantitative work on agriculture in India but also provided me the necessary fora. I learned more about India while working with Arvind Subramanian at the office of the Chief Economic Advisor in Delhi than in classrooms or books. It is here that I first got interested in agricultural

markets and the seeds of the first two essays were sown. Arvind's energy is infectious, and it is with him that I spent some of my most creative time.

Tridip Ray and Mausumi Das are the most selfless mentors and friends anyone could get. They always ensured that I have a home in Delhi. My knowledge about agriculture in India would be incomplete if not for conversations with Ajay Vir Jakhar and Mekhala Krishnamurthy.

Rohit Lamba is the reason I entered the hallowed corridors of Princeton, and he made sure that I did not fall when times were the hardest. Without Gaurav Sabharwal, my Princeton experience would be incomplete. Besides discussing vagaries of life on the third floor of Wallace, down Elm Dr. and over spirits, he has always challenged my intellect for the better. Skand Goel painstakingly read each and every word of the first two essays and provided critical feedback. Sanjay Raj Singh made sure that my economics was always right by asking the hard questions. With Skand and Sanjay, I have had a memorable journey from College to graduate school. It is with no one else that I would do it again.

I have learned more about economic geography from David Nagy than anyone else, and am very thankful for his friendship and ongoing tutelage. Dean Spears has been a watchful friend. He has always been around and provided valuable advice. Kababs, squash, and conversations with Dinsha Mistree ensured my happiness in Princeton. I owe a great deal of debt to my classmates Giulia Brancaccio, Emil Verner, Graham McKee, and Clara Santamaria, who created a terrific environment to learn in, were generous in sharing their knowledge, and made sure that I survived graduate school.

Lastly, I'd like to thank my family. This dissertation belongs to my mother, Lipika, as much as it does to me. She was my first teacher and always ensured, often at enormous costs to herself, that I had the resources and the mental space to accomplish this. My uncle, Sudip, has been a godsend and the iron pillar in my life. He did not so that I could. I am what I am because of my parents and I dedicate this dissertation to them.

I am profoundly grateful to my partner, Seher, for her patience, support, and wisdom. She has always read every word I have written, and like an alchemist brought out the intelligible in my gibber. Her belief that I would accomplish the pages that you are about to read made this happen.

S.C.

Princeton, New Jersey, USA.

June 2018.

For my parents, Lipika, Sudip, and Somir

Contents

Abstract iii

Acknowledgments v

1 Market Power and Spatial Competition in Rural India **1**

1.1 Introduction 1

1.2 Agricultural Trade in India 9

1.2.1 Empirical Context 9

1.2.2 The *Mandi* Institution: A Brief History 11

1.2.3 Geographical Placement of Markets 12

1.3 Data 13

1.4 Does more spatial competition increase farmers' prices? 15

1.4.1 Effect of within- and out-of-state competition 16

1.4.2 Causal Effects of Local Competition on Prices 18

References 21

1.A States in Sample 30

1.B Construction of Markets 31

1.C Placebo Tests 32

1.D Data Appendix 33

1.D.1 District Names 33

1.D.2 Organizing the Data on Farmer Prices and Retail 33

1.E Minimum Support Prices and Farmer Prices 35

2 General Equilibrium Effects of Spatial Competition **40**

2.1 Introduction . . . 40

2.1.1 Model . . . 41

2.1.2 Solving the Model . . . 44

2.1.3 Demand in Retail Markets . . . 46

2.2 Estimating the Model . . . 46

2.2.1 Estimating τ and δ . . . 47

2.2.2 Calibrating the production function and demand elasticities . . . 52

2.3 Counterfactual Analysis . . . 53

2.3.1 Changing Farmers' Threat Point . . . 54

2.3.2 Allow Farmers To Sell At A Different Market . . . 56

2.3.3 Letting Farmers Reoptimize the Use of Intermediate Inputs . . . 57

2.3.4 Incorporating a Demand Side Response . . . 58

2.3.5 Effect of a Reduction in Transport Costs . . . 60

2.4 Robustness and Sensitivity of Quantitative Results . . . 61

2.5 Concluding Remarks . . . 63

References . . . 64

2.A Proof of Theorem 1 . . . 85

2.A.1 Existence . . . 85

2.A.2 Uniqueness . . . 85

2.B Proof of Lemma 1 . . . 86

2.C Change in Production in Response to a Change in Prices . . . 87

2.C.1 A Note on Calibrating $\gamma + \nu$ . . . 89

2.D Algorithm to Compute Equilibrium Prices . . . 89

2.E Crops-State Combinations in Structural Estimation . . . 90

2.F Sensitivity to Allocation of Crop Choice to Grid Cells . . . 91

3 Escaping Malthus: Economic Growth and Fertility Change in the Developing World **92**

3.1 Introduction . . . 92

3.2 Data . . . 98

3.3 Analysis of Period Fertility . . . 102

3.3.1 Defining Time Horizons . . . 102

3.3.2 Methods . . . 103

3.3.3 Results . . . 106

3.3.4 Extensions . . . 110

3.4 Analysis of Cohort Fertility . . . 113

3.4.1 Methods . . . 113

3.4.2 Results . . . 114

3.5 Interpretation through the Lens of Unified Growth Theory . . . 115

3.6 Conclusion . . . 121

References . . . 123

3.A Supplementary Tables and Figures . . . 139

3.B Proofs . . . 157

Chapter 1

Market Power and Spatial Competition in Rural India

1.1 Introduction

Small farmers in low-income countries are among the poorest people in the world. Their revenues may be low partly because they are unproductive but also because they may receive low prices for what they produce. One reason farmers receive low prices could be the monopsony power of intermediaries or "middlemen," who are the main buyers of farmers' output in developing countries. A source of their market power is the limited ability of farmers to arbitrage between different intermediaries due to high costs of transportation or policy regulations. This channel of spatial competition between intermediaries, i.e. the ability of farmers to conduct spatial arbitrage, may be an important determinant of the income and welfare of farmers. I use microdata from India on locations of intermediary markets, prices, and data on agriculture production to study the importance of spatial competition between intermediary markets for farmer incomes and production.

Key to my approach are the Agriculture Produce and Marketing Committee (APMC) Acts of Indian states that restrict farmers to selling their output to government licensed in-

termediaries in government regulated markets of their own state. The analysis has two parts. In chapter 1 of this dissertation, I establish that spatial competition between intermediaries is an important determinant of the prices that farmers receive. There are two challenges in analyzing the relationship between spatial competition and prices: first, credibly measuring spatial competition and second, the endogeneity in the location of intermediaries. Besides its intrinsic interest, the Indian context helps me address both of these challenges. Monopsony rights of the APMC markets and their fixed spatial location allow me to measure competition between markets and relate it to price realizations. More importantly, the restriction on farmers' crossing state boundaries ensures a discontinuity at the border in the competition faced by markets. This allows me to establish a causal relationship between competition and prices.

In particular, I define spatial competition as the number of markets in the neighborhood of each market weighted by the inverse of their distance.[1] I form market pairs by matching markets that are in close proximity to each other but lie on different sides of a state border. Because farmers cannot sell across state borders, the competition faced by markets changes discontinuously at the border. However, other factors that influence prices, like geography and local demand, are likely to be similar and vary continuously across the border. Therefore, I recover a causal effect of increased competition by comparing the difference in prices that farmers receive in each market pair to the difference in competition faced by them. My estimates suggest an average increase of 6.4% in farmer prices due to one standard deviation increase in spatial competition.

In the above context, a policy that would increase spatial competition is the removal of interstate trade restriction on farmers. The effects of such a policy change cannot be directly inferred from the data, and evaluating them requires a structural model. Specifically, a reduced-form exercise will fail to capture at least three crucial elements of the change. First, removal of trade restrictions will not only affect the markets on the border, but will also cause

[1]This is a standard measure of market access in the trade literature. e.g. Harris (1954), Donaldson and Hornbeck (2016), and Allen and Atkin (2016).

ripple effects in the markets located in the interior of states as farmers eliminate arbitrage between markets. The magnitude of these effects will depend on transportation costs, the actual location of markets, local demand conditions, and cropping patterns. Second, a change in the prices received by farmers reflects only a part of the benefit. Higher prices in turn incentivize farmers to reoptimize the use of intermediate inputs and hence alter production. Third, changes in supply can lead to changes in retail prices of crops, which would feed back into the prices received by farmers.

In the second chapter of this dissertation, I develop a spatial model of trade in agricultural markets in India to estimate gains from removing interstate trade restrictions. The model uses the economic geography and flexibly captures the determinants of spatial competition. In the framework, post-harvest, farmers optimally choose a market to sell their output. At the market, they Nash-bargain with the associated intermediary,[2] who in turn sells the purchased goods in the local retail market. When farmers bargain with an intermediary at a market, they alternatively consider transporting their goods to another market nearby and selling at a higher price. Geography and policy create spatial variation in the farmers' outside options, and therefore, spatial variation in the degree of market power that intermediaries can exert.

I recover a structural relationship between the prices that the farmer receives at a particular market site, the prices in the neighboring markets, and the local retail price. This relationship depends on the Nash-bargaining parameter, geographic location of the markets, and the transport costs. I estimate the model parameters using a method of simulated moments procedure to match the overall spatial distribution of farmer prices. This procedure helps me address issues stemming from the simultaneity in the determination of prices at each location.

[2]In practice, there may be 15-20 intermediaries in a market, but empirical (Banerji and Meenakshi, 2004; Meenakshi and Banerji, 2005) and ethonographic (Krishnamurthy, 2011) works suggest collusion among intermediaries in these markets. Therefore, in the absence of data on the number of intermediaries within a market, I assume that each market is associated with a representative intermediary.

Next, I use the estimated model to conduct policy counterfactuals in which I remove the restriction on farmers' selling across state borders. This change increases spatial competition between markets. I quantify the associated benefits to farmers under four different scenarios. In the first exercise, I keep all decisions of the farmer fixed except the threat of the farmers' selling at a different market. Specifically, farmers go to the same market as before but while bargaining with the intermediary, can threaten to sell in a market outside the state.[3] This increases farmer prices by about 3% on average. For the farmers with above-median gain, the average increase in prices is about 7%. The maximum increase in farmer prices is 20%. Increased competition between markets close to state borders changes the outside option of the farmers trading in the markets adjacent to them. This mechanism ripples across space, and as a result even farmers in the interior of states experience an increase in prices. However, due to costly transportation, the price increases are larger in markets that are closer to state borders.

In the second exercise, I additionally allow farmers to optimize their choice of the market where they trade. This increases farmer prices by about 11% on average. For the farmers with above-median gain, the average increase in prices is about 21%. The key channel that generates the large increase in prices over the previous exercise is the geographical location of farmers relative to the location of markets. Consider farmers near state borders who are closer to markets of a neighboring state but may be far from markets in their own state. These farmers incur large transportation costs when border crossing restrictions are in place. With no restriction on crossing state boundaries, these farmers benefit not only from greater spatial competition but also from gaining access to closer markets, which reduces their transportation costs.

In a third exercise, I also allow farmers to reoptimize the use of intermediate inputs in response to getting higher prices for their output and estimate the associated change in production. I find that the crop output of farmers increases by about 9% on average. For

[3]In the model, this changes the threat point of the farmer in the Nash-bargaining problem but not his market choice.

the farmers with above-median gain, the average increase in crop output is about 15%. My estimates suggest this increases the national crop output by about 8% and the value of the national crop output, at farmer prices, by about 18%.

Since increased crop output can depress retail prices and feed back into farmer prices, in a fourth exercise, I allow the retail prices to adjust as a result of the change in supply. I estimate the change in prices and production in the long-run equilibrium of the model. Although a small fraction of farmers lose because of the downward adjustment in retail prices, the losses are small with respect to the overall gains. Farmer prices increase by about 9% on average. For the farmers with above-median change, the average increase in prices is about 19%. Among the farmers who lose, the average decline in prices is about 10%. The reason why the average increase in prices is similar to the earlier exercises, even after accounting for losses, is the following. Upon removal of border crossing restrictions, farmers from regions with lower initial prices transport their output and sell in regions with higher initial prices. Thus, regions with higher initial prices experience an increased supply and hence a reduction in retail prices. The opposite happens in regions with low initial prices, and retail prices adjust upwards. Compared to the scenario with no adjustment of retail prices, farmers in the high initial retail price regions lose on average. The effect on farmers in regions with low initial prices is ambiguous. In the scenario with no adjustment of retail prices, these farmers incurred transport costs to get access to better markets. When I allow retail prices to adjust, they can get higher prices in their local regions and save the transport cost. Therefore, they could potentially gain more compared to the scenario with no retail price adjustment.

The quantitative spatial model also allows me to show that spatial competition interacts in important ways with other development policies. The market access of farmers can be increased either by lowering transport costs or by letting them sell in markets in other states. Using the model, I show that there is a large interaction effect of the two policies. Farmers gain substantially more from the simultaneous reduction of transport costs and the

elimination of border restrictions to trade, than the sum of the gains from enacting each of the two policy changes alone. Moreover, the additional gains occur in locations with greater market power of intermediaries, where reductions in transport costs on their own would not have a large effect.

The quantitative magnitudes from the counterfactual exercises should matter for policymakers. First, they provide an estimate of the gains of a potential reform in India that the central government has been pushing the states toward—creating a National Agriculture Market (NAM) that lets farmers trade in any market in the country. Second, an understanding of the spatial distribution of market power provides insight on policy design that is likely to yield better results for raising farmer incomes (and concomitant poverty alleviation) in specific regions. For example, farmers in more competitive regions are likely to gain more from investments such as road construction. On the other hand, policymakers need to explicitly target competition between intermediaries in regions where they exert considerable market power. Third, while subsidies in the input and output markets have conventionally been the focus of policies meant to improve farmer incomes in developing countries, this work highlights the importance of physical output market sites and policies that restrict market access for farmer incomes. Subsidies not only place a fiscal burden on the exchequer but also distort market prices, whereas gains from implementing policies that increase competition[4] can essentially be achieved with little financial cost.

In summary, I make three contributions. First, I establish that spatial competition between intermediaries is an important channel for determining prices that farmers get. Second, I show that policies that increase spatial competition can have large effects on farmers' welfare which manifest through two channels; (a) a ripple effect of one market affecting the other; and (b) alteration of the production decisions of farmers. Third, I

[4]Many other policies can increase spatial competition, such as creating new markets or allowing farmers to directly sell to supermarkets. Quantifying gains from those policy changes will require additional information, such as estimates of the fixed cost of entry of new markets or the location and capacities of existing supermarkets. In the absence of such data, I choose to relax a realistic policy regulation, which allows me to study the same channel without having to worry about other margins.

show that spatial competition can have quantitatively large interaction effects with other development interventions such as road construction.

The work in chapters 1 and 2 relates to a number of different strands of literature. Most relevant is the work that considers the distribution of gains from trade in the presence of variable markups and intermediation. Although there is recent work that examines gains from trade in the presence of variable markups, the focus has been on producers rather than intermediaries with market power (see, for example, Melitz and Ottaviano, 2008; Edmond et al., 2015; De Loecker et al., 2016). In my framework, I introduce the dimensions of intermediation and spatially varying market power, and estimate gains from changes in the market power of intermediaries. In that respect, my work is complementary to Atkin and Donaldson (2015), who show that the market power of intermediaries is an important determinant of actual prices paid by consumers. However, the goal in Atkin and Donaldson (2015) is to separately identify trade costs from market power. I contribute by establishing the micro-foundations of a distinctive channel—that of bargaining with spatially varying threat points—through which intermediaries in remote locations may enjoy market power. This is essential to conducting the policy counterfactuals.

Although the core mechanism in my setup is not surprising and has been conjectured by others (e.g. Dillon and Dambro, 2016), the theoretical literature on intermediation and trade (Antras and Costinot, 2011; Bardhan et al., 2013; Krishna and Sheveleva, 2017) does not explicitly model it. I contribute to this literature by building a spatial model of bargaining and trade with intermediation where intermediaries enjoy market power because it is costly for the farmers in remote locations to access alternative buyers.

This work is also related to a growing body of quantitative economic geography literature that estimates welfare gains from integration. For example, Donaldson (2015b) and Donaldson and Hornbeck (2016) estimate welfare gains as a result of reduction in transport costs. Allen and Atkin (2016) explore the effect of trade on the link between integration and volatility of returns for Indian farmers. While this literature has largely focused on perfectly

competitive environments often building on the seminal work of Eaton and Kortum (2002), I show that gains from integration or reduction in trade costs are heterogeneous in space and the magnitudes relate to the spatial variation in market power of intermediaries.

The existing empirical literature that measures market power of intermediaries in agricultural markets (surveyed in Dillon and Dambro, 2016) provides mixed evidence, with some papers estimating sizable market power of intermediaries (Bergquist, 2017; Casaburi and Reed, 2016) and others finding that it is not so much (Fafchamps et al., 2006). This body of work has focused on testing competitiveness of particular agricultural markets without regard to their spatial locations. Papers either estimate pass-through rates at markets (e.g. Bergquist, 2017; Casaburi and Reed, 2016), directly measure margins of traders (Fafchamps and Minten, 2002), or account for the entry and exit of intermediaries (Fafchamps et al., 2005). My research shows that the interaction of economic geography with spatial competition can generate spatially varying market power for intermediaries. It, therefore, provides a possible rationale for isolated studies finding different estimates of the market power of intermediaries.

There are three related papers on agricultural markets in India. Banerji and Meenakshi (2004) and Meenakshi and Banerji (2005) analyze transactions-level data from markets in North India to identify collusion among traders. Mitra et al. (2017) estimate high trader margins in the state of West Bengal. They conclude that their results are inconsistent with long-term contracts between farmers and traders but consistent with a model of ex-post bargaining. While these papers provide useful insights about the working of these markets, I further contribute to this literature by using microdata to estimate the spatial distribution of the market power of intermediaries in most of India.

Finally, since the seminal work of McCallum (1995), an extensive literature has documented significant costs of crossing national borders (e.g. McCallum, 1995; Anderson and van Wincoop, 2003; Engel and Rogers, 1996). I show that significant border costs exist even within a large developing country. This work is complementary to Cosar et al. (2015),

that studies spatial competition between wind turbine manufacturers and the impact of the Denmark-Germany border on their production costs.

The rest of this chapter is organized as follows. In the next section, I provide the institutional background of agricultural trade and production in India. Section 1.3 outlines data sources. Section 1.4 presents reduced-form evidence for the main mechanism, i.e. spatial competition causing an increase in prices that farmers get.

In the second chapter of this dissertation, I shed light on spatial and aggregate implications of the intra-country trade restriction of agriculture goods in India. In section 2.1, I develop the quantitative framework that I will use to conduct policy counterfactuals. Section 2.2 discusses the structural estimation of the model. In section 2.3, I study the effect of increased competition due to removal of inter-state trade restrictions on farmer incomes and production. Section 2.4 discusses robustness and sensitivity of results. Finally, in section 2.5, I offer concluding remarks.

1.2 Agricultural Trade in India

In this section, I provide a brief description of agricultural trade in India, the current institutional setting, and a brief history of the regulated agricultural markets.

1.2.1 Empirical Context

Agriculture is an important sector of the Indian economy: in 2011, 54.6% of the total workforce was employed in agriculture, and the sector comprised 18.52% of India's total GVA.[5] Eight non-perishable crops—rice, wheat, maize, sorghum, barley, finger millet, pearl millet, and soybean—account for 70% of India's gross cropped area. In this work, I will restrict attention to these crops.

[5]Agriculture Statistics at a Glance 2016, Ministry of Agriculture & Farmer Welfare, Government of India.

There are two cropping seasons each year: fall or *kharif* (July to November during the south-west monsoon) and spring or *rabi* (November to March). Some regions also have a third summer crop between March and June. Rice, sorghum, maize, millet, and soybean are primarily grown in the *kharif* season, while wheat and barley are grown in the *rabi* season.

The median Indian farming household operates a small farm of 1.5 hectares and cultivates it with the help of family or village labor. Its net annual income[6] is approximately USD 365.[7] Usually, the farmers keep a small fraction of the final crop output for personal consumption and sell the rest to licensed intermediaries in a government-regulated marketplace (known as a *mandi*).

Thus, the institutional setting is comprised of three economic markets: (a) a market for intermediate inputs for farming; (b) the regulated market where farmers sell their output to government-licensed intermediaries; and (c) the retail markets where the intermediaries sell. The focus of my analysis is the regulated market and the transaction between the farmer and the intermediary.

Present-day trade in agricultural commodities in India is regulated by two key laws, the Essential Commodities Act (1955) of the central government, which prohibits hoarding, and the autonomous state-level APMC Acts. The APMC Acts mandate that after harvest the *first* sale and purchase of agricultural commodities produced in the state must be carried out in government-designated marketplaces, and buyers of agricultural output must obtain a license from the marketing committee of the marketplace. The laws de facto prohibit farmers from selling in marketplaces of a neighboring state (Chand, 2012). Thus, these Acts restrict the set of buyers of farmers' output to intermediaries within the state. Thereby these laws also reduce competition between marketplaces across state borders and provide a source of variation that can be exploited to estimate gains from increasing competition in the presence of spatially varying market power.

[6]Includes personal consumption valued at market prices.

[7]Economic Survey of India, Ministry of Finance, Government of India, 2016

According to officials in the central and state governments,[8] the enforcement of the border restriction is not hard. Typically, whenever farmers try to cross state borders, intermediaries protest with local law enforcement agencies who in turn punish a farmer to set an example. This does not preclude all leakages but it is well understood in public policy debates that this restriction hurts farmers. Moreover, the central government has been trying to push the state governments to reform these laws.[9] However, these efforts have had little success.[10]

As such, the regulated markets are the most important institutions for determining farmer incomes in India. The laws, as discussed, create licensed intermediaries who have monopsony rights over the farmers' output in the state. An intermediary has a license to a particular marketplace. Although multiple licenses are given per market, only a few are active and cartelization among intermediaries is common. Incumbent intermediaries also prevent new entrants (Chand, 2012). Once an intermediary has bought crops from farmers, he moves them up the supply chain and sells them to a large miller or in the retail markets. There are no legal restrictions on where the intermediaries can sell.

Given the importance of the regulated markets and the APMC Acts, I present a brief history of them next.

1.2.2 The *Mandi* Institution: A Brief History

The idea of the regulated agricultural marketplace in India has its genesis in a 1928 report by the British colonial government's Royal Commission on Agriculture in India. The chairman of this commission, Lord Linlithgow, had presided over a similar committee in 1922

[8] While interning at the Office of the Chief Economic Advisor of India in 2015–16, I interviewed senior bureaucrats at the Ministry of Finance, the Ministry of Agriculture, and in three state governments (Punjab, Karnataka and Madhya Pradesh) to understand the functioning of agricultural markets and the implementation of the laws.

[9] See, for example, "Will a national common market for farm products curb price rise?," in Livemint, July 10, 2014 (Click here), and "Agriculture reform: Government takes first step for a national farm market," in The India Express, July 3, 2015 (Click here), which describe how the current cross-border restrictions hurt farmers and the efforts of the central government to ameliorate them.

[10] For details, see "Farm Market Reforms: Not in NAMe alone" in The Indian Express, June 30th, 2016 (Click here).

that was supposed to investigate and comment on the "prevailing legal chaos" that raged around British agricultural markets. Four years later, in 1926, Linlithgow brought with him many of his unimplemented ideas derived from the British markets to his new job in India (Harris, 1984; Krishnamurthy, 2011). In his comprehensive 1928 report, he declared that the countrywide establishment of regulated markets "would confer an immense boon on the cultivating classes of India" (Royal Commission on Agriculture in India 1928).

This idea was entrenched in the post-colonial policy for agricultural development and regulated markets were seen as key to helping farmers realize a reasonable price in an environment where private trade was underdeveloped and controlled by mercantile power. Therefore, when the constitution of independent India gave legislative powers regarding agriculture to the states, each state of India adopted its own APMC Act in the 1960s.

Under the aegis of the APMC Acts, new markets were created, especially during the Green Revolution that started in the mid-1970s, when agricultural production started increasing exponentially. Between 1976 and 1991, agricultural production grew by 74%, and there was a concomitant 76% increase in the number of markets. The construction of *mandis* has plummeted since the late 1990s. Since 2001, very few new markets have been constructed (Chand, 2012).[11]

1.2.3 Geographical Placement of Markets

Most of these market were constructed between 1960 and 1985, during the peak of the green revolution in India, when grain production increased exponentially. According to the senior bureaucrats I interviewed in the state and central ministries of agriculture, historically, more markets were created in regions that had a greater yield or output per capita. In Appendix 1.B, I verify that controlling for size, cropped area, and population of districts in 1970–1985, the more productive districts have more markets today.

[11]In my data set, I do not observe the year of construction or age of a market. As such, I will assume that in my sample period the number of markets is constant. Chand (2012) reports the aggregate growth rate of markets after 2001 to be 0.7%, and hence, this will not be a major concern for me.

1.3 Data

Geospatial data on the location of markets and prices of commodities sold in them is necessary to assess the importance of spatial competition between intermediary traders in price determination. To capture the demand side, this should be matched geographically to local retail markets and must capture the local distribution of the population and retail prices. Further, credible quantification of production losses requires fine geospatial data on land productivity, land use, rainfall, and crop choice. Because such a data set is not publicly available, I have assembled a rich and novel microdata set for India covering the decade of 2005–2014 by combining data from various sources.

Intermediary Markets: My main data set is comprised of monthly data on the modal price[12] of eight major non-perishable commodities, specifically rice, wheat, maize, sorghum, barley, pearl millet, finger millet, and soybean, sold in any regulated rural agricultural market, along with the village names of the market. This is the price farmers get when they sell in these markets. I obtained this from the Ministry of Agriculture in India. I use Google Maps API to geocode the location of these markets.

I combine this with data on retail prices and production. These data are available only for administrative districts of India. There are 455 districts in the sample I consider, so they can substantially capture the spatial heterogeneity.

Retail Prices: I use monthly data on retail prices at the district level from the National Sample Survey (NSS) Schedule 3.01(R) - Rural Price Collection Survey (RPC) survey of the Central Statistical Organization of India. These data are collected from a fixed set of 603 villages spread across India and are available for the years 2005–2011, except 2008.

[12]These are producer prices,–i.e. the price that a farmer receives. The Indian government also has a Minimum Support Price (MSP) program and announces the procurement prices for 26 crops prior to the start of an agricultural season. Although legally, the government promises to buy farmers' output at the announced MSP whenever market price falls below it, reality is far removed. The MSP is only implemented for rice and wheat, and most procurement occurs in two states, Punjab and Haryana. Even in these states, the government actually relies on the intermediaries to procure the crops (Chatterjee and Kapur, 2017). Appendix Figures 1.E.1–1.E.10 plot the distribution of log prices alongside the MSPs for the crops. There is little evidence of price heaping just at or above the MSPs, as well as substantial mass below MSPs, suggesting any biases from excluding MSPs from the model is limited.

Production and Yields: To understand the cropping patterns of a district, I use data from the National Sample Survey (NSS) Schedule 33 - Situation Assessment Survey of Agricultural Households (Round 70) 2013. This is a large survey of rural agricultural households conducted twice a year, once in each cropping season, in 4529 villages covering 35000 households. Because the survey is representative at the district level, it provides me with a good estimate of the local cropping patterns.

In addition, I use estimates of the price elasticity of demand from Deaton (1997) and of local crop yields at the district level for the years 2005–2014 provided by the Ministry of Agriculture of India.

I further match the above data on prices, production, and consumption with finer data on land use, land elevation, and distribution of population. In particular, I obtained gridded data on land use from Princeton University's Geospatial Information Systems library, gridded data on land elevation produced by NASA's Shuttle Radar Topography Mission (SRTM),[13] gridded data on rainfall from Willmott and Matsuura (2001), and geocoded data on village population from the Census of India 2001. I compute district level monthly rainfall by taking an inverse-distance weighted average of all the grid points within the boundary of any district.

Sample: For my analysis I exclude the mountainous states of India because of the difficulty in measuring distances. I exclude the state of Bihar and all North Eastern states because they do not have an APMC Act, and hence there is no data on market–level prices. I also exclude some other territories and islands where agriculture is not practiced on any substantial scale. Therefore, my sample includes 15 states in mainland India covering 89% of the total cropped area and accounting for 90% of total production. The states are shown in Figure 1.A.1. In my data set, this results in a total of 2978 primary grain markets.

Table 1.1 presents summary statistics on prices at the market and retail levels. The average variance of farmer prices across markets within crop-month is 0.029. The variance

[13]National Aeronautics And Space Administration And The National Geospatial Intelligence Agency, 2000

in retail prices across regions, conditional on the crops' being produced in that region, is lower than the variance in farmer prices. The average variance in retail prices across districts within state-crop-month is 0.014, whereas the average variances in farmer prices across markets within state-crop-month is 0.022.

1.4 Does more spatial competition increase farmers' prices?

The analysis in this section is focused on the regulated marketplace where the transaction between the farmer and the government-licensed intermediary occurs.[14] There are two potential forms of local competition among buyers (i.e. the licensed intermediaries in markets): between and within market sites. In this chapter, I will focus on between-market competition because current empirical evidence suggests that intermediaries within a market collude.[15] This section provides two forms of reduced-form evidence that suggest: (a) local competition between market sites increases prices that farmers get for their output (hereinafter farmer prices); and (b) state borders attenuate competitive forces, i.e. the presence of more markets in the neighboring state does not influence farmer prices.

The hypothesized mechanism influencing farmer prices, which will be explicit in the quantitative model, hinges on a farmer's access to alternative buyers while negotiating with a buyer at a given market site. More markets in the vicinity increase the set of alternatives available to the farmer, increasing the competition faced by the present buyer. Consequently, the farmer is likely to be offered a better price. Therefore, greater competition is likely when alternative markets are closer and there are more of them. Moreover, since the APMC laws prevent farmers from crossing state borders, more markets in a neighboring state should not have an influence on farmer prices. I test for these implications below.

[14] See Section 1.2.1 for a discussion of the overall economic environment of agriculture trade in India.
[15] See, for example, Banerji and Meenakshi (2004) and Meenakshi and Banerji (2005).

1.4.1 Effect of within- and out-of-state competition

To explore the association between local competition and farmer prices, I regress farmer prices at a particular market site on a measure of local market density (or local competition faced by a market). Similar to the market access measure in Harris (1954) and Donaldson and Hornbeck (2016), I construct a local competition measure by taking a weighted sum of other markets near a particular market site but in the same state. The weights are the inverse of distances of the neighboring markets to the origin market.[16] For any market m,

$$\text{comp}_m = \sum_{j \in \mathcal{M} \setminus \{m\}} \left\{ \frac{1}{\text{distance}_{mj}} \right\} \mathbb{1} \{\text{state of } m = \text{state of } j\} \tag{1.4.1}$$

$\mathcal{M}$ is the set of all markets in the country. As competition (in any market m) is driven by a farmer's ease of access to alternative markets, the comp_m measure assigns a greater weight to a closer market. I also create an analogous competition measure for a market site from the markets not in the same state. Under the null hypothesis, this measure should have no association with prices received by farmers at any market site.

$$\text{comp}'_m = \sum_{j \in \mathcal{M} \setminus \{m\}} \left\{ \frac{1}{\text{distance}_{mj}} \right\} \mathbb{1} \{\text{state of } m \neq \text{ state of } j\} \tag{1.4.2}$$

My main specification takes the following form:

$$\log p^f_{cmdst} = \beta_0 + \beta_1 \text{comp}_m + \beta_2 \text{comp}'_m + \mathbf{X}'_{cdt}\beta_3 + \gamma_t + \gamma_c + \gamma_s + \epsilon_{cmdt}, \tag{1.4.3}$$

Here, p^f_{cmdst} is the price that a farmer receives at market site m located in district d, state s, for crop c, at time (month-year) t. All price observations are at market-crop-month level. $\mathbf{X}$ includes controls for district specific time (year) varying controls like crop yields, crop

[16] Donaldson and Hornbeck (2013) compute a weighted average of the total trade in a market. I do not have data on trade and therefore do not compute the trade-weighted measure. However, the number of proximate competitors, as a measure of competition, has also been used by Machiavello and Morjaria (2016).

area, population, and district-crop specific rainfall shocks. γ_t is a month-year fixed effect and controls for crop and district invariant unobservables such as macroeconomic shocks, large scale droughts, etc. γ_c controls for crop-specific factors such as crop-specific price levels and national tastes for particular crops. γ_s controls for state-specific and crop-time invariant effects such as state-level policies, relative incomes of states, and local tastes. Since cropping decisions and other shocks are likely to be spatially correlated, I compute robust standard errors clustered at the district level.

Variation in comp_m comes from geographical differences in the placement of markets. Because the local crop production to population ratio drives this variation, I control for local yields, local cropped area, and population in the regression (1.4.3). β_1 is identified from within-state variation in comp_m under the assumption that comp_m is uncorrelated with the residuals. As very little *mandi* construction occurred in my sample period, and I do not have data on the date of construction of markets, comp_m does not vary over time. Therefore, threats to identification can come only from spatially varying factors, and I have tried to control for as many of them as possible. I have not included district fixed effects because there are very few markets within a district, and therefore, I do not have enough within-district variation in spatial competition to identify its effect on prices.

Within-market competition is an omitted variable in all my regressions and can potentially confound my estimates if markets in local areas with a greater market density also have greater within-market competition. In the absence of any data on the number of buyers within a market or their transactions, I cannot control for it directly. However, within-market competition is unlikely to be important because existing evidence points towards collusion among intermediaries within agricultural markets, not only in India (Banerji and Meenakshi, 2004; Meenakshi and Banerji, 2005) but also in other parts of the developing world (Bergquist, 2017). In the Indian context, incumbent intermediaries have also actively tried to prevent entry of new traders (Chand, 2012).

Table 1.2 reports the regression results. Column 1 reports results from the base specification given in equation (1.4.3). The coefficient on competition, β_1, is significant and equals 0.0163. In Column 2, I also include crop-time fixed effects, which control for monthly world price shocks, and state-year fixed effects, which control for state-specific income levels and policy changes that vary over time. This does not change the estimate of β_1 much, which now equals 0.0146 and is still significant. The coefficient on out-state competition β_2 is close to zero and not significant in both columns.

The key message is that greater local competition increases farmer prices. A one standard deviation increase in competition increases prices by 2.7%. Moreover, the competition from markets in other states has no impact on prices. To get a sense of the overall magnitude of gains, note that if we removed all border restrictions to trade, the median increase in competition would be 1.6 standard deviations. This would increase prices in half the markets in India by at least 4.4%.

1.4.2 Causal Effects of Local Competition on Prices

Although the previous regression design shows that local competition increases farmer prices and borders reduce out-of-state market competition, one can still be concerned about other forms of unobserved heterogeneity. Therefore, to test this idea more concretely, I implement a border discontinuity design with market pairs. I match all markets, which are less than x kilometers apart but lie on different sides of a state boundary. Here, $x = 25, 30$, and 35. I regress the difference in prices for each market pair on the difference in their competition measures.

The basic idea is that other determinants of prices like demand, productivity via soil quality, and rainfall vary continuously across a state boundary. Indian state boundaries were primarily determined along linguistic or cultural lines rather than geography (Guha, 2007). Thus, the only determinant of prices that changes discontinuously across state borders is local competition because farmers are not allowed to sell their output in the markets of

other states.[17] Therefore, along a state boundary, the correlation between the difference in prices and the difference in market density between market pairs should help us identify causal effects of local competition on prices.

In the previous design, since I relied on within-state variation for identification, I could not fully address concerns about unobserved heterogeneity. The advantage of this design is that I can difference out unobserved factors other than competition that affect prices by choosing market pairs very close to each other.

I start with a levels specification as in (1.4.3) and then take a difference between all market pairs $(m - m')$ that are within the selected bandwidth but lie in different states.

$$\begin{aligned} \log p^f_{cmdt} &= \beta_0 + \beta_1 \text{comp}_m + \gamma_t + \gamma_c + \gamma_s + \epsilon_{cmdt} \\ \Delta \log p^f_{cmdt} &= \beta_1 \left(\Delta \text{comp}_m\right) + \gamma_{ss'} + \tilde{\epsilon}_{cmt} \end{aligned} \tag{1.4.4}$$

β_1 is the causal effect of competition on prices. $\gamma_{ss'}$ is a border fixed effect, specific to two states sharing a common border. This specification controls for unobservables like differences in state-specific policies, law enforcement, and tax regimes that do not vary along the common border but might potentially differ across state pairs. Because the market pairs are very close to each other, shocks can be spatially correlated. I report (in square brackets) standard errors corrected for spatial correlation.

Following Conley (1999), I allow for spatial dependence of an unknown form within a 100 km surface from the mid-point of the line segment joining the market pair. To be more precise, the variance-covariance matrix is estimated as a weighted sum of cross-products of the OLS moments ($E\left[x_i \epsilon_i\right] = 0$). I use two kernels for determining the weights: (a) the uniform kernel, which puts equal weight on all observation pairs within 100 km and 0

[17] Indian languages change very gradually over distance. Therefore, it is not the case that farmers in one state will not be able to communicate with farmers or intermediaries in the neighboring state.

otherwise; and (b) the Bartlett kernel, which gradually decreases the weight the further the observations are. For comparison, I also report robust standard errors in parentheses.[18]

I report the results in Table 1.3. The preferred estimates are reported in column (2) which has a bandwidth of 30 km. The marginal effect of competition on prices, β_1, equals 0.035. Choosing bandwidths lower than 30 km reduces the number of market pairs and then I do not have enough power to precisely estimate the effect of competition on prices. As reported in Column 1 of Table 1.3, where the bandwidth is 25 km, β_1 is equal to 0.025 but is not precisely estimated. Also, as can be seen from Column 3 of the same table, the estimate of β_1 settles down to about 0.035 as we increase the bandwidth further.[19]

Like before, there is a positive effect of local competition on farmer prices. However, the effect size is almost double compared to the estimate of 0.016 from the levels specification reported in Table 1.2. A farmer selling in a market that faces the 75th percentile of competition compared to one that faces the 25th percentile of competition gets a 3.5% higher price on average. A one standard deviation increase in competition would lead to a 6.4% increase in price on average. Appendix 1.C presents placebo tests to show that there are no discontinuities at state borders in other observable factors like potential crop productivity, land elevation, and distance to the nearest district headquarter that might influence farmer prices.

Another way of estimating this causal effect is to use a levels specification like (1.4.3) to mimic the border discontinuity design. β_1 could be biased in (1.4.3) if it is correlated with other omitted factors that also affect prices. However, if we limit our sample to markets very close to all state borders, then these omitted factors are likely to affect comp'_m in a manner similar to their effect on comp_m. Therefore, $\beta_1 - \beta_2$, should help us identify the "true" competition effect as β_2 purges β_1 from any contamination.

[18] I am restricted to computing cluster robust standard errors for the levels specification (1.4.3) because of the computational burden of computing Conley (1999) standard errors.

[19] Note that my choice of bandwidths is much smaller than similar border discontinuity designs as in Dell (2010).

I report these results in Columns 3 and 4 of Table 1.2, where we indeed find that the estimated effect of 0.032 is very close to the estimated effect of 0.035 in the border discontinuity design.

To summarize, in this section I have provided evidence for two aspects of spatial competition between intermediary markets. First, local competition, determined via density of markets, matters for farmer prices. Second, border restrictions have real effects in that they nullify the competitive forces from markets in neighboring states. As such, the quantitative effects of increased competition on prices, agricultural output from an induced change in use of intermediate inputs, and other welfare measures can be large if border restrictions on trade are removed. With farmers' being able to freely cross state boundaries, not only does competition increase among markets at the border but also potentially throughout the country via a ripple effect—since a price increase in one market influences the prices in adjacent markets and so on—and therefore affecting prices even in the interior markets.

Quantification of such ripple effects needs the aid of a spatial model of trade that uses the location of markets and the economic geography of the country and allows farmers to remove price arbitrages in space. In the next chapter, I set out such a model and study the spatial and aggregate implication of removing intra-country trade restrictions on Indian farmers.

References

Allen, Treb and David Atkin, "Volatility and the Gains from Trade," *NBER Working Paper*, 2016, (22276).

Anderson, James E. and Eric van Wincoop, "Gravity with Gravitas: A Solution to the Border Puzzle," *The American Economic Review*, 2003, *93* (1), 170–192.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro, "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *The Quarterly Journal of*

Economics, 2017, *132* (4), 1553–1592.

Antras, Pol and Arnaud Costinot, "Intermediated Trade," *The Quarterly Journal of Economics*, August 2011, *126* (3), 1319–1374.

Atkin, David and Dave Donaldson, "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs," *NBER Working Paper*, 2015, (21439).

Bai, Xue, Kala Krishna, and Hong Ma, "How You Export Matters: Export Mode, Learning and Productivity in China," *Journal of International Economics*, 2017, *104*, 122–137.

Banerji, A. and J. V. Meenakshi, "Buyer Collusion and Efficiency of Government Intervention in Wheat Markets in Northern India: An Asymmetric Structural Auctions Analysis," *American Journal of Agricultural Economics*, 2004, *86* (1), 236–253.

Bardhan, Pranab, Dilip Mookherjee, and Masatoshi Tsumagari, "Middlemen Margins and Globalization," *American Economic Journal: Microeconomics*, 2013, *5* (4), 81–119.

Bates, Robert H, *Markets and States in Tropical Africa: The Political Basis of Agricultural Policies*, Berkeley: University of California Press, 1981.

Bergquist, Lauren Falcao, "Pass-through, Competition, and Entry in Agricultural Markets: Experimental Evidence from Kenya.," *UC Berkeley Mimeograph*, 2017.

Casaburi, Lorenzo and Tristan Reed, "Competition and Interlinkages in Agricultural Markets: An Experimental Approach," *Working Paper*, 2016.

Chand, Ramesh, "Development Policies and Agricultural Markets," *Economic and Political Weekly*, 2012, *47* (52).

Chatterjee, Shoumitro and Devesh Kapur, "Six Puzzles in Indian Agriculture," *India Policy Forum 2016–17*, 2017, *13.*

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

Cooper, Russell and Huacong Liu, "Money or Grit? Determinants of Mismatch by Race and Gender," *NBER Working Paper*, 2016, (22734).

Cosar, A Kerem, Paul Grieco, and Felix Tintelnot, "Borders, Geography, and Oligopoly: Evidence from the Wind Turbine Industry," *Review of Economics and Statistics*, 2015, *97* (3), 623–637.

Deaton, Angus, *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*, World Bank Publications, 1997.

Dell, Melissa, "The Persistent Effects of Peru's Mining Mita," *Econometrica*, 2010, *78* (6), 1863–1903.

Dillon, B. and C. Dambro, "How Competitive are Food Crop Markets in Sub-Saharan Africa? A Review of the Evidence," *Working Paper, Evans School of Public Policy and Governance, University of Washington*, 2016.

Donaldson, Dave, "The Gains from Market Integration," *The Annual Review of Economics*, 2015, *7* (1), 619–647.

_ , "Railroads of the Raj: Estimating the Impact of the Transportation Infratructure," *American Economic Review, forthcoming*, 2015.

_ **and Richard Hornbeck**, "Railroads and American Economic Growth: A Market Access Approach," *The Quarterly Journal of Economics*, 2016, *131* (2), 799–858.

Eaton, Jonathan and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, *70* (5), 1741–1779.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu, "Competition, Markups, and The Gains from International Trade," *The American Economic Review*, 2015, *105* (10), 3183–3221.

Engel, Charles and John H. Rogers, "How Wide Is the Border?," *The American Economic Review*, 1996, *86* (5), 1112–1125.

Fafchamps, M, E Gabre-Madhin et al., "Agricultural markets in Benin and Malawi.," *African Journal of Agricultural and Resource Economics*, 2006, *1* (1), 67–94.

Fafchamps, Marcel and Bart Minten, "Returns to social network capital among traders," *Oxford Economic Papers*, April 2002, *54* (2), 173–206.

_ , Eleni Gabre-Madhin, and Bart Minten, "Increasing returns and market efficiency in agricultural trade," *Journal of Development Economics*, 2005, *78* (2), 406–442.

Gourieroux, Christian, Alain Monfort, and Eric Renault, "Indirect Inference," *Journal of Applied Econometrics*, 1993, *8* (S1).

Great Britain and Royal Commission on Agriculture in India, *Report of the Royal Commission on Agriculture in India.*, London: H.M. Stationery Office, 1928.

Guha, Ramachandra, *India after Gandhi: The History of The World's Largest Democracy*, New Delhi: Harper Collins, 2007.

Harris, Barbara, *State and market: State Intervention in Agricultural Exchange in a Dry Region of Tamil Nadu, South India*, New Delhi: Concept Pub. Co., 1984.

Harris, Chauncy D., "The Market as a Factor in the Localization of Industry in the United States," *Annals of the Association of American Geographers*, 1954, *44* (4), 315–348.

Honore, Bo E., Thomas Jorgensen, and Aureo de Paula, "Sense of Sensitivity," *Working Paper*, 2017.

Horn, Henrick and Asher Wolinsky, "Bilateral Monopolies and Incentives for Merger," *The RAND Journal of Economics*, 1988, pp. 408–419.

Kapur, Devesh and Mekhala Krishnamurthy, "Market Towns and the Dynamics of India's Rural and Urban Transformations," *Working Paper, International Growth Center*, 2014, (F-35007-INC-1).

Krishna, Kala and Yelena Sheveleva, "Wheat or Strawberries? Intermediated Trade with Limited Contracting," *American Economic Journal: Microeconomics*, August 2017, *9* (3), 28–62.

Krishnamurthy, M., "Harda Mandi: Experiencing Change in an Agricultural Market in Central India (1980-2010)." Ph.D. thesis, UCL (University College London) October 2011.

Loecker, Jan De, Pinelopi K Goldberg, Amit K Khandelwal, and Nina Pavcnik, "Prices, Markups, and Trade Reform," *Econometrica*, 2016, *84* (2), 445–510.

McCallum, John, "National Borders Matter: Canada-US Regional Trade Patterns," *The American Economic Review*, 1995, *85* (3), 615–623.

Meenakshi, J.V. and A. Banerji, "The Unsupportable Support Price: An Analysis of Collusion and Government Intervention In paddy Auction Markets in North India," *Journal of Development Economics*, 2005, *76* (2), 377 – 403.

Melitz, Marc J. and Gianmarco Ottaviano, "Market Size, Trade, and Productivity," *The Review of Economic Studies*, 2008, *75* (1), 295–316.

Ministry of Agriculture, Government of India, "Agricultural Statistics at A Glance 2016-17," 2016.

Ministry of Finance, Government of India, "Economic Survey of India 2014–15," 2015.

_ , "Economic Survey of India 2015–16," 2016.

Mitra, Sandip, Dilip Mookherjee, Maximo Torero, and Sujata Visaria, "Asymmetric Information and Middleman Margins: An Experiment with Indian Potato Farmers," *The Review of Economics and Statistics, forthcoming*, 2017.

National Aeronautics And Space Administration And The National Geospatial Intelligence Agency, "Shuttle Radar Topography Mission 30 Arc Second Finished Data," 2000.

Redding, Stephen J and Daniel M Sturm, "The Costs of Remoteness: Evidence from German Division and Reunification," *The American Economic Review*, 2008, *98* (5), 1766–1797.

Willmott, C. J. and K. Matsuura, "Terrestrial Air Temperature and Precipitation: Monthly and Annual Time Series (1950 - 2014)," 2001.

Tables and Figures

TABLE 1.1: SUMMARY STATISTICS

Log Farmer Prices (Rs. per kg)	
Mean	
Barley	2.24
Finger Millet	2.21
Maize	2.15
Paddy	2.25
Pearl Millet	2.17
Sorghum	2.33
Soybean	2.92
Wheat	2.46
Variance Across Markets, Within crop x month	0.029
Variance Across Markets, Within state x crop x month	0.022
Variance Across Months, Within market x crop x agricultural season	0.003
Log Retail Prices (Rs. per kg)	
Mean	
Barley	2.55
Finger Millet	2.29
Maize	2.18
Paddy	2.11
Pearl Millet	2.20
Sorghum	2.34
Soybean	3.33
Wheat	2.49
Variance across districts, within crop x month	0.053
Variance across districts, within crop x month (Conditional on production)	0.026
Variance across districts, within state x crop x month	0.029
Variance across districts, within state x crop x month (Conditional on production)	0.014
Variance Across Months, Within district x crop x agricultural season	0.002
Standard Deviation of comp (as defined in equation 1)	1.85

TABLE 1.2: FARMER PRICES AND LOCAL COMPETITION

Dependent Variable:	log price			
	All Markets		Border Markets within 30 km	
	(1)	(2)	(3)	(4)
comp (β_1)	0.0163	0.0146	0.0322	0.0292
	(0.0043)***	(0.0035)***	(0.0154)***	(0.0132)***
comp' (β_2)	-0.0005	-0.0007	-0.0057	-0.0056
	(0.0012)	(0.0009)	(0.0047)	(0.0041)
$\beta_1 - \beta_2$	0.0169	0.0154	0.0381	0.0348
	(0.0041)***	(0.0034)***	(0.0144)***	(0.0123)***
Observations	211963	211954	15838	15751
R-Squared	0.60	0.68	0.57	0.68
State, Crop FE	✓	✓	✓	✓
Month-Year FE	✓		✓	
State-Year, Crop-Month-Year FE		✓		✓

Notes: Ordinary least squares. Each observation is a crop-market-month-year. All regressions include controls for district-year specific crop-yields, crop area, crop-specific rainfall shocks, and local population. Robust standard errors clustered at the district level reported in parenthesis. Crops in sample are barley, finger millet, maize, paddy, pearl millet, sorghum, soybean and wheat. *p $<$ 0.1, ** p $<$ 0.05, *** p $<$ 0.01

TABLE 1.3: BORDER DISCONTINUITY REGRESSIONS

Dependent Variable:	Δ log price		
	Markets <25 km apart	Markets <30 km apart	Markets <35 km apart
	(1)	(2)	(3)
Δ comp (β_1)	0.025	0.035	0.036
Conley (Uniform)	[0.015]	[0.016]**	[0.012]**
Conley (Bartlett)	[0.015]	[0.017]**	[0.012]**
Robust	(0.011)***	(0.013)***	(0.009)***
Obs	3538	5292	9018
Mkt Pairs	108	173	284
State-Pair FE	✓	✓	✓

Notes: Observations are market-pair-crop-month. Column (a) reports regressions with markets pairs on state borders and (b) reports placebo regressions with market pairs not on state borders. Brackets contain Conley (1999) standard errors that allow for spatial and serial correlation of the error terms. Arbitrary spatial correlation in error terms is allowed within 100 km from the geographic midpoint of the market pairs. Serial correlation is allowed over 6 months, equivalent to one crop season in India. Robust standard errors are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix

1.A States in Sample



Figure 1.A.1: States in Sample

1.B Construction of Markets

In this section, I present conditional correlations between factors that drove geographical placement of markets. I regress the number of markets in a district as observed in my main data set on average cereal output, cereal yield, district area, cropped area, population between 1970-1985, and state fixed effects. This time period was during the peak of the green revolution when most new market construction occurred. District level data on production and population for this period come from the Village Dynamics in South Asia (VDSA) "meso" data set. We can see from Table 1.B.1 that conditional on district size, cropped area, and population, districts that had a greater cereal output in 1970-1985 also have a greater number of markets today.

TABLE 1.B.1:FACTORS INFLUENCING MARKET CONSTRUCTION

Dependent Variable:	Number of markets in 2005-2014	
cereal output	35.7487	
	(10.1235)	
cereal yield		0.3723
		(0.1018)
crop area	-0.0342	0.3315
	(0.1277)	(0.1198)
total area	0.4074	0.4083
	(0.0981)	(0.0979)
population	0.2400	0.2290
	(0.1062)	(0.1067)
N	251	251
R-Squared	0.67	0.67

Notes: All observations are at district level. All variables are in logs. All right hand side variables are district level averages over 1970-1985. All regressions have state fixed effects. OLS standard errors in parenthesis.

1.C Placebo Tests

In this section, I provide regression results to show that there are no discontinuities in other observable factors that might influence farmer prices. I run specifications like (1.4.4), but regress observable factors on competition instead of farmer prices.

In panel A, I regress log of potential yields of various crops on competition to show that there is no discontinuity in crop productivity at the border. In panel B, I regress log of land elevation and log on distance to the district headquarter on competition. Asher, Nagpal and Novosad (2017) shows that distance to district headquarter effects delivery of public goods and services to villages. Differences in public goods like roads can influence farmer prices and can be a concern if correlated with spatial competition. Column 2 of Panel B allays this concern.

PANEL A

	Δ log (agro-climatically attainable yield)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Rice	Wheat	Maize	Millet	Barley	Sorghum	Soybean
Δ comp	0.007	-0.016	-0.002	0.003	-0.003	0.002	0.009
	(0.020)	(0.029)	(0.006)	(0.005)	(0.026)	(0.006)	(0.008)
Obs	173	173	173	173	173	173	173
State-Pair FE	✓	✓	✓	✓	✓	✓	✓
Kernel	Uniform	Uniform	Uniform	Uniform	Uniform	Uniform	Uniform

Notes: Brackets contain Conley (1999) standard errors that allow for spatial and serial correlation of the error terms. Arbitrary spatial correlation in error terms is allowed within 100 km from the geographic midpoint of the market pairs. Bandwidth for market pairs is 30 km. Data on Agro-Climatically attainable yield from FAO GAEZ and averaged for a 50 km zone within the sate around each market. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

PANEL B

	(1)	(2)
	Δlog(elevation)	Δlog(dist to district HQ)
Δcomp	0.053	0.301
	(0.076)	(0.820)
Obs	173	173
State-Pair FE	✓	✓
Kernel	Uniform	Uniform

Notes: Brackets contain Conley (1999) standard errors that allow for spatial and serial correlation of the error terms. Arbitrary spatial correlation in error terms is allowed within 100 km from the geographic midpoint of the market pairs. Bandwidth for market pairs equals 30 km. Data on land elevation from NASA SRTM and averaged around a 50 km zone within the state for each market. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

1.D Data Appendix

1.D.1 District Names

To maintain consistency across time periods, I appropriate all data to 2001 census boundaries of Indian districts. That is, districts that got bifurcated after 2001 are assigned to their parent districts.

1.D.2 Organizing the Data on Farmer Prices and Retail

Each observation in the farmer prices data provides the market level mode price for a given month for each commodity that was sold in any market. There is a variable that has

information about the variety of the commodity. The village, district, and state of the market is also recorded.

I cleaned the data on farmer prices in the following steps. First, I winsorized the top and bottom 1% of the outlier observations. Unfortunately, most of the data (by the number of observations) is under the variety "other" and hence, this variable is not of much use. In the second step to maintain homogeneity of commodities, I dropped high-value varieties like "Super Fine Rice," "Basmati Rice," "Dehradoon Rice," "Sharbati Wheat," "Fine Wheat" etc. I also dropped Maize and Millets meant for "Cattle Fodder." Third, I corrected for Indian names of villages such that there is no duplication and then I geocoded the markets using the names of the state, district, and village in Google Maps API.

The unit of observation in the retail price data is a commodity-district-month-year. Data are recorded for three types of rice—fine, medium, and coarse. I use the data on medium and coarse varieties of rice and compute the median price across varieties. I fixed district names to match those in the farmer price data. Then I removed the outlier observations, by manually examining district-crop level data. I did not follow the 1% winsorizing rule since it would drop more observations than required. Then I computed district-crop-month level median prices, which I use as my final data on retail prices.

1.E Minimum Support Prices and Farmer Prices

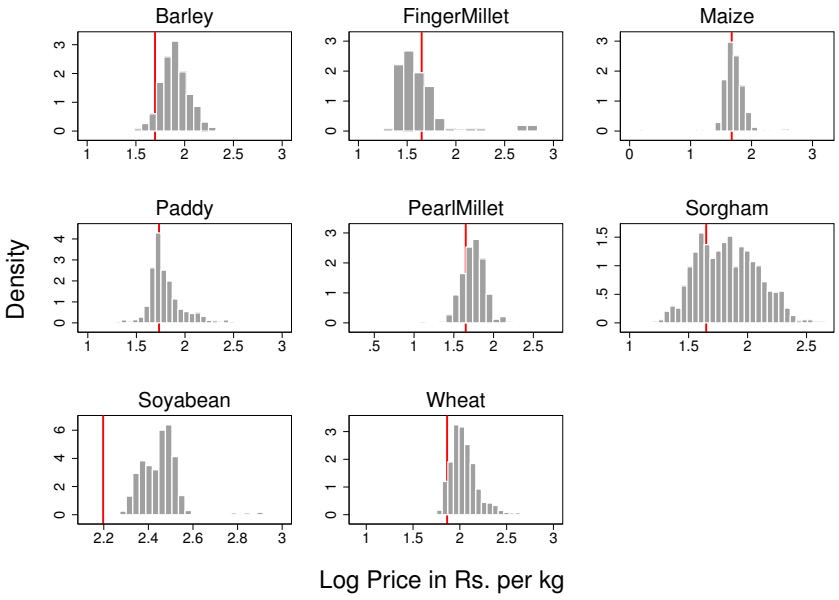


Figure 1.E.1: Distribution of Prices and MSPs in Crop Year 2004–05
Notes: This figure plots the distribution of log prices across markets for the sample crops in the 2004–05 crop year. Vertical lines show minimum support prices (MSPs).

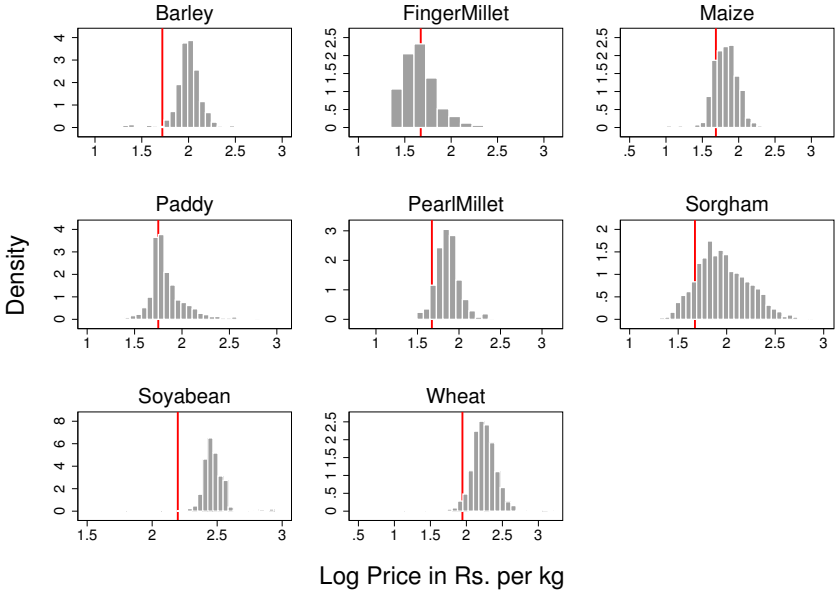


Figure 1.E.2: Distribution of Prices and MSPs in Crop Year 2005–06
Notes: This figure plots the distribution of log prices across markets for the sample crops in the 2005–06 crop year. Vertical lines show minimum support prices (MSPs).

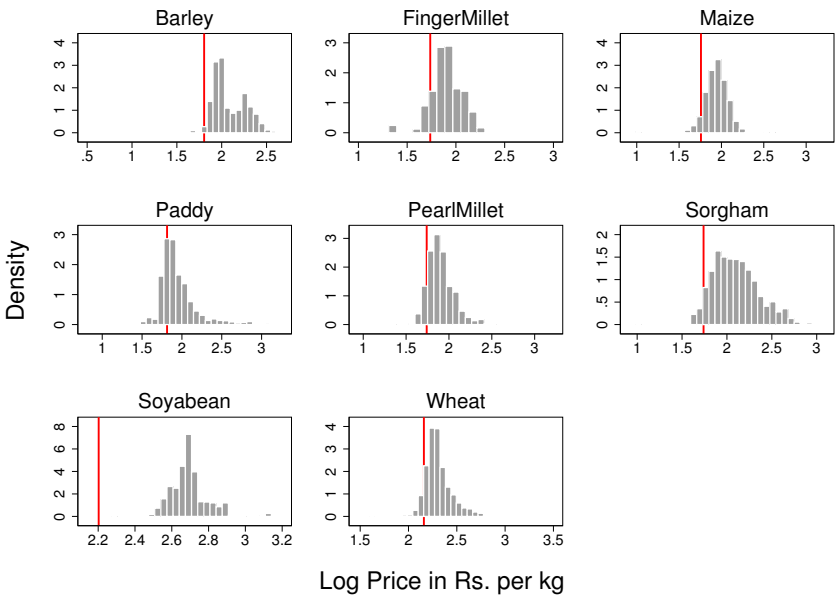


Figure 1.E.3: Distribution of Prices and MSPs in Crop Year 2006–07

Notes: This figure plots the distribution of log prices across markets for the sample crops in the 2006–07 crop year. Vertical lines show minimum support prices (MSPs).

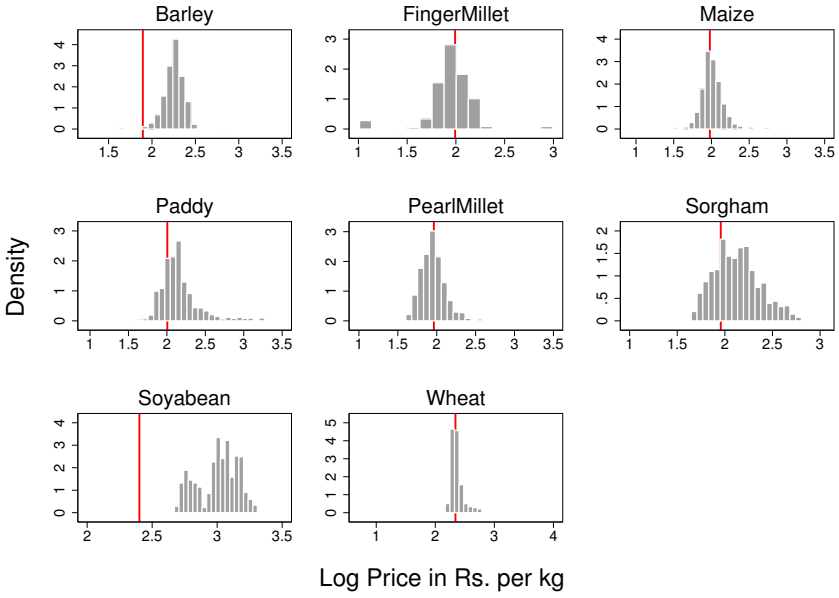


Figure 1.E.4: Distribution of Prices and MSPs in Crop Year 2007–08

Notes: This figure plots the distribution of log prices across markets for the sample crops in the 2007–08 crop year. Vertical lines show minimum support prices (MSPs).

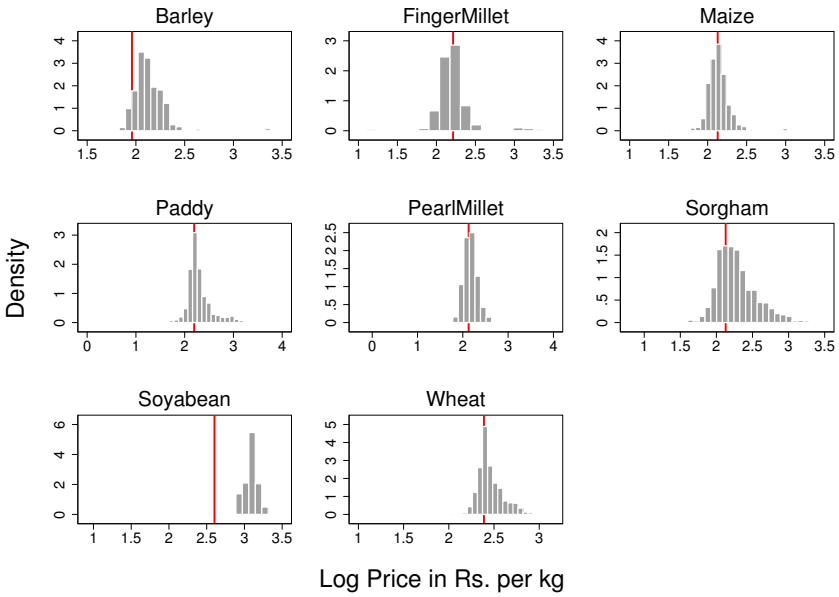


Figure 1.E.5: Distribution of Prices and MSPs in Crop Year 2008–09

Notes: This figure plots the distribution of log prices across markets for the sample crops in the 2008–09 crop year. Vertical lines show minimum support prices (MSPs).

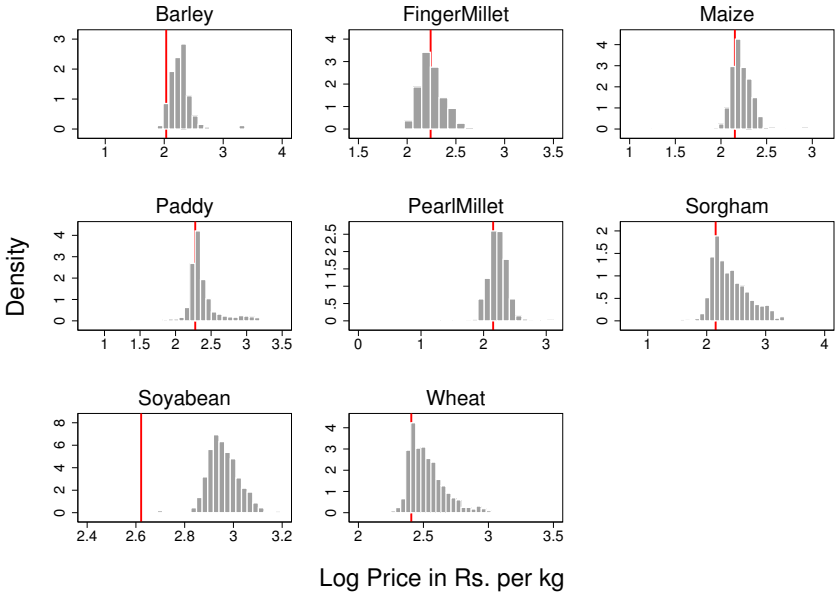


Figure 1.E.6: Distribution of Prices and MSPs in Crop Year 2009–10

Notes: This figure plots the distribution of log prices across markets for the sample crops in the 2009–10 crop year. Vertical lines show minimum support prices (MSPs).

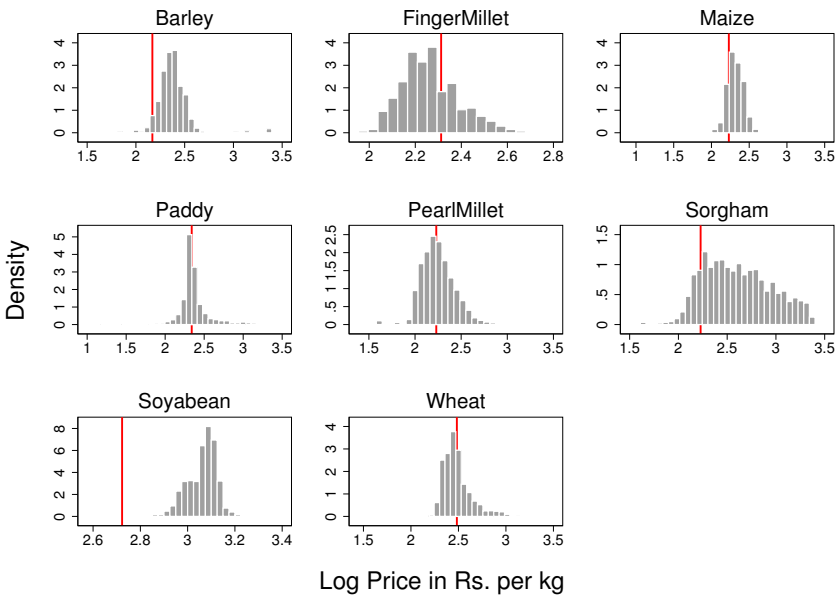


Figure 1.E.7: Distribution of Prices and MSPs in Crop Year 2010–11

Notes: This figure plots the distribution of log prices across markets for the sample crops in the 2010–11 crop year. Vertical lines show minimum support prices (MSPs).

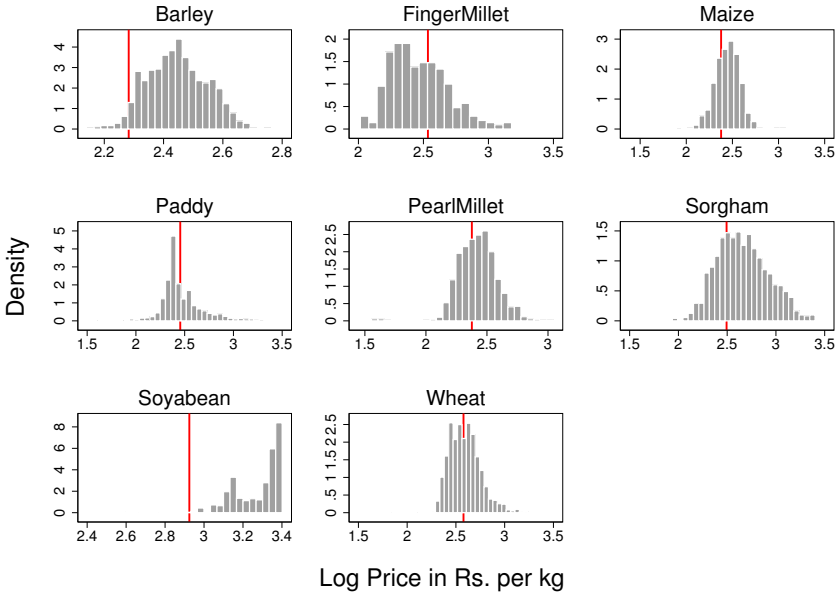


Figure 1.E.8: Distribution of Prices and MSPs in Crop Year 2011–12

Notes: This figure plots the distribution of log prices across markets for the sample crops in the 2011–12 crop year. Vertical lines show minimum support prices (MSPs).

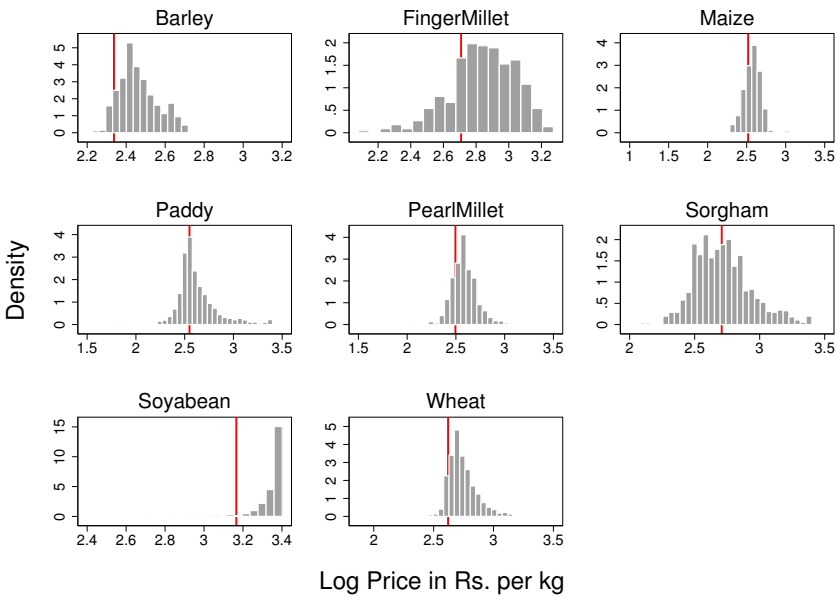


Figure 1.E.9: Distribution of Prices and MSPs in Crop Year 2012–13

Notes: This figure plots the distribution of log prices across markets for the sample crops in the 2012–13 crop year. Vertical lines show minimum support prices (MSPs).

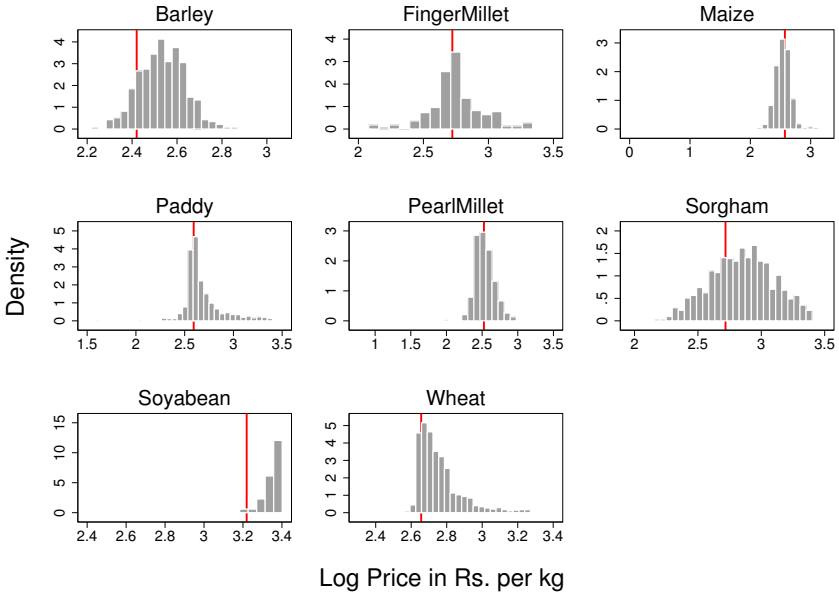


Figure 1.E.10: Distribution of Prices and MSPs in Crop Year 2013–14

Notes: This figure plots the distribution of log prices across markets for the sample crops in the 2013–14 crop year. Vertical lines show minimum support prices (MSPs).

Chapter 2

General Equilibrium Effects of Spatial Competition

2.1 Introduction

In this chapter, I develop a simple model of trade in agricultural markets in India that will flexibly capture the forces of spatial competition between markets and aid us in quantifying the income and productivity effects of border restrictions on trade. The basic structure of the model is as follows. The economy consists of many geographic regions. Within each region there are two types of agents: farmers and intermediaries. Each intermediary is identified with a market. The location of markets is determined exogenously by the government. Each farmer first chooses input levels. After realization of output, he chooses the market where he wants to sell his output. Transporting goods is costly. The price that a farmer receives at a given market is determined by Nash bargaining between the farmer and the intermediary after the farmer has arrived at the market with the output. Intermediaries sell all of the purchased output in the retail market at an exogenously given price. Figure 2.1 shows the timing of events.

I provide formal details of the model next.

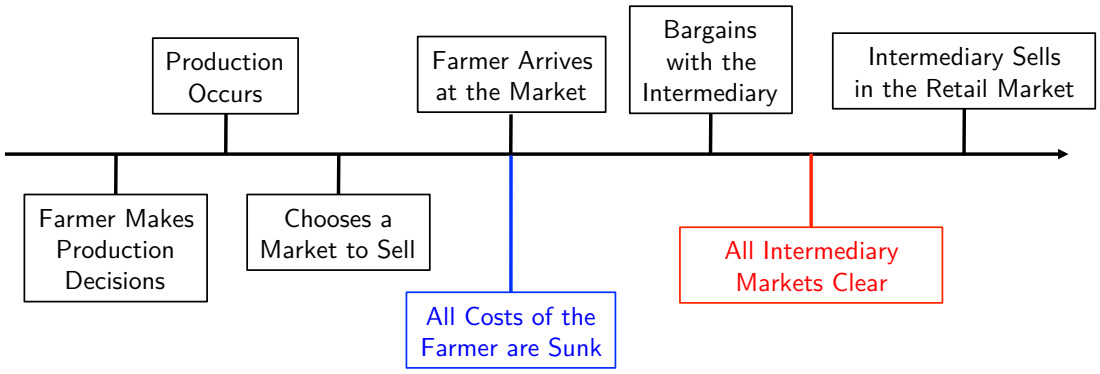


Figure 2.1: Timing of Events

2.1.1 Model

Geography

There are S regions in the economy. Regions in the model are analogous to the states of India. For now, and motivated by the regulatory structure in India that prohibits farmers from selling their output across different states, there will be no interactions between regions. In what follows, I will focus on a generic region and abstract from indexing the region.

The region has two types of agents: F farmers and M intermediaries. $\mathcal{F}$ and $\mathcal{M}$ are the sets of all farms and intermediaries in the region, respectively. Each farmer f is identified by a given plot of land, or farm, and each intermediary m is identified by a specific market. Farms and markets are distributed in space, so I will also refer to f and m as locations. Geography matters because it is costly to move goods across locations. I model these costs as iceberg trade costs and denote by τ_{xy} (> 1) the amount of goods that must be moved from location x in order for 1 unit to arrive at location y.

Farmers

There is a single crop in the economy.[1] Land and labor available to a given farm are fixed and denoted by h_f and l_f respectively.[2]

[1] I will incorporate multiple crops later but will continue to treat crop choice as exogenous.

[2] This is the farm of the median Indian farmer who owns about 1.5 hectares of land (Economic Survey of India, Ministry of Finance, Government of India, 2016) and mostly uses family or village labor for cultivation.

A farmer purchases $\left\{x_f^k\right\}_k$ units of K intermediate inputs and his output is given by:

$$y_f = \tilde{A}_f \left(h_f^{\gamma} l_f^{\nu} \prod_{k=1}^{K} \left(x_f^k\right)^{\alpha_k} \right), \quad \sum_{k=1}^{K} \alpha_k + \nu + \gamma = 1; \alpha_k > 0 \forall k, \nu, \gamma > 0, \tag{2.1.1}$$

$\tilde{A}_f$ is the total factor productivity for farm f. The market for inputs is perfectly competitive. w^k, the price of the kth input, is the same for all farmers in the region and is exogenously given. Farmers are small compared to the input, output, and the retail markets and therefore cannot influence prices directly.

Market Choice

Upon harvest, farmers optimally choose the market where they want to sell. Because farmers are subject to iceberg trade costs, the quantity of output actually reaching any market m from farm f is

$$q_{fm} = \frac{y_f}{\tau_{fm}}. \tag{2.1.2}$$

The price that farmers would get in market m is denoted by $p^f(m)$. Thus, farmers choose market m to sell their output such that their income $\frac{p^f(m) y_f}{\tau_{fm}}$ is maximized.

Price Determination and Intermediaries

Once a farmer f reaches a market m with his goods, all his costs are sunk. He bargains with the intermediary associated with market m over the price for a unit of his output. The outcomes are determined via Nash bargaining. The outside option of the farmer, $\underline{p}(m)$, is the value that he would receive by going to the best alternative market. The best alternative is given by

$$\underline{p}(m) = \max_{k \in \mathcal{M} - \{m\}} \left\{ \frac{p^f(k)}{\tau_{mk}} \right\}. \tag{2.1.3}$$

To the extent that large farmers in India may hire migrant labor, my estimates of gains in income and productivity should be thought of as lower bounds that ignore adjustment on the labor market.

The outside option for the intermediary is zero. If an intermediary m purchases q units of output from a farmer, he will sell this on the local retail market and receive a price p_m^r for each unit of output. A local retail market is the Indian administrative district in which the market is located. Each market is small compared to the local retail market and therefore the intermediary takes the retail price as given.[3]

If δ denotes the bargaining weight for the farmer, the Nash bargaining outcome is the solution to:

$$\max_{\lambda} \left(\lambda - \underline{p}\left(m\right) q_m\right)^{\delta} \left(p_m^r q_m - \lambda\right)^{1-\delta} \tag{2.1.4}$$

Here, λ is the farmer's income. Since $p_m^r q_m$ is the total surplus, $p_m^r q_m - \lambda$ is the income of the intermediary trader from this transaction. We should note that all farmers trading at a particular market m receive the same price because what matters is the distance of the next best alternative from a particular market and not the distance to that farmer's farm. Therefore, I omitted the subscript f in the bargaining problem.

Equilibrium

To solve for the spatial equilibrium, I employ the Nash-in-Nash (à la Horn and Wolinsky, 1988) solution concept—price negotiated at a market will be the Nash bargaining solution given that farmer-trader pairs in all other markets reach an agreement. The equilibrium for a region S is defined as follows:

Definition 1. Given parameters $\{\alpha_k\}_{k=1}^{K}, \nu, \gamma$ and δ, endowments $\left\{h_f, l_f, \tilde{A}_f\right\}_{f\in\mathcal{F}}$, transportation cost function τ, and retail and intermediate input prices $\{p_m^r\}_{m\in\mathcal{M}}$ and $\left\{w^k\right\}_{k=1}^{K}$ respectively, an equilibrium is a set of farmer prices $p^f\left(.\right)$, intermediate input choices $\left\{x_f^k\right\}_{k\in\mathcal{K},f\in\mathcal{F}}$ and the optimal market choice at each farm, denoted by $\left\{\mu\left(f\right)\right\}_{f\in\mathcal{F}}$, such that:

[3]For the main analysis, I read local retail prices from the data. In Section 2.3.4, I will endogenize retail prices and let them adjust in response to changes in local supply.

1. Farmers choose intermediate inputs to maximize their profits, to solve:

$$\max_{\left\{x_f^k\right\}_k} p^f\left(\mu\left(f\right)\right)\tilde{A}_f\left(h_f^{\gamma}l_f^{\nu}\prod_{k=1}^{K}\left(x_f^k\right)^{\alpha_k}\right)-\sum_{k\in K}w^k x_f^k, \tag{2.1.5}$$

2. Farmers optimally choose a market to sell their output, according to:

$$\mu\left(f\right)=\arg\max_{m\in\mathcal{M}}p^f\left(m\right)q_{fm}=\arg\max_{m\in\mathcal{M}}\left\{\frac{p^f\left(m\right)y_f}{\tau_{fm}}\right\}, \tag{2.1.6}$$

3. Farmers and intermediaries in any market Nash bargain over the total value of the farmer's output assuming that bargaining in all other markets have reached an agreement.

2.1.2 Solving the Model

We can now solve the model backwards from the timeline in figure 2.1. Conditional on the retail prices that intermediaries get, we first solve the bargaining problem between the farmers and the intermediaries.

Solution of the Bargaining Problem

The solution of the bargaining problem (2.1.4) gives us:

$$p^f\left(m\right)=\left(1-\delta\right)\underline{p}\left(m\right)+\delta p_m^r \tag{2.1.7}$$

where, $p^f\left(m\right)=\frac{\lambda}{q_m}$.

We will have one such equation for every market and all of them are interrelated through the outside option, $\underline{p}(m)$. Equilibrium farmer prices is then just the fixed point for the system of M equations in M endogenous variables defined in (2.1.7).[4]

Theorem 1. *Given retail prices* $\{p_m^r\}_{m\in\mathcal{M}}$*, there exists a unique fixed point* $p^{f*}(.) : \mathcal{M} \to \Re_+$*, that solves the system of equations given in (2.1.7).*

The proof can be found in Appendix 2.A. Intuitively, I can show that $p^f(.)$ is a contraction as $\delta \in (0, 1)$ and $\tau \geq 1$, and therefore there is a unique solution.

Once we know the equilibrium farmer price in any market $m \in \mathcal{M}$, denoted by $p^{f*}(m)$, the optimal market choice is given by (2.1.6). Subsequently, the input choices would be the result of a standard profit maximization problem (2.1.5), given the production function and intermediate input prices.

Lemma 1. *Given retail prices, removing restrictions on farmers from trading across regions will improve the prices that they receive.*

The formal proof of Lemma 1 is relegated to Appendix 2.B; the main intuition is as follows. In the system of equations (2.1.7) that determine equilibrium prices, the only variable that changes in the counterfactual is the set of outside available options to the farmer. In particular, $\mathcal{M}$ will include all markets in India as opposed to markets within the region (or that state). As the local retail prices are held fixed and we are taking the maximum over all outside options, the prices they receive in equilibrium cannot be strictly lower than before.

Lemma 1 guarantees that as long as retail prices do not change, no farmer will be worse off as a result of removal of trade restrictions. When farmers are allowed to trade across states, the degree of spatial competition between intermediary markets on both sides of the

[4]In the quantitative section, I will assume the problem to be independent across crops and time periods. Therefore, the problem with all crops and multiple periods involves finding a fixed point of a system of $M \times C \times T$ equations in $M \times C \times T$ variables, where T is the number of time periods.

border increases resulting in an increase in farmer prices in markets on both sides of the border. The only way farmers can lose in this setup is if due to an increase in supply, the retail prices fall enough to lower the total surplus and thus the prices farmers receive in markets.

2.1.3 Demand in Retail Markets

To complete the model, I define the demand side here. Each region above is further divided into subregions, indexed by d. Each subregion corresponds to an Indian administrative district and as discussed in Section 2.1.1 is the local retail market. The aggregate demand function for a subregion is given by

$$Q^d = Q\left(p_d^r\right). \tag{2.1.8}$$

Let η be the price elasticity of demand. All intermediaries m located in subregion d sell in the same subregion and hence, $p_m^r = p_d^r \; \forall m \in d$.[5] Finally, for retail markets to clear, total quantity demanded in a subregion d must equal total quantity sold in that subregion

$$\sum_{f \in \{\mu^{-1}(m) | m \in d\}} y_f = Q^d. \tag{2.1.9}$$

2.2 Estimating the Model

In order to take the model to data, I need to expand it to include multiple crops and time periods. Crops and time periods are assumed to be independent and therefore, in the quantitative section we will have $M \times C \times T$ equations of the form (2.1.7), where C is the number of crops and T is the number of time periods. To quantify the gains from increased

[5]This is an assumption. Equivalently, retail markets can be assumed to be perfectly integrated and then retail prices would differ across regions only due to transport costs. In the latter case, the intermediary will be indifferent between selling in different retail markets. Table 1.1 provides support for this by showing low variance of retail prices across different regions that produce a certain crop.

spatial competition due to removal of restrictions on farmers' trading across state borders, I structurally estimate the model presented in the previous section. First, I parametrize τ and then estimate δ and the parameters in the trade cost function jointly to match the spatial distribution of farmer prices. I conduct a sensitivity analysis to understand what features of the data inform the parameters. I set the production function parameters to match to land and labor shares in my data and I take demand elasticities from the literature.

2.2.1 Estimating τ and δ

I parametrize trade costs as a function of distance, and I estimate the parameters using data on the location of markets. I define the trade cost between two markets m and k, for crop c , and at time t as:

$$\tau_{mkct} = \begin{cases} 1 & m = k \\ 1 + A \cdot d_{mk} + \epsilon_{mct} & m \neq k \end{cases}, \tag{2.2.1}$$

$$\log \epsilon_{mct} \sim N(0, \sigma),$$

where, d_{mk} is the geodesic distance[6] between locations m and k and A is a scale parameter. The shock term, ϵ_{mct}, represents origin market-crop specific costs like broken roads, availability of a truck, or a strike among intermediaries, which are not observable to the econometrician but are known to the farmers. τ_{mkct} is set to infinity if markets m and k lie in different states to incorporate trade restrictions.

If I denoted crops by c and time period by t, then the key equilibrium equation of the model (2.1.7) becomes

$$p^f(m, c, t) = (1 - \delta) \max_{k \in \mathcal{M} - \{m\}} \left\{ \frac{p^f(k, c, t)}{1 + A \cdot d_{mk} + \epsilon_{mct}} \right\} + \delta p^r_{mct}, \tag{2.2.2}$$

[6]The shortest distance along the earth's surface.

From equation (2.2.2), we can see that prices in the different markets are interrelated. Therefore, the marginal likelihood function of price at one market is not independent of prices in the other markets. Hence, estimation cannot be carried out via a standard maximum likelihood or a method of moments procedure. To address this, I use a method of simulated moments procedure to estimate (δ, σ, A). In particular, I choose (δ, σ, A) to minimize the distance between moments of the data and their simulated counterparts.

Formally, the method of simulated moments chooses parameters $\Theta = (\delta, \sigma, A)$ to minimize the distance between simulated and data moments, solving:

$$\hat{\Theta} = \arg\min \left(\Psi^d - \Psi^s(\Theta)\right) W \left(\Psi^d - \Psi^s(\Theta)\right)',$$

where W is the optimal weighting matrix—the inverse of the variance-covariance matrix of the moments; Ψ^d are the data moments; and $\Psi^s(\Theta)$ are the simulated moments. I compute robust standard errors clustered at the state-season level and use numerical methods to find the required gradients.

Choice of Moments and Estimation

To identify moments that will inform the estimation, I choose an auxiliary regression model that mimics the equilibrium equation (2.2.2) as closely as possible.[7]

$$\ln p^f_{mct} = \tilde{\beta}_0 + \tilde{\beta}_1 \left(\max_{k \in \mathcal{M} - \{m\}} \ln p^f_{kct} \right) + \tilde{\beta}_2 \log d_{mk^*} + \tilde{\beta}_3 \ln p^r_{mct} + \gamma_c + \xi_{mct}, \tag{2.2.3}$$

Here, p^f_{mct} and p^r_{mct} are the prices farmers get and the retail prices in market m for crop c at time t respectively. d_{mk^*} is the distance from market m to market k^*, which has

[7]Following Gourieroux et al. (1993), identification of the auxiliary regression is not necessary for consistent estimation of the structural parameters. It just needs to capture key moments of the data that can inform identification of the structural parameters. One reason not to use only the border regression (1.4.4) as an auxiliary model is that it does not capture the relationship between farmer prices, distances to other markets, and retail prices. These moments are important for identification of the model parameters.

the highest price in the neighborhood, and γ_c is a crop fixed effect. I match five moments from the auxiliary regression: $\tilde{\beta}_0, \tilde{\beta}_1, \tilde{\beta}_2, \tilde{\beta}_3$, and the Mean Squared Error of the auxiliary regression. In addition, I also match the distribution of farmer prices using the mean and the variance of farmer prices, and the fraction of data in different parts of the distribution. Let us denote the k-th percentile of the distribution of farmer prices in the actual data by pc_k, and simulated farmer prices by $\tilde{p}^f_{mct}$. Then the moment condition is defined as:

$$\psi(\Theta) = \mathbb{1}\left[\tilde{p}^f_{mct} \in [pc_k, pc_{k+\zeta}]\right] - \zeta,$$

where ζ is the chosen bin size. I use four bins: 5th–25th, 25th–50th, 50th–75th and the 75th–95th percentiles.

While in principle all parameters are jointly estimated to match the distribution of prices, some moments are more important for identifying some parameters. Intuitively, the Nash-bargaining parameter, δ, is identified from the correlation between retail and farmer prices, $\tilde{\beta}_3$, and that between farmer price in a market and the maximum price in the neighborhood, $\tilde{\beta}_1$. Given δ, the scale parameter on distance, A, adjusts to match the level of farmer prices and is therefore identified of, $\tilde{\beta}_3 - \tilde{\beta}_1$, and the mean of farmer prices. Finally, the variance of the shock term, σ, is identified from the residual variance in prices captured by the MSE of the auxiliary regression and the distribution of prices. However, because the shock term enters inside the max function in the denominator, it is the smallest shocks that matter for identification. As the log normal distribution is bounded below by zero, this parameter is not estimated very precisely and this is reflected by the standard errors. Finally, $\tilde{\beta}_0$ reflects the scale in the auxiliary regression and therefore informs all parameters in the model.

To carry out the estimation, I set up the data in the following manner. First, because retail prices are only available at the level of an Indian administrative district, I assign the same retail price to all markets in a given district-time period. Second, data on retail prices

are available only between 2005 and 2011, barring 2008[8]; therefore I focus on the data in these years for estimation. Third, to reduce the computational burden, I use the median prices at the crop-market-season level for estimation.[9] There is little variation in retail prices of a district or farmer prices at a market within a season (see Table 1.1) and thus there is no consequential loss of information. However, it helps me to create a balanced panel as retail prices in certain months are erroneously reported or are missing. Fourth, I exclude the states for a crop where it has never been produced. To infer this, I use the information in the crop calendar and production statistics reported by the Ministry of Agriculture of India in its annual publication, Agriculture Statistics at a Glance, for the years in my sample. I provide the details in Appendix 2.E. This leaves me with data on 8 crops in 2978 markets for 15 time periods.

Table 2.1 reports the estimated parameters with robust standard errors clustered at the state-time level. The bargaining power of the farmer, δ, is estimated to be 0.69. The estimated parameters of the trade costs correspond to a deterministic 16% transport cost over a distance of 100 km. In Section 2.4, I also consider other functional forms for trade cost such as a power function, a quadratic, and the exponential function. In general, the data admit any functional form for trade costs that is concave in distance.

Sensitivity of Parameters

In this subsection, I compute four different statistics to understand the sensitivity of the parameters to different aspects of the data. First, following Andrews et al. (2017), I compute the semi-elasticity of the probability limit of each estimator with respect to the moments scaled by their standard deviations. Second and third, following Honore et al. (2017), I compute the elasticity of each estimator with respect to the diagonal elements of the weighting matrix and the elasticity of the variance co-variance matrix of the estimator

[8]2008 was a drought year. Therefore, I am not worried about sample selection. If anything, the data would have been outliers.

[9]I use three crop seasons–fall, spring, and summer.

with respect to the variance co-variance matrix of the moments. Fourth, following Cooper and Liu (2016), I compute the elasticity of the moments with respect to the parameters. All the estimates are reported in Tables 2.2 and 2.3 and the results are also shown in Figures 2.2–2.5.

The first statistic helps us understand to which moments the estimators are most sensitive. In particular, if a moment is not exactly equal to 0 but equal to some positive number ρ, i.e. $E(\psi(x_i, \theta_0)) = \rho$, then I compute the derivative of the probability limit of $\hat{\theta}$ with respect to ρ. I scale the moments by their respective standard deviations. The results are presented in Figure 2.2. A general takeaway is that the coefficients $\tilde{\beta}_1, \tilde{\beta}_2$ and $\tilde{\beta}_3$ aid in identification of all the parameters.

The second statistic helps us ask what would happen to the estimate of a parameter if I gave a little more weight to a moment, holding all other weights in the weighting matrix fixed. This is calculated as the derivative of the probability limit of the estimator with respect to the kkth element of the weighting matrix. The results in Table 2.3 (and Figure 2.3) are qualitatively very similar to the previous statistic. An additional insight here is that σ is sensitive to the weight on almost all moments and thus the estimate of σ is very sensitive to the choice of the weighting matrix.

The third statistic corresponds to the thought experiment: How much precision would we lose if we had a different data set in which a certain moment is subject to a little additional noise? This differs critically from the previous two statistics in that we allow the weighting matrix to change optimally as more noise is added. In my case, the results presented in Figure 2.4 provide useful insurance against noise in the retail prices. Because data on retail prices is aggregated at the level of the district, it is measured with noise. Panel 1 of Figure 2.4 tells us that δ is not much affected by such noise in the retail prices. In fact, none of the parameters are.

The fourth statistic is a direct measure of sensitivity of the moments with respect to model parameters. In contrast to the other statistics, this does not take into account optimal

weighting at all and asks the opposite question: If data were generated from the same model with different parameter values, which moments are likely to change? It is again comforting to note that changing δ does not change $\tilde{\beta}_3$, the coefficient on log retail prices in the auxiliary regression, as much. What stands out in the results presented in Figure 2.5 is that the bottom of the price distribution is very sensitive to δ. This makes sense because the lowest farmer prices are likely to occur in very isolated markets such that distances to other markets are very large. This would make the simulated prices in the bottom tail very sensitive to the scale parameter on distance, shocks to transport costs, and retail prices. This has little implication for estimation per se, given that the optimal weight matrix will ensure that only relevant information is extracted from this moment.

Goodness of Fit

To assess how well the model fits the data within the sample, I regress log of actual prices to log of prices generated from the model. A linear regression without an intercept term has a coefficient of 1.065 with an R-square of 0.59. The fit is shown in Figure 2.6. Figure 2.7 shows a density plot of model-generated prices and actual prices and Table 2.4 presents different moments of the price distribution. We can notice that the model does a good job of capturing the distribution of prices. In Figure 2.7, the humps in the density of actual data are because of different crop specific means. This is not exactly replicated in the model generated data as I do not have any crop-specific parameters in the model. However, I take it to be a success of the model that I match the distribution of prices well even without incorporating crop-specific parameters.

2.2.2 Calibrating the production function and demand elasticities

To calibrate $\gamma + \nu$, the share of land and labor, I use the first order conditions from the farmers' profit maximization problem. Because of fixed factors, I need data on input costs and farm profits, which I obtain from the Situation Assessment Survey of Agricultural

Households in India 2012–13.[10] I find that $\gamma + \nu = 0.57$. In section (2.4), I present results for other values of $\gamma + \nu$.

I take estimates of the price elasticity of demand, η, in India from Deaton (1997). For rice $\eta = -1.19$, wheat $\eta = -1.29$, and other crops $\eta = -3.31$.

2.3 Counterfactual Analysis

The model I have laid out allows me to compute changes in farmer prices, production, and revenues as a result of more competition between intermediary markets. In this section, I use the removal of the restriction on farmers' selling in markets of other states as a convenient experiment to increase spatial competition between intermediary markets. The key mechanism is the following. Consider two markets on the Madhya Pradesh-Maharashtra border marked in Figure 2.8. The farmer living close to the market in Madhya Pradesh gets low prices because competition is low in his local region as the next market is quite far away. Further, this farmer cannot cross the border and sell in Maharashtra. When we remove restrictions on farmers' selling across state borders, these two markets start competing for the farmer's output. This increases prices in both these markets. The farmer, wherever he chooses to sell, is better off. Moreover, these two markets are the "threat points" to other markets nearby. As the prices increase in the border markets, they also increase in other markets and this ripple effect increases prices even in the interior of the states. This is the direct benefit to farmers via an increase in prices. In response to an increase in output prices, farmers also adjust their use of intermediate inputs, and their incomes improve further as a result of increased production. Finally, as a result of changes in production and changes in the market-site where farmers choose to sell, retail prices may adjust and feed back into the prices farmers get.

[10] For details, see Appendix 2.C.1.

In the model, this removal of trade restrictions on farmers is equivalent to changing the trade cost between two locations in different states from infinity to what it would be under equation (2.2.1). To study the role of the different channels mentioned in the previous paragraph, I conduct four exercises relaxing one constraint at a time. In the first exercise, I do not allow retail prices to adjust. I keep all decisions of the farmer fixed except their threat to sell at a market outside the state (Section 2.3.1). In addition to allowing farmers to offer a different threat, in the second exercise, I also allow them to choose a different market to trade in (Section 2.3.2). In a third exercise, I let farmers also adjust their intermediate input choices to study the change in their crop output (Section 2.3.3). In the final exercise, I let retail prices at the district level adjust in response to changes in supply to understand medium-run implications for farmer prices and output (Section 2.3.4).

I also use the model to understand the implications of reductions in transport costs and examine the quantitative magnitudes of the interaction of such public investment with trade liberalization for farmers (Section 2.3.5).

Note that all these exercises require data on crop choices of farmers. For this, I use data from the Situation Assessment Survey of Agricultural Households 2012–13, and I match them with data on retail prices from 2011–12. These are the two closest years for which I have detailed data on crop choice and retail prices.[11] Since regional cropping patterns in India have been stable over the last decade, this is not a major concern for interpreting my results.

2.3.1 Changing Farmers' Threat Point

In this exercise, I use my model and estimated parameters to study what happens to farmer prices when we remove the restriction on farmers' selling in markets of other states but keep all other decisions fixed. We can think of this as a hypothetical exercise where farmers go

[11]The Central Statistical Organization of India should soon release data on retail prices from 2012–13, at which point I will update this analysis.

to the same market where they went when restrictions were in place. However, they can threaten intermediaries with going to a market in another state when bargaining with them.

The changes in farmer prices at the markets are presented in Figures 2.9 and 2.10. Figure 2.9 plots the relationship between the change in prices in each market in the counterfactual scenario and the change in competition for each market, where competition is defined as in the reduced-form regressions using (1.4.1). The bottom of Figure 2.9 plots the empirical distribution of the change in competition as a result of removal of the interstate trade restriction. There is a monotonic relationship between price increase and increase in competition. The median increase in competition is 2.9, which means that half of all markets experience at least a 2.8% increase in price.

Figure 2.10 plots the estimated relationship between gain in price at a market and the distance to the nearest border. The bottom of the figure plots the empirical distribution of the location of markets with respect to state borders. We can see that the biggest gains are in markets closest to state borders and the gains decline the farther the markets are. However, due to the ripple effect, even markets in the interior of some states realize an increase in competition and therefore an increase in price.

To understand what this means for farmers living in different locations, I divide India into grid cells of 5 arc minute by 5 arc minute.[12] I use data on land use to exclude cells that are uncultivated like deserts, large water bodies, mines, marshes etc. The remaining cells are hypothetical farms in India for which I will compute changes in prices and production as a result of increased spatial competition. In order to determine the price the farm gets, I also need to know what crops this farm grows. I use data from the Situation Assessment Survey of Agricultural Households 2012–13 to estimate the district level land shares for each crop in each crop season, i.e. the crop choice of the district. I then assign the crop choice of the district to each cell that lies in a majority in that district. I compute the price received

[12]This corresponds to a cell having land area of approximately of 10 km by 10 km. Although, this is much larger than a typical farm, it is not a major concern because I consider a production function which is homogeneous of degree 1, and cropping patterns in India are similar within 10 km distances. Limitations to computing power does not allow for consideration of finer grid cells.

by a farm as the Laspeyres price index. I use yield estimates from the Situation Assessment Survey of Agricultural Households to estimate the quantities produced on each farm.

The results for the fall season are presented in Figure 2.11; Figure 2.12 presents the results for the spring season. The average increase in farmer prices is 2.74% and 2.65% in the fall and spring seasons respectively. The average increase in prices for the top 25% of farmers with the largest gains is 7.36% in the fall season and 5.84% in the spring season. The maximum increase in prices is 9% and 13% in the fall and the spring season respectively. These increases in farmer prices are purely due to an increase in spatial competition because we do not allow farmers to change any other decision.

2.3.2 Allow Farmers To Sell At A Different Market

In the second exercise, in addition to providing a greater threat to intermediaries while bargaining, I also allow farmers to choose a different market to trade in. Their new market choice could be another market in the same state or a neighboring state. As we will see, the largest increases in farmer prices will come from changing market choice.

Figures 2.13 and 2.14 present the changes in prices that farmers get as a results of increased competition because farmers can sell in markets of other states. We can see that farmers are significantly better off in many regions of India. The price increases are generally larger near state borders and lower in the interior of the states. The average increase in prices is 12.6% in the fall and 10.7% in the spring. The distribution of price increases is, however, skewed. The top 50% of the farmers who experience an increase in prices get 23.37% and 20.04% higher prices on an average in the fall and the spring. Consistent with Lemma 1, no farmer loses.

In the fall season, we notice bigger gains in central India (in particular, in the state of Madhya Pradesh). To explore the mechanism further, I will zoom into the Madhya Pradesh and Maharashtra border.[13] This region mainly grows soybean in the fall. Figures 2.15 and

[13]Refer to Figure 1.A.1 for state names.

2.16 plot the log of soybean prices that farmers get in this region before and after the policy change respectively. We can see that price levels before the policy change are higher in Maharashtra than in Madhya Pradesh. This is so because near the state border, the density of markets is higher in Maharashtra.

After the policy change, markets on the border compete and the price in them increases. Through a ripple effect of one market affecting the other, prices increase even in the interior of the states. As can be seen in Figure 2.16, prices equalize on both sides of the state border. Of course, because prices were lower in Madhya Pradesh to begin with, the percentage increase there is higher (see Figure 2.17).

Prices increase by nearly 60% in Madhya Pradesh near its border with Maharashtra for the following reason. Farmers in Madhya Pradesh close to the border are generally far away from markets in their own state. Therefore, with border restrictions, they pay a substantial transport costs to trade in these markets. When states open up to trade, these farmers get access to markets in Maharashtra, which are much closer. This saves farmers' transport costs, leading to even larger gains.

2.3.3 Letting Farmers Reoptimize the Use of Intermediate Inputs

The Economic Survey of India 2015–16 documents a huge variation in the use of intermediate inputs. Farmers in more productive regions are usually more input-intensive. I highlight a possible reason behind this pattern in this exercise. I allow farmers to change their use of intermediate inputs in response to the change in prices we observe in the counterfactual exercises of the previous section. From the model, we can show that for farm f, the new output $\left(y'_f\right)$ to the old output (y_f) ratio is a function of the ratio of the new to old prices. In particular[14]

$$\frac{y'_f}{y_f} = \left(\frac{p^{f\prime}}{p^f}\right)^{\frac{1-(\gamma+\nu)}{\gamma+\nu}},$$

[14]For a formal proof see Appendix 2.C.

here γ and ν are the share of land and labor in production, and primes represent the new output and prices. Since $\gamma + \nu < 1$, the percentage increase in production is smaller than percentage increases in prices. The results presented in Figures 2.18 and 2.19 use the same crop choices as before. To measure quantity produced at each farm, I compute the Laspeyres quantity index. The average increase in production is 9.1% and 7.75% in the fall and the spring respectively. The average increase in production for the top 50% of the farmers is 16.9% and 14.46% in the fall and the spring respectively.

Farmers in less competitive regions use fewer intermediate inputs as a result of low output prices. Therefore, the benefit to farmers from the increase in competition comes from two sources—an increase in prices and an increase in quantities. This leads to an increase in the total revenue of farmers. The aggregate increase in the value of production is 17.3% in the fall season and 20.9% in the spring season.

2.3.4 Incorporating a Demand Side Response

The gains described above could be overestimates because an increase in production could depress retail prices, which would feed back into farmer prices. To understand this channel better, I use estimates of demand elasticities from Deaton (1997) and an iterative algorithm[15] to find the new equilibrium after border restrictions to trade are removed.

Figures 2.20 and 2.21 present the changes in prices incorporating adjustment of retail prices for the fall and the spring respectively.[16] Although there are regions in the interior of some states where farmers lose, most farmers gain. The average increase in farmer prices is 9.62% in the fall season and 9.56% in the spring season. The average increase in prices for the farmers with above-median changes is about 20%. For the farmers who lose, the average decline in prices is about 10%.

[15] See Appendix 2.D for details.

[16] Compare to Figures 2.13 and 2.14, that present changes in farmer prices when we do not allow retail prices to adjust.

The average price increases here are only marginally lower than in the exercise when we kept retail prices fixed despite a fall in prices of some farmers. The main reason is the following. When I remove the restriction on farmers to trade across state borders, supply, on average, increases in regions with high initial retail prices. The supply increases because these regions attract more farmers, who in turn, produce even more as they get paid more. This adjusts the retail prices downwards, and farmers from such regions lose on average. Since supply reduces in regions with low initial retail price, as farmers sell in other regions, there is an upward adjustment of retail prices on average. In equilibrium, the net effect on farmers from these regions is ambiguous. Earlier, when retail prices were held fixed, in response to removal of border restrictions these farmers incurred transport costs to get access to higher prices (Section 2.3.2). Now, these farmers can get higher prices in their local neighborhood. These prices are not as high as in the counterfactual before, but the farmers save their transport costs. Hence, in principle, some farmers can gain more as a result of the adjustment of retail prices.

The results from the exercises above are summarized in Table 2.5. I have shown that increased competition due to the removal of trade restrictions substantially increases the prices that farmers receive. The average farmer gets about a 10% higher price for his output. The gains are skewed as the average increase in price for the top 50% of farmers is about 19%. Approximately 6% of the increase is purely due to competition, while a majority of it comes because farmers also reoptimize their market choice and save on transport costs. Increase in price is only part of the gain that farmers get as this is reinforced by an increase in production as farmers adjust their use of intermediate inputs. As a result, the total value of production increases by about 19.90%. Finally, I also allay concerns that an increase in output might reduce retail prices so much that farmers in most regions end up losing.

2.3.5 Effect of a Reduction in Transport Costs

The exercises that I have presented heretofore reflect the direct gains from competition via increased prices for farmers. Spatial competition can also interact with other development policies in interesting ways, and I can use my model to shed some light on this aspect. Poor roads and lack of access to good vehicles increase transportation costs in India, and they are especially high for farmers (Kapur and Krishnamurthy, 2014). However, over the last 15 years, the central government has invested around $40 billion in the construction of rural roads (Chatterjee and Kapur, 2017). Between 2001 and 2015, approximately 185,000 habitations have been connected by all-weather roads. The aim in this section is to show that farmers can realize substantially greater benefits of such public investments in the presence of greater competition among intermediaries.

I consider reducing A by 40%, and for illustrative purposes, I present results only for the fall season. Figure 2.22 presents the increase in farmer prices due to reduction in transport costs in the fall season in the presence of border restrictions. Figure 2.23 presents the additional increase in farmer prices due to more spatial competition in the absence of border restrictions to trade. I computed this by calculating the increase in farmer prices due to simultaneously reducing the transport costs and removing the border restrictions, and then subtracting the increase in prices just due to the reduction in transport costs (computed in Figure 2.22) and the increase just due to the elimination of border restrictions to trade (computed in Figure 2.13). Farmer prices increase by an additional 8% on average, because of increased spatial competition.

We can draw two lessons. First, a reduction in transport costs, even in the presence of border restrictions, leads to increases in farmer prices in the interior of states as farmers can travel farther, and this changes their threat point in the bargaining problem. Second, the additional gains from the increase in spatial competition occur in those regions where transport cost reduction was not effective in improving farmer prices, e.g., in the states of

Odisha, Madhya Pradesh, and Rajasthan.[17] Therefore, transport infrastructure polices can be thought of as complements to policies that increase competition.

2.4 Robustness and Sensitivity of Quantitative Results

In this section, I argue that my results are robust to two potential misspecifications in the setup above. First, it could be possible that farmers in villages do not sell in markets but to village intermediaries who aggregate from all farmers in a village and in turn sell to market intermediaries. Second, there are costs associated with activities of the intermediaries that are missing in the model. Additionally, I show that my quantitative estimates of increases in farmer prices and output are not very sensitive to the way I assign crop choices to farm-cells, the choice of $\gamma + \nu$, and the functional form of trade cost.

It could be possible that farmers in some regions do not travel to the markets themselves but sell in villages to village intermediaries who in turn sell in the markets to market intermediaries. In this case, all the quantitative results remain unchanged if we allow for a constant pass-through from the village intermediary to the actual farmer. In particular, the village intermediaries perform the same role that farmers play in the current model and pass a share of their earnings to the farmers. Therefore, p^f is now the price that the village intermediaries receive as they engage in Nash-bargaining (defined in 2.1.4) with the market intermediaries. Let us further suppose that these village intermediaries pass a fixed share to the actual farmers, κ^f. I can allow this pass-through rate, κ^f, from the village intermediaries to the farmers to vary by each farm. However, we have to assume that it does not change in the counterfactual scenarios. Since all my results are computed in changes, they remain the same.

In my model, the intermediaries buy from farmers and sell in retail markets and have no associated costs. In the presence of costs to intermediaries, the actual prices that

[17] For state names, see Figure 1.A.1

intermediaries get must be lower than the retail prices measured in my data. δ is insensitive to noise in retail prices following the arguments I presented in Section 2.2.1. Therefore, to match the level of farmer prices, A would be lower than is estimated. This would make trade costs lower and the competition channel more important. Thus my results can be viewed as a lower bound on the actual gains.

To show that my results are not sensitive to the allocation of crop choice to farm-cells, in Figures 2.F.1 and 2.F.2 I present the counterfactual increases in prices that farmers receive assuming that farms grow the main crop (by crop sown area) of the district. The results look very similar to Figures 2.13 and 2.14, because in the data, the main crop of the district occupies on average 80% of the cropped area of the district. The mean increase in farmer prices is about 11%, and the average increase in prices for the 50% of farmers with the largest gains is about 20%.

In Table 2.6, I present estimates of increases in production and revenues for higher values of $\gamma + \nu$. The increase in output is decreasing in $\gamma + \nu$. For $\gamma + \nu$ between 0.57 and 0.63, the average increase in output of farmers is between 6% and 9%. The average increase in output for 50% of the farmers with the largest gains is between 12% and 19%.

As far as different functional forms for trade cost is concerned, in Table 2.7, I present results if trade cost is assumed, following Redding and Sturm (2008), to be $\tau_{mk} = (1 + A \cdot d_{mk})^{\frac{1}{3}} + \epsilon_{mct}$. Since the trade costs in this case increases at a slower rate with distance than with the linear trade costs, the associated price increases are larger. The average increase in farmer prices is about 20%, and the average increase for the 50% of farmers with the largest gains is about 35%.

The data seem to be rejecting functional forms that are convex in distance. I estimated a quadratic function for trade costs: $\tau_{mk} = 1 + A \cdot d_{mk} + B \cdot d_{mk}^2 + \epsilon_{mct}$. The estimation procedure resulted in $B = 0$, rejecting this functional form in favor of the linear functional form. Similarly, the estimation procedure results in $A = 0$ for an exponential functional form $\tau_{mk} = \exp(A \cdot d_{mk}) + \epsilon_{mct}$.

2.5 Concluding Remarks

The median annual farmer income in India is USD 365 (Ministry of Finance, Government of India 2016). Using unique data on the location of intermediary markets and farmer prices, I have shown that spatial competition between intermediaries is an important determinant of the prices that farmers in India get for what they produce. High transport costs and policies that limit the ability of farmers to arbitrage between different intermediaries cause the market power of intermediaries to vary in space. Farmers who live in regions where there is more competition between intermediaries receive higher prices for their output. A one standard deviation increase in competition increases farmer prices by 6.4%.

I also show that increasing competition in one region spreads through the rural economy via a ripple effect. In particular, I consider the effect of reforming the laws of Indian states that prohibit farmers from selling in markets of other states. I show that such a policy not only increases incomes of farmers who live close to state borders but also of those living in the interior. I also show that increased competition further increases farmer incomes because of the increased output of farmers as they optimize the use of intermediate inputs. I find that the average increase in prices and output for the farmers with above-median gains is about 21% and 15% respectively. Further, the value of national crop output increases by about 18%. Moreover, the results indicate that isolated studies in agricultural markets can indeed find varying estimates of market power of intermediaries because they are partly driven by spatial competition.

Finally, I document that there are large complementarities between market integration and spatial competition. Adding to the large literature on how reduction in trade costs affects market integration and gains from trade (summarized in Donaldson, 2015a), I show that welfare benefits of market integration and trade cost reduction are substantially more in presence of greater competition.

Although I have focused on one institutional feature in one particular developing country, I believe that these findings have broader importance. Similar institutions and policies

exist elsewhere. Monopsony powers of agricultural marketing boards or intermediaries is a common feature across developing countries (e.g. see Bates, 1981). Policies that restrict trade through intermediaries exist even outside agriculture. For example, until 2004 medium-sized firms in China were required by law to export via government approved intermediaries (Bai et al., 2017). The underlying mechanism of spatial competition will apply to all such environments and future work could study its importance in other settings.

References

Allen, Treb and David Atkin, "Volatility and the Gains from Trade," *NBER Working Paper*, 2016, (22276).

Anderson, James E. and Eric van Wincoop, "Gravity with Gravitas: A Solution to the Border Puzzle," *The American Economic Review*, 2003, *93* (1), 170–192.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro, "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *The Quarterly Journal of Economics*, 2017, *132* (4), 1553–1592.

Antras, Pol and Arnaud Costinot, "Intermediated Trade," *The Quarterly Journal of Economics*, August 2011, *126* (3), 1319–1374.

Atkin, David and Dave Donaldson, "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs," *NBER Working Paper*, 2015, (21439).

Bai, Xue, Kala Krishna, and Hong Ma, "How You Export Matters: Export Mode, Learning and Productivity in China," *Journal of International Economics*, 2017, *104*, 122–137.

Banerji, A. and J. V. Meenakshi, "Buyer Collusion and Efficiency of Government Intervention in Wheat Markets in Northern India: An Asymmetric Structural Auctions Analysis," *American Journal of Agricultural Economics*, 2004, *86* (1), 236–253.

Bardhan, Pranab, Dilip Mookherjee, and Masatoshi Tsumagari, "Middlemen Margins and Globalization," *American Economic Journal: Microeconomics*, 2013, *5* (4), 81–119.

Bates, Robert H, *Markets and States in Tropical Africa: The Political Basis of Agricultural Policies*, Berkeley: University of California Press, 1981.

Bergquist, Lauren Falcao, "Pass-through, Competition, and Entry in Agricultural Markets: Experimental Evidence from Kenya.," *UC Berkeley Mimeograph*, 2017.

Casaburi, Lorenzo and Tristan Reed, "Competition and Interlinkages in Agricultural Markets: An Experimental Approach," *Working Paper*, 2016.

Chand, Ramesh, "Development Policies and Agricultural Markets," *Economic and Political Weekly*, 2012, *47* (52).

Chatterjee, Shoumitro and Devesh Kapur, "Six Puzzles in Indian Agriculture," *India Policy Forum 2016–17*, 2017, *13*.

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

Cooper, Russell and Huacong Liu, "Money or Grit? Determinants of Mismatch by Race and Gender," *NBER Working Paper*, 2016, (22734).

Cosar, A Kerem, Paul Grieco, and Felix Tintelnot, "Borders, Geography, and Oligopoly: Evidence from the Wind Turbine Industry," *Review of Economics and Statistics*, 2015, *97* (3), 623–637.

Deaton, Angus, *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*, World Bank Publications, 1997.

Dell, Melissa, "The Persistent Effects of Peru's Mining Mita," *Econometrica*, 2010, *78* (6), 1863–1903.

Dillon, B. and C. Dambro, "How Competitive are Food Crop Markets in Sub-Saharan Africa? A Review of the Evidence," *Working Paper, Evans School of Public Policy and Governance, University of Washington*, 2016.

Donaldson, Dave, "The Gains from Market Integration," *The Annual Review of Economics*, 2015, *7* (1), 619–647.

_ , "Railroads of the Raj: Estimating the Impact of the Transportation Infratructure," *American Economic Review, forthcoming*, 2015.

_ **and Richard Hornbeck**, "Railroads and American Economic Growth: A Market Access Approach," *The Quarterly Journal of Economics*, 2016, *131* (2), 799–858.

Eaton, Jonathan and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, *70* (5), 1741–1779.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu, "Competition, Markups, and The Gains from International Trade," *The American Economic Review*, 2015, *105* (10), 3183–3221.

Engel, Charles and John H. Rogers, "How Wide Is the Border?," *The American Economic Review*, 1996, *86* (5), 1112–1125.

Fafchamps, M, E Gabre-Madhin et al., "Agricultural markets in Benin and Malawi.," *African Journal of Agricultural and Resource Economics*, 2006, *1* (1), 67–94.

Fafchamps, Marcel and Bart Minten, "Returns to social network capital among traders," *Oxford Economic Papers*, April 2002, *54* (2), 173–206.

_ , **Eleni Gabre-Madhin, and Bart Minten**, "Increasing returns and market efficiency in agricultural trade," *Journal of Development Economics*, 2005, *78* (2), 406–442.

Gourieroux, Christian, Alain Monfort, and Eric Renault, "Indirect Inference," *Journal of Applied Econometrics*, 1993, *8* (S1).

Great Britain and Royal Commission on Agriculture in India, *Report of the Royal Commission on Agriculture in India.*, London: H.M. Stationery Office, 1928.

Guha, Ramachandra, *India after Gandhi: The History of The World's Largest Democracy*, New Delhi: Harper Collins, 2007.

Harris, Barbara, *State and market: State Intervention in Agricultural Exchange in a Dry Region of Tamil Nadu, South India*, New Delhi: Concept Pub. Co., 1984.

Harris, Chauncy D., "The Market as a Factor in the Localization of Industry in the United States," *Annals of the Association of American Geographers*, 1954, *44* (4), 315–348.

Honore, Bo E., Thomas Jorgensen, and Aureo de Paula, "Sense of Sensitivity," *Working Paper*, 2017.

Horn, Henrick and Asher Wolinsky, "Bilateral Monopolies and Incentives for Merger," *The RAND Journal of Economics*, 1988, pp. 408–419.

Kapur, Devesh and Mekhala Krishnamurthy, "Market Towns and the Dynamics of India's Rural and Urban Transformations," *Working Paper, International Growth Center*, 2014, (F-35007-INC-1).

Krishna, Kala and Yelena Sheveleva, "Wheat or Strawberries? Intermediated Trade with Limited Contracting," *American Economic Journal: Microeconomics*, August 2017, *9* (3), 28–62.

Krishnamurthy, M., "Harda Mandi: Experiencing Change in an Agricultural Market in Central India (1980-2010)." Ph.D. thesis, UCL (University College London) October 2011.

Loecker, Jan De, Pinelopi K Goldberg, Amit K Khandelwal, and Nina Pavcnik, "Prices, Markups, and Trade Reform," *Econometrica*, 2016, *84* (2), 445–510.

McCallum, John, "National Borders Matter: Canada-US Regional Trade Patterns," *The American Economic Review*, 1995, *85* (3), 615–623.

Meenakshi, J.V. and A. Banerji, "The Unsupportable Support Price: An Analysis of Collusion and Government Intervention In paddy Auction Markets in North India," *Journal of Development Economics*, 2005, *76* (2), 377 – 403.

Melitz, Marc J. and Gianmarco Ottaviano, "Market Size, Trade, and Productivity," *The Review of Economic Studies*, 2008, *75* (1), 295–316.

Ministry of Agriculture, Government of India, "Agricultural Statistics at A Glance 2016-17," 2016.

Ministry of Finance, Government of India, "Economic Survey of India 2014–15," 2015.

_ , "Economic Survey of India 2015–16," 2016.

Mitra, Sandip, Dilip Mookherjee, Maximo Torero, and Sujata Visaria, "Asymmetric Information and Middleman Margins: An Experiment with Indian Potato Farmers," *The Review of Economics and Statistics, forthcoming*, 2017.

National Aeronautics And Space Administration And The National Geospatial Intelligence Agency, "Shuttle Radar Topography Mission 30 Arc Second Finished Data," 2000.

Redding, Stephen J and Daniel M Sturm, "The Costs of Remoteness: Evidence from German Division and Reunification," *The American Economic Review*, 2008, *98* (5), 1766–1797.

Willmott, C. J. and K. Matsuura, "Terrestrial Air Temperature and Precipitation: Monthly and Annual Time Series (1950 - 2014)," 2001.

Tables and Figures

TABLE 2.1: STRUCTURAL ESTIMATION RESULTS

PARAMETER	POINT ESTIMATE	CLUSTER ROBUST SE
δ	0.69	0.0025
σ	0.00039	0.0048
A	0.0016	0.0007
J-Statistic	30.00	

Notes: Standard Errors Clustered at State-Time Level

TABLE 2.2: SENSITIVITY MEASURES

	Andrews, Gentzkow & Shapiro (2017)			Cooper & Liu (2016)		
	δ	σ	A	δ	σ	A
β_0	-1.57	2.12e+05	114.87	327.31	4.42e-4	-4.85
Max P(β_1)	0.23	-32110	-16.77	-122.77	4.30e-4	2.20
Distance (β_2)	0.08	-10652	-5.91	72.16	1.31e-4	-1.04
Retail Price (β_3)	0.17	-21988	-12.64	29.90	7.20e-5	-0.41
Mean Sq Error	1.75e-2	-2924.4	-1.27	0.07	-1.98e-5	-0.01
Mean	1.55e-5	-33.02	-1.17e-3	-15.97	-1.96e-6	0.14
Variance	-6.24e-6	0.83	4.55e-4	-6.91	-9.97e-7	-0.011
5^{th}-25^{th} P'ctile	-2.20e-3	-334.66	0.15	-210.49	-0.03	7.19
25^{th}-50^{th} P'ctile	7.89e-2	-1478.7	-0.57	-2.60	1.14e-3	-0.004
50^{th}-75^{th} P'ctile	1.2e-2	-1872.5	-0.90	-40.07	1.83e-3	1.24
75^{th}-90^{th} P'ctile	8.79e-3	-1379.2	-0.64	-9.97	-1.38e-4	0.11

TABLE 2.3: SENSITIVITY MEASURES (HONORÉ, JØRGENSEN & DE PAULA 2017)

Elasticity of	Estimator wrt Diagonal of the Weighting Matrix			Variance of Estimator wrt Variance of Moments		
	δ	σ	A	δ	σ	A
β_0	37.09	-1.08E+06	-2426.7	0.45	1768.7	151.65
Max P(β_1)	-50.30	1.66E+06	3447.3	129.09	1230	236.05
Distance (β_2)	-8.33	2.42E+05	545.19	91.23	7.14	7.02
Retail Price (β_3)	-20.12	5.90E+05	1321.8	7.46	2.54	0.94
Mean Sq Error	-1.63e-3	40.28	0.10	1.37	1.34	0.41
Mean	-4.74	1.59E+05	328.67	1.59	40.85	4.71
Variance	-2.34	93086	177.29	0.075	27.41	1.80
5^{th}-25^{th} P'ctile	-0.45	12957	29.627	0.56	0.20	0.007
25^{th}-50^{th} P'ctile	-3.62	17806	169.4	97.46	13.80	20.11
50^{th}-75^{th} P'ctile	-0.44	-7490.6	12.92	43.80	41.19	13.38
75^{th}-90^{th} P'ctile	-0.36	12010	25.10	15.76	14.04	4.82

TABLE 2.4: GOODNESS OF FIT

Moment of Farmer Prices	Data	Model
Mean	2.18	2.31
Variance	0.13	0.14
25^{th} Percentile	1.90	2.04
Median	2.19	2.25
75^{th} Percentile	2.39	2.54

TABLE 2.5: SUMMARY OF COUNTERFACTUAL EXERCISES

Change in Farmer Prices (in %)	Median		Mean		Mean For the Top 50%		Mean For the Top 25%	
	Fall	Spring	Fall	Spring	Fall	Spring	Fall	Spring
#1: Change the treat point	1.75	2.50	2.74	2.65	4.97	4.50	7.36	5.84
#2: #1+Market Choice	8.35	5.82	12.60	10.72	23.37	20.04	32.67	29.69
#3: #2+Adjust Retail Price	5.98	6.21	9.62	9.56	20.32	19.49	31.70	28.74

TABLE 2.6: SENSITIVITY TO $\gamma+\upsilon$

	Mean Increase in Output (%)			
	All Farmers		Top 50%	
$\gamma+\upsilon$	Fall	Spring	Fall	Spring
0.57	9.10	7.75	16.10	14.46
0.60	7.91	6.73	14.68	12.55
0.63	7.00	5.96	12.99	11.11

TABLE 2.7: SENSITIVITY TO FUNCTIONAL FORM FOR TRADE COST

				Mean Increase in Prices (%)			
				All Farmers		Top 50%	
Functional Form	δ	σ	A	Fall	Spring	Fall	Spring
1+A.distance	0.69	0.00039	0.0016	12.6	10.72	23.37	20.04
$(1+A.distance)^{0.33}$	0.69	0.00058	0.0057	21.5	21.75	37.04	33.36

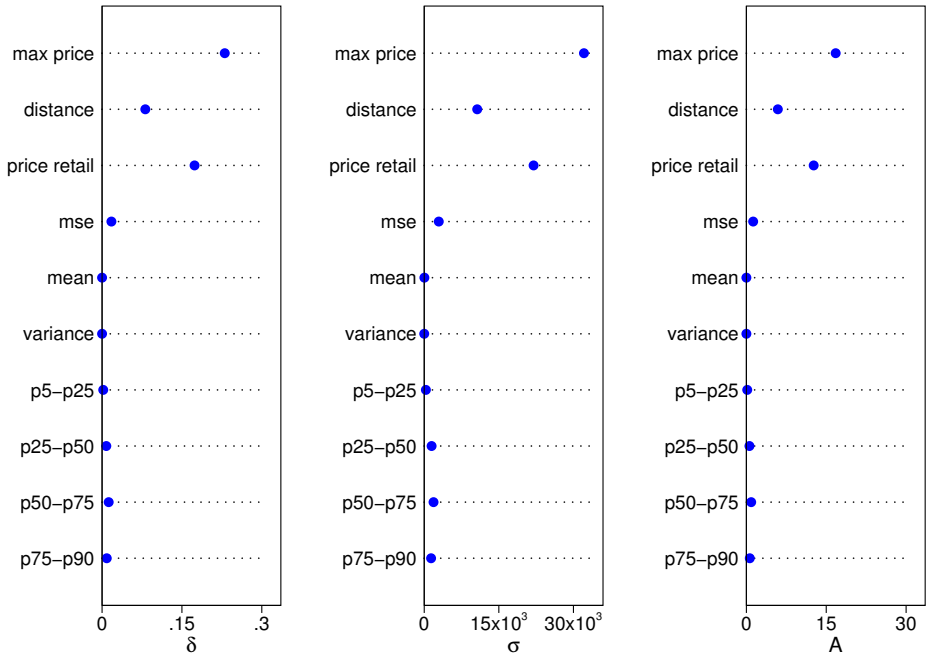


Figure 2.2: Semi-Elasticity of the Prob Limit of the Estimator wrt Standardized Moments.

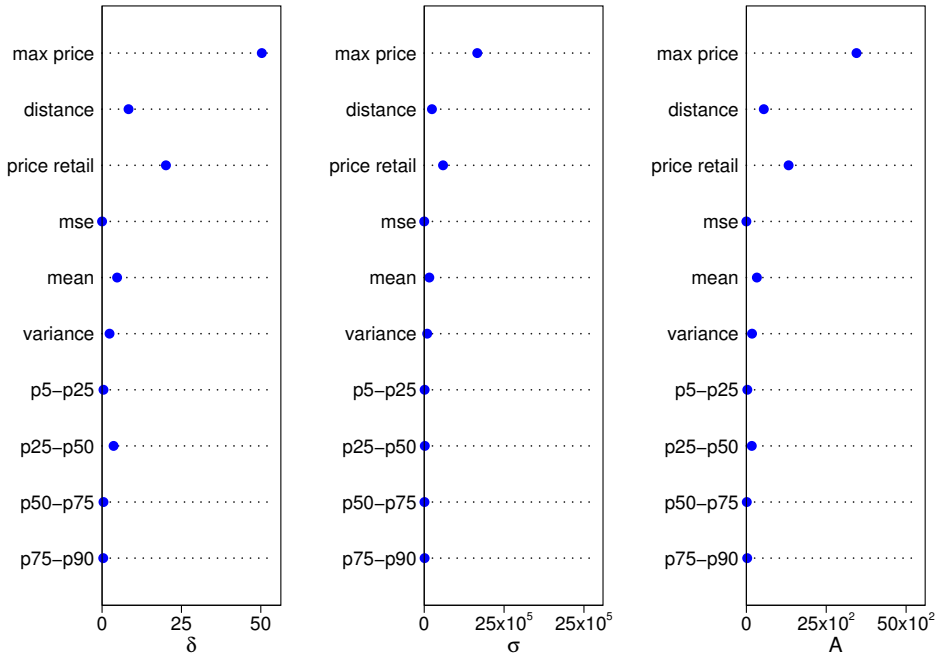


Figure 2.3: Elasticity of the Estimator wrt the Weight on a Moment.

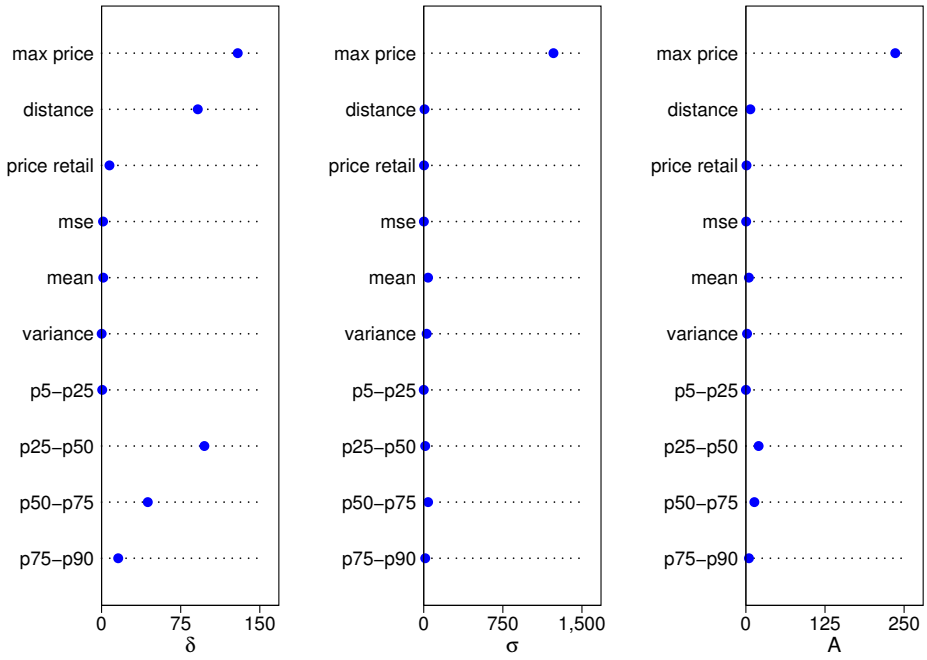


Figure 2.4: Elasticity of the Variance of Estimator wrt Variance of Moments.

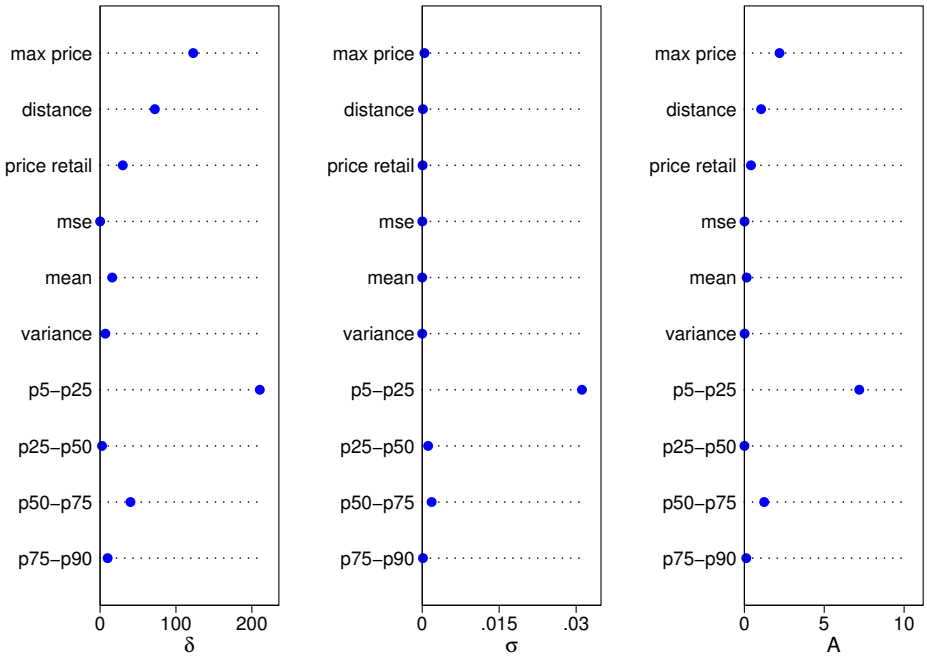


Figure 2.5: Elasticity of the Moment wrt the Parameter.

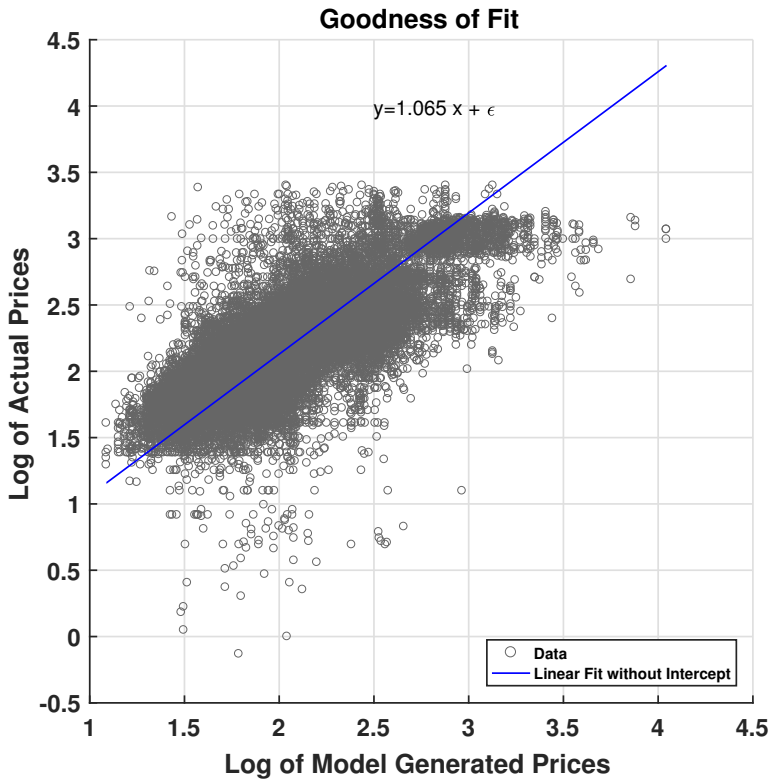


Figure 2.6: Linear Fit: Data vs. Model.

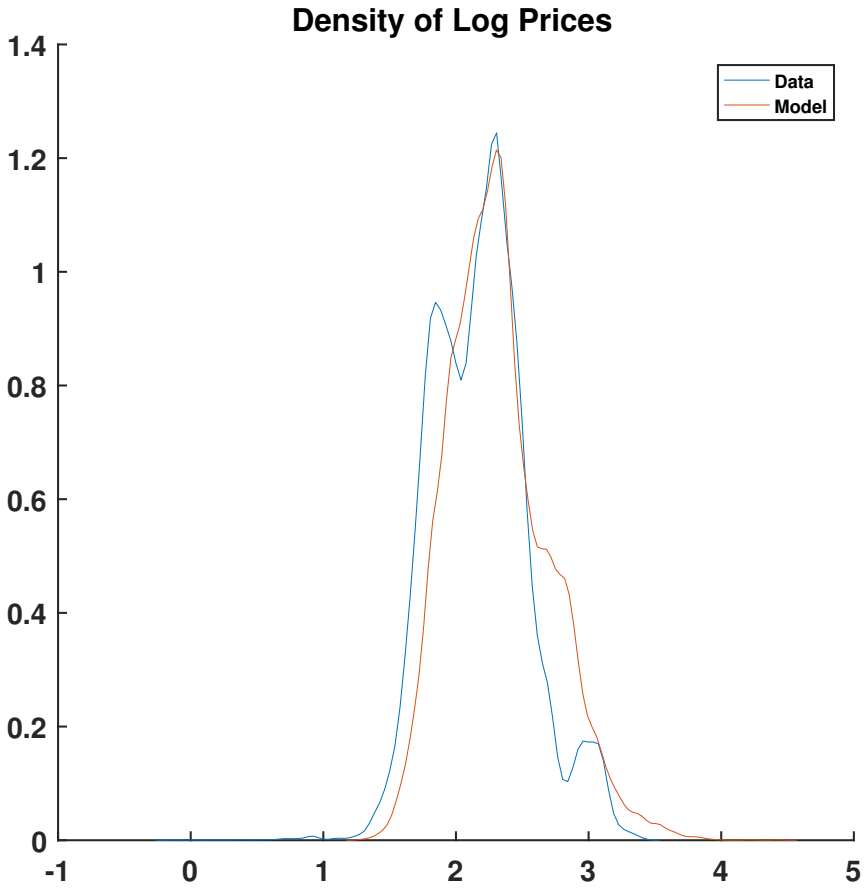


Figure 2.7: Empirical Density of Prices: Data vs. Model.

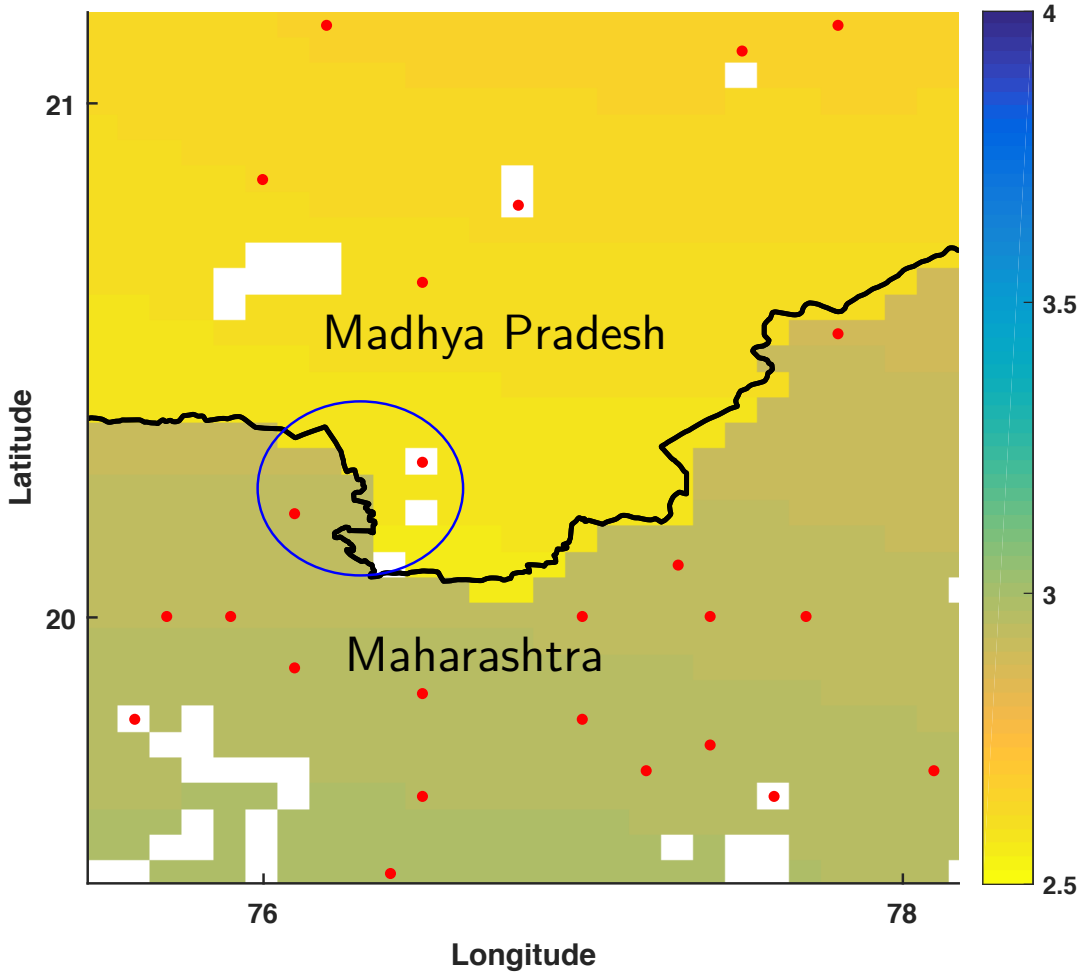


Figure 2.8: Illustrative Example.

Notes: The red dots are primary grain markets. The border markets in the example are encircled in blue. The figure also plots the log of soybean prices (Rs. per kilo) received by farmers.

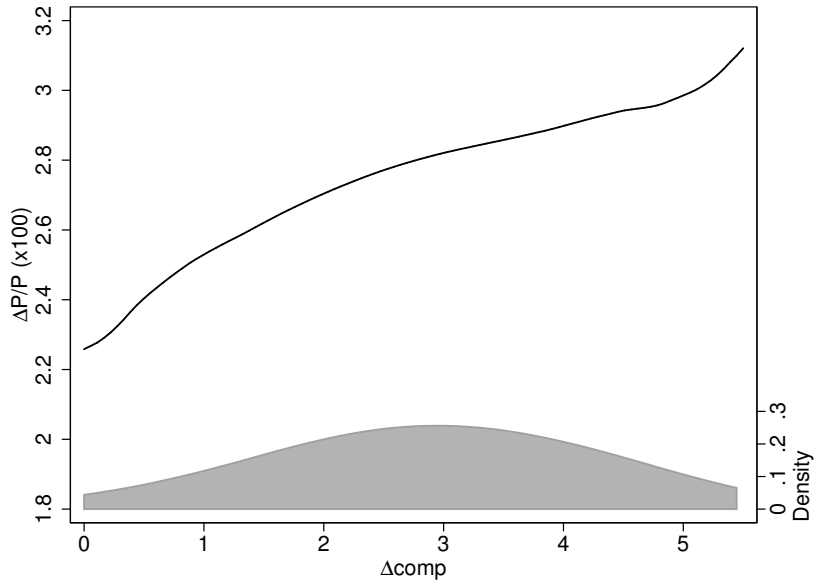


Figure 2.9: Change in Price at Each Market vs. Change in Competition.

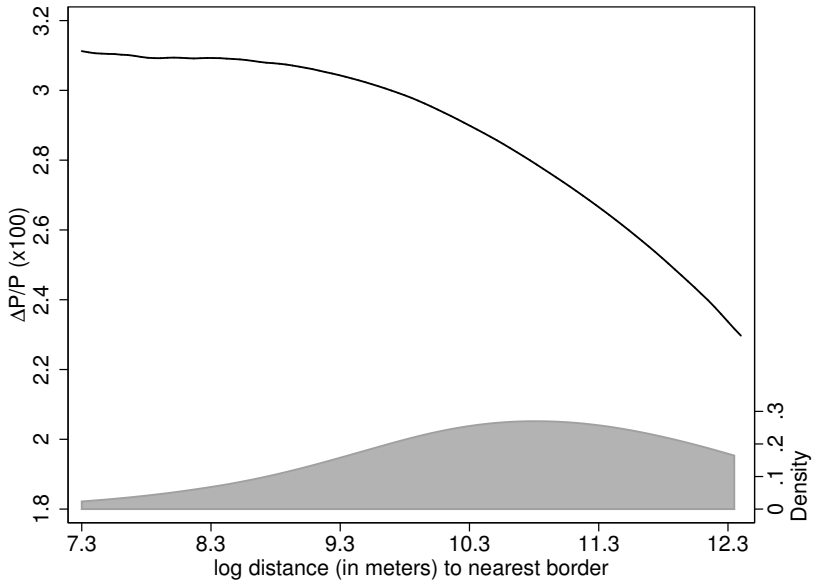


Figure 2.10: Change in Price at Each Market vs. Distance to Nearest Border.

Notes: The solid line in Figure 2.9 is the estimated relationship between change in price at a market and the change in competition, Δcomp, due to removal of the interstate trade restriction. Competition is measured as in the reduced-form regressions (defined in equation 1.4.1). The solid line in Figure 2.10 is the estimated relationship between change in price at a market and the distance from the market to the nearest state border. Estimated relationships using local linear regressions of bandwidth 1. Domain is 1st to 99th Percentile. Bottom of both figures plot the empirical distribution of the x-axis variable.

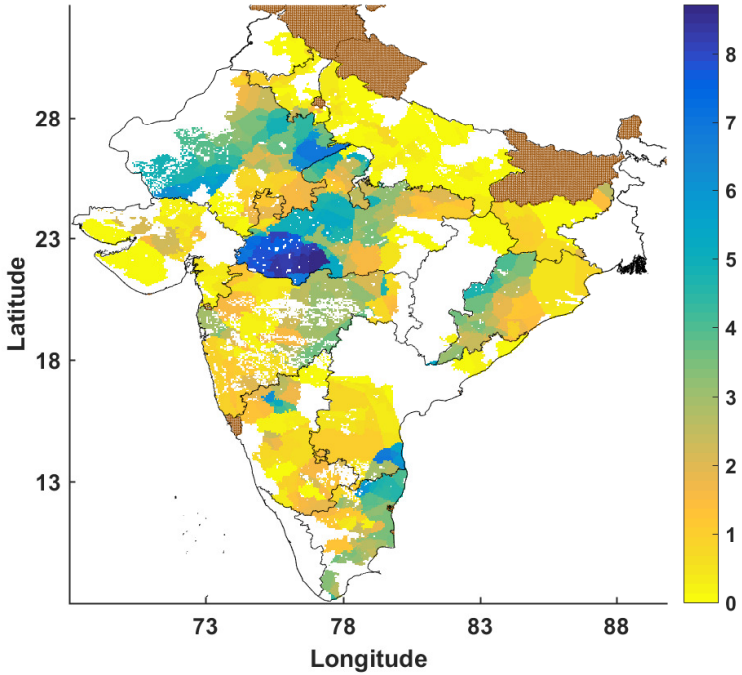


Figure 2.11: % Increase in Prices in the Fall Holding Market Choice of Farmers Fixed.

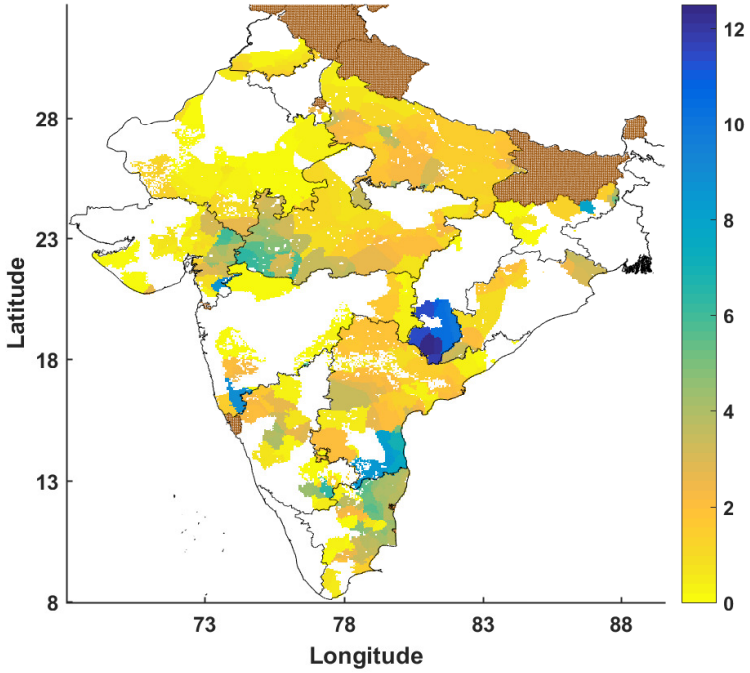


Figure 2.12: % Increase in Prices in the Spring Holding Market Choice of Farmers Fixed.

Note: Brown color represents no data and white represents zero gains.

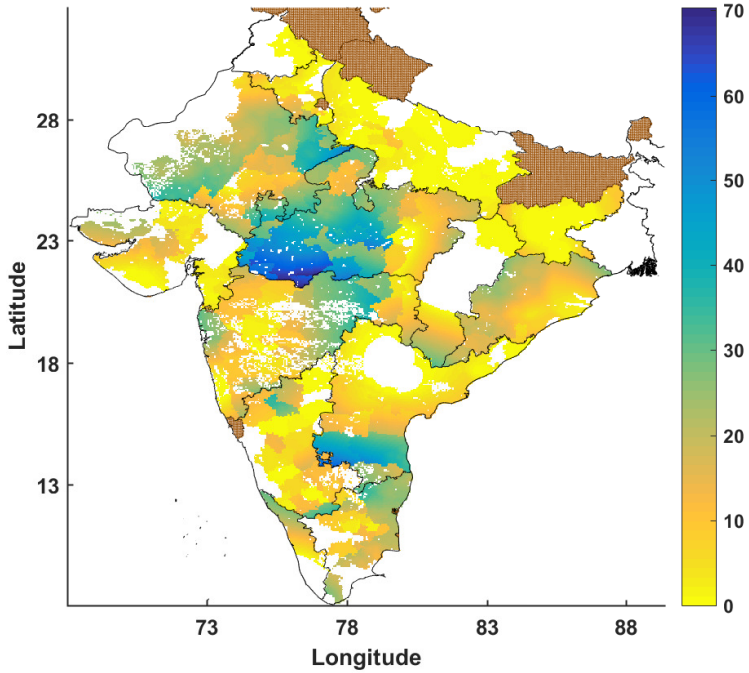


Figure 2.13: % Increase in Prices allowing for Optimal Market Choice—Fall Season.

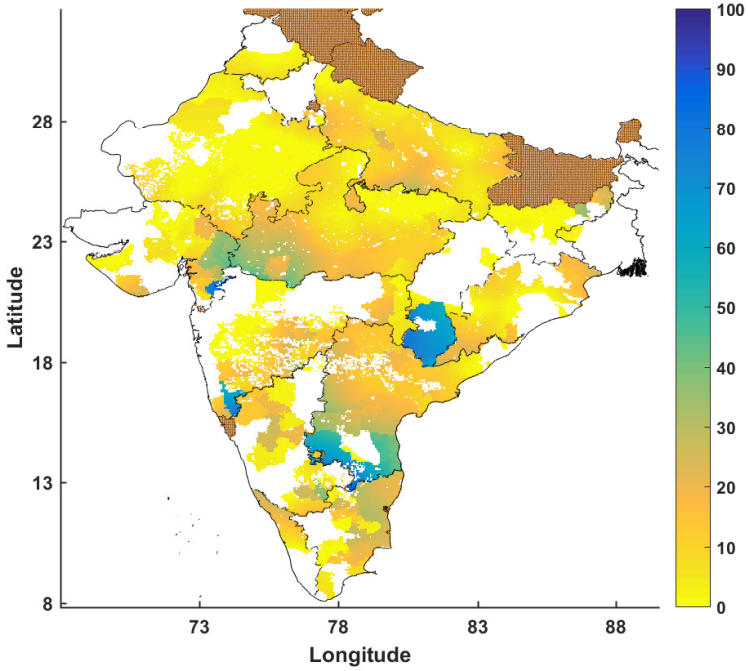


Figure 2.14: %Increases in Prices allowing for Optimal Market Choice—Spring Season.

Notes: Brown color represents no data and white represents zero gains.

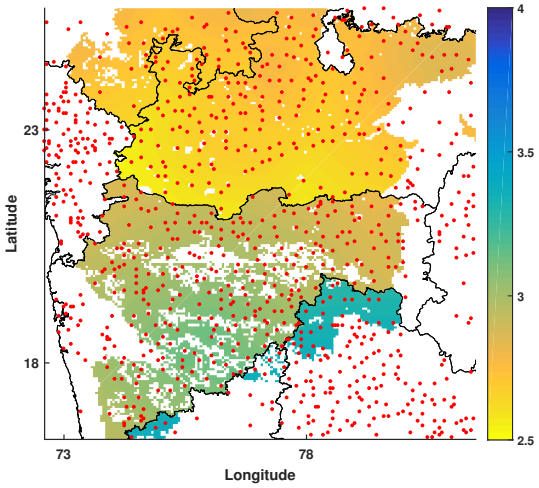


Figure 2.15: Log of Soybean Prices (Rs. per kilo) Before Policy change.

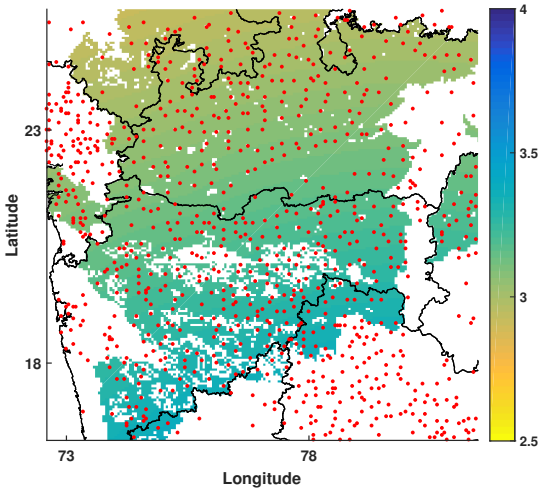


Figure 2.16: Log of Soybean Prices (Rs. per kilo) After Policy change.

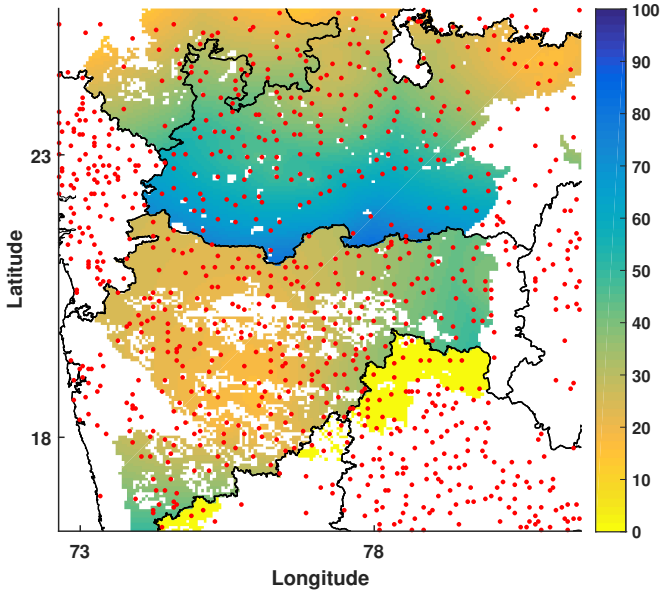


Figure 2.17: %Change in Soybean Prices in Maharashtra and Madhya Pradesh.

Note: The red dots are primary grain markets.

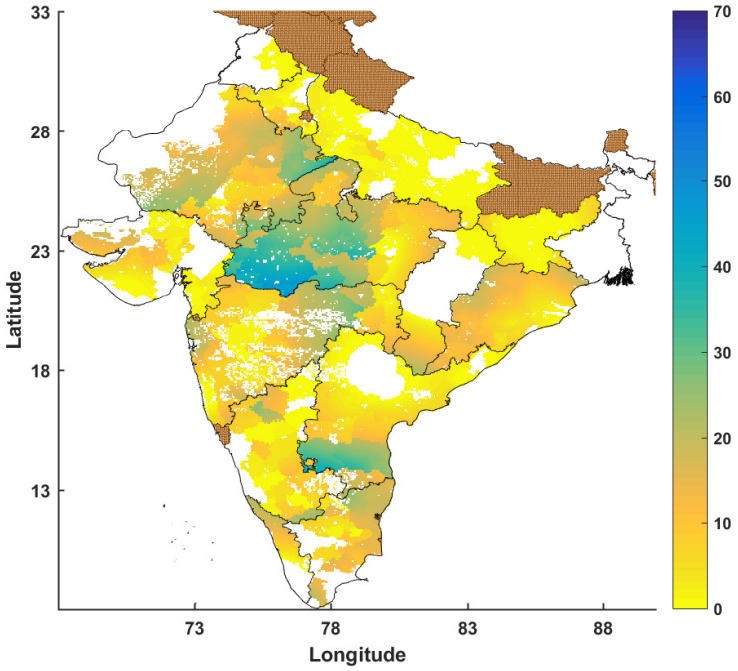


Figure 2.18: % Increase in Production in the Fall Season.

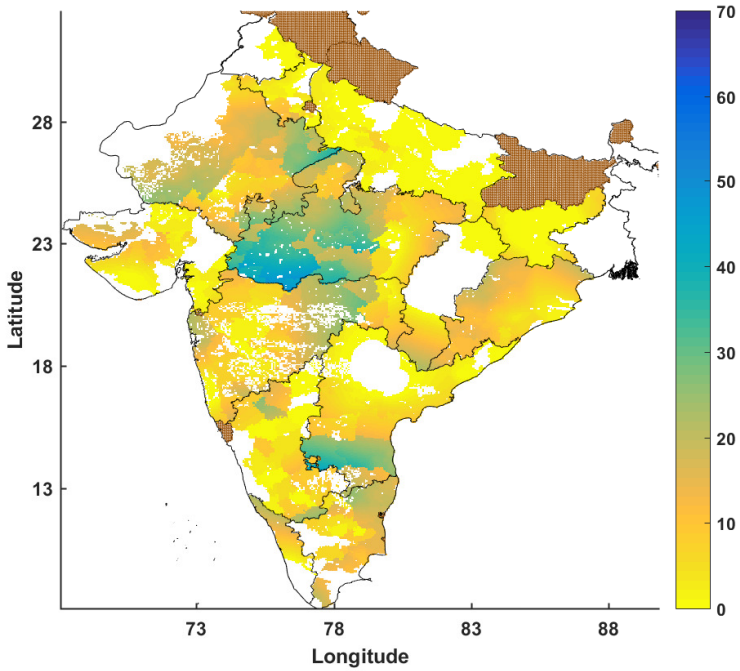


Figure 2.19: % Increase in Production in the Spring Season.

Notes: Brown color represents no data and white represents zero gains.

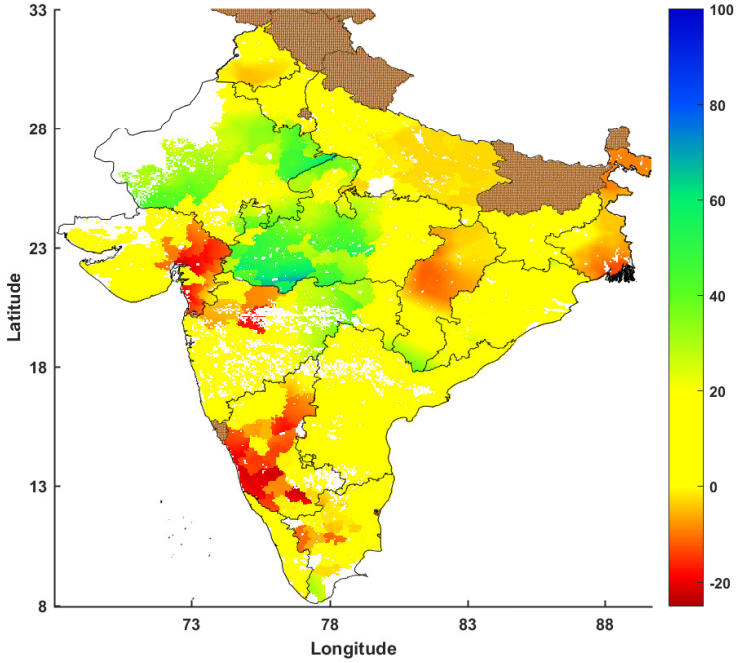


Figure 2.20: % Increase in Prices in the Fall Season Incorporating a Demand Side Response.

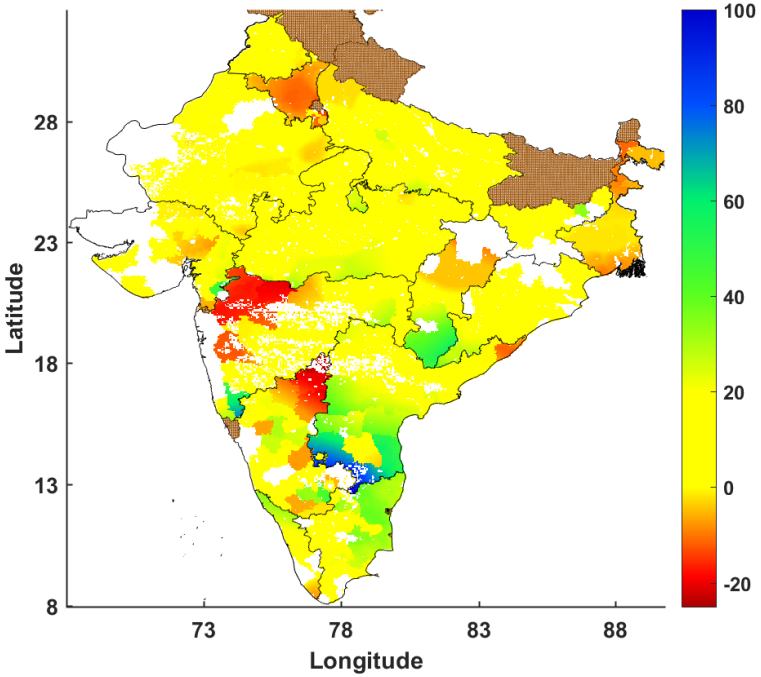


Figure 2.21: % Increase in Prices in the Spring Season Incorporating a Demand Side Response.

Notes: Brown color represents no data and white represents zero gains.

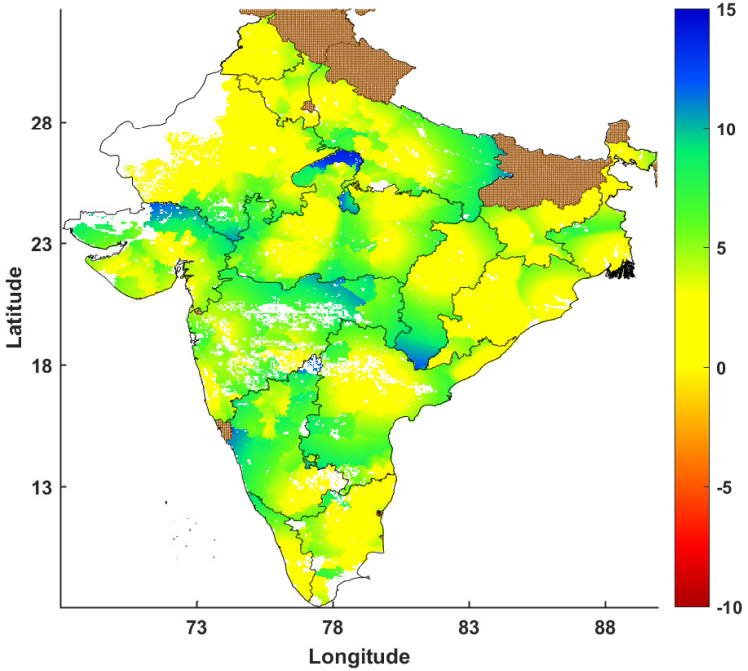


Figure 2.22: % Increase in Prices due to Reduction in Trade Costs with Border Restrictions.

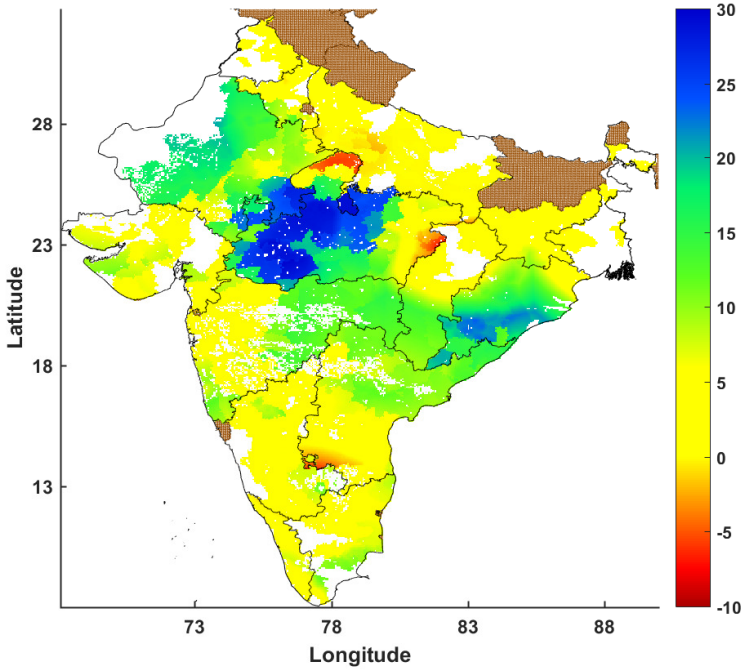


Figure 2.23: Additional Changes in Prices (in %) because of a Reduction in Trade Costs and Elimination of Border Restrictions.

Notes: Brown color represents no data and white represents zero gains. Estimates are for the fall season. To compute the changes in the Figure (2.23), I first compute changes in prices due to reduction in trade cost and elimination of border restrictions. From this I subtract the gains in prices just due to reduction in trade costs (Figure 2.22) and just due to elimination of border restrictions to trade (Figure 2.13).

Appendix

2.A Proof of Theorem 1

To show that there is unique fixed point of (2.1.7):

$$p^f(m) = (1-\delta) \max_{k \in \mathcal{M}-\{m\}} \left\{ \frac{p^f(k)}{\tau_{mk}} \right\} + \delta p_m^r \tag{2.A.1}$$

Define $p^f : \mathcal{M} \to [0, \max_m p_m^r]$. Note that $\mathcal{M} = \{1, 2, ..., M\}$, $\tau : \mathcal{M} \times \mathcal{M} \to [1, \infty]$, $\delta \in (0,1)$ and $p_m^r \in \Re_+$.

2.A.1 Existence

Since the max function is continuous, the RHS of (2.A.1) is a continuous function. Therefore, Brouwer's Fixed Point Theorem guarantees existence of a $p^{f*}(.)$ that solves the system of equations defined in (2.A.1). Further, since $\delta \in (0,1)$ and $p_m^r \in \Re_+$, it must be that $p^{f*}(.) > 0$.

2.A.2 Uniqueness

Define $T(f)(s) = (1-\delta) \max_{s' \in \Gamma(s)} \left\{ \frac{f(s')}{\tau(s,s')} \right\} + \delta p_s^r$. I will verify Blackwell's sufficient conditions to show that $T(.)$ is a contraction.

1. Monotonicity: Let $f(s) \geq g(s)$ for all s.

$$\begin{aligned} (1-\delta) \max_{s' \in \Gamma(s)} \left\{ \frac{f(s')}{\tau(s,s')} \right\} + \delta p_s^r &\gtreqless (1-\delta) \max_{s' \in \Gamma(s)} \left\{ \frac{g(s')}{\tau(s,s')} \right\} + \delta p_s^r \\ \max_{s' \in \Gamma(s)} \left\{ \frac{f(s')}{\tau(s,s')} \right\} &\gtreqless \max_{s' \in \Gamma(s)} \left\{ \frac{g(s')}{\tau(s,s')} \right\} \end{aligned}$$

Since $f(s) \geq g(s) \, \forall s$, let us say $s^* = \arg\max_{s'} \frac{g\left(s'\right)}{\tau\left(s,s'\right)}$, then $\frac{f(s^*)}{\tau(s,s^*)} \geq \frac{g(s^*)}{\tau(s,s^*)}$. It must be true that

$$\max_{s' \in \Gamma(s)} \left\{ \frac{f\left(s'\right)}{\tau\left(s,s'\right)} \right\} \geq \frac{f\left(s^*\right)}{\tau\left(s,s^*\right)} \geq \frac{g\left(s^*\right)}{\tau\left(s,s^*\right)} = \max_{s' \in \Gamma(s)} \left\{ \frac{g\left(s'\right)}{\tau\left(s,s'\right)} \right\}$$

Therefore,

$$T(f)(s) \geq T(g)(s)$$

2. Discounting: Let $a \in \mathbb{R}_+$.

$$\begin{aligned} T(f+a)(s) &= (1-\delta) \max_{s' \in \Gamma(s)} \left\{ \frac{f\left(s'\right)+a}{\tau\left(s,s'\right)} \right\} + \delta p_s^r \\ &\leq (1-\delta) \max_{s' \in \Gamma(s)} \left\{ \frac{f\left(s'\right)}{\tau\left(s,s'\right)} \right\} + \delta p_s^r + (1-\delta) a \\ &= T(f)(s) + (1-\delta) a \end{aligned}$$

The second line follows because $\tau(.,.) \geq 1$.

2.B Proof of Lemma 1

Let $\mathcal{M}$ denote the set of markets in a region. When border restrictions are removed, let the set of markets available to farmers in this region increases to $\mathcal{N} \supset \mathcal{M}$. Denote equilibrium farmer prices before and after border restrictions are removed by $p^1(t)$ and $p^2(t)$ respectively, where $t \in \mathcal{M}, \mathcal{N}$.

Suppose $\exists m \in \mathcal{M}$, such that $p^2(m) < p^1(m)$. This implies

$$(1-\delta)\max_{k\in\mathcal{N}-\{m\}}\left\{\frac{p^2(k)}{\tau_{mk}}\right\}+\delta p_m^r<(1-\delta)\max_{j\in\mathcal{M}-\{m\}}\left\{\frac{p^1(j)}{\tau_{mj}}\right\}+\delta p_m^r$$
$$\iff \max_{k\in\mathcal{N}-\{m\}}\left\{\frac{p^2(k)}{\tau_{mk}}\right\}<\max_{j\in\mathcal{M}-\{m\}}\left\{\frac{p^1(j)}{\tau_{mj}}\right\}$$
$$\iff \max_{k\in\mathcal{M}-\{m\}}\left\{\frac{p^2(k)}{\tau_{mk}}\right\}<\max_{j\in\mathcal{M}-\{m\}}\left\{\frac{p^1(j)}{\tau_{mj}}\right\}.$$

The last line follows from $\mathcal{M}\subset\mathcal{N}$. Because of uniqueness it must be that

$$\max_{k\in\mathcal{M}-\{m\}}\left\{\frac{p^1(k)}{\tau_{mk}}\right\}<\max_{j\in\mathcal{M}-\{m\}}\left\{\frac{p^1(j)}{\tau_{mj}}\right\}.$$

This implies that

$$p^1(k^*)<p^1(k^*),$$

which is a contradiction. Therefore, it must be that $p^2(m)\geq p^1(m)\,\forall m\in\mathcal{M}$.

2.C Change in Production in Response to a Change in Prices

Let p denote prices that farmers get and w^k the price of the kth intermediate input. The production function is given in (2.1.1). The profit maximization problem for a farm can be written as follows where I omit the f subscript:

$$\max_{x^k}\quad p\tilde{A}\left(h^\gamma l^\nu\prod_{k=1}^{K}\left(x^k\right)^{\alpha_k}\right)-\sum w^k x^k$$

The First Order Condition is:

$$py\alpha_k=w^k x^k \tag{2.C.1}$$

Therefore:

$$\frac{\alpha_k}{\alpha_j} = \frac{w^k x^k}{w^j x^j}$$
$$x^j = \left(\frac{\alpha_j}{w^j}\right)\left(\frac{w^k x^k}{\alpha_k}\right)$$

From the first order condition (2.C.1), we can write:

$$p\left(\frac{\alpha_k}{x^k}\right)\left(\tilde{A}h^\gamma l^\nu \prod_{j=1}^{K}\left(x^j\right)^{\alpha_j}\right) = w^k$$
$$\tilde{A}h^\gamma l^\nu p\left(\frac{\alpha_k}{x^k}\right)\left(\prod_{j=1}^{K}\left(x^j\right)^{\alpha_j}\right) = w^k$$
$$\tilde{A}h^\gamma l^\nu p\left(\frac{\alpha_k}{x^k}\right)\left(\frac{w^k x^k}{\alpha_k}\right)^{1-\gamma-\nu}\left(\prod_{j=1}^{K}\left(\frac{\alpha_j}{w^j}\right)^{\alpha_j}\right) = w^k$$
$$\tilde{A}h^\gamma l^\nu p\left(\frac{\alpha_k}{x^k}\right)^{\gamma+\nu}\left(\prod_{j=1}^{K}\left(\frac{\alpha_j}{w^j}\right)^{\alpha_j}\right) = \left(w^k\right)^{\gamma+\nu}$$
$$\tilde{A}h^\gamma l^\nu p\left(\frac{\alpha_k}{w^k}\right)^{\gamma+\nu}\left(\prod_{j=1}^{K}\left(\frac{\alpha_j}{w^j}\right)^{\alpha_j}\right) = \left(x^k\right)^{\gamma+\nu}$$
$$\Upsilon p^{\frac{1}{\gamma+\nu}} = x^k$$

where, Υ is a constant.

Now, new output to old output ratio is:

$$\begin{aligned}\frac{y'}{y} &= \frac{\prod_{k=1}^{K}\left(x^k\right)'^{\alpha_k}}{\prod_{k=1}^{K}\left(x^k\right)^{\alpha_k}} \\ &= \prod_k\left(\frac{\left(x^k\right)'}{x^k}\right)^{\alpha_k} \\ &= \prod_k\left(\frac{p'}{p}\right)^{\frac{\alpha_k}{\gamma+\nu}} \\ &= \left(\frac{p'}{p}\right)^{\frac{1-(\gamma+\nu)}{\gamma+\nu}}\end{aligned}$$

2.C.1 A Note on Calibrating $\gamma + \nu$

From the first order condition (2.C.1),

$$\begin{aligned} py \sum_{k=1}^{K} \alpha_k &= \sum_{k=1}^{K} w^k x^k \\ \iff py\,(1-\gamma-\nu) &= \sum_{k=1}^{K} w^k x^k \end{aligned}$$

Therefore, denoting wage rate by w^l and rental rate of land by w^h, expenditure on fixed factors must be:

$$w^l l + w^h h + \pi = (\gamma + \nu)\, py,$$

and therefore

$$\gamma + \nu = \frac{w^l l + w^h h + \pi}{w^l l + w^h h + \sum_{k=1}^{K} w^k x^k + \pi}.$$

2.D Algorithm to Compute Equilibrium Prices

1. Start with a distribution of retail prices $\{p_0^r\}$, associated farmer prices, and crop output $\{q_0^r\}$.

2. Remove border restrictions to trade and compute the new farmer prices and the anticipated output. Given the output and the demand elasticities, find the anticipated retail prices $\{p_A^r\}$.

3. Guess the new retail prices $\left\{p_{guess}^r\right\}$ as a fraction of the anticipated retail prices.

4. Find the associated farmer prices and quantities.

5. Find the actual change in retail prices $\{p_1^r\}$. If the actual retail prices are equal to the guess, then stop.

6. Otherwise update $p_{guess}^r = p_o^r + 0.2 \times \left(p_{guess}^r - p_1^r\right)$. Go to Step 4.

2.E Crops-State Combinations in Structural Estimation

For structural estimation, I include data from the following states. This choice is meant to exclude states that do not produce a particular crop. To determine this I make use of Agriculture Statistics at a Glance published each year by the Ministry of Agriculture, Government of India.

Crop	States Excluded
Rice	Rajasthan
Wheat	Andhra Pradesh, Chhattisgarh, Kerala, Odisha, Tamil Nadu
Barley	All except Haryana, Madhya Pradesh, Punjab, Rajasthan, Uttar Pradesh
Maize	Chhattisgarh, Jharkhand, Kerala, Odisha, Punjab, West Bengal
Sorghum	Chhattisgarh, Kerala, Odisha, Tamil Nadu, West Bengal
Soybean	All except Madhya Pradesh, Maharashtra, Karnataka, Rajasthan
Finger Millet	Haryana, Kerala, Punjab, Rajasthan, Uttar Pradesh
Pearl Millet	Chhattisgarh, Jharkhand, Kerala, Odisha, Punjab, West Bengal

2.F Sensitivity to Allocation of Crop Choice to Grid Cells

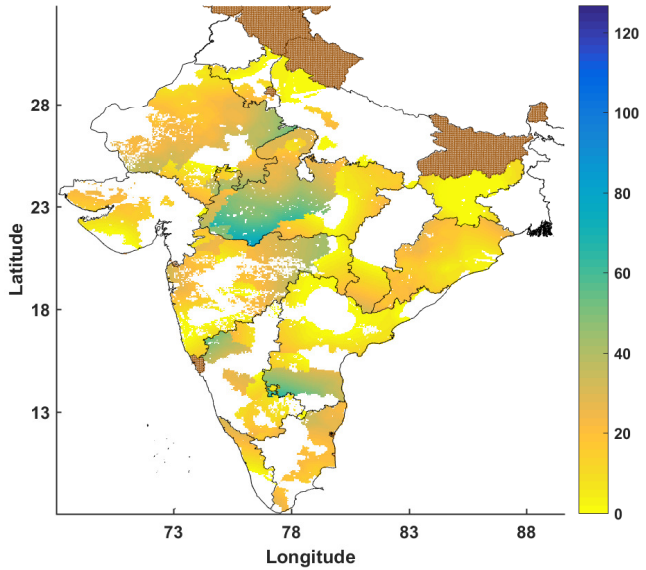


Figure 2.F.1: % Increase in Prices - Fall Season
For the main crop grown in the district.

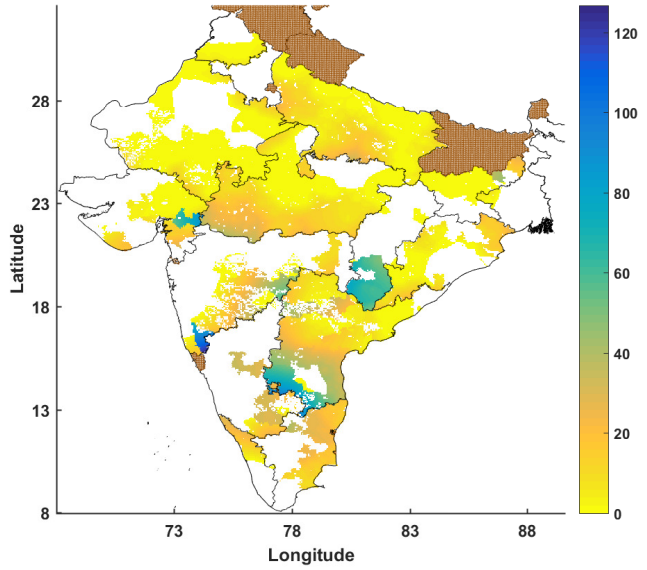


Figure 2.F.2: %Increases in Prices - Spring Season
For the main crop grown in the district.

Notes: Brown color represents no data and white represents zero gains.

Chapter 3

Escaping Malthus: Economic Growth and Fertility Change in the Developing World*

3.1 Introduction

The 1960s were a time of grave popular, political, and academic concern about global overpopulation (Connelly 2008). The world's population was growing faster than ever before, led by the developing world, where its growth rate exceeded that in the developed world threefold. Yet a Malthusian crisis was averted as fertility more than halved in the developing world, from six children per woman to fewer than three today, while both GDP per worker and GDP per capita more than doubled.[1] Despite a large body of economic theory positing a link between economic growth and fertility (see Galor 2011 for a review),

*Co-authored with Tom Vogl, Assistant Professor of Economics and International Affairs, Princeton University. Email: tvogl@princeton.edu.. A version of this chapter is forthcoming in the American Economic Review with the same title.

[1]Here we follow the UN in defining Europe, the Western Offshoots and Japan as More Developed Regions, and the rest of the world as Less Developed Regions. The fertility claim is based on the UN *World Population Prospects*, while the labor productivity claim is based on the Penn World Table. Section 2 introduces both data sources.

existing empirical research provides few basic facts on how they related to each other during this puzzling episode in world population history. Did the simultaneous rise in income and drop in fertility across the developing world signify a break from long-standing Malthusian dynamics in which population rose with productivity, or have these dynamics persisted in the background due to the continued importance of land as a factor of production (Hansen and Prescott 2002)?

This chapter seeks to empirically characterize the relationship between economic growth and fertility change in the developing world since 1950, with an eye toward matching the facts with leading economic theories of growth and demographic change. We argue that careful attention to time horizons and the lifecycle is crucial to developing a coherent account of the relationship. Fertility may respond differently to growth occurring over different time horizons, and these responses may vary over the lifecycle. For instance, long-run growth may alter returns (e.g., to child investment) or prices (e.g., women's wages) in ways that short-run fluctuations do not, while short-run fluctuations may have liquidity and intertemporal substitution effects that are absent or different in the long run. And in both cases, variation in growth can affect both the lifetime number of births—in demographic parlance, the *quantum* of fertility—and their timing—the *tempo*. Existing research pays limited attention to these issues, particularly as they relate to the process of economic development. We fill this gap by using rich microdata to disentangle fertility responses to fluctuations and long-run growth over the lifecycle.

Table 3.1 details how we extend the existing literature, separating (for both existing findings and our own) fluctuations from long-run growth and overall findings from decompositions of tempo and quantum effects. Existing evidence on fluctuations is most complete in currently rich countries, where a large literature finds that birth rates are procyclical both at present (Sobotka et al. 2011) and hundreds of years ago (Galloway 1988; Lee 1997). Most studies do not track subsequent fertility over the lifecycle, but Currie and Schwandt (2014) find in the US that downturns at particular ages decrease lifetime fertility, implying a mix

of tempo and quantum effects. In the contemporary developing world, research on select countries finds fewer births after economic crises, but its generalizability across countries is unclear, as is its lifecycle interpretation.[2] Moving to the longer run, two stylized facts motivate the theoretical literature on growth and fertility (Galor 2011): (1) the coincidence of fertility decline with the emergence of modern growth in historical time series, and (2) the inverse association between per capita income and fertility across countries today. Although these patterns suggest a link between growth and fertility decline, panel analyses of the relationship are rare and inconclusive, so its magnitude and contributing factors are poorly understood.[3] Even more poorly understood is the relationship between long-run growth and the timing of fertility over the lifecycle.[4]

To study growth-fertility linkages at different time and lifecycle horizons, we combine macro-economic data with the reproductive histories of 2.3 million women from 255 World Fertility Surveys (WFS) and Demographic and Health Surveys (DHS), covering 81 low- and middle-income countries. The survey data allow us to avoid standard cross-country demographic databases, which rely heavily on interpolation, smoothing, and demographic modeling (United Nations 2015; World Bank 2015): problematic for studying fluctuations—which may be smoothed out—and age heterogeneity—which is typically restricted by demographic models. We also use the survey data to calculate cohort measures of fertility and investigate within-country heterogeneity—both difficult to do with cross-country databases—and to better match the timing of conception (rather than birth).

[2]See National Research Council (1993); Tapinos et al. (1997); Lindstrom and Berhanu (1999); and Adsera and Menendez (2011), which with 18 Latin American countries provides the broadest geographic coverage.

[3]Bongaarts and Watkins (1996) find that growth in the Human Development Index (HDI) predicts decline in the total fertility rate (TFR), 1960-90, but do not separate the components of the HDI. Schultz (1997) finds that, conditional on a range of covariates, economic growth was unrelated to TFR change, 1972-88, with no further exploration of the result. Over the 20^{th} century, Herzer et al. (2012) find that growth predicts fertility decline, while Murtin (2013) finds a non-monotonic relationship, but neither distinguishes fluctuations from long-run growth. All draw on standard cross-country datasets and do not consider the lifecycle.

[4]In the West, fertility control depended on the marriage timing during the Malthusian era (Hajnal 1965) and on increased spacing and earlier stopping (Knodel 1987; Bean et al. 1990) during the fertility transition. In contemporary developing countries, early marriage is becoming rarer (Mensch et al. 2005), and birth spacing has increased (Casterline and Odden 2016). None of this research touches on the role of economic growth.

Table 3.1: Summary of Existing and New Evidence on Fertility Responses to Growth

	Overall	**Tempo (timing) vs. quantum (lifetime)**
Fluctuations	**Existing evidence:** fertility procyclical in MDCs and pre-industrial Europe; limited systematic evidence on LDCs.	**Existing evidence:** in US, downturns delay births, reduce lifetime fertility if experienced in the 20s; no evidence on LDCs.
	New evidence: fertility procyclical in LDCs; driven by downturns, stronger among less-educated women.	**New evidence:** downturns delay births in LDCs, reduce lifetime fertility if experienced after age 30.
Long-run growth	**Existing evidence:** in historical time series, growth takeoffs coincide with fertility decline; incomplete correlational evidence on LDCs.	**Existing evidence:** varied importance of starting, spacing, and stopping strategies; little direct evidence on relation to growth.
	New evidence: growth associated with fertility decline in LDCs; related to rising enrollment, not sectoral comp., death, female educ., FLFP.	**New evidence:** growth associated with fertility delay in LDCs; age-specific rates fall more quickly before 40, slower after.

Note: MDC/LDC = more/less developed country. FLFP = female labor force participation.

With these data in hand, we can provide a geographically and temporally broader account of growth-fertility links at various time and lifecycle horizons than previously possible. We carry out three analyses to add new evidence to each cell of Table 3.1: two on the flow of fertility at the population level (i.e., period fertility) and one on the stock of fertility at the cohort level (i.e., cohort fertility). First, we examine how the annual rate of starting a successful pregnancy (i.e., conceiving a future liveborn child) responds to growth fluctuations in the short run, assessing whether immediate effects are subsequently offset and whether fertility responds differently to booms and busts. Second, we estimate how the long-run rate of economic growth relates to the long-run rate of fertility decline across the lifecycle. Third, we look across cohorts within a country to ask how completed fertility varies with macroeconomic conditions experienced along the lifecycle.

Our findings support the view that time horizon and lifecycle heterogeneity is important. When we study annual fluctuations, we find that fertility is procyclical at all ages, with statistical significance at prime parenting ages (20-34, when fertility is highest). The procyclicality is driven by downturns and is stronger for less-educated women, suggesting a role for liquidity constraints that prevent poor households from smoothing through recessions. In contrast, when we study average rates of change over 20 or more years, we find that fertility declines and delays with economic growth, consistent with price effects unique to the long run. Faster-growing economies see faster declines in prime-age fertility and slower declines in older-age fertility, with the former dominating for total fertility. The asymmetry between short- and long-run growth is substantial: among 25-29 year olds, a 1-log point growth fluctuation is associated with a spike of 0.6 pregnancies per 1000 women, whereas the same magnitude increase in the long-run growth rate is associated with an annual decline of 0.4 pregnancies per 1000.[5] When we compare cohorts within a country, we find that growth experienced early in the reproductive period is unrelated to lifetime fertility, while growth in the 30s leads to higher lifetime fertility. We also show that many

[5]We analyze growth in GDP per adult to avoid reverse causality; results are similar when we use overall GDP.

of our findings would be obscured in standard cross-country fertility data, while others are altogether impossible to study, underscoring the benefits of combining hundreds of surveys. In a simple theoretical framework, we demonstrate that these findings are consistent with the parental choice problem from the literature on long-run growth and the demographic transition, extended to include a lifecycle with liquidity constraints.

While all of the analyses are correlational, they differ substantially in their statistical interpretation. The short-run and cohort analyses use demanding regression specifications that isolate within-country variation, while the long-run analysis must rely on cross-country variation in rates of change. But notably, the long-run results are not explained by the initial levels of GDP per adult and population density; nor by changes in adult female education, adult female labor force participation, the sectoral composition of value added, urbanization, infant mortality, conflict, and democracy. We find only one trend that explains part of the long-run association: rising school enrollment (among teens, not mothers). Insofar as this trend suggests rising returns to human capital investment, our results are consistent with human capital-based theories of unified growth (Galor and Weil 2000).

Our use of hundreds of survey datasets places us in an exciting literature at the intersection of growth economics and development economics, which combines large quantities of microdata from across the developing world to shed light on macroeconomic aspects of the development process (Young 2013; Kleven and Landais 2016; Aaronson et al. 2017; Lagakos et al. 2017). Advances in data availability make it possible to characterize population moments with more precision and (geographic and temporal) generalizability than ever before. Besides the direct contribution to knowledge on growth and the demographic transition, our results also relate to the recent literature that estimates wealth or income effects on fertility using variation from natural resource (Black et al. 2013; Brueckner and Schwandt 2015) or housing (Lovenheim and Mumford 2013; Dettling and Kearney 2014) booms. While the question of whether children are 'normal' is related to ours, substitution and liquidity effects are key to understanding our setting, making it a separate contribution. Conceptually and

methodologically, our analysis is more similar to research on the link between economic growth and mortality change—which is typically negative at all time horizons in developing countries (Deaton 2007; Baird et al. 2011)[6]—and on the cumulative mortality effects of economic shocks over the lifecycle (Cutler et al. 2016).

The chapter clarifies how to place the developing world's postwar experience in the timeline of long-run growth and demographic change. In the Malthusian era, productivity growth increased living standards in the short run and population in the long run (through higher fertility and lower mortality), with the latter effect limiting the duration of the former effect (Lee 1997; Ashraf and Galor 2011). The patterns here represent a marked departure from those long-standing dynamics; instead of rising to offset productivity growth, fertility falls with growth in the long run, augmenting its per capita effects. This departure has implications for cost-benefit analyses of development policies, which are more effective without Malthusian population offset. In a complementary analysis of data from developed countries, we find similar short- but not long-run results, suggesting that our results pertain to the phase after an economy begins the escape from the Malthusian trap but before it attains high living standards and low fertility.

3.2 Data

We use survey microdata to generate fertility aggregates but draw on standard databases to measure economic growth. This section describes both data sources and explains how we use them to construct our analysis datasets. It then introduces several other databases we use in extensions.

Macroeconomic Data To measure growth, we obtain data on GDP from the Penn World Table (PWT) v. 8.1, matching it with data on population and age structure from the United

[6] An exception is Colombia, where infant health is countercyclical (Miller and Urdinola 2010), as in the United States (Dehejia and Lleras-Muney 2004).

Nations. The central independent variable is the logarithm of GDP per adult age 15-64 for country c in year t, $GDPpa_{ct}$, where the age range for the denominator is chosen to minimize concern about the endogeneity of population size due to fertility and mortality. To ease interpretation, this variable is multiplied by 100, so that the results are quantified in log points, which approximately reflect percentage points. When analyzing levels, we adjust for purchasing power parity (PPP); when analyzing growth rates, we adjust for inflation but use national prices, following the recommendation in Johnson et al. (2013).

Fertility Microdata To measure fertility, we assemble data from all publicly available, standard format WFS and DHS surveys that are nationally representative for all women and can be merged with our macroeconomic dataset, leading to a sample of 81 countries. Appendix 3.A Table 3.A1 lists the number of surveys for each of these countries, which were all classified as low- or middle-income at the time of the surveys (World Bank 2015). Respondents provided full birth histories, listing all of their children ever born, with information on birth date and survival status.[7] These data allow us to track fertility behavior over time and the lifecycle, although they are sometimes subject to reporting errors (Schoumaker 2014), a matter we discuss further below. Reporting errors take the form of both omitted births (which are likely to have been in the distant past or to have involved deceased children) and displaced births (either forward or backward in time). Some surveys only interviewed women who had ever been married or had completed schooling; in those cases, we only use data on women who at the time of the survey belonged to an age group in which the rate of ever-marriage or school completion exceeded 95%.

For all analyses, we collapse the individual-level data into country-year-age or country-cohort cells, allowing us to weight countries in a consistent way across a range econometric models. We categorize age and birth cohort in single years. To estimate the fertility rate

[7]The DHS from El Salvador 1985 and Nigeria 1999 have well-known deficiencies in their birth histories (Casterline and Odden 2016). For these surveys, we do not use the birth histories but do use data on lifetime fertility.

for each cell, we pool data from all surveys in the same country and rescale the survey weights to reflect each survey's sample size contribution to the cell, excluding cells with fewer than 30 observations (<5% of cells).[8] We generate two types of fertility rates: period, summarizing fertility outcomes by age in a given year, and cohort, summarizing the lifetime fertility outcomes of women born in the same year.

Period Dataset For the analysis of period fertility, we study the age-specific conception rate, CR_{cta}: the number of conceptions per 1000 women aged a in year t from country c. Because we do not have information on miscarriages or abortions, we focus only on conceptions that resulted in a live birth; because we do not have information on gestational age at delivery, we assume that conception took place 9 months before the date of birth.[9] As such, we do not directly analyze childbearing behavior but instead use approximate dates of conception for fetuses that survived gestation: a limitation, given that economic conditions may affect the risk of fetal death.[10] To help distinguish the short and long run, we focus on country-age combinations with conception rates and macroeconomic data spanning at least 20 years. As reported in the first two columns of Table 3.2, this sample definition gives rise to 58,992 distinct cells defined by country, year, and age, with conception rates based on the fertility histories on 2.3 million women from 65 countries.[11] Pooling all ages 15-44, age-specific conception rates have a mean of 199 per 1000. The annual change in log GDP per adult averages 1.0, with a standard deviation of 5.8. Female education averages 4 years; the urbanization rate, 38%; and GDP per adult, 4,476 international dollars, adjusted for purchasing power parity.

[8]Because reporting errors may be more likely for longer recall periods, we apply a Bartlett kernel to the rescaled survey weights in a robustness exercise, down-weighting births that occurred long before the survey.

[9]We count multiple births as coming from a single conception and allow for the possibility that a woman may conceive twice in one year.

[10]The term 'conception rate' is thus a slight abuse of terminology, but one that follows Currie and Schwandt (2014) in their work on unemployment and fertility. The focus on conception (rather than birth) timing also follows them.

[11]Every sample in Table 3.2 has fewer than 81 countries because the period analysis omits surveys with unrepresentative or low quality birth histories, while the cohort analysis omits cohorts that lack complete macroeconomic histories. Appendix 3.A Table 3.A2 lists the countries included in each sample.

Cohort Dataset For the analysis of cohort fertility, we study the completed fertility rate, CFR_{cj}: the number of children per 1000 women from country c and birth cohort j. We only include women over 45 at the time of the survey, treating their fertility as complete. Our main cohort analyses are based on all children ever born, although we show that we obtain similar results when we only count children who survived to the date of the survey. As reported in the last column of Table 3.2, data are available on 935 country-cohort cells from 62 countries, containing 212k women over 45. The completed fertility rate per 1000 women averages 5951 children ever born and 4862 surviving children; an individual woman experiencing these rates would bear 6 children, of whom 1 would die before she reached her late 40s. Compared with the period sample, the cohort sample is characterized by a higher average age because it excludes women below age 45. Average educational attainment is also lower because earlier cohorts received less education; other characteristics have similar means.

Additional Data Sources For covariates, analyses of heterogeneity, alternative measures of fertility, and comparisons with developed countries, we draw on additional aggregate data sources. We obtain alternative fertility data from the World Development Indicators (WDI) and United Nations (UN); information on contraceptive use, population density, and the sectoral composition of value added from the UN; school enrollment and labor force statistics from the WDI and International Labor Organization (ILO); democratization scores from the Polity IV project; and conflict indicators from the UCDP/PRIO Armed Conflict Dataset. Average female education, the infant mortality rate, and the urbanization rate are estimated from WFS/DHS microdata. For comparison with developed countries, we use the Human Fertility Database, which assembles high-quality vital registration data.

3.3 Analysis of Period Fertility

Our analysis of period fertility focuses on how changes in fertility across the age distribution vary with short- and long-run economic growth. We begin by laying out how we distinguish between short- and long-run patterns empirically, followed by the main results for both horizons. We then delve further into the results—assessing non-linearity, lag structure, heterogeneity, and alternative covariates—first for the short run and then for the long run.

3.3.1 Defining Time Horizons

A key issue is how to define "short run" and "long run." To allow the data to speak to this issue, we run a series of first-difference regressions in which we vary the length of the difference. Because we are holding age fixed as time changes, we are effectively studying cross-cohort changes in age-specific fertility. For each 5-year age group A from $[15, 19)$ to $[40, 44)$, we run:

$$CR_{cta} - CR_{c,t-\Delta,a} = \beta^A \left(Y_{ct} - Y_{c,t-\Delta}\right) + \alpha_a^A + \varepsilon_{cta}^A \tag{3.3.1}$$

where $Y_{ct} = 100 \times \ln\left(GDPpa_{ct}\right)$. Because the distribution of single-year ages within each 5-year age group varies across countries and over time, we include a single-year age effect α_a^A, thus allowing fertility levels to trend differently for each single-year age within the 5-year age group. We estimate equation (3.3.1) using a range of values for the length of the difference Δ, from 1 to 30 years. Figure 3.1 displays the results, for each age group plotting estimates of the coefficients against the length of the difference.[12]

Figure 3.1 makes clear that economic growth and fertility change have different relationships over different time horizons. In the annual first difference ($\Delta = 1$), all age groups have positive coefficients, indicating procyclical fertility. But all but one age group immediately begin trending downward with rising Δ, becoming negative by $\Delta = 14$ and leveling off at

[12] Appendix 3.A Figure 3.A1 shows similar patterns if we leave out the single-year age effects or only use data from country-age cells with conception rates spanning at least 20 or at least 30 years.

about $\Delta = 20$. In other words, for all age groups less than 40, economic growth is negatively associated with fertility change in the long run. Above 40, economic growth is positively related with fertility change at all time horizons.

These patterns have two further implications. First, the annual first-difference coefficients estimated with equation (3.3.1) actually reflect a mix of short- and long-run associations, depending on the relative contributions of transitory fluctuations and long-run growth to the variance of the annual growth rate. Second, the sharp drop of the coefficients for the first five age groups beyond $\Delta = 1$, as well as their leveling at about $\Delta = 20$, suggest 1 and 20 years as reasonable definitions of the short and long run.

3.3.2 Methods

Based on the patterns in Figure 3.1, we divide our study of period fertility into a short-run analysis of annual fluctuations and a long-run analysis of average annual changes over periods of at least 20 years. To distinguish time horizons as clearly as possible, we restrict both analyses to country-age cells that span at least 20 years. For the short-run analysis, we modify equation (3.3.1) to fully disentangle the short- and long-run relationships by including a country effect, which absorbs the country's average changes in log GDP per adult and the conception rate. For completeness, we also include a year effect, addressing any spurious global trends. The first-difference specification then becomes:

$$\Delta CR_{cta} = \beta^A g_{ct} + \lambda_c^A + \tau_t^A + \alpha_a^A + \varepsilon_{cta}^A \tag{3.3.2}$$

where ΔCR_{cta} is the change in the conception rate from the previous year, and g_{ct} is the annual change in $100 \times \ln(GDPpa_{ct})$, which approximates the growth rate. λ_c^A, τ_t^A, and α_a^A are the country, year, and single-year age effects, which in first differences serve to control

for correlated level trends.[13] The coefficient β^A isolates how fluctuations of the growth rate from its long-run country average affect changes in the conception rate, net of year- and age-specific factors. A 1-log point growth fluctuation raises the change in the conception rate by β^A.

For the long-run analysis, we seek to estimate a long-difference version of equation (3.3.1). The standard approach would relate simple changes in log GDP per adult to simple changes in fertility over an interval of 20 years. However, because our country-year-age conception rates are noisy estimates from surveys, we deviate from this standard approach in order to leverage as much information as possible on the rate of long-run fertility change. Instead of the long difference, we analyze the average annual rates of change in the two variables, $\overline{g}_{ca}$ and $\overline{\Delta CR}_{ca}$, over periods of at least 20 years. To estimate these quantities using as much information as possible, we regress $100 \times \Delta \ln(GDPpa_{ct})$ and CR_{cta} on year within each country-age cell, using the slope of the trend as the estimated average annual rate of change.[14] We then run the following regression for each 5-year age group A:

$$\overline{\Delta CR}_{ca} = \beta^A \overline{g}_{ca} + \alpha_a^A + \varepsilon_{ca}^A \tag{3.3.3}$$

As before, the single-year age effect α_a^A absorbs any age-related factors common across countries. By collapsing the country-year-age observations into country-year cells, we remove the time dimension from our panel, so the year effect τ_t^A drops out. Similarly, because

[13] Equation (3.3.2) can be obtained from differencing a level specification with country (μ_c^A), year ($\tilde{\tau}_t^A$), and age (ω_a^A) fixed effects, as well as country (λ_c^A) and age (α_a^A) linear trends:

$$CR_{cta} = \beta^A (100 \times \ln(GDPpa_{ct})) + \mu_c^A + \tilde{\tau}_t^A + \omega_a^A + \lambda_c^A t + \alpha_a^A t + \tilde{\varepsilon}_{cta}^A$$

On differencing, μ_c^A and ω_a^A drop out, while λ_c^A, $\tau_t^A \equiv \Delta\tilde{\tau}_t^A$, and α_a^A become country, year, and age effects. However, serial correlation, non-stationarity, and the need for PPP adjustment make the level specification unattractive.

[14] An alternative approach would take the mean of observed annual changes g_{ct} and ΔCR_{cta}, but this approach does not use all available information because of gaps in the data. For example, if data were collected only in 1970, 1971, 1990, and 1991, then the mean of the two observed annual changes would ignore developments during 1971-1990. Another alternative approach would take the annualized long difference from 1970 to 1991, but this approach loses precision because both variables are stochastic, and conception rates are measured with sampling error. Figure 3.A2 provides a case study for further intuition.

equation (3.3.3) primarily analyzes variation in ΔCR_{cta} and g_{ct} that was absorbed by the country effect in equation (3.3.2), we omit λ_c^A from the long-run regression. Here, β^A represents the cross-country association of long-run economic growth with long-run fertility change, net of age-related factors. A 1-log point faster annualized rate of long-run growth is associated with a β^A higher annual rate of change in the conception rate.

For both equations (3.3.2) and (3.3.3), we summarize the age group results by reporting the implied result for the total conception rate (TCR) per 1000 women, defined as the expected number of conceptions in a hypothetical cohort of 1000 women who experience current age-specific conception rates at every age from 15 to 44:

$$\beta^{TCR} = 5\left(\sum_A \beta^A\right) \tag{3.3.4}$$

This summary measure is the sum of the age group coefficients, multiplied by 5 to account for the length of each age group. While heterogeneity over the age distribution is key to our investigation, β^{TCR} provides an overall measure of the association between economic growth and fertility change.

To assess the roles of other aggregate variables in explaining any relationship between economic growth and fertility change, we report estimations of equations (3.3.2) and (3.3.3) with and without a main set of covariates. In the extended models, we control for variables available for all country-years in our dataset: the initial levels of population density and log PPP-adjusted GDP per adult, as well as the annual change (for the short-run analysis) or average annual rate of change (for the long-run analysis) in female education, urbanization, infant mortality, and armed conflict. Female education and urbanization are averaged at the country-year-age level, while infant mortality is measured at the country-year level to minimize noise.[15] We consider covariates available only for subsamples later.

[15]To avoid endogeneity concerns, the short-run analysis relies on mortality rates among infants conceived in the previous year. We include an indicator for missing mortality data to accommodate the first cell in any country-age series. Results do not changed if we drop these cells instead.

For weighting and variance estimation, we make conservative choices that clarify interpretation. Sample sizes in individual WFS and DHS surveys range from fewer than 5,000 to more than 100,000 women, suggesting possible efficiency gains to weighting by cells size, but we choose to weight cells equally to ease interpretation of the results. We also cluster standard errors by country, allowing for arbitrary error covariance within country while imposing independence across countries.

3.3.3 Results

Figures 3.2–3.4 present the main results from the analysis of period fertility. To aid interpretation of the age-group-specific regression estimates, Figure 3.2 first summarizes the level of fertility and its annual rate of change across age groups. Average conception rates follow an inverted u-shape in age, peaking at 261 per 1000 among 20-24 year olds. The TCR per 1000 women is 5483, so that a woman experiencing these age-specific conception rates over her lifecycle would expect 5.5 conceptions. Despite this high level, fertility was falling throughout the sample period for all age groups. On average, conception rates in all age groups declined by 1 to 3 points per year, with a -67 annual change in the TCR, corresponding to a decadal reduction of two-thirds of a conception per woman.

These rates of fertility change serve as dependent variables in Figures 3.3–3.4, which present the regression results graphically; for additional reference, Appendix 3.A Table 3.A3 presents them numerically. Figure 3.3 reports the short-run coefficients from equation (3.3.2), with and without covariates. In both models, all short-run coefficients are positive, indicating procyclical fertility, with statistical significance in the prime parenting ages. The inclusion of covariates—the lagged levels of GDP per adult and population, as well as changes in conflict, female education, urbanization, and infant mortality—does not meaningfully alter the coefficients or their significance levels. 25-29 year olds show the largest short-run response to growth fluctuations; their coefficient of 0.56 in the base model implies that a one log point increase in GDP per adult raises the number of conceptions by roughly $\frac{1}{2}$ per 1000 women

in the age group. Moving to neighboring age groups, the coefficients decline more than would be proportional to the level of fertility. This finding is consistent with the theoretical framework's prediction that older parents (who are closer to menopause) are less willing to forego births during a recession, although it is only suggestive evidence. Combining all age groups, the TCR increases by 8.8 per 1000 in response to a log point positive growth fluctuation.

Additional robustness checks appear in Appendix 3.A Figure 3.A3, which plots estimates of β^A across age groups for a range of alternative short-run models. One weights cells by their size; another reweights observations within each cell to give more weight to fertility outcomes with shorter recall periods (using a Bartlett kernel); another omits country, year, and age effects; and three others add country-specific linear, quadratic, and cubic time trends. The alternative weighting schemes and trend specifications he weighted model and the trend models deliver results very similar to those reported in Figure 3.3.[16] However, in the model with no country, year, or age effects, the coefficients at prime parenting ages (20-34)—while still statistically significant—shrink by roughly one-quarter, while the coefficient in the 40-44 age group grows by the same proportion.

That the omission of country, year, and age effects modifies coefficients across the age distribution in different directions is easily reconciled by the analysis of long-run rates of change, where the results are nearly opposite the short-run estimates. As shown in Figure 3.4, long-run economic growth and long-run fertility change are negatively correlated at prime ages but positively correlated at older ages. A comparison of women around age 30 with women in their early 40s provides the starkest contrast. In the base model, among 25-29 and 30-34 year olds, a 1-point faster average annual rise in log GDP per adult is associated with a 0.43-point *faster* average decline in conception rates: roughly equal and opposite in sign from the short-run coefficients in Figure 3.3. Among 40-44 year olds, the same increase in long-run growth is associated with a 0.14-point *slower* average decline in conception rates.

[16]The robustness to downweighting fertility outcomes that occurred long ago suggests that our results are not an artifact of women failing to recall infants who were born alive but died in the neonatal period.

Because the declines are concentrated in the middle of the reproductive period, these results suggest increased spacing, rather than later starting or earlier stopping. In the extended model, the estimates are robust to controlling for the initial levels of GDP per adult and population, as well as the average rates of change in conflict, female education, urbanization, and infant mortality.

On net, the offsetting coefficients at different ages imply a long-run TCR coefficient of -6.1 in the base model, so that overall, faster long-run economic growth is associated with more rapid fertility decline. The average annual rate of change in log GDP per adult has a standard deviation of 1.8, so a one standard deviation increase in long-run growth is associated with fertility declining at a decadal rate of one conception for every nine women. Conditional on the single-year age effects, the R^2 is 0.09 for all age groups pooled and 0.16 for the 25-29 age group, implying that economic growth can account for a meaningful share of fertility change in developing countries over the long run.

Further specification checks are reported in Appendix 3.A Figure 3.A4, which plots age-group-specific coefficients from a range of alternative long-run models. Changing the minimum long-run time horizon from 20 years to 15 or 25 does not change the estimates; nor does reweighting conception rates using a Bartlett kernel. We also obtain similar results when we use the average of observed annual changes instead of the slope of the trend, as well as when we use GDP instead of GDP per adult. This final result confirms that our results are not driven by reverse causality.

In both the short- and long-run analyses, the log-linear specifications may mask theoretically relevant non-linearities. To examine this possibility for the short-run, we discretize the distribution of $100 \times \Delta \ln (GDPpa_{ct})$ into six bins and then run a semi-parametric version of equation (3.3.2) that replaces the continuous variable g_{ct} with bin indicators. Figure 3.5 presents summary estimates for the TCR; age-group-specific estimates appear in Appendix 3.A Figure 3.A5. An asymmetry emerges: conceptions fall sharply in deep recessions but do not rise in rapid expansions. Relative to the base category (0-5 log points), a recession of

more than 10 log points decreases the total conception rate by 171 per 1000 women: nearly one-fifth of a child per women.[17] To shed light on the functional form governing the long-run relationship, we estimate local linear regressions of the average annual rate of change in the conception rate on the average annual rate of rate of change in log GDP per adult. Figure 3.6 reports summary estimates for the TCR; age-group-specific estimates appear in Appendix 3.A Figure 3.A6. The long-run results do not deviate substantially from linearity.[18]

Because women may offset past fertility adjustments, we also estimate a distributed-lag version of equation (3.3.2). Figure 3.7 reports summary estimates for the TCR with four lags; age-group-specific estimates appear in Appendix 3.A Figure 3.A7. Some but not all of the short-run response is offset through subsequent adjustments to childbearing. In response to a 1-point growth fluctuation, TCR rises by 11 at first, falls by 8 in the following year, and then fluctuates by smaller amounts. Summing across lag coefficients, the cumulative effect of a fluctuation on TCR shrinks to 3 one year after the fluctuation but settles at a significant 6-8 children per 1000 women thereafter.[19] While Figure 3.7 identifies interesting dynamics in the fertility response to aggregate fluctuations, two caveats are worthy of note. First, much of the immediate offset reflects the inability of currently or recently pregnant women to conceive. Second, because we study the year-to-year change in the conception rate at a fixed age, rather than the change for a fixed birth cohort, this exercise does not map exactly onto the evolution of fertility over time for a particular woman. The cohort analysis in Section 4 will address both of these issues, shedding light on the lifetime fertility effects of economic fluctuations at particular ages.

[17] As shown in the histogram at the bottom of Figure 3.5, recessions of this magnitude are rare but not unprecedented, with 3% of the sample (1,644 cells) in this category.

[18] The estimated regression functions are negative for all age groups in Appendix 3.A Figure 3.A6, confirming that the long-run results relate primarily to the pace of fertility decline, rather than increase.

[19] In Appendix 3.A Figure 3.A7, most age groups exhibit a similar lag structure to the TCR. An exception is the 30-34 year old age group, which displays weak offset behavior: a pattern relevant for the cohort results in Section 4.

3.3.4 Extensions

This section summarizes a number of extensions that shed additional light on mechanisms underlying our results and their relation to theory. Appendix 3.A reports all results.

Short Run

Heterogeneity Fertility responses to growth shocks may be heterogeneous with respect to both individual and aggregate characteristics. To shed light on heterogeneity within countries, Appendix 3.A Table 3.A4 studies how four average characteristics of mothers change over the business cycle: age, education, urban residence, and ever-marriage.[20] The average education of mothers falls during recessions, implying that poorer, less-educated women are more responsive to growth fluctuations. Other average characteristics do not vary over the business cycle.[21] The marriage null result is surprising, given the institution's historical role in fertility limitation (Malthus 1798; Hajnal 1965; Wrigley 1981), so Appendix 3.A Table 3.A5 confirms that conception rates are procyclical both inside and outside marriage, and that neither the rate nor hazard of first marriage varies significantly with growth fluctuations. Across countries, Appendix 3.A Table 3.A6 finds no significant variation in the short-run TCR coefficient by the lagged levels of GDP per adult, contraceptive prevalence, average education, or urbanization, nor with the female labor force share. At the same time, Appendix 3.A Table 3.A7 finds more procyclicality in Africa and Latin America than in Asia.

Comparison with developed countries How do the fertility responses documented here compare with those in developed countries, the focus of the literature on fertility and the

[20]To minimize changes in sample composition from cells with no births, we run this analysis at the country-year level. We control for changes in the age structure and average characteristics of all women in the country-year cell, so the coefficient on g_{ca} captures how the composition of conceiving mothers varies with growth fluctuations, over and above any association with the composition of women at risk for conceiving.

[21]Table 3.A4 also examines sex ratios, given Trivers and Willard's (1973) hypothesis that male fetuses are more vulnerable to the mother's condition. Growth fluctuations do not affect sex ratios, suggesting little role for fetal death.

business cycle (Sobotka et al. 2011)? Because data on conception rates by year in developed countries are not readily available, Appendix 3.A Table 3.A8 analyzes birth rates from our WFS/DHS microdata and from the Human Fertility Database (HFD), a compilation of natality data from populations with high-quality vital registration systems. Estimates of equation (3.3.2) reveal that developing country fertility is more procyclical in absolute terms but less procyclical in relative terms.[22]

Comparison with standard aggregate data To shed light on the gains from using microdata, Appendix 3.A Table 3.A8 reruns the analysis using total fertility rates from the World Development Indicators (WDI), a popular cross-country dataset. For developed countries, estimates from the WDI are similar, but for developing countries, they are insignificant and close to zero, likely because the WDI's fertility data are overly smooth for countries with low-quality vital registration systems. Researchers using this popular cross-country dataset would have incorrectly concluded that fertility is far more procyclical in richer, lower-fertility countries.

Long Run

Additional covariates Although controlling for trends in conflict, female education, urbanization, and infant mortality did not alter results in Figure 3.4, other relevant covariates were omitted because they were not available for all country-years. For the long-run analysis, however, yearly measurements are less important. Appendix 3.A Table 3.A9 controls for the average annual rate of change in each of four covariates that are not available for the whole sample but may shed light on mechanisms: secondary school enrollment, female labor force participation, the sectoral composition of value added, and the extent of democracy. When we control for average rate of change in secondary enrollment (from WDI), the

[22]These regressions relate changes in birth rates to the weighted average of current and lagged changes in log GDP per adult, assigning weight $\frac{1}{4}$ to the current change and $\frac{3}{4}$ to the lagged change, roughly matching the conception period for the current year's births.

coefficients on average growth rise substantially, and the TCR coefficient shrinks by roughly half.[23] Therefore, trends in contemporaneous secondary school enrollment—which reflects the desirability of schooling, *not* the education of mothers—can partly explain our long-run results. In contrast, we find no evidence that trends in female labor force participation, the sectoral composition of value added, or democratization explain the results.[24]

Heterogeneity Because the long-run analysis is effectively cross-sectional, comparing countries with different long-run growth rates, the small number of countries in our sample limits our ability to explore heterogeneity. One question we can address involves the role of Africa. As Appendix 3.A Figure 3.A8 shows in non-parametric plots, prime-age fertility falls and older-age fertility rises with long-run growth both inside and outside Africa. Still, the full sample results are partly explained by Africa lagging behind the rest of the developing world in both growth and fertility decline.[25]

Comparison with developed countries As a companion to Appendix 3.A Table 3.A8, Appendix 3.A Table 3.A10 compares long-run results for developing and developed countries countries. Here we find a stark difference between the WFS/DHS and the HFD. Faster long-run growth is not associated with greater fertility decline in high-income, low-fertility populations, suggesting that the fertility-reducing substitution effects of long-run growth are stronger during the development process.

[23]In the WDI, gross enrollment ratios are available from many more countries in our sample than net enrollment ratios, so we use the former. A well-known problem with gross enrollment ratios is that they can by biased by grade repetition, and indeed, more than one-third of the country-years in our dataset have primary school ratios in excess of 100. We therefore rely on secondary school enrollment ratios, which never exceed 100 in our sample.

[24]However, the labor force participation rate is measured noisily due to differing definitions across censuses and surveys, and it is a poor proxy for wages, which are more relevant to theories of fertility change (Galor and Weil 1996). We use the rate for women over 15, as assembled by Olivetti (2014) from ILO databases.

[25]If we include an Africa indicator in equation (3.3.3), the coefficients shrink but remain significant for key age groups: from -0.430 [0.113] to -0.227 [0.115] for 25-29 year olds and from 0.136 [0.051] to 0.123 [0.053] for 40-44 year olds.

Comparison with standard aggregate data Appendix 3.ATable 3.A10 further compares our long-run results with those from a standard aggregate fertility dataset. Here, we draw on fertility rates from the UN, which are available only in 5-year intervals (compared to the annual data from the WDI) but are disaggregated into 5-year age groups. Unsurprisingly, interpolation and smoothing in the UN data do not bias the long-run results for the total fertility rate. However, consistent with restrictions on age heterogeneity in demographic models, the UN data perform poorly at the oldest ages: the 40-44 coefficient in developing countries has the wrong sign. Here again, standard cross-country data miss an important nuance in the relationship between growth and fertility change.

3.4 Analysis of Cohort Fertility

If women fully offset short-run responses before the end of childbearing, then the observed procyclicality will not affect lifetime fertility; growth fluctuations will alter the tempo but not the quantum of fertility. Full offset may be more likely for fluctuations early in the lifecycle; younger women may have more time than older women to make up for lost childbearing opportunities. To investigate these issues, this section changes the unit of analysis to the country-cohort cell (women born in the same country and year), relating a cohort's completed fertility to its experience of economic growth over the lifecycle.

3.4.1 Methods

We follow Currie and Schwandt (2014) by regressing the cohort's completed fertility rate on average economic conditions experienced in each age interval A from $[15, 19)$ to $[40, 44)$, a location (in our case, country) fixed effect λ_c, and a cohort fixed effect δ_j:

$$CFR_{cj} = \sum_A \beta^A \overline{g}^A_{cj} + \lambda_c + \delta_j + \varepsilon_{cj} \tag{3.4.1}$$

where $\overline{g}_{cj}^{A}$ is the average annual change in log GDP per adult over age interval A, measured in log points. The β^A coefficients capture how completed fertility responds to within-country, within-cohort differences in economic growth experienced over the lifecycle. The isolation of within-country variation is key; the country fixed effect λ_c absorbs cross-country variation in long-run economic growth, so the β^A coefficients are identified by fluctuations. However, the underlying variation is not the same as that in the short-run period analysis; 5-year growth may reflect deeper business cycle variation with greater liquidity effects.[26] The aggregation increases precision and tractability.

3.4.2 Results

Figure 3.8 graphically displays how a cohort's completed fertility rate relates to its experience of economic growth over the lifecycle; Appendix 3.ATable 3.A11 presents the results numerically. Up to age 30, fluctuations have no relation with lifetime fertility, consistent with full offset of short-run responses. Offset opportunities appear to diminish thereafter, with the results indicating permanent effects of fluctuations in the 30s. Net of the long-run growth rate, a 1 log point increase in the average annual growth rate experienced during 30-34 or 35-39 raises completed fertility by roughly 30-40 children per 1000 women, regardless of whether we count all children ever born or only those that survived until the survey date.[27] Appendix 3.A Table 3.A11 shows that the estimated coefficients shrink slightly but remain significant if we control for average education and share urban.

Notably, the magnitudes in Figure 3.8 exceed what the short-run effects would imply if they were permanent. This accumulating effect appears inconsistent with the offset patterns documented in the short-run impulse response function in Figure 3.7. Research on the US

[26] Along these lines, it would be illuminating to study the sequence of growth over the 5-year interval; alternating positive and negative shocks may have different quantum effects than a series of positive shocks followed by a series of negative shocks. Unfortunately, our dataset is underpowered for an examination of this issue

[27] Because the count of children ever born requires respondents to recall deceased children who were born long ago, recall error may bias it toward the count of surviving children.

has found similar patterns of short-run effects accumulating over the lifecycle (Currie and Schwandt 2014), although the key margin in that context is childlessness, which does not play an important role here.[28] Three points may help explain this puzzle. First, in the short-run model with lags, 30-34 is the age group with the weakest offset pattern. Second, as discussed in Section 3, the immediate offset in Figure 3.7 may mechanically reflect a woman's inability to conceive during and soon after pregnancy, rather than any behavioral response. Third, fertility may respond non-linearly to a sustained and deep recession, which may be better reflected in a 5-year average than a single-year growth measure. Beyond these three points, the 95% confidence intervals contain values consistent with each other.

3.5 Interpretation through the Lens of Unified Growth Theory

Our empirical results demonstrate varied fertility responses to growth at different time horizons and lifecycle stages. How far can the theoretical literature on long-run growth and the demographic transition go in interpreting these findings? That literature's standard overlapping generations model posits parents choosing the quantity (and sometimes quality) of their children in a single period. We explore a simple extension of the standard model, especially emphasizing Galor and Weil (2000), to include a lifecycle with a finite childbearing period and liquidity constraints.[29]

The literature on which we build treats fertility entirely as a choice variable; we do the same, even if some of our findings leave room for biological mechanisms. For instance, although two results—the stability of sex ratios over the business cycle and the robustness of our main findings to controlling for infant mortality—suggest that infecundity from

[28] Childlessness rates are low in our sample cohorts, averaging 4%, and are unrelated to within-country variation in cohort experiences of economic growth.

[29] See also Galor and Weil (2000); Galor and Moav (2002); Hazan and Berdugo (2002); De La Croix and Doepke (2003); Doepke (2004); de La Croix (2013); and Cervellati and Sunde (2015).

malnutrition is not the main cause of procyclical fertility, it may still contribute. The discussion below seeks solely to square our results with existing economic theory, so it ignores biological mechanisms. For tractability, we also follow the literature in assuming continuous rather than binary fertility.

Setup Period utility is separable over parental consumption c_t, the number of children n_t, and their mean human capital h_t:

$$U(c_t, n_t, h_t) = u_c(c_t) + u_n(n_t) + u_h(h_t), \tag{3.5.1}$$

where the sub-utility functions $u_x(\cdot)$ are increasing, concave, and twice continuously differentiable, with $\lim_{x\downarrow 0} u'_x(x) = \infty$ and $\lim_{x\uparrow\infty} u'_x(x) = 0$. Parents live $t = 1, \cdots, T$ periods and maximize expected lifetime utility, discounted by factor β. They start their lives with assets A_0 and then receive stochastic wages w_t (with a period time endowment of 1) and unearned income y_t in subsequent periods.[30] In each period, they allocate assets and potential income to consumption, the quantity and quality costs of children, and savings (at gross return R), but they cannot borrow ($A_t \geq 0$).

Parents set a birth rate $b_t \in [0, 1]$ in each period, starting their lives with no children and accumulating them according to $n_t = n_{t-1} + b_t$ until they reach menopause at age $M < T$. During each year of childhood up to age K, a child costs $\tau \in (0, 1)$ units of time and κ units of the consumption good, plus any education spending e_t to produce human capital. Education spending is transformed into human capital by a twice continuously differentiable human capital production function $h(e_t; \bar{g})$, which also depends on the long-run growth rate of technology $\bar{g}$. Because we are not primarily interested in the allocation of education within the family, we simplify by assuming that parents plan a single education level e for all of their children in period 0 of the model, before the first period of the lifecycle. Then from

[30] Because child costs depend on w_t but not y_t, one can think of these variables as women's and men's wages.

period 1 to period M, parents make a sequence of consumption and birth decisions, followed by a sequence of consumption decisions till death at T.

Optimization In light of the finite horizon, we work backward: characterizing the consumption sequence first, the birth sequence next, and education spending last. As shown in Appendix 3.B, the first-order conditions to the lifecycle problem lead to a consumption Euler equation:

$$u_c'(c_t) = \beta R E_t\left[u_c'(c_{t+1})\right] + \lambda_t, \tag{3.5.2}$$

where λ_t is the Lagrange multiplier on the borrowing constraint. When the borrowing constraint does not bind, parents set the current marginal utility of consumption to the discounted expected marginal utility of consumption in the next period. When it does bind, they fall short of consuming enough in the current period to satisfy this condition, with a positive multiplier filling the gap.

Although the Euler equation does not directly involve fertility, the consumption smoothing motive is key to understanding the timing of births over the business cycle. This point becomes apparent upon inspection of a separate first-order condition, which equates the marginal benefit of consumption with the discounted marginal benefits of childbearing:

$$u_c'(c_t) = \frac{u_n'(n_{t-1} + b_t) + E_t\left[\Upsilon\right] + \mu_t^0 - \mu_t^1}{\tau w_t + \kappa + e}, \tag{3.5.3}$$

where

$$\Upsilon = \sum_{s=t+1}^{T} \beta^{s-t} u_n'(n_s) \frac{\partial n_s}{\partial b_t} - \sum_{s=t+1}^{t+K} \beta^{s-t} \nu_s \left(\tau w_s + \kappa + e\right),$$

and μ_t^0 and μ_t^1 are the multipliers on the constraints that $0 \leq b_t \leq 1$, while ν_t is the multiplier on the period t budget constraint. The numerator reflects the current marginal benefit and future marginal benefits (net of marginal costs) of childbearing; the denominator reflects the current marginal cost.

Assuming an interior solution, the education spending plan satisfies the first order condition:[31]

$$u'_h\left(h\left(e;g\right)\right)h_e\left(e;g\right)\left(\tfrac{1-\beta^T}{1-\beta}\right)=E_0\left[\sum_{t=1}^{M+K}\sum_{k=0}^{K}\beta^{t-1}\nu_t b_{t-k}\right], \tag{3.5.4}$$

The left-hand side reflects the discounted lifetime marginal utility of children's mean human capital. The right-hand side reflects the discounted lifetime marginal cost of the education plan (the number of children in the household) multiplied by the marginal utility of income (the Lagrange multiplier on the period budget constraint).

Mapping to the Data This simple theoretical framework captures the short- and long-run dynamics we document in the data. For the short run, equations (3.5.2)-(3.5.3) provide insight into the effects of transitory wage and income fluctuations. Because the right-hand side of equation (3.5.3) is divided by the marginal cost of childbearing—which includes w_t—a transitory wage cut incentivizes the shifting of births from the future to the present. If T is large relative to M, then serially independent fluctuations in w_t before menopause have minimal effects on expected lifetime income, so this intertemporal substitution effect is likely to dominate any income effect when the borrowing constraint does not bind. In this case, fluctuations in y_t are also unlikely to affect the birth rate. When the constraint binds, however, a transitory depression in wages or incomes decreases current consumption by equation (3.5.2), which then also incentivizes reduced childbearing by equation (3.5.3).[32] Thus, at least with serially independent shocks, a borrowing constraint is key to generating procyclical fertility.[33] The role of liquidity is consistent with the non-linearity and hetero-

[31]For simplicity, we assume that the period income is always high enough to cover educational expenses.

[32]Existing research on fertility over the business cycle has also noted offsetting liquidity and intertemporal substitution effects, albeit generally without fully specifying a model (e.g., Ward and Butz 1980; Adsera and Menendez 2011).

[33]Aguiar and Gopinath (2007) find that developing economies are characterized by permanant, not transitory, growth shocks, which would lead to a mix of short- and long-run mechanisms in the response to shocks. But Figure 3.5 shows that deep recessions to be key to the procyclicality of births; these recessions are autocorrelated only to the second lag.

geneity we find in the data: borrowing constraints may bind more during deep recessions, and the less-skilled may be especially likely to hit zero assets.

The durability of children adds several nuances in equation (3.5.3). The lagged number of children appears inside $u'_n(\cdot)$, placing a ceiling on the marginal utility of children for all $t > 1$. No such ceiling exists for the marginal utility of consumption, implying that births decline to zero more rapidly than consumption when parents are borrowing constrained. Also because the lagged number of children appears inside $u'_n(\cdot)$, parents may offset past adjustments in childbearing. For example, if a negative shock forced borrowing-constrained parents to forego births in period $t-1$, then the marginal gains from childbearing are high in period t. This offset is consistent with our distributed-lag estimates as well as our cohort estimates for fluctuations early in the reproductive period. Offset becomes impossible after menopause, so fluctuations toward the end of the reproductive period are more likely to have permanent effects on the number of children, as we also observe in our cohort estimates. When offset is incomplete ($\frac{\partial n_s}{\partial b_t} > 0$ for $s > t$), the benefit of current childbearing includes the marginal utility of children in the future. As a result, parents approaching menopause may tolerate greater declines in consumption to finance childbearing, a prediction consistent with the decline in the magnitude of procyclicality after age 30.

The insights about fluctuations from equations (3.5.2)-(3.5.3) do not carry to long-run productivity growth. We are agnostic about how long-run growth affects the parents' budget constraint, but as a starting point, it is useful to assume that $\bar{g}$ reflects the expected growth rate *across generations*, as in long-run growth models, and does not affect wages or incomes in the parents' lifetimes. Galor and Weil (2000) assume that education increases human capital at a decreasing rate ($h_e > 0$, $h_{ee} < 0$) while long-run growth depletes human capital but makes education more productive ($h_{\bar{g}} < 0$, $h_{e\bar{g}} > 0$). In this case, higher $\bar{g}$ raises the left-hand side of equation (3.5.3), which leads parents to raise e. The increase in e pushes up the denominator in equation (3.5.3), so the optimal number of children declines. Among liquidity-constrained parents, rising e may also delay or increase spacing between births to

allow accumulation of assets to pay child costs, so both the decline and the delay we observe in the data are consistent with the theory. That controlling for school enrollment makes the long-run coefficients less negative (see Appendix 3.A Table 3.A9) reinforces the hypothesized role of returns to child investment.

If $\bar{g}$ affects parents' budget constraint, then these predictions become less sharp, with the effect of long-run growth on fertility additionally depending on the balance of income effects from higher w_t and y_t and substitution effects from higher w_t. Many economic theories of the demographic transition emphasize the link between long-run growth and rising women's wages (Schultz 1985; Galor and Weil 1996), which in our framework can be seen as a shift from y_t to w_t in the composition of household income. By equation (3.5.3), such a shift incentivizes fertility reduction. But we find suggestive evidence that women's work is not a primary driver of our long-run results. Rising female labor force participation does not explain our long-run coefficients, although rising participation is not the same as the wage (i.e., the opportunity cost of children), which is unavailable for most of our sample. Goldin (1995) points out that female labor force participation is high early in the development process, but wages are low and work is compatible with childcare: for example, work on the family farm close to home. A rising opportunity cost of children may be better reflected in the size of the service sector, which employs women outside the home at higher wages. But the share of services in the economy also fails to explain our long-run coefficients.[34]

One complication for the child investment theory is that it relates long-run economic growth to the the *level* of fertility, while our analyses examine *changes* in fertility. But if parents adapt to the new economic environment slowly, then long-run growth will be associated with gradual fertility decline. In this sense, our long-run results may reflect how sustained growth leads parents to gradually update their expectations regarding the return to human capital investment. Expectations are also key to understanding our short-run results;

[34]However, sectoral composition is not a perfect proxy for the the ease of combining work and family. The latter may improve within the service sector as an economy grows, for example because of more generous parental leave policies, more childcare availability, or changing social norms.

couples avoid conception during recessions presumably due to concerns about liquidity after birth, not at the time of conception.

The discussion has so far ignored the factors that generate Malthusian fertility dynamics in models of long-run growth. Galor and Weil (2000) include a subsistence consumption constraint and also assume $h(0;\bar{g}) > 0$, admitting an optimum with $e = 0$. Both assumptions lead to corner solutions that are useful for understanding regimes in which fertility and population rise with productivity growth in the long run (Lee 1997; Ashraf and Galor 2011). Our long-run results do not fit this characterization, so we have focused on a regime in which subsistence and education spending constraints do not bind: the era of demographic transition and sustained economic growth.

3.6 Conclusion

Over the last half-century, the developing world saw rising living standards and falling fertility, but empirical evidence on the link between the two is surprisingly sparse. Combining hundreds of survey datasets, this chapter sheds new light on growth-fertility relationships, with careful attention to time horizons and lifecycle dynamics. Three main empirical results emerge. First, fertility is procyclical in the short run, falling during recessions. Second, fertility declines and delays with long-run economic growth. Third, across birth cohorts within a country, higher economic growth late in the reproductive period predicts higher completed fertility. These results are broadly consistent with an extension of long-run growth models with endogenous fertility to include a lifecycle with liquidity constraints.

The short-run procyclicality is consistent with evidence on fertility responses to economic fluctuations both in historical, pre-industrial populations and in contemporary, industrialized populations (Lee 1997; Sobotka et al. 2011). Distributed lag models and cohort analyses suggest that economic fluctuations affect both the tempo (timing) and quantum (lifetime cumulation) of fertility. The weight of the evidence suggests a role for liquidity constraints,

but beyond this implication, the mechanism behind procyclicality is not clear. One possibility is that couples take intentional steps to reduce conception risk during recessions, by using modern contraception or traditional birth control strategies like withdrawal, the rhythm method, or abstinence. Because we only measure conceptions that resulted in live birth, abortion may play a role. But mechanisms beyond conscious choice may also be at work. Stress may decrease coital frequency among cohabiting couples, and migration for labor market opportunities may temporarily split couples (Timaeus and Graham 1989). Crisis-related malnutrition may also reduce fecundity or *in utero* survival (Bongaarts 1980), although we fail to find evidence of this mechanism in our tests.

The short-run patterns stand in stark contrast to the relationship between long-run trends in income and fertility, which is negative on average but heterogeneous across age groups. The main takeaway is that some force that accompanies long-run economic growth leads to faster declines in childbearing, as reflected in the negative long-run coefficient for the total conception rate, and also to increased birth spacing, as reflected in the positive long-run coefficient for 40-44 years olds. This force is related to rising secondary school enrollment, but *not* declining child mortality, rising adult female education or labor force participation, structural transformation, or democratization. Theories positing that long-run economic growth raises the return to child investment (Galor and Weil 2000) may therefore go a long way in explaining the long-run results.

As for Malthus (1798), his theory performs poorly for the episode we study. Fertility declined with sustained productivity growth, reinforcing a rise in living standards and therefore contradicting the core of his argument. The procyclicality of fertility might appear more supportive, but Malthus never considered liquidity, which is likely a key force behind it. He was perhaps more prescient in viewing fertility control as a practice of "civilized nations" (see Chapter 4), which one could charitably (though not exclusively) interpret as positing that it becomes more prevalent with economic development.

While our results help clarify the relationship between aggregate income growth and fertility change in developing countries, they raise interesting questions about mechanisms and about how fertility's relation to economic growth varies with the underlying source of that growth. They also leave open the question of whether and how fertility affects growth, which has long concerned researchers and policymakers (Coale and Hoover 1958); recent findings suggest such effects are real but modest in size (Ashraf et al. 2013; Miller and Babiarz 2016).[35] Methodologically, they highlight the importance of careful measurement, showing how one can use large amounts of retrospective survey data to improve on standard cross-country datasets.

References

Aaronson, Daniel, Rajeev Dehejia, Andrew Jordan, Cristian Pop-Eleches, Cyrus Samii, and Karl Schulze, "The Effect of Fertility on Mothers' Labor Supply over the Last Two Centuries," NBER Working Paper 23717 2017.

Adsera, Alicia and Alicia Menendez, "Fertility Changes in Latin America in Periods of Economic Uncertainty," *Population Studies*, 2011, *65* (1), 37–56.

Aguiar, Mark and Gita Gopinath, "Emerging Market Business Cycles: The Cycle is the Trend," *Journal of Political Economy*, 2007, *115* (1), 69–102.

Ashraf, Quamrul and Oded Galor, "Dynamics and Stagnation in the Malthusian Epoch," *American Economic Review*, 2011, *101* (5), 2003–2041.

Ashraf, Quamrul H., David N. Weil, and Joshua Wilde, "The Effect of Fertility Reduction on Economic Growth," *Population and Development Review*, 2013, *39* (1), 97–130.

[35]The effect of fertility on economic growth is beyond the scope of this chapter, so we ruled out first-order effects by studying growth in GDP per adult rather than GDP per capita. We also obtained similar results when we studied growth in GDP, without dividing by population.

Baird, Sarah, Jed Friedman, and Norbert Schady, "Aggregate Income Shocks and Infant Mortality in the Developing World," *Review of Economics and Statistics*, 2011, *93* (3), 847–856.

Bean, Lee, Geraldine Mineau, and Douglas Anderton, *Fertility Change on the American Frontier: Adaptation and Innovation*, Berkeley: University of California Press, 1990.

Black, Dan, Natalia Kolesnikova, Seth Sanders, and Lowell Taylor, "Are Children Normal?," *Review of Economics and Statistics*, 2013, *95* (1), 21–33.

Bongaarts, John, "Does Malnutrition affect Fecundity? A Summary of Evidence," *Science*, 1980, *208* (4444), 564–569.

_ **and Susan Cotts Watkins**, "Social Interactions and Contemporary Fertility Transitions," *Population and Development Review*, 1996, pp. 639–682.

Brueckner, Markus and Hannes Schwandt, "Income and Population Growth," *Economic Journal*, 2015, *125* (589), 1653–1676.

Casterline, John B. and Colin Odden, "Trends in Inter-Birth Intervals in Developing Countries 1965–2014," *Population and Development Review*, 2016, *42* (2), 173–194.

Cervellati, Matteo and Uwe Sunde, "The Economic and Demographic Transition, Mortality, and Comparative Development," *American Economic Journal: Macroeconomics*, 2015, *7* (3), 189–225.

Coale, Ansley J. and Edgar M. Hoover, *Population Growth and Economic Development*, Princeton, NJ: Princeton University Press, 1958.

Connelly, Matthew, *Fatal Misconception: The Struggle to Control World Population*, Cambridge, MA: Belknap Press, 2008.

Croix, David De La and Matthias Doepke, "Inequality and Growth: Why Differential Fertility Matters," *American Economic Review*, 2003, *93* (4), 1091–1113.

Currie, Janet and Hannes Schwandt, "Short-and Long-term Effects of Unemployment on Fertility," *Proceedings of the National Academy of Sciences*, 2014, *111* (41), 14734–14739.

Cutler, David M., Wei Huang, and Adriana Lleras-Muney, "Economic Conditions and Mortality: Evidence from 200 Years of Data," NBER Working Paper 22690 2016.

de La Croix, David, *Fertility, education, growth, and sustainability*, Cambridge: Cambridge University Press, 2013.

Deaton, Angus, "Global Patterns of income and Health: Facts, interpretations, and Policies," *UN WIDER Annual Lecture 10*, 2007.

Dehejia, Rajeev and Adriana Lleras-Muney, "Booms, Busts, and Babies' Health," *Quarterly Journal of Economics*, 2004, *119* (3), 1091–1130.

Dettling, Lisa J and Melissa S Kearney, "House Prices and Birth Rates: The Impact of the Real Estate Market on the Decision to Have a Baby," *Journal of Public Economics*, 2014, *110*, 82–100.

Doepke, Matthias, "Accounting for fertility decline during the transition to growth," *Journal of Economic growth*, 2004, *9* (3), 347–383.

Galloway, Patrick R., "Basic Patterns in Annual Variations in Fertility, Nuptiality, Mortality, and Prices in Pre-Industrial Europe," *Population Studies*, 1988, *42* (2), 275–303.

Galor, Oded, *Unified Growth Theory*, Princeton University Press, 2011.

_ **and David N. Weil**, "The Gender Gap, Fertility, and Growth," *American Economic Review*, 1996, *86* (3), 374–387.

_ **and** _ , "Population, Technology, and Growth: From Malthusian Stagnation to the Demographic Transition and Beyond," *American Economic Review*, 2000, *90* (4), 806–828.

_ **and Omer Moav**, "Natural Selection and the Origin of Economic Growth," *Quarterly Journal of Economics*, 2002, *117* (4), 1133–1191.

Goldin, Claudia, "The U-Shaped Female Labor Force Function in Economic Development and Economic History," in T. P. Schultz, ed., *Investment in Women's Human Capital and Economic Development*, Chicago: University of Chicago Press, 1995, pp. pp. 61–90.

Hajnal, John, "European Marriage Patterns in Perspective," in D.V. Glass and D.E.C. Eversley, eds., *Population in History: Essays in Historical Demography*, London: Edward Arnold, 1965, pp. pp. 101–143.

Hansen, Gary D and Edward C Prescott, "Malthus to Solow," *American Economic Review*, 2002, *92* (4), 1205–1217.

Hazan, Moshe and Binyamin Berdugo, "Child Labour, Fertility, and Economic Growth," *Economic Journal*, 2002, *112* (482), 810–828.

Herzer, Dierk, Holger Strulik, and Sebastian Vollmer, "The long-run determinants of fertility: one century of demographic change 1900–1999," *Journal of Economic Growth*, 2012, pp. 1–29.

Johnson, Simon, William Larson, Chris Papageorgiou, and Arvind Subramanian, "Is Newer Better? Penn World Table Revisions and Their Impact on Growth Estimates," *Journal of Monetary Economics*, 2013, *60* (2), 255–274.

Kleven, Henrik and Camille Landais, "Gender Inequality and Economic Development: Fertility, Education, and Norms," 2016.

Knodel, John, "Starting, Stopping, and Spacing During the Early Stages of Fertility Transition: The Experience of German Village Populations in the 18th and 19th Centuries," *Demography*, 1987, *24* (2), 143–162.

Lagakos, David, Ben Moll, Tommaso Porzio, Nancy Qian, and Todd Schoellman, "Life-Cycle Wage Growth Across Countries," *Journal of Political Economy*, 2017.

Lee, Ronald D., "Population Dynamics: Equilibrium, Disequilibrium, and Consequences of Fluctuations," in Mark R. Rosensweig and Oded Stark, eds., *Handbook of Population and Family Economics*, Amsterdam: Elsevier, 1997, pp. pp. 1063–1115.

Lindstrom, David P. and Betemariam Berhanu, "The Impact of War, Famine, and Economic Decline on Marital Fertility in Ethiopia," *Demography*, 1999, *36* (2), 247–261.

Lovenheim, Michael F. and Kevin J. Mumford, "Do Family Wealth Shocks Affect Fertility Choices? Evidence from the Housing Market," *Review of Economics and Statistics*, 2013, *95* (2), 464–475.

Malthus, Thomas Robert, *An Essay on the Principle of Population*, J. Johnson, 1798.

Mensch, Barbara, Susheela Singh, and John Casterline, "Trends in the Timing of First Marriage Among Men and Women in the Developing World," in Cynthiaa Lloyd, Jere Behrman, Nelly Stromquist, and Barney Cohen, eds., *The Changing Transitions to Adulthood in Developing Sountries: Selected Studies*, Washington: National Academies Press, 2005, pp. pp. 118–71.

Miller, Grant and B Piedad Urdinola, "Cyclicality, Mortality, and the Value of Time: The Case of Coffee Price Fluctuations and Child Survival in Colombia," *Journal of Political Economy*, 2010, *118* (1), 113–155.

_ **and Kimberly Singer Babiarz**, "Family Planning Program Effects: Evidence from Microdata," *Population and Development Review*, 2016, *42* (1), 7–26.

Murtin, Fabrice, "Long-Term Determinants of the Demographic Transition, 1870–2000," *Review of Economics and Statistics*, 2013, *95* (2), 617–631.

National Research Council, *Demographic Effects of Economic Reversals in Sub-Saharan Africa*, Washington: National Academies, 1993.

Olivetti, Claudia, "The Female Labor Force and Long-Run Development: The American Experience in Comparative Perspective," in Leah Platt Boustan, Carola Frydman, and Robert A. Margo, eds., *Human Capital in History: The American Record*, Chicago: University of Chicago Press, 2014, pp. pp. 161–197.

Schoumaker, Bruno, *Quality and Consistency of DHS Fertility Estimates, 1990-2012*, Rockville, Maryland: ICF International, 2014.

Schultz, T. Paul, "Changing World Prices, Women's Wages, and the Fertility Transition: Sweden, 1860-1910," *Journal of Political Economy*, 1985, *93* (6), 1126–1154.

_ , "Demand for Children in Low Income Countries," in Mark R. Rosenzweig and Oded Stark, eds., *Handbook of Population and Family Economics*, Vol. 1, Elsevier, 1997, pp. pp. 275–347.

Sobotka, Tomáš, Vegard Skirbekk, and Dimiter Philipov, "Economic Recession and Fertility in the Developed World," *Population and Development Review*, 2011, *37* (2), 267–306.

Tapinos, Georges Photios, Andrew Mason, and Jorge Bravo, *Demographic Responses to Economic Adjustment in Latin America*, Oxford University Press, 1997.

Timaeus, Ian and Wendy Graham, "Labor Circulation, Marriage, and Fertility in Southern Africa.," in Ron Lesthaeghe, ed., *Reproduction and Social Organization in Sub-Saharan Africa*, Berkeley: University of California Press., 1989.

Trivers, Robert L. and Dan E. Willard, "Natural Selection of Parental Ability to Vary the Sex Ratio of Offspring," *Science*, 1973, *179* (4068), 90–92.

United Nations, *World Population Prospects: The 2015 Revision, Methodology of the United Nations Population Estimates and Projections*, New York: UN, 2015.

Ward, Michael and William Butz, "Completed Fertility and its Timing," *Journal of Political Economy*, 1980, *88* (5), 917–940.

World Bank, "Country Classification: A Short History," `http://go.worldbank.org/U9BK7IA1JO` 2015.

Wrigley, Edward A, *Marriage, Fertility, And Population Growth In Eighteenth-Century England*, New York: St. Martins Press., 1981.

Young, Alwyn, "Inequality, the Urban-Rural Gap and Migration," *Quarterly Journal of Economics*, 2013, *128* (4), 1727 – 1785.

Tables and Figures

Table 3.2: Summary Statistics

	Period analysis		Cohort
	Full sample	20-year sample	analysis
	(1)	(2)	(3)
Fertility rates (per 1000)			
Age-specific conception rate	191 (95)	199 (91)	
Completed fertility rate, ever-born			5951 (1284)
Completed fertility rate, surviving			4862 (870)
Macroeconomic conditions			
Country GDPpa, PPP	4711 (4218)	4476 (4087)	4229 (3292)
Change in log GDPpa, log pts.	1.0 (6.1)	1.0 (5.8)	0.7 (1.6)
Cell characteristics			
Average years of education	4.3 (3.1)	4.0 (2.9)	3.6 (2.4)
Percent urban at survey	40 (21)	39 (21)	38 (21)
Number of women	2,374,019	2,280,022	242,886
Number of cells	67,050	58,988	935
Number of countries	76	65	62

Notes: Period sample consists of country-year-age cells; cohort sample consists of country-cohort cells. Ages and cohorts are categorized in single years. "Conception rate" only includes conceptions that resulted in live birth. "GDPpa" is gross domestic product per adult age 15-64. For the cohort sample, macroeconomic conditions are first averaged over each cohort's reproductive lifecycle and then summarized across cohorts.

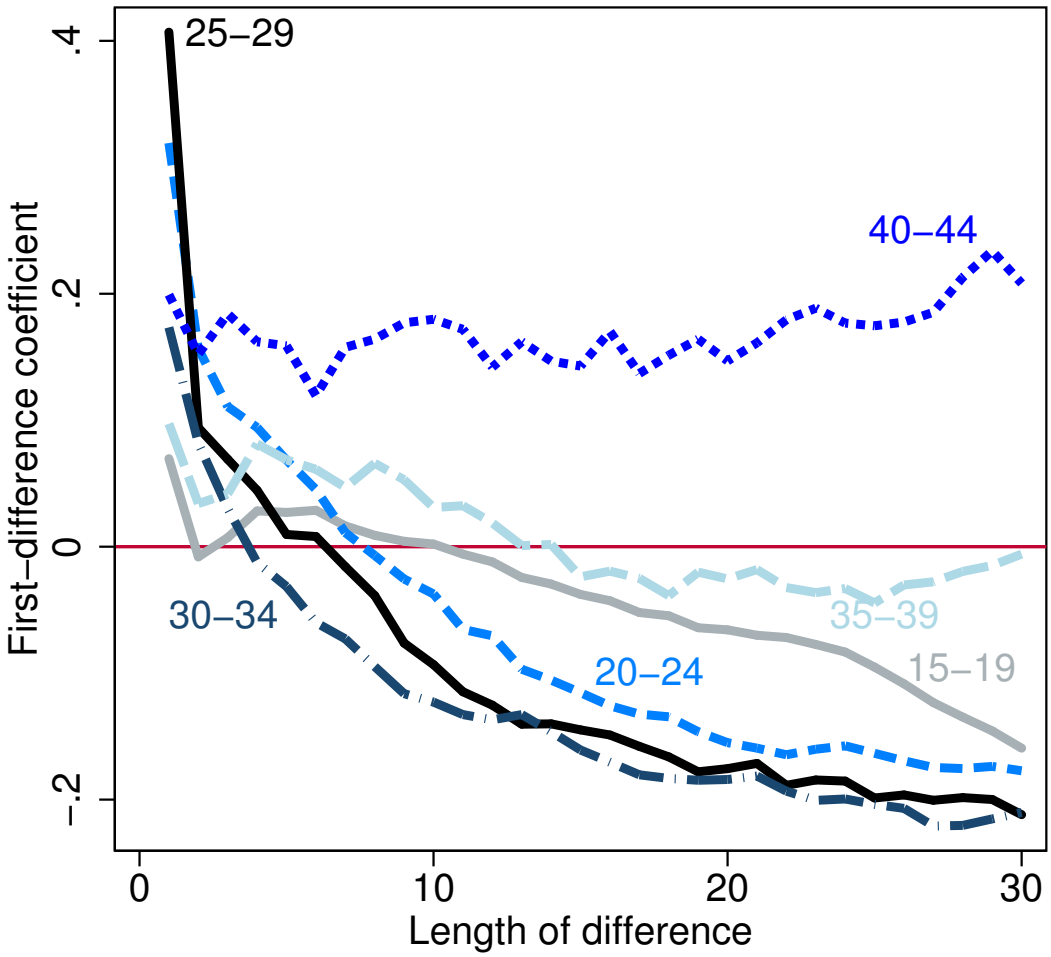


Figure 3.1: Economic Growth and Fertility Change over Varying Time Horizons
Notes: For each 5-year age group, the figure plots coefficients from regressions of the change in the conception rate from year t – Δ to year t on the change in 100 × log GDP per adult over the same period, controlling for single-year age indicators. Separate regressions were run for each integer value of Δ from 1 to 30.

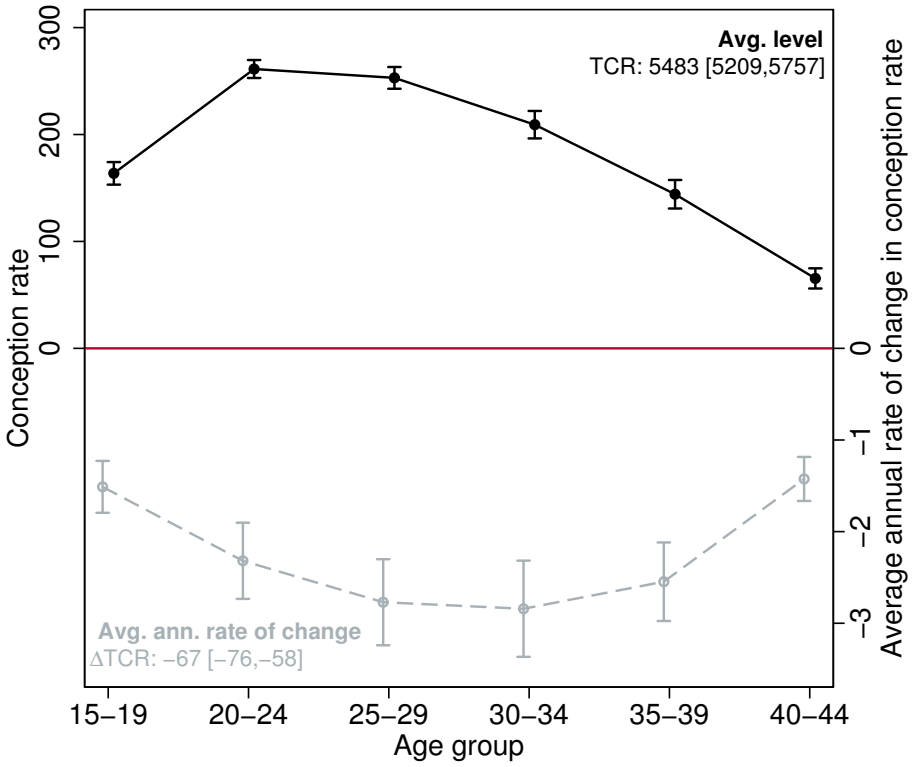


Figure 3.2: Age-Specific Conception Rate Levels and Changes

Notes: Means and 95% confidence intervals for the level of the age-specific conception rate (across 58,992 country-year-age cells) and its annual rate of change (across 1,595 country-age cells). "TCR" refers to the total conception rate per 1000; estimates equal 5 times the sum of age-group-specific estimates. Confidence intervals reflect standard errors clustered by country.

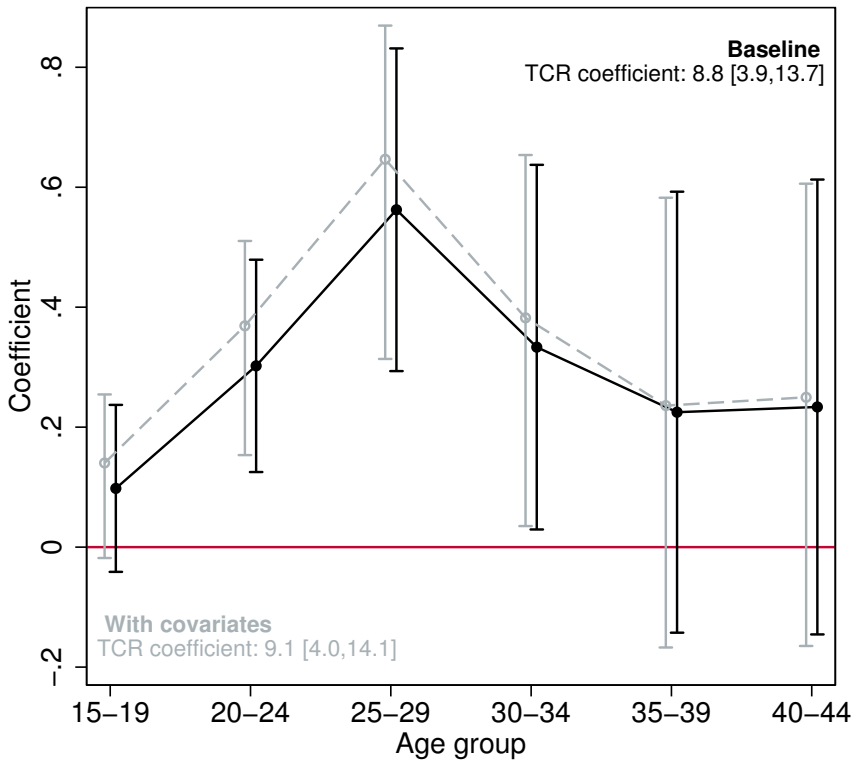


Figure 3.3: Short-Run Estimates

Notes: Coefficients and 95% confidence intervals from regressions of annual changes in the age-specific conception rate on annual changes in 100 × log GDP per adult, controlling for country, year, and single-year age indicators. In the gray plot, the model also controls for the lagged levels of GDP per adult (PPP) and population density; changes in conflict, female education, urbanization, and infant mortality; and an indicator for missing mortality information (less than 3% of all cells). "TCR" refers to the total conception rate per 1000; estimates equal 5 times the sum of age-group-specific estimates. Confidence intervals reflect standard errors clustered by country.

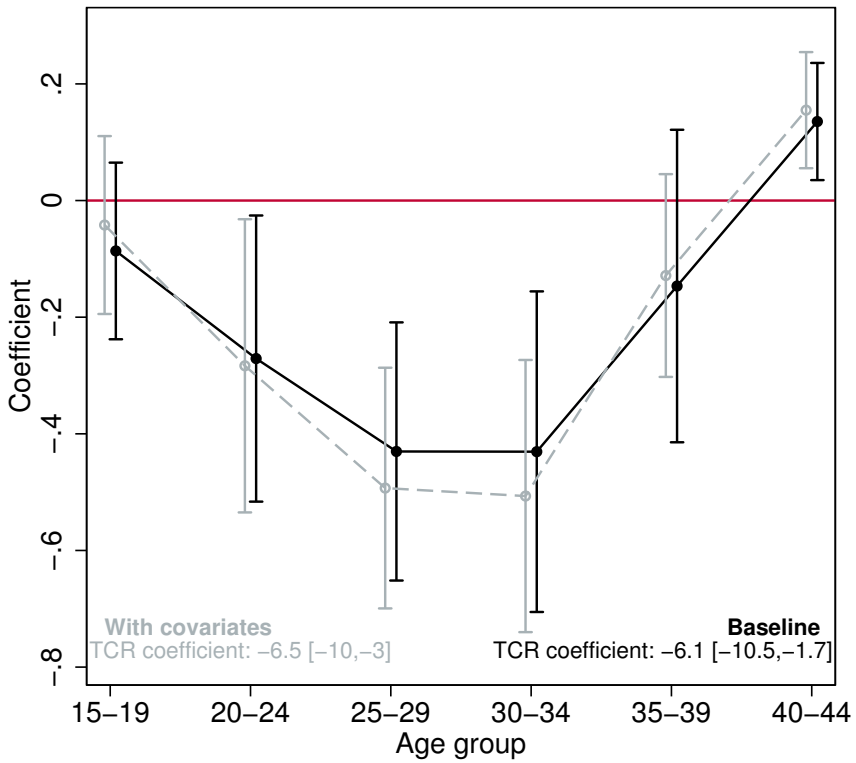


Figure 3.4: Long-Run Estimates

Notes: Coefficients and 95% confidence intervals from regressions of the average annual rate of change in the age-specific conception rate on average annual rate of change in 100 × log GDP per adult, controlling for single-year age indicators. In the gray plot, the model also controls for the lagged levels of GDP per adult (PPP) and population density, as well as annual rates of change in conflict, female education, urbanization, and infant mortality. "TCR" refers to the total conception rate per 1000; estimates equal 5 times the sum of age-group-specific estimates. Confidence intervals reflect standard errors clustered by country.

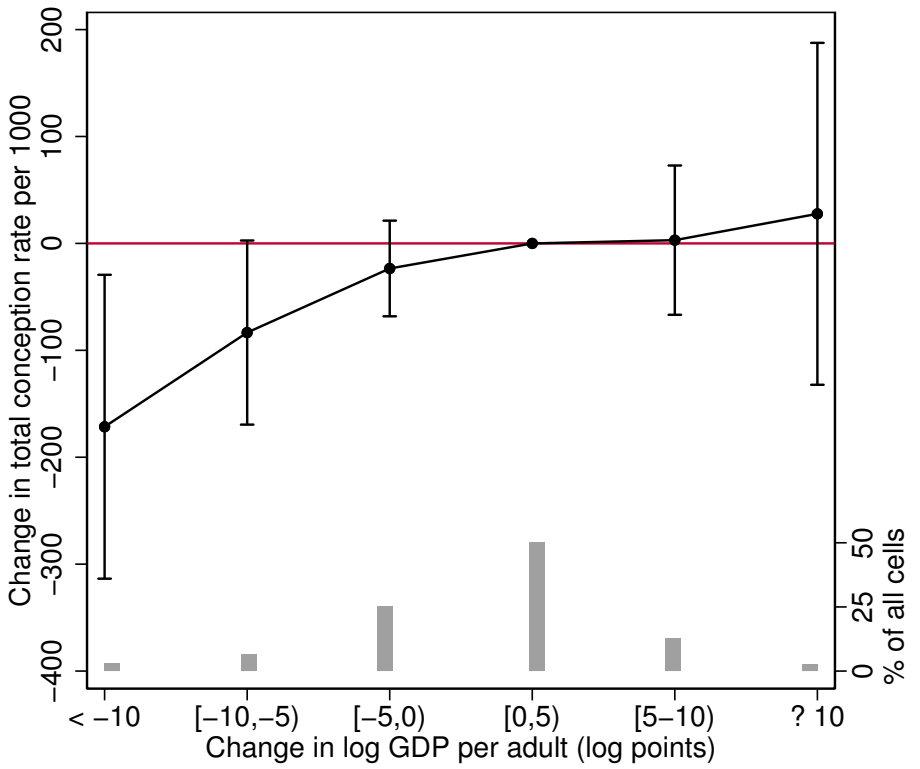


Figure 3.5: Non-Linear Short-Run Estimates

Notes: Total conception rate estimates and 95% confidence intervals based on age-group-specific regressions of annual changes in the age-specific conception rate on binned annual economic growth, controlling for country, year, and age fixed effects. Each estimate reflects the sum of the age-group-specific estimates, multiplied by 5. Omitted category is [0,5). Confidence intervals reflect standard errors clustered by country. For reference, a histogram of the binned growth variable appears at the bottom.

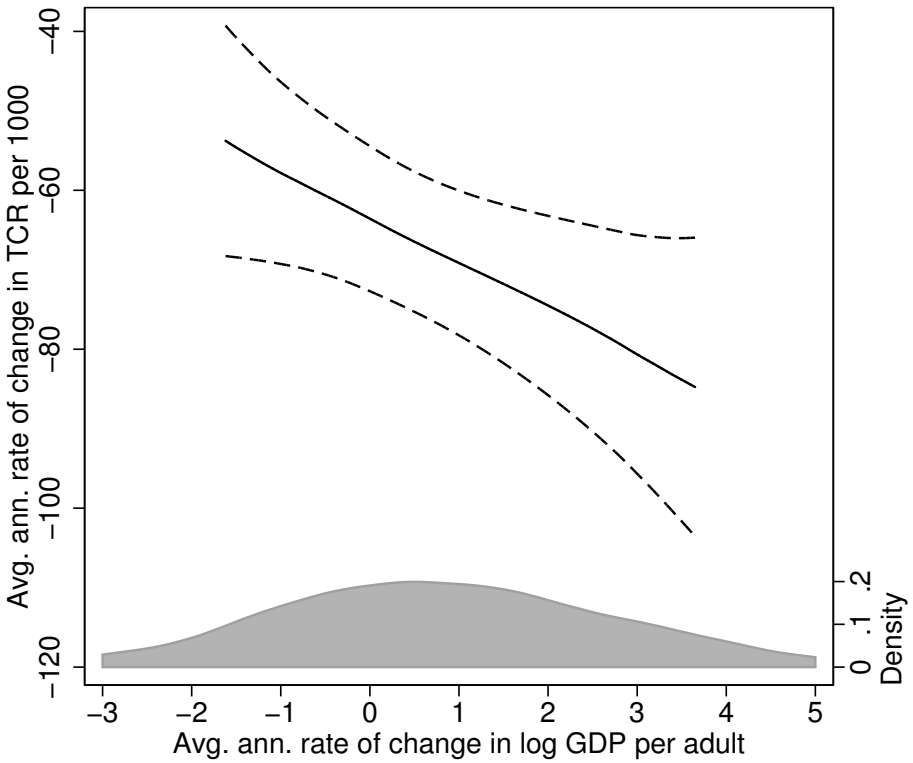


Figure 3.6: Non-Linear Long-Run Estimates

Notes: Total conception rate estimates and 95% confidence intervals based on age-group-specific local linear regressions, bandwidth = 2. The domain of each local linear regression runs from the age group's 5th to 95th percentile of the average annual rate of change in log GDP per adult. The regression function plotted is obtained by summing the age-group-specific estimates (for the domain in which they overlap) and multiplying by 5. Confidence intervals are based on standard errors block-bootstrapped by country. For reference, a kernel estimate (bandwidth = 1) of the density of long-run growth appears at the bottom.

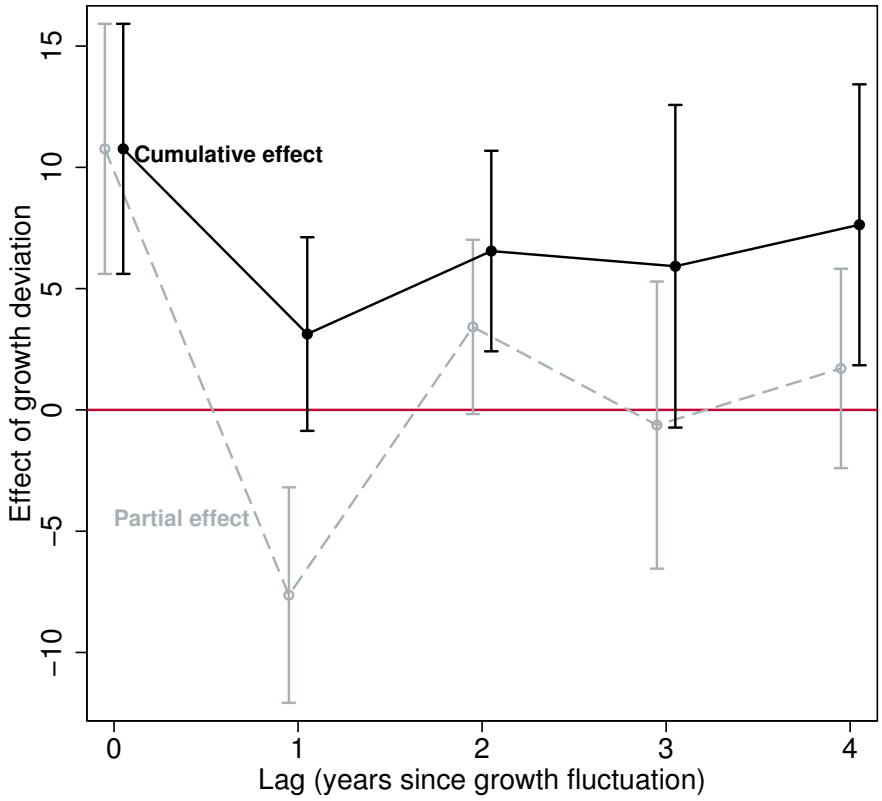


Figure 3.7: Distributed Lag Model

Notes: Total conception rate estimates and 95% confidence intervals based on an age-group-specific distributed lag model of annual changes in the age-specific conception rate on current and lagged annual changes in 100 × log GDP per adult, controlling for country, year, and age fixed effects. Confidence intervals based on standard errors clustered at the country level. Sample includes observations that have both lagged conception rates and lagged growth rates.

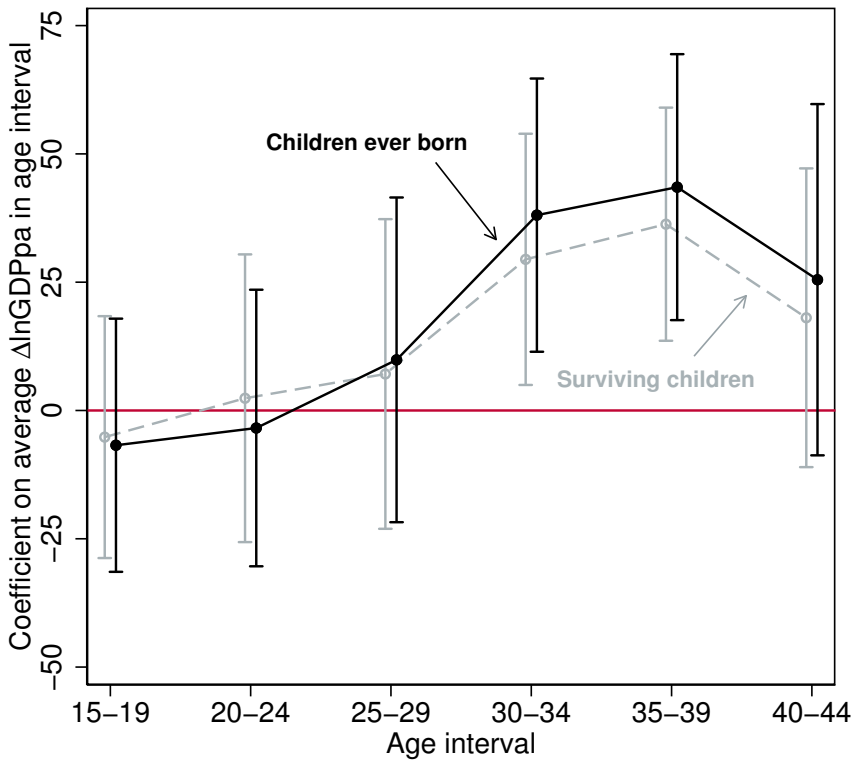


Figure 3.8: Economic Growth over the Lifecycle and Completed Fertility
Notes: Coefficients and 95% confidence intervals from regressions of the number of children ever born (black) or the number of surviving children (grey) per 1000 women on average annual growth conditions during each age interval, controlling for country and cohort indicators. Cohorts were over age 45 when surveyed.

Appendix

3.A Supplementary Tables and Figures

Table 3.A1: Number of WFS/DHS Surveys per Country

Albania	1	Ghana	7	Pakistan	4
Armenia	3	Guatemala	2	Panama	1
Azerbaijan	1	Guinea	3	Paraguay	2
Bangladesh	8	Honduras	2	Peru	9
Benin	5	India	3	Philippines	5
Bolivia	5	Indonesia	8	Rwanda	5
Brazil	2	Jamaica	1	Sao Tome and Principe	1
Burkina Faso	4	Jordan	5	Senegal	8
Burundi	2	Kazakhstan	2	Sierra Leone	1
Cambodia	4	Kenya	7	South Africa	1
Cameroon	5	Korea, Rep.	1	Sri Lanka	1
Central African Republic	1	Kyrgyz Republic	2	Swaziland	1
Chad	2	Lesotho	4	Syria	1
Colombia	7	Liberia	3	Tajikistan	1
Comoros	2	Madagascar	4	Tanzania	5
Congo, Dem. Rep.	2	Malawi	4	Thailand	1
Congo, Rep.	2	Maldives	1	Togo	3
Costa Rica	1	Mali	3	Trinidad and Tobago	2
Cote d'Ivoire	4	Mauritania	1	Tunisia	2
Dominican Republic	8	Mexico	2	Turkey	4
Ecuador	2	Moldova	1	Uganda	5
Egypt	8	Morocco	4	Ukraine	1
El Salvador	1	Mozambique	3	Uzbekistan	1
Ethiopia	3	Namibia	4	Venezuela	1
Fiji	1	Nepal	5	Vietnam	1
Gabon	2	Niger	4	Zambia	5
Gambia	1	Nigeria	5	Zimbabwe	5

Table 3.A2: Country Composition of the Analysis Samples

Country	Period		Cohort	Country	Period		Cohort
	Full sample	Analysis sample			Full sample	Analysis sample	
Albania	✓	✓	✓	Madagascar	✓	✓	✓
Armenia	✓			Malawi	✓	✓	✓
Azerbaijan	✓			Maldives	✓		✓
Bangladesh	✓	✓	✓	Mali	✓	✓	✓
Benin	✓	✓	✓	Mauritania	✓		
Bolivia	✓	✓	✓	Mexico	✓	✓	✓
Brazil	✓	✓	✓	Moldova	✓		
Burkina Faso	✓	✓	✓	Morocco	✓	✓	✓
Burundi	✓	✓	✓	Mozambique	✓	✓	✓
Cambodia	✓	✓	✓	Namibia	✓	✓	✓
Cameroon	✓	✓	✓	Nepal	✓	✓	✓
Central African Rep..	✓	✓	✓	Niger	✓	✓	✓
Chad	✓	✓	✓	Nigeria	✓	✓	✓
Colombia	✓	✓	✓	Pakistan	✓	✓	✓
Comoros	✓	✓	✓	Panama	✓	✓	
Congo, Dem. Rep.	✓	✓	✓	Paraguay	✓	✓	✓
Congo, Rep.	✓	✓	✓	Peru	✓	✓	✓
Costa Rica	✓	✓		Philippines	✓	✓	✓
Cote d'Ivoire	✓	✓		Rwanda	✓	✓	✓
Dominican Republic	✓	✓	✓	Sao Tome and Principe	✓	✓	✓
Ecuador	✓	✓	✓	Senegal	✓	✓	✓
Egypt	✓	✓	✓	Sierra Leone	✓	✓	✓
El Salvador			✓	South Africa	✓	✓	✓
Ethiopia	✓	✓	✓	Sri Lanka			✓
Fiji	✓			Swaziland	✓	✓	✓
Gabon	✓	✓	✓	Syria	✓		
Gambia	✓	✓		Tajikistan	✓	✓	
Ghana	✓	✓	✓	Tanzania	✓	✓	✓
Guatemala	✓	✓	✓	Thailand			✓
Guinea	✓	✓	✓	Togo	✓	✓	✓
Honduras	✓	✓	✓	Trinidad and Tobago	✓	✓	✓
India	✓	✓	✓	Tunisia	✓	✓	
Indonesia	✓	✓	✓	Turkey	✓	✓	✓
Jamaica	✓	✓		Uganda	✓	✓	✓
Jordan			✓	Ukraine	✓		
Kazakhstan	✓			Uzbekistan	✓		
Kenya	✓	✓	✓	Venezuela	✓	✓	
Korea, Rep.	✓			Vietnam			✓
Kyrgyz Republic	✓	✓		Zambia	✓	✓	✓
Lesotho	✓	✓	✓	Zimbabwe	✓	✓	✓
Liberia	✓	✓	✓				

Table 3.A3: Economic Growth and Conception Rates in the Short- and Long-Run

	Mean of conception rate per 1000 in...		Short run regressions		Long run regressions	
	Levels	Changes	Basic	Extended	Basic	Extended
	(1)	(2)	(3)	(4)	(5)	(6)
Ages 15-19	164	-1.5	0.10	0.12	-0.09	-0.03
			[0.07]	[0.07]	[0.09]	[0.08]
Ages 20-24	261	-2.3	0.30	0.33	-0.27	-0.25
			[0.09]	[0.09]	[0.13]	[0.13]
Ages 25-29	253	-2.8	0.56	0.59	-0.43	-0.46
			[0.14]	[0.14]	[0.11]	[0.11]
Ages 30-34	209	-2.8	0.33	0.34	-0.43	-0.48
			[0.16]	[0.16]	[0.14]	[0.12]
Ages 35-39	144	-2.5	0.22	0.21	-0.15	-0.11
			[0.19]	[0.19]	[0.14]	[0.09]
Ages 40-44	65	-1.4	0.23	0.22	0.14	0.16
			[0.19]	[0.20]	[0.05]	[0.05]
TCR	5483	-67	8.77	9.07	-6.15	-5.82
			[2.50]	[2.59]	[2.23]	[1.81]
# cells	58,992	1,595	56,926	56,926	1,595	1,595

Notes: Point estimates and standard errors associated with Figures 2-4. Columns (3)-(4) regress the annual change in the age-specific conception rate on the annual change in 100 $\times$ log GDP per adult, controlling for country, year, and single-year age effects; columns (5)-(6) regress the average annual rate of change in the age-specific conception rate on average annual rate of economic growth, controlling for single-year age effects. "Extended" models also control for the initial level of GDP per adult (PPP) and population density; and the change or trend in female education, urbanization, infant mortality, and conflict. Column (4) also an indicator for missing mortality information (3% of all cells). "Conception rate" only includes conceptions that resulted in live birth; "TCR" refers to the total conception rate per 1000; estimates equal 5 times the sum of age-group-specific estimates. Brackets contain standard errors clustered by country.

Table 3.A4: Cyclicality in the Composition of Births

		Average characteristics of...				
		Mothers				Children
	Concep. rate	Age	Education	% urban	% ever mar.	% male
	(1)	(2)	(3)	(4)	(5)	(6)
Δ log GDPpa	0.24	-0.0001	-0.0022	-0.013	-0.009	0.012
× 100	[0.07]	[0.001]	[0.0009]	[0.013]	[0.014]	[0.016]
Outcome mean	201	23	3.5	35	92	51
Outcome SD	(52)	(3)	(2.5)	(19)	(10)	(3)
# cells	2,831	2,831	2,831	2,831	2,831	2,831

Notes: Regressions of annual changes in average characteristics on annual changes in 100 × log GDP per adult, controlling for country and year fixed effects, as well as changes in age composition, average years of education, percent urban, and percent married among all women in each cell. "Conception rate" only includes conceptions that resulted in live birth. Brackets contain standard errors clustered by country.

Table 3.A5: Can Marriage Explain the Procyclicality of Conceptions?

	Conception rate			Marriage	
	Overall	Pre-marital	Post-marital	Rate	Hazard
	(1)	(2)	(3)	(4)	(5)
Δ log GDPpa	0.23	0.14	0.33	-0.05	-0.13
× 100	[0.07]	[0.09]	[0.16]	[0.06]	(0.12)
Outcome level mean	201	97	270	52	195
Outcome level SD	(52)	(36)	(78)	(23)	(73)
Number of cells	2831	2830	2831	2831	2830

Notes: Regressions of the changes in outcomes on annual changes in 100 × log GDP per adult, controlling for country and year fixed effects, as well as changes in the age composition of each cell. All rates are per 1000. Columns (2) and (5) have smaller sample sizes because 1 cell has no never-married women. Brackets contain standard errors clustered by country.

Table 3.A6: Aggregate Heterogeneity in the Procyclicality of Conceptions

	Lagged GDPpa, PPP	Lagged contraceptive prevalence	Female labor force share in 1990	Lagged average years of education	Lagged share urban at survey
	(1)	(2)	(3)	(4)	(5)
Coefficient below	8.46	8.58	9.95	8.83	8.80
variable's median	[3.19]	[3.32]	[3.62]	[3.42]	[3.37]
Coefficient above	9.68	10.22	8.33	9.24	7.86
variable's median	[2.65]	[3.24]	[3.15]	[1.94]	[1.78]
p-value for difference	0.741	0.725	0.725	0.911	0.779
Number of cells	56,926	48,092	56,926	56,926	56,926

Notes: Total conception rate coefficients based on regressions of annual changes in the age-specific conception rate on annual changes in 100 $\times$ log GDP per adult, controlling for country, year, and age fixed effects. Coefficients are estimated by 5-year age group and then summed and multiplied by 5 to obtain TCR coefficient. Brackets contain standard errors clustered by country. Sample sizes vary because data on some of the aggregate variables are not available for the full sample. "GDPpa" is GDP per adult, from the Penn World Table; contraceptive prevalence is the estimated share of women of childbearing age using modern contraceptives, from the UN; female labor force share is the percent of the labor force aged 15-64 that is female, from the WDI; average years of education and share urban are estimates from WFS/DHS survey data.

Table 3.A7: Regional Heterogeneity in Procyclicality

		p-values: coefficients equal within pair		
		Africa	C/W Asia	S/SE Asia
	(1)	(2)	(3)	(4)
Africa	9.43			
	[2.96]			
Central/Western Asia	4.53	0.01		
	[2.96]			
South/Southeast Asia	-2.46	0.18	0.44	
	[8.43]			
Latin America/Caribbean	10.81	0.69	0.09	0.12
	[2.33]			
p-value: all coefficients equal	0.15			
Number of cells	56,926			

Notes: Total conception rate coefficients based on full-sample regressions of annual changes in the age-specific conception rate on annual changes in 100 $\times$ log GDP per adult interacted with region indicators, controlling for country, year, and age fixed effects. An additional (unreported) interaction term is included for the group of five countries (Albania, Fiji, Korea, Moldova, Ukraine) that did not fit into these regional classifications. We do not interact the year and age effects with region indicators to conserve statistical power. Analyses are run by 5-year age group; age group associations are summed and multiplied by 5 to obtain TCR association. Brackets contain standard errors clustered by country.

Table 3.A8: Comparison of Procyclicality Results with Other Datasets

	Country-years in the WFS/DHS Mean 2005 GDPpa, PPP = 5,239 Mean GDPpa growth = 0.91			Country-years in the HFD Mean 2005 GDPpa, PPP = 46,993 Mean GDPpa growth = 2.41		
	WFS/DHS		WDI	HFD		WDI
	Mean rate	Regression	Regression	Mean rate	Regression	Regression
	(1)	(2)	(3)	(4)	(5)	(6)
Ages 15-19	138	0.163		28	0.173	
		[0.078]			[0.062]	
Ages 20-24	258	0.145		103	0.397	
		[0.097]			[0.135]	
Ages 25-29	261	0.475		126	0.179	
		[0.108]			[0.086]	
Ages 30-34	224	0.340		83	0.189	
		[0.107]			[0.066]	
Ages 35-39	161	0.358		35	0.141	
		[0.151]			[0.029]	
Ages 40-44	80	0.155		8	0.030	
		[0.277]			[0.008]	
Total fertility rate per 1000	5601	8.19 [2.63]	0.21 [0.24]	1920	5.55 [1.14]	6.44 [1.31]
Num. of cells	57,126	55,479	2,460	23,310	23,130	760

Notes: "WFS" = World Fertility Survey; "DHS" = Demographic and Health Survey; "HFD" = Human Fertility Database; "WDI" = World Development Indicators. Coefficients from regressions of annual changes in the age-specific fertility rate on the weighted average of current and lagged annual changes in 100 × log GDP per adult, with weight 0.25 on the current change and weight 0.75 on the lagged change. In the WFS/DHS and HDI, unit of observation is a country-year-age cell, and the dependent variable is the age-specific birth rate; analyses are run by 5-year age group and include country, year, and age fixed effects. Total fertility rate estimates equal 5 times the sum of age-group-specific estimates. In the WDI, unit of observation is a country-year cell, and the dependent variable is the total fertility rate; analyses are adjusted for country and year indicators. Brackets contain standard errors clustered by country. Sample includes all WFS/DHS and HFD cells that can be matched with macroeconomic data from the Penn World Table and total fertility rate data from the WDI, excluding cells with < 30 obs and from country-age combinations spanning < 20 yrs. WFS/DHS countries are listed in Table A1; HFD countries include Austria, Belarus, Bulgaria, Canada, Czech Republic, Estonia, Finland, France, Germany, Hungary, Iceland, Japan, Lithuania, Netherlands, Norway, Portugal, Russia, Slovakia, Slovenia, Sweden, Switzerland, Ukraine, United Kingdom, and the United States. We omit Japanese data for 1966, when birth rates dropped 25% due to superstition surrounding the year of the fire horse.

Table 3.A9: Alternative Long-Run Covariates

	Secondary school gross enrollment rate		Sectoral composition of value added		Female lab. force participation		POLITY IV score	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Ages 15-19	-0.06	-0.02	-0.06	-0.05	-0.11	-0.11	-0.14	-0.09
	[0.10]	[0.11]	[0.08]	[0.10]	[0.10]	[0.09]	[0.07]	[0.07]
Ages 20-24	-0.43	-0.36	-0.33	-0.37	-0.27	-0.27	-0.38	-0.38
	[0.16]	[0.17]	[0.13]	[0.17]	[0.14]	[0.15]	[0.11]	[0.11]
Ages 25-29	-0.42	-0.33	-0.46	-0.52	-0.42	-0.44	-0.48	-0.49
	[0.14]	[0.17]	[0.12]	[0.13]	[0.15]	[0.15]	[0.11]	[0.13]
Ages 30-34	-0.29	-0.12	-0.35	-0.42	-0.42	-0.45	-0.42	-0.41
	[0.14]	[0.15]	[0.15]	[0.20]	[0.17]	[0.17]	[0.14]	[0.16]
Ages 35-39	-0.08	0.06	-0.08	-0.05	-0.39	-0.41	-0.12	-0.08
	[0.15]	[0.16]	[0.15]	[0.19]	[0.16]	[0.16]	[0.13]	[0.13]
Ages 40-44	0.20	0.28	0.14	0.17	0.07	0.05	0.14	0.14
	[0.07]	[0.07]	[0.06]	[0.08]	[0.08]	[0.09]	[0.05]	[0.05]
TCR	-5.35	-2.42	-5.69	-6.16	-7.74	-8.12	-7.00	-6.57
	[2.92]	[3.09]	[2.50]	[3.16]	[3.00]	[2.96]	[2.16]	[2.30]
Covariate?		✓		✓		✓		✓
# cells	1,297	1,297	1,424	1,424	1,261	1,261	1,532	1,532

Notes: Regressions of the average annual rate of change in the age-specific conception rate on the average annual rate of economic growth. Each pair of columns restricts to the subsample with non-missing information on the average annual rate of change in the specified covariate. The even-numbered columns report models that include an age-specific coefficient on the average annual rate of change in the covariate. "TCR" refers to the total conception rate per 1000; estimates equal 5 times the sum of age-group-specific estimates. Brackets contain standard errors clustered by country.

Table 3.A10: Comparison of Long-Run Results with Other Datasets

	Country-ages in the WFS/DHS		Country-ages in the HFD	
	WFS/DHS	UN	HFD	UN
	(1)	(2)	(3)	(4)
Ages 15-19	-0.116	-0.008	0.191	0.36
	[0.069]	[0.090]	[0.160]	[0.152]
Ages 20-24	-0.286	-0.354	-0.181	0.511
	[0.126]	[0.109]	[0.482]	[0.263]
Ages 25-29	-0.476	-0.562	-0.673	-0.184
	[0.116]	[0.148]	[0.314]	[0.247]
Ages 30-34	-0.423	-0.529	-0.073	-0.329
	[0.132]	[0.158]	[0.177]	[0.234]
Ages 35-39	-0.220	-0.332	-0.030	-0.275
	[0.142]	[0.107]	[0.176]	[0.228]
Ages 40-44	0.152	-0.076	-0.024	-0.130
	[0.069]	[0.066]	[0.083]	[0.121]
Total fertility rate per 1000	-6.84	-9.30	-3.95	-0.23
	[2.27]	[2.52]	[3.58]	[3.35]
Num. of cells	1601	317	510	96

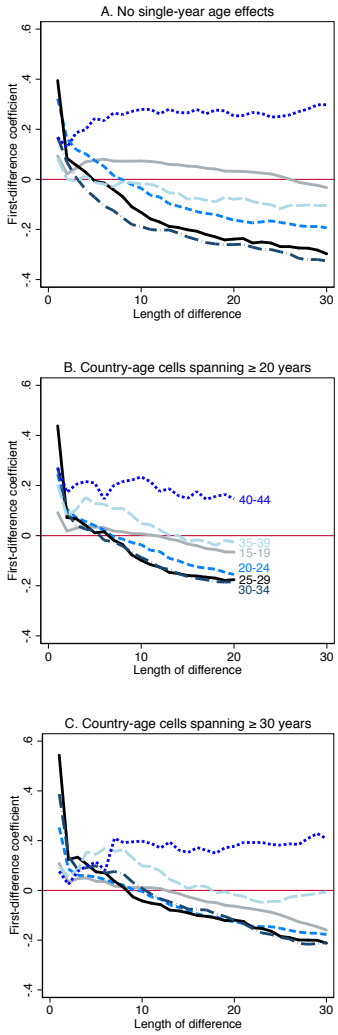
Notes: "WFS" = World Fertility Survey; "DHS" = Demographic and Health Survey; "HFD" = Human Fertility Database; "UN" = United Nations World Population Prospects, 2015 Revision. Coefficients from regressions of the average annual rate of change in the conception rate on the average annual rate of economic growth. The unit of observation is a country-age cell, and the dependent variable is the average annual rate of change in age-specific birth rate. Total fertility rate estimates equal 5 times the sum of age-group-specific estimates. Brackets contain standard errors clustered by country. WFS/DHS countries are listed in Table A1; HFD countries include Austria, Bulgaria, Canada, Finland, France, Germany, Hungary, Iceland, Japan, Netherlands, Norway, Portugal, Sweden, Switzerland, Taiwan, United Kingdom, and the United States. We omit Japanese data for 1966, when birth rates dropped 25% due to superstition surrounding the year of the fire horse.

Table 3.A11: Economic Growth over the Lifecycle and Completed Fertility

	Children Ever Born		Surviving Children	
	(1)	(2)	(3)	(4)
Avg. change in 100 × log GDPpa during ages...				
15-19	-7	-15	-5	-12
	[13]	[10]	[12]	[10]
20-24	-3	-15	2	-8
	[14]	[12]	[14]	[13]
25-29	10	1	7	-1
	[16]	[15]	[16]	[13]
30-34	38	26	29	18
	[14]	[12]	[12]	[11]
35-39	43	35	36	28
	[13]	[12]	[12]	[10]
40-44	25	34	18	26
	[17]	[12]	[15]	[12]
Cohort avg. ed.		-226		-226
		[54]		[54]
Cohort % urban		-4.7		-4.7
		[5.3]		[5.3]
Cohort FE	✓	✓	✓	✓
Country FE	✓	✓	✓	✓
Fertility measure	Ever-born	Ever-born	Surviving	Ever-born
Num. cells	935	935	935	935

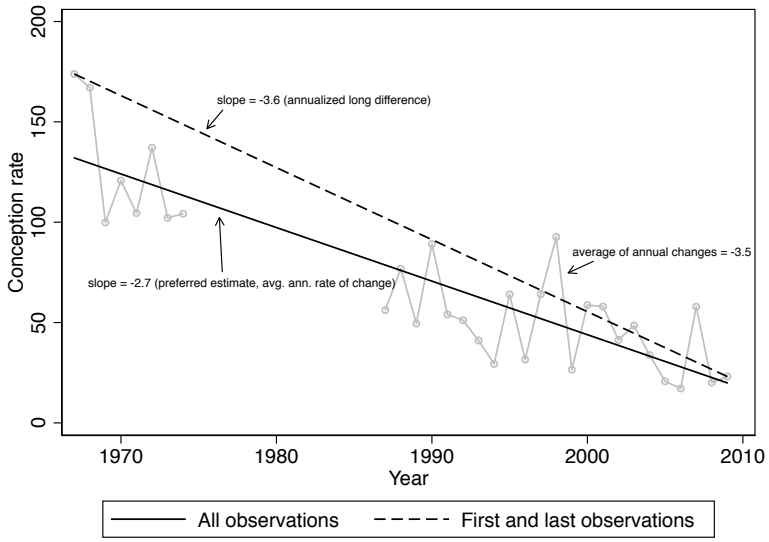
Notes: Sample includes single-age cohorts over age 45 when surveyed. Dependent variable is the number of children per 1000 women. Brackets contain SEs clustered by country.

Figure 3.A1: First-Difference Models with Varying Time Horizons, Constant Samples



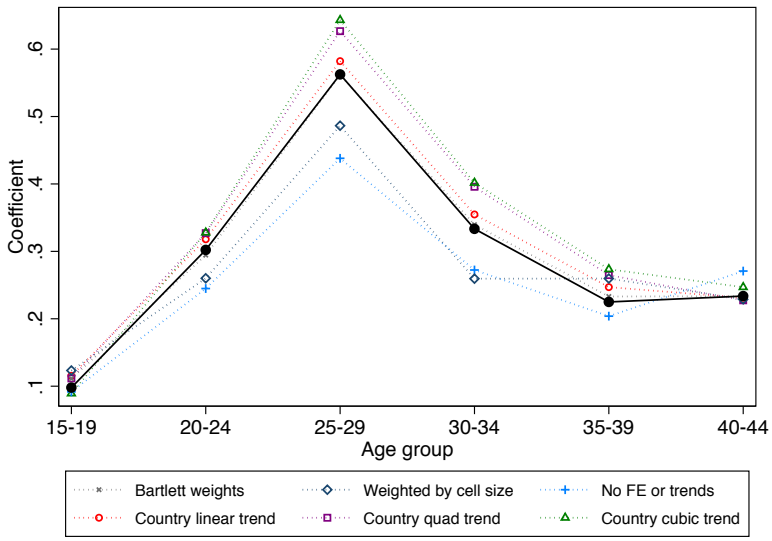
Notes: Reproduces Figure 1 without age effects or using samples that do not change for different time horizons. For each 5-year age group, each panel plots coefficients from regressions of the change in the conception rate from year t – Δ to year t on the change in 100 × log GDP per adult over the same period, controlling for single-year age indicators.

Figure 3.A2: Estimating the Average Annual Rate of Fertility Change: 40 year olds, Nepal



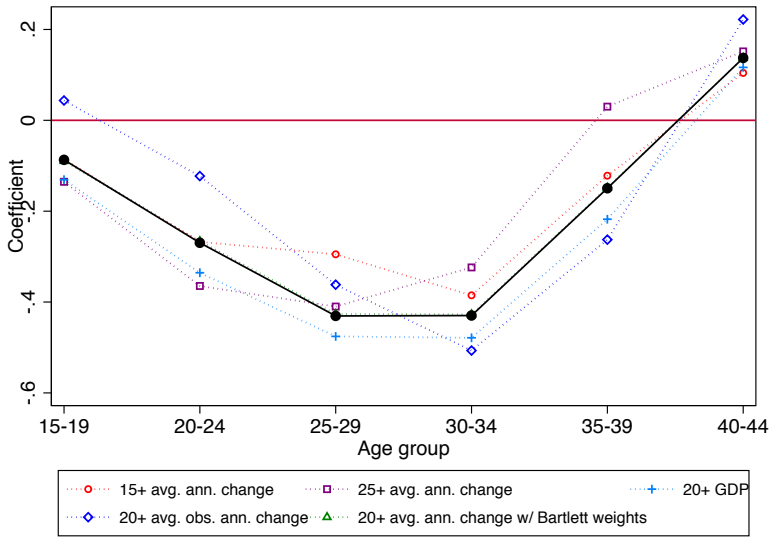
Notes: The figure plots estimated conception rates at age 40 over time in Nepal, with two trend lines, one estimated using all years and one estimated using just the first and last year of the series. We use the slope of the "all observations" trend as our estimate of the average annual rate of change because the other trend line (which is equivalent to the annualized long difference) uses less data, and the average of annual changes ignores trends during the data gap in the late 1970s and early 1980s.

Figure 3.A3: Alternative Short-Run Models



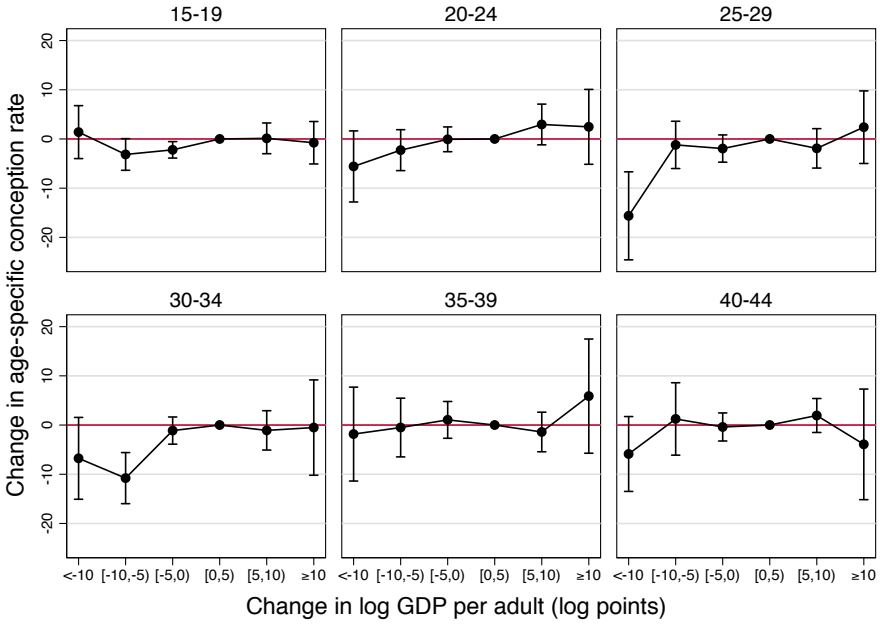
Notes: Age-group-specific coefficients from regressions of the change in the conception rate on the change in log GDP per adult. The thick black plot represents the coefficients from the short-run model reported in Figure 3. "Bartlett" uses a Bartlett kernel to downweight longer recall periods. The remaining models add country-specific polynomials in time to the baseline model.

Figure 3.A4: Alternative Long-Run Models



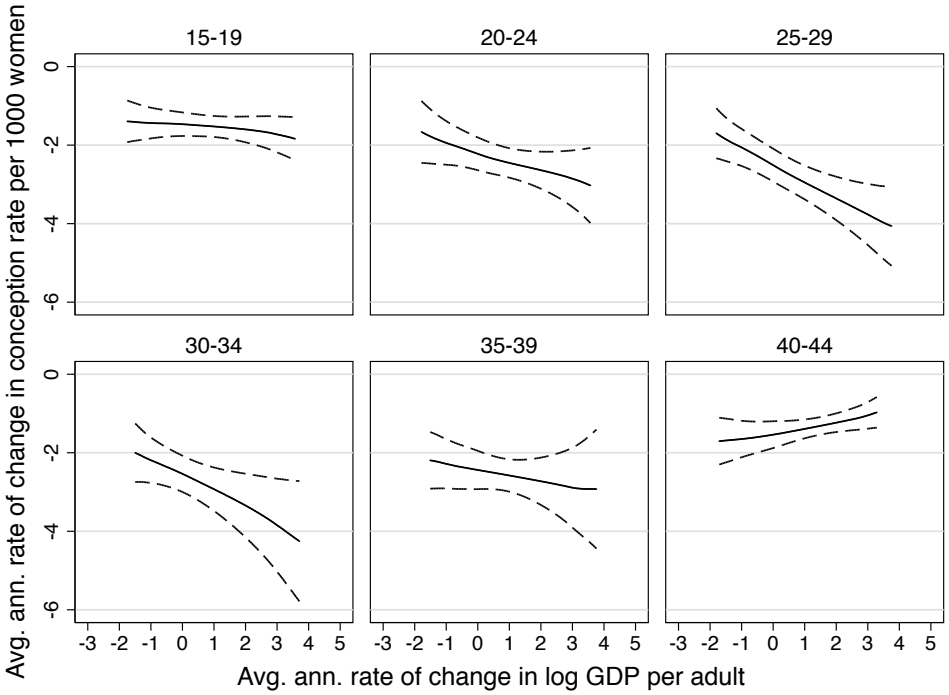
Notes: The figure compares results from different methods of computing the average annual rate of change. The thick black plot represents the coefficients from the long-run model reported in Figure 4. "15+" and "25+" use alternative minimum time horizons (15 and 25 years) to estimate the slope of the annual trend. "20+ avg. obs. ann. change" uses the average of observed annual changes (leaving out gaps in the panel) instead of the slope of the annual trend. "20+ Bartlett weights" downweights observations with longer recall periods, and "20+ GDP" uses GDP instead of GDP per adult.

Figure 3.A5: Non-Linear Estimates by Age Group, Short Run



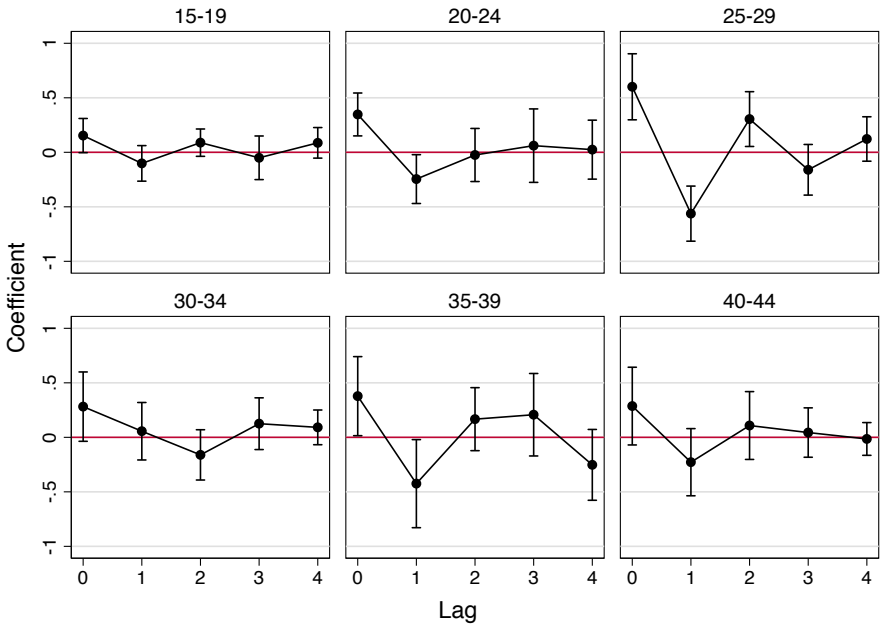
Notes: Semi-parametric results summarized in Figure 5, here shown by age group with 95% confidence intervals. Coefficients from regressions of annual changes in the age-specific conception rate on binned annual economic growth, controlling for country, year, and age fixed effects. Confidence intervals are clustered at the country level.

Figure 3.A6: Non-Linear Estimates by Age Group, Long Run



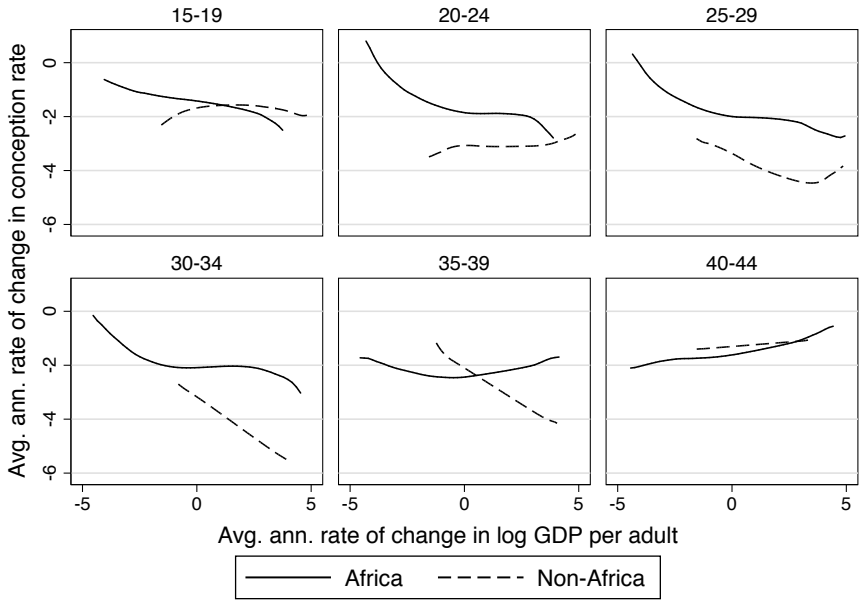
Notes: Non-parametric results summarized in Figure 6, here shown by age group with 95% confidence intervals. Local linear regressions; the dependent variable is the estimated trend in conception rates within a country-age cell, while the independent variable is the estimated trend in log GDP per adult in the same cell. Bandwidth equals 2, and confidence intervals are block bootstrapped at the country level.

Figure 3.A7: Distributed Lag Models by Age Group



Notes: Distributed lag model summarized in Figure 7, here shown by age group with 95% confidence intervals. Coefficients from regressions of annual changes in the age-specific conception rate on current and lagged annual changes in 100 $\times$ log GDP per adult, controlling for country, year, and age fixed effects. Confidence intervals are clustered at the country level. Sample includes observations that have both lagged conception rates and lagged growth rates.

Figure 3.A8: Non-Linear Estimates by Age Group and Region, Long Run



Notes: Replicates Figure A6, splitting the sample into African and non-African countries. Local linear regressions; the dependent variable is the estimated trend in conception rates within a country-age cell, while the independent variable is the estimated trend in log GDP per adult in the same cell. Bandwidth equals 2, and confidence intervals are block bootstrapped at the country level.

3.B Proofs

In the theoretical framework of Section 5, each household first chooses e to maximize:

$$U = E_0 \sum_{t=1}^{T} \beta^{t-1} \left\{ u_c\left(c_t\right) + u_n\left(n_t\right) + u_h\left(h_t\right) \right\}$$

and then chooses $\{c_t, b_t, A_{t+1}\}_{t=1}^{T}$ to maximize:

$$U = \sum_{t=1}^{T} E_t \beta^{t-1} \left\{ u_c\left(c_t\right) + u_n\left(n_t\right) + u_h\left(h_t\right) \right\}$$

subject to:

$$\begin{aligned}
n_t &= n_{t-1} + b_t \\
n_0 &= 0 \\
b_t &\in [0,1] \\
c_t &= w_t \left(1 - \tau \sum_{k=0}^{K} b_{t-k} \right) - \kappa \sum_{k=0}^{K} b_{t-k} - e \sum_{k=0}^{K} b_{t-k} + \left((1+r) A_t - A_{t+1} \right) \\
h_t &= \mathbf{1}[n_t > 0] h\left(e\right) \\
A_0 &\quad \text{given} \\
A_{T+1} &= 0 \\
A_{t+1} &\geq 0
\end{aligned}$$

The current value formulation of the Lagrangian is:

$$\mathcal{L} \equiv \sum_{t=1}^{T} \beta^{t-1} E_t \left\{ \begin{array}{l} u_c\left(c_t\right) + u_n\left(n_t\right) + u_h\left(h_t\right) + \lambda_t A_{t+1} + \left(\mu_t^0 - \mu_t^1\right) b_t + \\ \nu_t \left[w_t \left(1 - \tau \sum_{k=0}^{K} b_{t-k} \right) - \kappa \sum_{k=0}^{K} b_{t-k} \right. \\ \left. - e \sum_{k=0}^{K} b_{t-k} + \left((1+r) A_t - A_{t+1} \right) - c_t \right] \end{array} \right\}$$

where ν_t is the Lagrange multiplier on the period budget constraint, λ_t is the Lagrange multiplier on the borrowing constraint, and μ_t^0 and μ_t^1 are the Lagrange multipliers on births being between 0 and 1, respectively.

The first order conditions for consumption (in periods 1 to T) and births (in periods 1 to M) are:

$$u_c'(c_t) = \nu_t$$

$$-\nu_t + \lambda_t + \beta\left[(1+r)\, E_t \nu_{t+1}\right] = 0$$

$$\sum_{s=t}^{T} \beta^{s-t} E_t u_n'(n_s) \frac{\partial n_s}{\partial b_t} - \sum_{s=t}^{t+K} \beta^{s-t} E_t \nu_s \left\{\tau w_s + \kappa + e\right\} + \mu_t^0 - \mu_t^1 = 0$$

which imply equation (7):

$$u_c'(c_t, n_t, h) = \beta(1+r)\, E_t\left[u_c'(c_{t+1}, n_{t+1}, h)\right] + \lambda_t$$

and equation (8):

$$u_c'(c_t) = \frac{u_n'(n_{t-1} + b_t) + \sum_{s=t+1}^{T} \beta^{s-t} E_t\left\{\Psi\right\} + \mu_t^0 - \mu_t^1}{\tau w_t + \kappa + e}$$

where

$$\Psi = u_n'(n_s) \frac{\partial n_s}{\partial b_t} - \mathbf{1}_{\{s-t \leq K\}} \nu_s (\tau w_s + \kappa + e)$$

The first order condition for education (in period 0) is:

$$\sum_{t=1}^{T} \beta^{t-1} E_0\left\{u_h'(h(e;g))\, h_e(e;g)\right\} = E_0\left\{\sum_{t=1}^{M+K} \sum_{k=0}^{K} \beta^{t-1} \nu_t b_{t-k}\right\}$$

The left-hand side has no uncertainty, so we can remove the expectations sign and, noting that $\sum_{t=1}^{T} \beta^{t-1} = \sum_{t=0}^{T-1} \beta^{t} = \frac{1-\beta^{T}}{1-\beta}$, rewrite as equation (9):

$$u_h' \left(h \left(e; g \right) \right) h_e \left(e; g \right) \left(\frac{1-\beta^{T}}{1-\beta} \right) \quad = \quad E_0 \left\{ \sum_{t=1}^{M+K} \sum_{k=0}^{K} \beta^{t-1} \nu_t b_{t-k} \right\}$$