Essays on Road Infrastructure, Market Integration, and Development

Hundanol A. Kebede

A Dissertation presented to the Graduate Faculty
of the University of Virginia in Candidacy for the Degree of Doctor of Philosophy

Department of Economics
University of Virginia

April, 2020

Kerem Cosar ____________

Sheetal Sekhri ____________

John McLaren ____________

James Harrigan ____________

Peter Debaere ____________

Abstract

In the first chapter of this dissertation, I estimate the welfare gains from rural roads in Ethiopia. I develop a Ricardian trade model to explain how rural roads affect welfare in agricultural villages. The model captures two contrasting effects of rural roads on welfare. On one hand, decreases in trade costs lead to reallocation of village farmland to its comparative advantage crops and enable farmers to obtain higher prices for these crops. On the other hand, as these crops constitute significant part of the household consumption basket, the increase in prices decreases welfare. I take these predictions of the model to a very rich agricultural data. I show that, following decreases in trade costs due to massive rural road expansion, villages reallocate more farmland toward their comparative advantage crops and receive better prices for these crops, thus receiving higher nominal revenue. Despite increases in village price index, the gain in nominal income dominates, resulting in net welfare gain, on average. The size of this welfare gain depends on the crop composition of village consumption vis-a-vis production.

In the second chapter, I suggest a new approach to test separability– whether farm household's production and consumption decisions can be separately analyzed – and explore its link to market integration. The existing separability tests looked at the link between on-farm labor demand and household demographic characteristics. One problem with this approach is that on-farm labor demand is likely to be poorly measured in the context of self-employing agricultural households. I suggest an alternative test that is based on the link between household land allocation across different crops and the household's consumption tastes. I first estimate households' tastes for crops from their preference functions, and show that these tastes significantly dictate the households' land allocation across crops. The extent to which crop tastes dictate land allocation decreases with improvement in market integration due to construction of new rural roads.

In the third chapter, I study how frequent booms and busts in international coffee prices affect welfare of coffee producers in Ethiopia. I use unique panel data on

household consumption and land utilization to show that decreases in international coffee prices are passed on fully to consumption expenditure in households that allocate all farmland to coffee production. The decline in consumption has real consequences on child health and mortality. I find that children born in coffee producing districts during low coffee price periods are more likely to be underweight, wasted, stunted, and anemic, compared to their peers born in non-coffee districts; and their under-five mortality rate is higher by 1.8 percentage points.

Declaration

I hereby declare that this dissertation represents my own work which has been done after registration for the degree of Doctor of Philosophy at Department of Economics at University of Virginia, and has not been previously included in a work submitted to any other institution for a degree, diploma or other qualifications.

Signiture: ______________________

Date: ______________________

Acknowledgements

I would like to thank all members of my Dissertation Committee: Kerem Cosar, Sheetal Sekhri, John McLaren, James Harrigan and Peter Debaere. I thank all participants of the Trade and Development workshops for their helpful comments and encouragements throughout writing these papers. I also like to thank Chris Gist and Drew MacQueen for their technical supports in ArcGIS.

Introduction

This dissertation is organized into three chapters. The first two chapters of the dissertation study the effects of rural road infrastructure using micro data from Ethiopia while the third chapter studies the effects of large swings in international commodity prices on welfare of households who rely on production of these commodities.

In the first chapter, I estimate the welfare gains from rural roads. I develop a Ricardian trade model that captures the key mechanisms through which roads affect village welfare and estimate the predictions of the model using micro panel data on prices, land utilization and crop productivity at village level. I estimate welfare gains from the roads and how much of this welfare gain is explained by my trade model.

In the second chapter, I suggest a new test of separability – whether farm household's production and consumption decisions can be separately analyzed – and explore how construction of rural roads, by improving market integration, lead to decoupling of household production and consumption decisions. The separability test is derived from a simple theoretical insight that if household production decision is independent of its consumption preferences, the household's tastes for different crops would not affect household land allocation across crops. I implement this test using household level panel data on crop production and consumption fro Ethiopia.

In the third chapter, I study the effect of booms and busts in the international international coffee prices on welfare of producers in Ethiopia. I show that coffee price shocks significantly affect household consumption, and that this has substantial consequences for child malnutrition and child mortality.

Contents

List of Figures

List of Tables

Chapter 1

The gains from market integration: The welfare effects of new rural roads in Ethiopia

1.1 Introduction

A common feature of rural areas in developing countries is that they lack all-weather roads interconnecting the villages or connecting the villages to market centers in towns and cities. This makes villages isolated economies operating in a near-autarky environment. As a result, product markets are too thin for farmers to rely on, and farmers adopt subsistence farming where they self-produce most of the crops they need for consumption, instead of specializing in few crops and buying the rest from market.[1] This implies significant efficiency loss in land utilization compared to a scenario where villages specialize in crops to which their agro-climatic condition is most suited.

The welfare effects of road infrastructure that decreases trade costs in rural villages is complicated because the production and consumption decisions of farm households are intertwined, and roads affect both of these decisions. While decreases

[1] A manifestation of this is that a farm household in a developing country usually grows several crops, even though the average land holding of a household is less than two hectares. In Ethiopia for instance, the median land holding is about a hectare and a median farmers grows 5 crops.

in trade costs might lead to efficiency in resource allocation by inducing reallocation of land and labor towards comparative advantage crops and might enable farmers obtain better prices for these crops, it may also lead to significant increase in costs of consumption since these crops account for significant fraction of the consumption basket of the villagers. In this chapter, I study how construction of new roads affect village welfare using micro data on agricultural production and prices, and a large-scale rural road expansion project in Ethiopia as a source of variation to trade costs.

To understand the key mechanisms through which decreases in trade costs affect welfare in rural economies, I develop a multi-crop multi-location Ricardian model of trade with heterogeneous land quality within and across villages. On the demand side, a village maximizes utility by choosing optimal quantities of each crop to consume, given local crop prices. On the production side, the village makes a decision on how to optimally allocate its limited land across the potential crops, given local crop prices and the productivity of its village in different crops. Villages also engage in costly trade among each other, very similar to how countries trade in Eaton and Kortum (2002).

The model provides a number of sharp predictions on the mechanisms through which improvement in road infrastructure affects village welfare. First, decreases in trade costs lead to increases in the relative prices of the villages' comparative advantage crops (CA-crops). Second, decreases in trade costs lead to reallocation of farmland from a village's comparative disadvantage crops (CD-crops) towards the village's CA-crops. Third, the size of welfare gain from roads depends on the crop composition of the village's consumption basket vis-a-vis the fraction of village land allocated to these crops. For example, the model predicts that a village that has CA in cash-crops should benefit more from decreases in trade costs than a village that has CA in cereals because the latter faces significant increase in the relative costs of its consumption basket following decreases in trade costs. These predictions of the model can be tested easily given panel data on village level crop prices, land allocation, and a shock to transport infrastructure.

Next, I use the model structure to derive a sufficient statistic for the welfare effect of new roads which enables me account how much of the total welfare gain from roads estimated from a reduced-form regression is explained by my model. This sufficient statistics resembles the sufficient statistic for welfare gain from decreases in international trade costs in workhorse trade models, e.g., Arkolakis et al. (2012), and Eaton and Kortum (2002), with a straightforward modification to account for geographic variation in productivity of land in crops in my model. Interestingly, however, no trade flow data is required to construct the empirical measure of the sufficient statistic – it can be inferred from the villages' land allocations across crops and the model parameters.

I take these prediction of the model to a very rich village level panel data on crop prices, land allocation, and roads. The agricultural data covers about 2,000 nationally representative rural villages (locally named as *Kebele*s, which are the lowest administrative units) and all crops. The price data covers about 500 nationally representative rural villages and almost all agricultural products and inputs. I use a massive rural road construction project called Universal Rural Road Access Program (URRAP) that took place between 2013 and 2015 as a source of variation to trade costs. The objective of this program was to improve market integration by connecting rural villages to the nearest preexisting all-weather road or to the nearest town, and the program led to doubling of the total road length in the country between 2013 and 2015. About of half of my sample villages in each of agricultural and price data got road connection under this program.

At the core of my empirical exercises is defining a village's comparative advantage crop(s). This requires information on yield estimate for each crop in each village. My agricultural survey data includes village level crop yield estimated by trained enumerators.[2] However, there are two caveats to this yield estimate. First, yield estimates are provided only for the crops that are actually produced in the village.

[2]The enumerators use a method called crop cut where they take a sample of plots each with area of 4 square meter and conduct crop cut to obtain estimate of yield.

Second, such yield estimates are influenced by seasonal fluctuations to climatic factors such as rainfall. I overcome the second problem by averaging yield estimates across pre-URRAP program years within a village. As for the first problem, I assume a yield of zero for crops that are not produced in a village. As a robustness check, I use FAO-GAEZ data on agro-climatically attainable yield of crops in each village. This data uses a number of agro-ecological, soil and climatic factors, and sophisticated agronomic models to provide yield estimate at 5 arc-minute resolution. While the FAO-GAEZ data overcomes the above two problems with the yield estimates provided in my survey, it misses some of widely grown endemic crops in Ethiopia such as Enset and Teff. Given the yield estimates, I define a village's comparative advantage crops using the following procedure. First, I calculate a village's yield relative to national average for each crop. Next, I rank crops, within each village, based on their yield relative to national average. I define crops in the top 20% of ranking based on relative yield as my baseline comparative advantage crops. I relax this baseline threshold to top 30%, top 40%, etc. to see sensitivity of my results.

There are three main results in this chapter. First, the relative prices of CA-crops increased by at least 4% in villages that got new road connections compared to villages that didn't get new roads. As a result, overall price index increased by about 3% in these villages. Second, the fraction of farmland allocated to CA-crops increased by about 6% in villages that got road connection, relative to those that did not. These results are robust to alternative definitions of CA-crops and alternative estimation of yield (FAO-GAEZ data vs. survey data). Third, reduced-form estimation of the effect of the road expansion implies 12.5% increase in village welfare between 2012 and 2016. Using the sufficient statistics for welfare gain computed from the model, I conclude that about 7.5% of this welfare gain is attributed to the mechanism suggested in the model.

This chapter's novel contribution is that it studies the welfare effect of low-cost gravel roads connecting rural villages (that were previously inaccessible by any modern means of transport system) to the nearest market centers. Previous studies on the welfare effect of roads focused on trunk roads, high-ways, and railroads which connect different regions of a country and are much more expensive to build. The chapter also gives due emphasis to the underlying mechanisms through which roads affect village welfare.

An emerging literature studies the gains from intra-national market integration, particularly in the agricultural sector, using many-location many-good Ricardian trade models with heterogeneous factors of production (Costinot and Donaldson 2016, Donaldson 2018, Sotelo 2018, Allen and Atkin 2018, and Adamopoulos 2018). The most closely related papers to this chapter are Adamopoulos (2018), Donaldson (2018), and Sotelo (2018). Adamopoulos (2018) finds 13.6% increase in aggregate agricultural yield following road expansion and upgrading that reduced trade costs between Ethiopian districts and location of national grain market centers. The main difference with this chapter is that this chapter focuses on the effect of *rural* roads that connect village centers with district capitals, instead of a decrease in trade cost between district capitals and Addis Ababa or other major urban centers. Donaldson (2018) develops a multi-sector multi-region Ricardian model in which land is treated as homogeneous within a region to study the gains from the railway expansion in colonial India. Sotelo (2018) introduces heterogeneous land quality to study how falling trade costs due to (counter-factual) paving of roads increases agricultural productivity and welfare in Peru. This chapter builds on these two papers for the theoretical part and contribute to this literature by estimating the effect of low cost rural road construction on crop prices, land allocation, and welfare.

This chapter also relates to broad literature in development economics on how rural roads improve livelihood of households in developing countries (Gebresilasse 2018, Shamdasani 2018, Shrestha 2018, and Asher and Novosad 2019). Asher and Novosad (2019) exploit strict implementation rule of India's massive rural road expansion project called Pradhan Mantri Gram Sadak Yojana (Prime Minister's Village Road Program, or PMGSY) to identify the program's causal effect using fuzzy regression discontinuity design. They find that the roads' main effect is to facilitate the movement of people out of agriculture, with little or no effect on agricultural income and consumption. However, this chapter relies on proxies, instead of direct measures, for agricultural outcomes due to lack of data at fine geographic level. This chapter uses large household-level agricultural survey and price surveys at detailed geography to construct real agricultural income and consumption. Shamdasani (2018) studies the effect of large road-building program in India and finds that remote farmers who got access to road diversified their crop portfolio by starting to produce non-cereal hybrids, adopted complementary inputs and improved technologies, and hired more labor. Gebresilasse (2018) studies how rural roads complement with agricultural extension program, a program that trains farmers on how to use best agricultural practices and technology adoption, to increase farm productivity in Ethiopia. Shrestha (2018) finds that a 1% decrease in distance to roads due to expansion of highways resulted 0.1–0.25% increase in the value of agricultural land in Nepal. This chapter uses a theoretical model structure to identify the mechanisms through which rural roads affect welfare in village economies.

The rest of the chapter is organized as follows. In section 2, I present the data, identification issues, and give some descriptive statistics that motivate the model presented in Section 3. Sections 4 takes the key predictions of the model to the data. Section 5 presents estimation of key model parameters and the welfare gain. Section 6 concludes the chapter.

1.2 Data

1.2.1 Sources

Agricultural production data: I primarily use the Agricultural Sample Survey (AgSS), which is the largest annual agricultural survey in the country covering over 40,000 farm households in about 2200 villages. While this dataset goes back as far as 1995, villages were resampled every year until 2010 which makes tracking a village overtime difficult. Staring from 2010, Central Statistical Agency (CSA) kept the sample of villages fixed but took a random sample of about 20 farmers per village every year. This dataset includes detailed production information: areas of land covered by each crop, application of fertilizer and other inputs, and quantities of harvest. Moreover, every three-year starting from the year 2009/2010, CSA also gathered crop utilization information, i.e., the fraction of crop production used for own consumption, the fraction sold, the fraction used to pay wages, the fraction used for seeds, etc, for all crops.

Consumption data: To estimate some preference parameters of the model (the elasticity of substitution between crops), I need consumption information. I use Ethiopian Socioeconomic Survey (ESS) data which is an exceptionally detailed panel data of about 4,000 nationally representative farm households for the years2011, 2013 and 2015. The main advantage of the ESS dataset is it includes consumption information disaggregated by crops.[3] A big caveat of this data set is that it covers households in only about 330 villages.

Price data: The main price data is the Agricultural Producer Price Survey (AgPPS), which is a monthly survey of farm-gate prices at a detailed geography (villages) for almost all crops and many other agricultural produces.[4] This data covers over 500 representative villages which can be tracked over period since 2010. I also use the Retail Price Survey (RPS), which is a monthly survey of prices of almost all crops and non-agricultural commodities in major urban centers throughout the country. This dataset covers over 100 urban centers across all administrative zones of the country. Importantly, the agricultural products covered in both datasets overlap almost fully.

[3]The consumption information is based on a seven-day recall of basic consumption items, which are predominantly crops. However, household's crop utilization information also gives how much of each crops produced is consumed within the household.

[4]CSA claims that the prices in this survey can be considered as *farm-gate* price because they are collected at the lowest market channel where the sellers are the producers themselves, i.e., no intermediaries involved.

Rainfall and agro-climatic data: I use FAO/GAEZ agro-climatically attainable yield for low/intermediate input use to construct villages' crop suitability, as a robustness check to yield measures in AgSS data. This data covers about 19 crops. However, it misses some of the endemic crops that are widely grown in Ethiopia such as Enset and Teff. As a result, I use this data only for robustness exercise. The rainfall data comes from Climate Hazards Group InfraRed Precipitation with Station data (CHIRPS), which provides rainfall dataset starting from 1981. CHIRPS incorporates 0.05 resolution satellite imagery with station data to create gridded rainfall time series for trend analysis and seasonal drought monitoring. It is widely used to monitor drought in East Africa (Funk et al., 2015).

Road data: I use administrative data on the entire road-network in the country. This data includes the attributes of the roads (such as surface type), the role of the road (trunk road, link road, etc), and ownership (federal government, regional government, etc). In this chapter I use the massive rural road expansion under URRAP as a source of variation to villages' access to road/market. Over the period 2011-2015, the Ethiopian government gave exclusive focus to the URRAP and constructed over 62,413 kms of new all-weather roads connecting village centers to the nearest road or the nearest town, which ever is shorter. Figure 2.1 shows map of the road network before and after URRAP.

The main objective of the URRAP was to improve villages' access to product and input markets. The program increased the overall road density per 1000 square-km from 44.4 in 2010 to 100.4 in 2015 (Ethiopian Road Authority, 2016). Though the URRAP was launched in 2011, very few roads were commenced in the years 2011 and 2012, which are officially considered as capacity building years. Almost all the rural roads constructed under this program were started and completed between 2013-2015.

1.2.2 Identification of the effect of roads

There are two main issues that have to be dealt with to identify the causal effect of roads on village outcomes. The first is the concern of selection bias – villages are selected for the program based on some demographic, geographic, social, and economic factors.[5] The second is heterogeneity in treatment intensity and potential spillover effects. That is, villages that get connected to a dense network may gain more from the road than those that get connected to sparse network. Moreover, road connection in a village may have spillover effects to other villages that are not directly connected. When a village is connected to the preexisting road network or to the nearest urban center, all its neighbors also have improved access to market via the connected village. As a result, non-connected villages may not serve as control groups in identification of the effects of road connection.

I address the potential selection bias by using a Matching-Based Difference-in-Differences (MB-DID) strategy. That is, I first obtain a matched sample of treated and non-treated villages based on their observable characteristics that might be relevant for selection of villages for URRAP. I then conduct DID estimation based on these matched sample of treated and non-treated villages. Combining matching with DID strategy is a powerful approach to address the selection problem. The matching step enables me to compare treated villages with non-treated villages that have similar observed characteristics and hence similar treatment probability. The DID strategy on these matched samples helps me to washout unobserved time-invariant village characteristics that may confound the treatment effect.

To identify relevant village characteristics for matching treated and non-treated villages, I use information from officials at Ethiopian Roads Authority (ERA). They suggest that the main factors determining whether a village would be selected for URRAP in a particular year are: (1) the village's distance to preexisting road

[5] Unfortunately there was no official guideline as to which villages should be selected for the URRAP in a given year. Even though the project was fully funded by the federal government, implementation of URRAP was completely decentralized to regional governments. Within each regional government, districts propose list of villages that should get a road during a particular year and the regional governments approve villages based the available regional budget.

network, (2) population density of the village, and (3) the terrain and landscape of the village. Distance to preexisting roads is crucial because movement of machineries and other construction materials to the construction sites, by itself, requires roads that are passable by vehicles. Population density is relevant both for political consideration and the project's labor input requirements.[6] Finally, terrain and landscape significantly affects the road construction costs. Villages that require many bridge constructions or cutting and digging of hills are usually less favorable due to cost considerations.

Based on these insights from the officials, I use the following list of variables to match treated and non-treated villages: distance to nearest town, distance to preexisting road network, population size, average slope of land in the village, average elevation in the village, and average rainfall over 1990-2010 period. I use Digital Elevation Model (DEM) data and ArcGIS tools to calculate average slope and elevation of each village.

To address the heterogeneity in treatment intensity and spillover effects, I use market access measure derived from general equilibrium trade models (see Donaldson and Hornbeck 2016) that are calculated using the entire road network and the distribution of population across Ethiopian villages. Changes in market access capture treatment benefit from both direct and indirect connectivity, and properly account for the density of the network to which a village is connected. See Appendix A for details in the construction of market access measure. The constructed market access measure increases both for villages that are directly connected and those that are not by 47%, on average, but it increases more for the directly connected villages by about 40%.

1.2.3 Evidences on improvement in market integration

In this subsection, I present some evidences on how the road construction under URRAP improved market integration among rural villages.

[6] Most of the labor input for the URRAP roads are contributed by local residents, about three-quarters of which is a free labor.

Farmers face considerable barrier to trade: ESS data includes direct questions about transport costs. I use this survey to obtain estimation of transport costs. The ad valorem trade cost (transport cost per value of transaction) on vehicle is very high (the median is 6.5% and the mean is 11.4%). The size of this cost is comparable to international trade costs estimated by Hummels (2007) for US and New Zealand import, although in my data the distance traveled is just 12 kilometers.

URRAP decreased trade costs: The main objective of URRAP roads was to integrate rural villages to market centers (Ethiopian Road Authority, 2016). If URRAP roads really integrated rural villages to local market centers, we would see the price gap between the rural villages and the market centers decreasing for villages that got road connection relative to villages that did not get roads. I test this by looking at the difference in crop prices between zone capitals and the villages within the zones using the two rich price surveys, AgPPS and RPS. I run the following regression:

$$lnP^k_{zmt} - lnP^k_{zvmt} = \alpha_1 Post_t + \alpha_2(Post_t * URRAP_v) + \gamma_v + \gamma^k_m + \gamma_t + \varepsilon^k_{zvmt}$$

where k denotes crop, v is village, z is zone capital, m is month, t is year, $Post$ equals zero for all month-years before URRAP and one for all month-years after URRAP; $URRAP_v$ is a dummy variable representing whether a village got URRAP road between 2012 and 2015; and γ^k_m is crop-month fixed effect which captures possible seasonality of crop prices.

The result is reported in Table 1.1. It shows that road connection significantly decreased the urban-rural price gap. The first column pools all 56 crops for which data is available on both urban and rural prices. It shows that trade cost, as proxied by the ratio of urban to rural prices, decreased by about 3% for villages that got road connection, relative to villages that did not get road connection. In column 2, the estimation is restricted to perishable products, vegetables and fruits. The estimated decrease in trade cost for these products is more than twice the estimate for all crops: trade cost for vegetables and fruits decreased by about 8%. This is not surprising because trading such products is difficult when there is no road passable by vehicle connecting a village to the urban center due to their perishability. In the last column, the sample is restricted to observations in which urban prices are higher than rural prices, which is what one would expect if villages are net exporters of crops to urban centers.[7] The gap between these two prices are plausibly capturing trade costs, which decrease by about 2.4%.

URRAP decreases the correlation between local prices and yields: One key indicator of an integrated market is that local prices are less sensitive to local supply. Under autarky, prices are relatively lower (higher) for the goods in which a region has a comparative advantage (disadvantage). Market integration weakens this inverse relationship between local prices and local comparative advantage. I run the following generalized difference-in-differences regression to investigate this:

$$
\begin{aligned}
lnP^k_{vt} &= \alpha_1 lnA^k_v + \alpha_2(Post_t * URRAP_v) + \alpha_3(lnA^k_v * Post_t * URRAP_v) \\
&\quad + \gamma_v + \gamma_k + \gamma_t + \varepsilon^k_{vt}
\end{aligned}
$$

where P^k_{vt} is price of crop k in village v, A^k_v is a village's productivity in crop k which is proxied by GAEZ potential yield for the crop.

The result is presented in table A.2. Panel A uses binary treatment dummy. The first column shows that there is a negative relationship between local prices of a crop and local yield of the crops. Column 3 shows that this negative relationship is significantly weakened when a village gets road connection. The elasticity of village price to village yield is 2.7% for a village with no road connection and a road connection decreases this estimate to 1.8%.[8] Panel B of table A.2 reports the corresponding estimation result using market access measure instead of binary treatment dummy. The result clearly shows that in villages that see an increase in

[7] Note that about 80% of observations (67,147 out of 82,944) conform with this expectation.

[8] Alternatively, a positive α_3 would imply that road connectivity increases the prices of crops in which a village has a comparative advantage.

their market access, the negative correlation between crop price and yield becomes significantly weaker.

Informed by these evidences on the effect of roads on market integration, in the next section, I develop a multi-sector multi-village Ricardian trade model to analyze the mechanisms through which road constructions affect village welfare.

1.3 Theoretical framework

The model builds on Donaldson (2018) and Sotelo (2018). Consider an economy composed of V villages indexed by $v = 1, ..., V$, each village represented by a representative household. A village derives utility from consumption of K homogeneous crops indexed by $k = 1, ..., K$ that can be potentially produced or purchased.

Preferences: The village spends all its income on crops and its preference over different crops is given by

$$U_v = \left(\sum_k (q_v^k)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}$$

where σ is elasticity of substitution between crops, U_v is utility of the village, and q_v^k is the quantity of crop k consumed by the village.

Technology: Similar to Sotelo (2018) and Allen and Atkin (2018), I assume that the farmer's technology is constant returns to scale. I also assume, for simplicity of exposition, that land is the only input.[9] Each village has L_v amount of land, which is divided into a continuum of plots of size one indexed by $\omega \in \Omega_v$, where Ω_v is the set of plots in village v such that $\int_{\Omega_v} \omega d\omega = L_v$. Each of the plot is potentially different in how well it is suited to growing different crops, which I denote as $z_v^k(\omega)$. Assuming that a given plot can only be used to grow one crop at a time (plots cannot

[9] The model can easily be extended to include labor without altering any of the analysis in this section but at a cost of introducing new notations. Hence, I abstract from introducing labor in this section.

be divided), the production function is given as:

$$y_v^k(\omega) = z_v^k(\omega)$$

where $y_v^k(\omega)$ is quantity of crop k per unit of plot.

A representative farmer in each village draws $z_v^k(\omega)$ independently for each plot-crop from a Fréchet distribution with the following cumulative distribution function:

$$F_v^k(z) = Pr(Z_v^k < z) = \exp(-(A_v^k)^\theta z^{-\theta})$$

where A_v^k is the location parameter for the distribution of crop-suitability of land across the set of plots in village v, Ω_v. A_v^k can be interpreted as the average productivity of village v in crop k. For villages with agro-climatic conditions that are impossible to produce crop k, A_v^k is set to zero. θ is the (inverse) measure of dispersion in the productivity of land in a village, and it is constant across villages and crops.

Trade: Villages operate in a perfectly competitive crop market. There is an iceberg trade cost of $\tau_{vv'}^k \geq 1$ between villages v and v' in crop k. Motivated by the result in table 1.1, which shows that spatial price variation differs across crops, trade costs are assumed to vary across crops to reflect that some crops, such as vegetables, are more costly to trade (e.g., perishable) than others such as cereals. I assume that $\tau_{vv}^k = 1, \forall k$, and impose the standard assumption of triangle inequality in trade costs, $\tau_{vv'}^k \times \tau_{v'v''}^k \geq \tau_{vv''}^k, \forall k$.

I assume no arbitrage condition, so that for any two villages v and v' equilibrium crop prices satisfy $p_{vv'}^k = \tau_{vv'}^k p_{vv}^k$ where $p_{vv'}^k$ is price in village v' of crop k originating from village v, p_{vv}^k is price in village v of crop k originating from the same village v.

Distribution of prices: Let r_v is the rental rate of plots in village v, which is determined in equilibrium. The unit cost of production of crop k in village v is $c_v^k = \frac{r_v}{Z_i^k}$, which is stochastic because it is a function of stochastic productivity Z_v^k.

As a result, the price at which village v supplies crop k to village v', $p^k_{vv'} = \frac{r_v}{Z^k_i}\tau^k_{vv'}$, is stochastic.

Using the distribution of Z^k_i, we obtain the following distribution of the prices of crop k that village v' is offered by another village v:

$$G^k_{vv'}(p) = 1 - \exp\Big(- (A^k_v)^\theta (r_v \tau^k_{vv'})^{-\theta} p^\theta \Big)$$

Because crops supplied by different villages are homogeneous, village v' buys each crop k from the village that supplies the crop at the lowest price. Thus, the distribution of the price of crop k that is actually paid in village v' is the distribution of the lowest prices across all other villages and is given by:

$$\begin{aligned} G^k_{v'}(p) &= 1 - \prod_{v=1}^{V} (1 - G^k_{vv'}(p)) \\ &= 1 - \exp\Big(- p^\theta \sum_{v=1}^{V} (A^k_v)^\theta (r_v \tau^k_{vv'})^{-\theta} \Big) \end{aligned} \tag{1.1}$$

The probability that village v is the cheapest supplier of crop k to village v' (the probability that village v's productivity draw, adjusted for trade costs and rental rates, is the highest compared to all other potential villages trading with v') is:

$$\begin{aligned} \pi^k_{vv'} &= \Pr\Big[p^k_{vv'} \leq \min_n \{p^k_{nv'}\} \Big] \\ &= \frac{(A^k_v)^\theta (r_v \tau^k_{vv'})^{-\theta}}{\sum_v (A^k_v)^\theta (r_v \tau^k_{vv'})^{-\theta}} \end{aligned}$$

which is increasing in the average productivity of village v in crop k, A^k_v and decreasing in the trade cost $\tau^k_{vv'}$ and the rental rate in village v, r_v, relative to other villages.

The probability that a village will be the cheapest supplier of a crop to itself is

$$\begin{aligned} \pi^k_{vv} &= \Pr\Big[p^k_{vv} \leq \min_{\{n \neq v\}} \{P^k_{nv}\} \Big] \\ &= \frac{(A^k_v)^\theta r_v^{-\theta}}{\sum_n (A^k_n)^\theta (r_n \tau^k_{nv})^{-\theta}} \end{aligned}$$

A village is more likely to self-produce crop k if the village is more productive in the crop relative to other villages and/or if there is high trade costs between the village and other villages.

Following Donaldson (2018), we can obtain the expected value of this price distribution:

$$p_v^k = \Gamma\Big(\sum_{v'=1}^{I}(A_{v'}^k)^\theta(r_{v'}\tau_{v'v}^k)^{-\theta}\Big)^{\frac{-1}{\theta}} \tag{1.2}$$

where Γ is a gamma function $\Gamma(1+1/\theta)$. I assume that farmers in village v make their production and consumption decisions based on this expected price. Given the CES preferences, the price index faced by village v is given by:

$$P_v = \left(\sum_k p_v^{k\,1-\sigma}\right)^{\frac{1}{1-\sigma}} \tag{1.3}$$

where p_v^k is given by equation 1.2.

Equilibrium land allocation: I assume that there is a competitive land rental market. Each village decides how to allocate its farmland across different crops given prices p_v^k and the suitability of the village land for various crops. Revenue maximization implies that each plot of land is allocated to a crop that yields the highest return:

$$R_v(\omega) = \max_k\{p_v^k z_v^k(\omega)\} \tag{1.4}$$

where $R_v(\omega)$ is revenue from plot ω. Together with the Fréchet distribution, this implies the following land allocation rule:

$$\eta_v^k = \frac{(p_v^k A_v^k)^\theta}{(\Phi_v)^\theta}, \quad \text{where} \quad \Phi_v = \left(\sum_{l=1}^{K}(p_v^l A_v^l)^\theta\right)^{\frac{1}{\theta}} \tag{1.5}$$

where η_v^k is the fraction of land in village v allocated to crop k. It increases with the price of the crop and the average productivity of the the village in the crop, relative to all other crops. The following proposition summarizes the key mechanism through which roads affect village welfare in this model:

Proposition 1. *Decreases in trade costs lead to reallocation of farmland to a village's comparative advantage crops, resulting in more specialization of villages.*

Proof. The elasticity of land share of a crop in a village to the village's productivity is given as $\frac{d\ln\eta_v^k}{d\ln A_v^k} = \theta(1-\eta_v^k)$. Differentiating with respect to trade costs, we obtain $\frac{d^2\ln\eta_v^k}{d\tau_{vv'}^k d\ln A_v^k} = -\theta\frac{d\eta_v^k}{dp_v^k}\frac{dp_v^k}{d\tau_{vv'}^k}$. The term $\frac{d\eta_v^k}{dp_v^k}$ is always positive (see equation 1.5). Consider two villages v and v' and suppose the price of crop k in village v' is normalized, so that the price in village v is defined relative to the price in village v'. This implies that $p_v^k = \tau_{vv'}^k$ if k is CD-crop in village v (i.e., $p_v^k > p_{v'}^k = 1$) or $p_v^k = \frac{1}{\tau_{vv'}^k}$ if k is CA-crop in village v (i.e., $p_v^k < p_{v'}^k = 1$). Thus the term $\frac{dp_v^k}{d\tau_{vv'}^k}$ has a positive sign if crop k is a CD-crop in village and a negative sign otherwise.[10] This implies that as trade costs decrease, villages reallocate more land to their CA-crops. □

The intuition is simple. As trade costs decrease, a village's CA-crops ('export' crops) become relatively more expensive and CD-crops ('import' crops) become relatively cheaper at local markets. This makes growing CA-crops relatively more attractive and growing CD-crops relatively less attractive, which induces reallocation of land to these CA-crops.

Revenue per plot and equilibrium rental rate: In appendix B, I derive the conditional distribution of land productivity $\mathcal{Q}_v^k(z) \equiv \mathcal{P}\Big(Z_v^k(\omega) < z | \omega \in \Omega_v^k\Big)$, i.e., the distribution of productivity of a plot conditional on the plot being used for crop k, which gives the following distribution function:

$$\mathcal{Q}_v^k(z) = \exp\Big(-\Big(\frac{\Phi_v}{p_v^k}\Big)^\theta z^{-\theta}\Big)$$

which is Fréchet with the expected value of $\frac{\Phi_v}{p_v^k}$.

Suppose crop k is the crop that maximizes revenue from plot ω so that optimal revenue from plot ω is given by $R_v(\omega) \equiv p_v^k q_v^k(\omega) = p_v^k z_v^k(\omega)$. The conditional

[10]Recall that, from the no-arbitrage and $\tau_{vv'}^k \geq 1$ conditions, $p_v^k = p_{v'}^k\tau_{vv'}^k$, if $p_v^k > p_{v'}^k$ or $p_v^k = p_{v'}^k/\tau_{vv'}^k$, if $p_v^k < p_{v'}^k$

distribution of revenue from a plot conditional on the plot being used for crop k, $\mathcal{P}\Big(R_v(\omega) < R|\omega \in \Omega_v^k\Big)$, is also Fréchet with the expected value of Φ_v because revenue is just the productivity term multiplied by a non-stochastic price p_v^k. Moreover, given the assumption of competitive land rental market, rental rate per plot is equal to revenue per plot. Thus, the conditional distribution of rental rate per plot is the same as the conditional distribution of revenue per plot (note from equation 1.4 that $r_v(\omega)|\omega \in \Omega_v^k = p_v^k z_v^k(\omega)$ which has a Fréchet distribution with parameter Φ_v).

A sufficient statistic for the welfare gain: Because land is the only factor of production in the model, average real rental rate per plot, which is also equal to average real revenue per plot, can be used as a measure of village welfare: $W_v \equiv \frac{r_v}{P_v} = \frac{R_v}{P_v}$ where r_v is average rental rate per plot, R_v is average revenue per plot, and P_v is the village price index. Given data on quantities of each crop produced in a village, village prices, and model parameters needed to construct village price index, one can construct village real revenue $\frac{R_v}{P_v}$ and compare its changes over time in villages that get new road connection against villages whose road status did not change in a difference-in-differences strategy. However, this DID estimate does not tell us how much the change in welfare due to roads is explained by the mechanism suggested in this model.

The simplicity of the current trade model allows me to derive a sufficient statistic for the welfare gain from road. This sufficient statistic is crucial to shed light on how much of the welfare gain from road obtained in the reduced-form DID estimation is attributed to the mechanism that our trade model captures. Recall that the probability that village v is the cheapest supplier of crop k to village v' is equal to $\pi_{vv'}^k = (A_v^k)^\theta (r_v \tau_{vv'}^k)^{-\theta} (p_{v'}^k)^\theta$. Evaluating this expression at $v = v'$ and solving for p_v^k we obtain $p_v^k = r_v (A_v^k)^{-1} {\pi_{vv}^k}^{1/\theta}$. Plugging this into CES price index we obtain $P_v = r_v \Big[\sum_k \Big((A_v^k)^{-1} (\pi_{vv}^k)^{\frac{1}{\theta}}\Big)^{1-\sigma}\Big]^{\frac{1}{1-\sigma}}$ which implies:

$$\ln\Lambda_v \equiv \ln\frac{r_v}{P_v} = \frac{1}{\sigma - 1}\ln\sum_k {A_v^k}^{\sigma-1} {\pi_{vv}^k}^{\frac{1-\sigma}{\theta}} \tag{1.6}$$

where Λ_v denotes the sufficient statistic for the welfare effects of roads. This expression is similar to the sufficient statistic expressions for welfare gain from a decrease in trade cost given in Arkolakis et al. (2012), except that the expenditure share on own product π^k_{vv} is weighted by local productivity of crops. This modification is due to the productivity distribution is assumed to vary across locations for each crop (compared to the same technology parameters assumed for all goods in Arkolakis et al. (2012)).

The main result in this model that enables me to obtain the sufficient statistic of the welfare gain from roads without observing any trade flows is Proposition 2.

Proposition 2. *The probability that a village is the cheapest supplier of crop k to itself, π^k_{vv}, is proportional to the fraction of the village farmland allocated to the crop, i.e., $\pi^k_{vv} = \kappa\eta^k_v$.*

Proof. The proof is straight forward. Recall that the probability that village v is the cheapest supplier of crop k to village v' is given by $\pi^k_{vv'} = \kappa(A^k_v)^\theta(r_v\tau^k_{vv'})^{-\theta}(p^k_{v'})^\theta$. Evaluating this expression at $v = v'$ we obtain $\pi^k_{vv} = (A^k_v)^\theta(r_v)^{-\theta}(p^k_v)^\theta$, which can be rearranged to give $\pi^k_{vv} = \kappa\frac{(A^k_v p^k_v)^\theta}{r^\theta_v} = \kappa\frac{(A^k_v p^k_v)^\theta}{\sum_l (A^l_v p^l_v)^\theta} = \kappa\eta^k_v$, where the second equality uses the expression for the average rent from the distribution of equilibrium rent, $r_v = \Phi_v$. $\kappa = \Gamma(1 + 1/\theta)^{-\theta}$, where $\Gamma(.)$ is a gamma function. □

Proposition 3. *The size of the welfare gain from road depends on the fraction of land allocated to different crops in a village vis-a-vis the consumption composition of the village.*

Recall that welfare is given by $\frac{r_v}{P_v} = \left(\sum_{k=1}^K (p^k_v A^k_v)^\theta\right)^{\frac{1}{\theta}} / \left(\sum_k p^{k\,1-\sigma}_v\right)^{\frac{1}{1-\sigma}}$. Taking logs and differentiating this with respect to p^k_v gives:

$$\mathrm{dln}W_v = \sum_k \left(\eta^k_v - s^k_v\right)\mathrm{dln}p^k_v \tag{1.7}$$

This equation shows that two forces govern the welfare effects decreases in trade costs. The effect of change in trade costs of a crop on village welfare depends on whether or not the crop is CA-crop in the village and the fraction of farmland allocated to the crop relative to the share of the crop in the village expenditure $\eta_v^k - s_v^k$. Whether a crop is CA-crop determines the sign of $d\ln p_v^k$ – which is negative if crop k is CA-crop in the village and negative otherwise. Thus, decreases in trade costs of CA-crop increases village welfare if $\eta_v^k > s_v^k$ and decreases village welfare if $\eta_v^k < s_v^k$. The opposite holds for a decrease in trade costs of CD-crops. The intuition is simple. The net welfare effect of change prices of village CA-crops depends on, loosely speaking, whether the village is net seller or net buyer of the crops. If a village is net seller of the CA-crop, then $\eta_v^k >> s_v^k$ and the decrease in trade costs of the crop has a positive effect on welfare. If the crop is instead a CD-crop, the net welfare effect of a decrease in trade costs will be positive only if the village is net buyer of the crop, $\eta_v^k << s_v^k$.

A testable implication of this proposition is that a village that specializes in cereals gains less from road connection compared to a village that specializes in non-cereal crops. This is because, while both groups of villages gain from the increase in the prices of their CA-crops, the villages specializing in cereals experience an increase in their consumption expenses relative to the villages that specialize in non-cereals.

In order to construct the empirical measure of sufficient statistic for welfare gain from roads, we need to obtain estimates for the parameters of the model: the elasticity of substitution σ, and the measure of homogeneity of plots in a village θ.

1.4 Testing the model predictions

1.4.1 New roads and reallocation of land

The theoretical results in section 2.3 suggest that decreases in trade costs would lead to reallocation of farmland towards a village's CA-crops. I test this directly. In order to identify a village's CA-crop(s), I primarily use village level crop yield estimates provided in AgSS. This yield estimate is conducted by trained enumerators who use a method called crop cut where they take a sample of plots of area of 4 square meters and conduct crop cut to obtain estimate of yield. However, there are two caveats to this yield estimate. First, yield estimates are provided only for the crops that are actually produced in the village. Second, such yield estimates are influenced by seasonal fluctuations in climatic factors such as rainfall and crop diseases. I overcome the second problem by averaging yield estimates across pre-URRAP program years within a village. As for the first problem, I assume a yield of zero for crops that are not produced in a village.

As a robustness check, I use FAO-GAEZ data on agro-climatically attainable yield of crops in each village. This data uses a number of agro-ecological, soil and climatic factors, and sophisticated agronomic models to provide yield estimate at 5 arc-minute resolution. While the FAO-GAEZ data overcomes the above two problems with the yield estimates provided in AgSS data, it covers only a partial list of crops produced in Ethiopia. In particular, it misses some of widely grown endemic crops in Ethiopia such as Enset and Teff. Nevertheless, I use this alternative yield estimate for a robustness check. The correlation between the FAO-GAEZ yield estimates and the AgSS yield estimates, based on crops that exist in both datasets, is remarkably high – about 0.8, and all my results are robust to which yield estimates used.

Given the yield estimates A_v^k, I define a village's comparative advantage crops using the following procedure. First, I calculate a village's yield relative to national average for each crop, $\frac{A_v^k}{A^k}$. Next, I rank crops, within each village, based on their yield relative to national average $\frac{A_v^k}{A^k}$. I define crops in the top 20% of ranking based on relative yield as my baseline comparative advantage crops. I relax this baseline threshold to top 30%, top 40%, etc. to see sensitivity of my results. One issue is some villages grow only a handful of the major crops considered for analysis, and using the above procedure would end up classifying all or most of the crops grown in a village as CA-crops in villages that grow few crops.[11] For instance, if a village grows only 5 of, say, 25 crops in the data and we define CA crops as top 20% in

[11] About 4% of the villages grow five or less crops and the maximum number of crops grown in a given village is 19, out of the 25 major crops.

relative yield, all the 5 crops grown in the village would be classified as CA crops and as a result we would not discern any land reallocation because we are defining the set of all crops grown in the village as CA-crops. To overcome this problem, I keep only crops that are grown in a village in at least one of pre- or post-road years and rank these crops in their relative productivity $\frac{A_v^k}{A^k}$ with in the village. In this approach, for a village that grows only 5 of the 25 crops, the CA-crop is the crop that is at the top in ranking of $\frac{A_v^k}{A^k}$ within the village.

I estimate the following regression at a village level:

$$\eta_{vt}^k = \alpha_1(Post_t * Road_v) + \alpha_2 CA_v^k + \alpha_3(Post_t * Road_v * CA_v^k) + \beta\mathbf{X} \\ + \gamma^k + \gamma_v + \gamma_t + \varepsilon_{vt}^k$$

where η_{vt}^k denotes the share of land allocated to crop k in village v in year t. CA_v^k is dummy variable indicating whether crop k is among the village's CA-crops, and X is a vector of village characteristics such as the population density and rainfall. I also include crop fixed effects to account for mean variation in land intensity across crops. The regression includes village fixed effect to account for time invariant village characteristics that may confound our result and year fixed effect to account for any year specific factor shared across villages. I also estimate similar regression by using log market access derived from general equilibrium trade model in stead of the binary treatment dummy.

For the main analysis, I drop tree-crops and crops that are very rarely produced and have insignificant land share. The issue with tree-crops such as coffee, banana, orange, etc. is that farmers are unlikely to switch between these crops and others in the short and medium run due to their long gestation period. This brings down the number of crops in the analysis to 25. However, I show that my results are robust to including all crops.

As mentioned in section 2.2, I use a matching-based DID estimation strategy to minimize selection bias. That is, I first obtain a matched sample of treated and non-treated villages based on a set of observed village characteristics before conducting DID estimation. Particularly, I use the *gmatch* command in STATA for its handiness in panel data setting. For each treated village, the *gmatch* algorithm finds, a non-treated village(s) that has/have the closest observed characteristics or propensity score. I conduct a DID estimation on these matched samples of treated and non-treated villages. Figure 1.3 shows the histogram of propensity score by treatment status. The figure clearly shows that the region of common support is large as very few non-treated villages lie outside the common support. Table 1.3 reports the balancing of the matching variables. All the t-statistics are insignificant and the bias percentage is small.

In what follows, I report results for both matching-based DID estimation and DID estimation (without matching), but most of my discussions are based on results for matching-based DID estimation.

Table 1.4 presents the main results. The table clearly shows that villages that got road connection reallocated more land to their CA-crops. Panel A uses binary treatment dummy. The first column shows that area of land allocated to the top 20% crops (in relative yield rankings) increased by about 6.3% (0.01/0.159), relative to the land share of these crops in 2012, following road connection. The fourth columns show that the estimated reallocation of land towards CA-crops decreases. This should be the case because as we relax the cutoff for definition of CA-crops we classify more and more crops as CA-crops even if the village is not strongly more productive in those crops relative to national average. Panel B shows similar results using market access measure instead of binary treatment dummy. We see that the fraction of land allocated to CA-crops increases significantly for villages that have seen an increase in their market access, and that this effect becomes weaker as we relax the cutoff for CA-crop definition. Moreover, the estimated reallocation effects are roughly the same magnitude as those in Panel A. Table 1.5 reports the corresponding results using DID estimation (without matching). These results are very similar to those in table 1.4.

I explore two robustness exercises. First, I include all the 45 crops reported in AgSS data for which complete information is available. Table A.1 shows the estimation result. We see that the estimated increases in the share of land allocated CA-crops is slightly smaller than those in table 1.4 in absolute terms. However, relative to pre-URRAP average land share of these crops the effect is slightly larger because the average land share of CA-crops decreases significantly when less common crops are included. In short, the fraction of land allocated to CA-crops increased by 9.3% (0.009/0.097). In the second robustness exercise, I use FAO-GAEZ yield measure. This data includes only 19 crops.[12] Table A.2 reports the results. The estimated (percentage) increases in fraction of land allocated to CA-crops is about 19%, which larger in magnitude compared to those in tables 1.4 and A.1. This is mostly driven by the selection of crops – GAEZ data predominantly includes cereals.

Overall, these results clearly show that villages reallocate farmland towards their CA-crops following improvement in market integration.

1.4.2 New roads and crop prices

As trade costs between a village and its trading partners decrease due to new road construction, the relative prices of the village's CA-crops increase (or the relative prices of the village's CD-crops decrease). This prediction of the model can be directly tested, given data on village level prices of crops before and after the program and our CA-crop definition in the previous subsection. I use the following regression to test this prediction:

$$
\begin{aligned}
\ln p^k_{vmt} &= \alpha_1(Post_t * URRAP_v) + \alpha_2 CA^k_v + \alpha_3(Post_t * URRAP_v * CA^k_v) \\
&\quad + \gamma_v + \gamma_{mt} + \gamma^k_m + \varepsilon^k_{vt}
\end{aligned}
$$

where $\ln p^k_{vmt}$ denotes the log price of crop k in village v in month m of year t. CA^k_v is dummy variable indicating whether crop k is among the village's CA-crops. I

[12] Some of these crops are rarely produced. Only 11 of the 19 crops are produced across many villages. As a result, in regressions using this data, I drop crop fixed effects because including them increases multicollinearity with CA-crop dummy variables.

include crop-month fixed effects to account for seasonal fluctuation in crop prices.

I use the AgPPS data for the same 25 major non-tree crops in the main analysis. This data covers about 450 villages (after cleaning for missing information). The results for matching-based DID estimation are reported in table 1.6. Panel A uses binary treatment dummy whereas Pane B uses log market access. Both specification clearly show that relative prices of CA-crops increased in villages that got road connection, relative to villages that did not get road connection. Relative prices of CA-crops increased by 4%. Table 1.7 reports results for DID estimation (without matching), and results are very similar to those in 1.6. As a robustness check, I use CA-crops defined based on GAEZ yield estimates. Table A.3 reports the results. The results are qualitatively similar to those in 1.7, but quantitatively we see larger increase in the prices of CA-crops, an increase of 9%. Part of this is driven driven by selection of crops.

Table 1.8 reports the resulting effect on village (monthly) price index. The village price index is computed using equation 1.3 and the elasticity of substitution estimated in section 1.5. The results in table 1.9 show that the village price index also increased significantly, by about 3%. This not surprising given that the consumption basket of villagers are likely to be biased towards locally produced CA-crops. Table 1.9 reports the DID estimation results without matching. As discussed in the theoretical section, the welfare effect of decreases in trade costs is ambiguous because it may lead to an increase in cost of consumption baskets if consumption baskets are biased towards local CA-crops. The data clearly supports this conjecture as village price indices, driven by increases in prices of CA-crops, increased significantly following road connection.

1.5 Estimation of parameters and welfare gain

1.5.1 Estimation of model parameters

Estimation of σ: I use the expression for the expenditure share of a crop and ESS data on household consumption disaggregated by crops to estimate the σ. I estimate the following equation:

$$\ln s_{it}^k = \alpha_0 - (\sigma - 1)\ln p_{it}^k + \delta_1 E_{it} + \delta_2 \mathrm{H} + \gamma_t + \gamma_r^k + \gamma_z + \varepsilon_{it}^k \tag{1.8}$$

where the expenditure share equation is augmented by including household annual expenditure E (a proxy for income), household size H, year fixed effects γ_t, region-crop fixed effects γ_r^k to account for regional variation in crop tastes, and a set of dummy variables for agro-ecological zones.

There is a well known endogeneity concern in estimation of equation 1.8 because unobserved factors could be correlated with both prices and the expenditure share of crops or measurement error in price could attenuate the coefficient of price. To address this, I use GAEZ yield measure to instrument for prices. GAEZ yield measure is a valid instrument because locations that have higher crop productivity have lower prices for the crop. The first-stage F-stat is 24.6. I obtain the coefficient of price of -0.3 (s.e. 0.07), which implies sigma of 1.3. While this estimate looks too small at first glance, it is in fact in line with other studies estimating the elasticity of substitution across food varieties in low income countries. For instance, Behrman and Deolalikar (1989) find $\hat{\sigma} \approx 1.2$. But other studies such as Sotelo (2018) find larger estimate of about 2.3 using data from Peru.[13]

Estimation of θ: I follow Sotelo (2018) in the estimation of θ. I rely on AgSS village level crop yield estimate that is constructed based on a random sample of crop cut. To purge out the noise in yield estimate and fluctuations due to whether conditions, I take the average across all years (2010-2016) to obtain a time invariant

[13] Sotelo (2018) admits that his estimate of σ is the upper bound when compared to similar studies from developing countries.

measure of yield for a crop in a village. I assume that the true crop productivity in a village A_v^k is related to the AgSS's village yield measure Y_v^k in the following equation:

$$A_v^k = \delta^k Y_v^k \exp(-u_v^k)$$

where $\exp(-u_v^k)$ is a random noise, and δ^k is crop-specific constant. Plugging this for A_v^k in the land share equation and taking logs gives:

$$\ln\Big(P_v^k Y_v^k\Big) = \frac{1}{\theta}\eta_v^k + \ln\Phi_v - \ln\delta^k + u_v^k$$

The empirical counterpart of this is:

$$\ln\Big(P_{vt}^k Y_v^k\Big) = \frac{1}{\theta}\eta_{vt}^k + \gamma_v + \gamma^k + \gamma_t + u_{vt}^k$$

where γ_v and γ^k are village and crop fixed effects respectively.

I obtain a value of $\hat{\theta} = 2.7$ for productivity dispersion, which is larger than the estimate of Sotelo (2018) around 1.7, but smaller than that of Donaldson (2018) who reports a mean of about 7.5 across the 17 crops in his data. However, Donaldson (2018) estimates θ from a gravity equation using data on trade flows between regions in colonial India.

1.5.2 Estimation of welfare gain

Recall that the measure of welfare W_{vt} in the model is real revenue $\frac{R_{vt}}{P_{vt}}$, which is also equal to real rental rate $\frac{r_{vt}}{P_{vt}}$. Given data on crop production and prices before and after the program, I can directly construct village revenue as $R_{vt} = \sum_k p_{vt}^k y_{vt}^k$. However, my price data (AgPPS data) does not cover all the villages for which agricultural production data is available (AgSS villages). I impute missing prices as follows. For each village in AgSS data that got road connection under URRAP, I find the nearest village in AgPPS data that also got road under URRAP. Similarly, for each village in AgSS data that did not get road under URRAP, I find the nearest village in the AgPPS data that did not get road under URRAP. I use these imputed prices to calculate village revenue. The price index is also calculated using the same data and the elasticity of substitution estimated in the previous subsection: $P_{vt} = (\sum_k {p_{vt}^k}^{1-\sigma})^{\frac{1}{1-\sigma}}$.

To estimate the welfare gain from road connectivity, I run the following regression:

$$\ln W_{vt} = \alpha_1 Post_t + \alpha_2 (Post_t * URRAP_v) + \delta \mathbf{X} + \gamma_v + \varepsilon_{vt} \tag{1.9}$$

where X includes a vector of village characteristics such as rainfall. I also run similar regression using market access measure instead of binary treatment.

Finally, in order to estimate how much of the estimated welfare gain in the above reduced-form estimation is attributed to the trade mechanism suggested in the model, I use the sufficient statistics derived from model and mediation analysis. First, I construct the sufficient statistics using equation 1.6 and proposition 2 as $\ln\Lambda_{vt} = \frac{1}{\sigma-1}\ln \sum_k A_v^{k\sigma-1} {\eta_{vt}^k}^{\frac{1-\sigma}{\theta}}$, given data on village yield A_v^k and land share of crop η_{vt}^k, and the values of parameters θ and σ estimated in the previous subsection. Then, I rerun equation 1.9 by including $\ln\Lambda_{vt}$ on the right-hand side:

$$\ln W_{vt} = \beta_1 Post_t + \beta_2 (Post_t * URRAP_v) + \beta_3 \ln\Lambda_{vt} + \delta \mathbf{X} + \gamma_v + \varepsilon_{vt} \tag{1.10}$$

The logic behind this mediation analysis is as follows. Because $\ln\Lambda_{vt}$ captures the welfare gain from road that is mediated through the trade mechanism suggested in our model, once this term is included on the right-hand side of 1.10 the coefficient of road access (β_2 in equation 1.10) captures only the welfare effect of roads that is mediated through non-trade mechanisms. Thus, α_2 captures the total welfare gain from road, β_2 captures part of this welfare gain that is attributed to non-trade mechanisms, and the difference $\alpha_2 - \beta_2$ captures the welfare gain attributed to the trade mechanism.[14] If roads affected real revenue *only* through the mechanism suggested in the model, then β_2 in equation 1.10 should be small and statistically

[14]See this excellent note on how mediation analysis works and the procedures to implement it http://web.pdx.edu/~newsomj/semclass/ho_mediation.pdf.

insignificant. If instead, β_2 is statistically significant but significantly smaller that α_2, then the trade mechanism captures only a portion of the effect of roads on real revenue.

Table 1.10 reports the estimation results for equations 1.9 and 1.10 using the matching-based DID strategy. Panel A uses dummy variable as treatment measure while panel B uses market access approach. Column 1 reports results for estimation of 1.9. From panel A, we see that real agricultural income per a hectare of land (a measure of welfare) increased by 12.5% for treated villages compared to non-treated villages over the period of 2012 to 2016. This is the total welfare gain from URRAP roads, which can be attributed to a number of mechanisms. From column 2, we see that the sufficient statistic measure of welfare derived from the model $\ln\Lambda_{vt}$ increases by 16.2% for the treated villages compared to non-treated villages. This is important for the validity of the mediation analysis because it proves that the roads significantly increase measure of welfare in the model. Column 3 reports estimation result for equation 1.10. Once the model predicted welfare $\ln\Lambda_{vt}$ is included in the right-hand side, the coefficient of $Post_t * URRAP_v$ decreases to 0.05 and loses its significance. Comparing the coefficient of $Post_t * URRAP_v$ in columns 1 and 3 implies that 7.5% (12.5.9%-0.05%), out of the total 12.5% welfare gain, is attributed to the trade mechanism. Panel B tells similar story using market access approach. However, the coefficient of $Post_t * URRAP_v$ in column 3 remains statistically significant (this is due to large variation in market access as opposed to the binary treatment dummy). Nevertheless, panel B also shows that significant fraction of the overall welfare gain is attributed to the trade mechanism.

In table 1.11, I report the DID estimation with out matching. Unlike in the previous subsection where matching treated and non-treated villages did not have significant effect on the estimation results, here the non-matching based results are significantly larger than the matching-based results, implying upward selection bias. Perhaps, the reason why such selection bias was not prevalent in the previous section where we estimated the effect of road connection on land reallocation and prices of CA-crops is the triple-difference specification where treatment is interacted with dummy variables for CA-crops.

Overall, the results in table 1.10 show that the URRAP roads have led to significant increases in village welfare, despite increases in village price index, implying that nominal income increased by significantly larger. Moreover, improvement in trade opportunity is the main mechanism through which the roads improved welfare.

1.5.3 Heterogeneity in welfare gain: cereal vs. non-cereal villages

One of the key testable implications from the model is that the welfare gain from roads is heterogeneous and depends on the fraction of land allocated to different crops in a village vis-a-vis the expenditure share of these crops in village consumption. In particular, villages that specialize in cereals should gain less from roads because they experience an increase in the prices of their consumption baskets while non-cereal producing villages experience the opposite.

Table 1.12 presents the empirical evidence supporting this result. I interact the fraction of village farmland allocated to cereals with the treatment dummy (in the first panel) and the market access measure (in the second panel). To facilitate interpretation, the land share of cereal is standardized. The table shows that the welfare gain from roads is decreasing in the fraction of land allocated cereals, i.e., villages that allocate higher fraction of land to cereal gain significantly less than those that allocate lower fraction of land to cereals.

1.5.4 Remoteness and welfare gain from new roads

An important policy question is how the gain from infrastructure expansion, such road connection, is distributed among villages of different characteristics. Would remote villages gain significantly more or less compared to villages that are located near roads or population centers. Theoretically, the answer is ambiguous. On one hand non-remote villages would face competition from remote villages that may now have improved access to market centers. On the other hand, the decrease in trade costs might not be large enough for remote villages to engage in trade with distant markets but significant enough for villages near markets to engage in trade. In this subsection, I use villages' distance from nearest town (population center with more than 20,000 population), distance to major trunk roads (which includes inter-state roads, roads connecting the zones to capital city and roads that connect to neighboring countries), and distance to pre-URRAP road networks as measure of remoteness.

Table 1.13 reports the estimation result. To facilitate interpretation, the distance measures are standardized so that the coefficients are interpreted as the effect of one standard deviation increase in distance. Panel A reports results using the binary treatment indicator and Panel B uses log market access. Overall the result shows that remote villages gain significantly less than villages near population centers or roads. Using the results in column 1 of Panel A, villages that are one standard deviation farther than towns gain 14.5% less than villages with average distance from towns (the average and standard deviation of distance from town are about 13km and 9km, respectively). Using the market access approach gives an even precise and robust evidence across the different measures of remoteness used.

1.6 Conclusions

In this chapter, I estimate the welfare gain from a massive rural road expansion in Ethiopia. I develop a Ricardian trade model with multi-crop multi-location feature where land productivity is allowed to vary both within and across villages to study the key mechanisms through which roads affect village welfare. On the demand side, a representative farmer in each village maximizes utility by choosing optimal quantities of crops to consume, given prices. On the production side, the representative farmer decides how to allocate its limited land across potential crops, given local prices and local productivity of crops. The village also engage in costly trade with other villages.

The model gives sharp predictions about the effects of decreases in trade costs on village land allocation across crops and village prices of crops. However, the model implies that the net welfare effect of roads on village welfare is ambiguous. While a decrease in trade cost leads to increases in village income by inducing reallocation of crops towards village CA-crops and enables farmers receive better prices for these crops, it also leads to increase in the cost of consumption baskets for the villagers. The net effect on welfare depends on the strength of these two contrasting forces, which in turn depends on the composition of the villages' CA-crops vis-a-vis their consumption baskets.

I directly test the model's predictions using micro data on agricultural production, crop prices, and geospatial data on the entire road network before and after the road expansion program. First, the road expansion led to significant increases in the relative prices of CA-crops by 4% for villages that got road connection relative those that did not. This led to about 3% increase in the overall village price index for the treated villages relative to non-treated ones. Second, the road expansion led to reallocation of farmland towards a village's CA-crops for the treated villages relative to non-treated ones. The fraction of land allocated to CA-crops increased by about 6% in the treated villages relative to non-treated villages.

Finally, I estimate the welfare gain from the road expansion in a reduced-form regression and decompose the estimated total welfare gain from road into the trade mechanism captured in my model and other potential mechanisms not in the model. To do so, I derive a sufficient statistic for the welfare effect of roads from the model. This sufficient statistic is similar to those in workhorse Ricardian trade models, but it can be inferred only from data on land allocation and the model parameters, which I directly estimate. I estimate a total welfare gain of 12.5% from the road expansion between 2012-2016, and 7.5% of this is attributed to the improvements in the trade opportunities for the villagers.

While the chapter is silent on what other mechanisms explain the remaining 5%, several other mechanisms through which roads can improve welfare can be imagined. In particular, future studies may explore how road connectivity might affect accessibility of health services and agricultural inputs, such as fertilizer and extension services, which are critical for agricultural productivity. Anecdotal reports show that construction of rural roads have significantly improved maternal and child health by facilitating faster access to health stations using modern transport.

Figure 1.1: Rural road expansion under URRAP

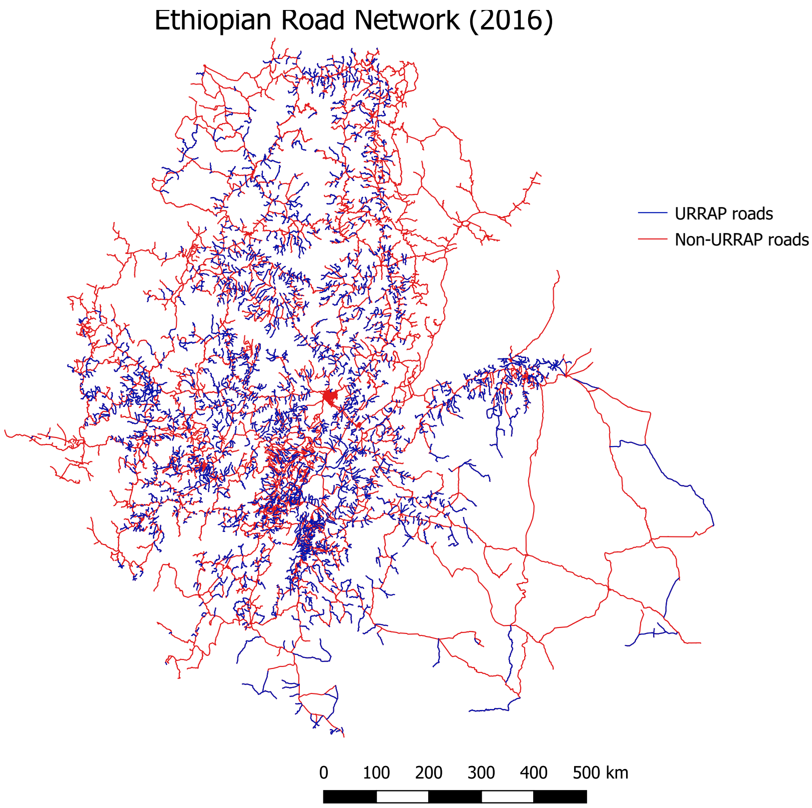

Figure 1.2: Completed URRAP roads (pictures are taken from Oromia Roads Authority).

Figure 1.3: Common support of propensity score matching

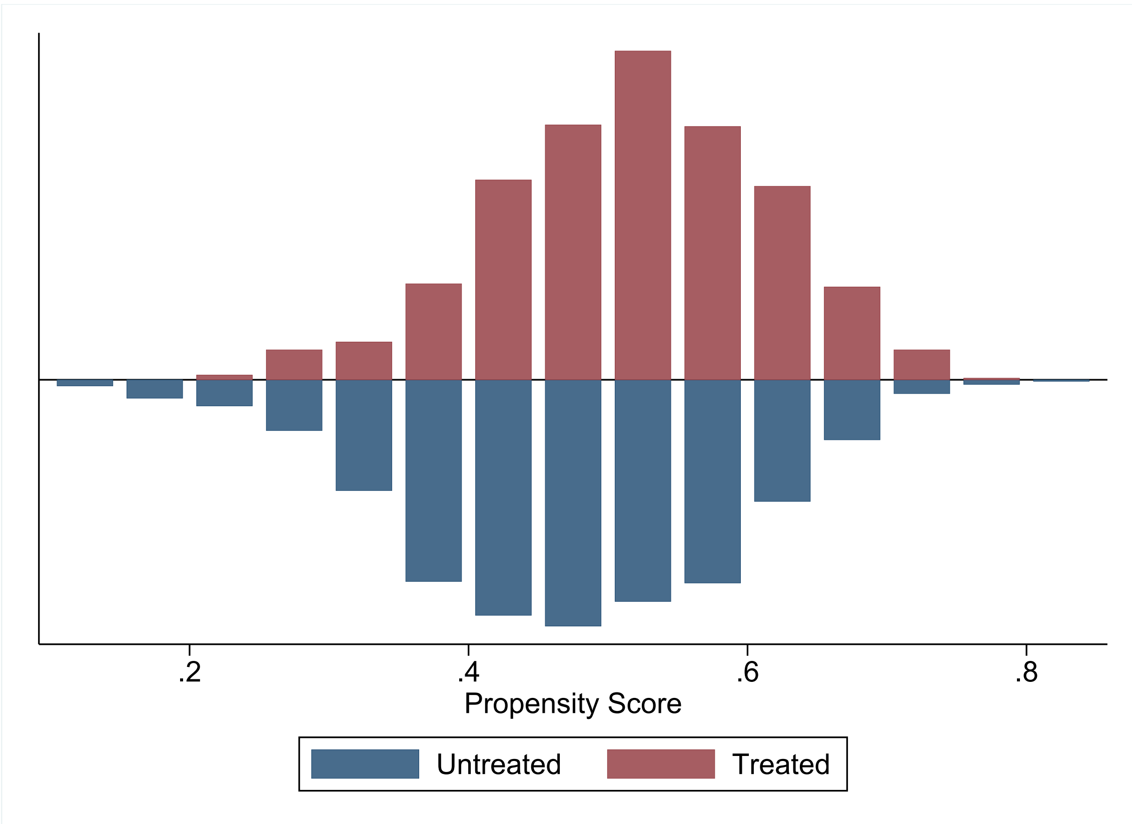

Table 1.1: URRAP road access and trade costs

	Dependent Variable: log(Price in Zone Capital/Price in village)		
	All crops	Vegetables and Fruits	Cases where dep. var >0
Post*URRAP	-0.031**	-0.079*	-0.024*
	(0.016)	(0.044)	(0.013)
N	82944	24468	67147
R^2	0.378	0.360	0.493

Notes: Standard errors are clustered at village level. This table is based on AgPPS and RPS datasets. The regression includes 422 villages, 57 urban centers, and 56 crops. All regressions include village, crop-month, and year fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.2: Rural roads and the link between local prices and local yield: the dependent variable is crop-village level prices in 2012 and 2015.

	(1)	(2)	(3)
	Panel A: Binary Treatment		
LogYield	-0.036***		-0.027***
	(0.003)		(0.003)
Post*URRAP		0.017	-0.083***
		(0.023)	(0.023)
LogYield*Post*URRAP			0.009***
			(0.003)
N	59270	59270	59270
R^2	0.752	0.739	0.776
	Panel B: Market access approach		
LogYield	-0.036***		-0.099***
	(0.003)		(0.026)
LogMarketAccess		-0.026**	-0.043**
		(0.012)	(0.020)
LogYield*LogMarketAccess			0.006**
			(0.003)
N	59270	59270	59270
R^2	0.790	0.780	0.795

Notes: Standard errors are clustered at village level. The regression includes 277 villages, and 20 crops. All regressions include crop and year fixed effects, and log rainfall as a control. The last column includes village fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.3: Balancing of variables for Average Treatment Effect on Treated (ATT)

	Treated	Control	% bias	t-stat	p-value
Distance to nearest town (meters)	10596	10417	2.2	0.57	0.571
Distance to nearest trunk road (meters)	4551.8	4280.5	6.2	1.47	0.143
Distance to preexisting road network (meters)	1320.5	1221.5	3.6	0.93	0.351
Population	5340.9	5458.1	-3.1	-0.63	0.528
Average slope (degrees)	9.9697	10.018	-0.9	-0.20	0.840
Average altitude (meters)	1946.2	1959.1	-2.3	-0.54	0.587
Rainfall (mm)	1183	1161.1	6.6	1.50	0.133

Notes: Population and rainfall correspond to the period before URRAP.

Table 1.4: Road construction and reallocation of farmland towards CA-crops: Matching-Based DID estimation

	Definition of CA-crops:			
	Top 20%	Top 30%	Top 40%	Top 50%
	Panel A: Binary treatment			
CA-crop	0.041***	0.049***	0.055***	0.056***
	(0.003)	(0.003)	(0.003)	(0.003)
Post*URRAP	-0.004**	-0.005**	-0.006***	-0.005**
	(0.002)	(0.002)	(0.002)	(0.002)
CA-crop*Post*URRAP	0.010**	0.010**	0.010**	0.007**
	(0.005)	(0.004)	(0.004)	(0.004)
N	28030	28030	28030	28030
adj. R^2	0.310	0.318	0.324	0.324
	Panel B: Market access approach			
CA-crop	-0.061	-0.047	-0.009	0.028
	(0.039)	(0.034)	(0.032)	(0.029)
LogMarketAccess	-0.003	-0.004*	-0.004	-0.002
	(0.002)	(0.002)	(0.002)	(0.002)
CA-crop*LogMarketAccess	0.011***	0.010***	0.007**	0.003
	(0.004)	(0.003)	(0.003)	(0.003)
N	28030	28030	28030	28030
adj. R^2	0.310	0.318	0.324	0.324
Mean land share (CA-crops) in 2012	0.159	0.162	0.160	0.153
Mean land share (all) in 2012	0.118	0.118	0.118	0.118

Notes: Standard errors are clustered at village level. This regression analysis includes 25 major non-tree crops. Each column represents different cutoff for definition of CA-crops. In column 1, CA-crops are crops in the top 20% of within village ranking of crops based on yield relative to national average. In column 2 CA-crops are top 30% in the ranking, and so forth. All regressions include village, crop and year fixed effects. Panel A uses binary treatment dummy. Panel B uses market access measure constructed from the entire road network before and after the program, and pre-program spatial population distribution. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.5: Road construction and reallocation of farmland towards CA-crops: DID estimation

	Definition of CA-crops:			
	Top 20%	Top 30%	Top 40%	Top 50%
	Panel A: Binary treatment			
CA-crop	0.043***	0.051***	0.056***	0.056***
	(0.003)	(0.003)	(0.002)	(0.002)
Post*URRAP	-0.003*	-0.003*	-0.004*	-0.004**
	(0.002)	(0.002)	(0.002)	(0.002)
CA-crop*Post*URRAP	0.009*	0.007*	0.006*	0.006*
	(0.005)	(0.004)	(0.004)	(0.003)
N	33821	33821	33821	33821
adj. R^2	0.324	0.332	0.338	0.337
	Panel B: Market access approach			
CA-crop	-0.058	-0.022	0.019	0.040
	(0.036)	(0.032)	(0.029)	(0.028)
LogMarketAccess	-0.002	-0.002	-0.001	-0.001
	(0.002)	(0.002)	(0.002)	(0.002)
CA-crop*LogMarketAccess	0.011***	0.008**	0.004	0.002
	(0.004)	(0.003)	(0.003)	(0.003)
N	33821	33821	33821	33821
adj. R^2	0.324	0.332	0.338	0.337
Mean land share (CA-crops) in 2012	0.159	0.162	0.160	0.153
Mean land share (all) in 2012	0.118	0.118	0.118	0.118

Notes: Standard errors are clustered at village level. This regression analysis includes 25 major non-tree crops. Each column represents different cutoff for definition of CA-crops. In column 1, CA-crops are crops in the top 20% of within village ranking of crops based on yield relative to national average. In column 2 CA-crops are top 30% in the ranking, and so forth. All regressions include village, crop and year fixed effects. Panel A uses binary treatment dummy. Panel B uses market access measure constructed from the entire road network before and after the program, and pre-program spatial population distribution. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.6: New road construction and prices of CA-crops: Matching-Based DID estimation

	Definition of CA-crops:			
	Top 20%	Top 30%	Top 40%	Top 50%
	Panel A: Binary treatment			
CA-crop	-0.042***	-0.046***	-0.053***	-0.048***
	(0.013)	(0.011)	(0.011)	(0.012)
Post*URRAP	0.024	0.024	0.023	0.042*
	(0.020)	(0.020)	(0.021)	(0.023)
CA-crop*Post*URRAP	0.040*	0.027	0.020	-0.013
	(0.023)	(0.020)	(0.018)	(0.020)
N	35154	35154	35154	35154
adj. R^2	0.798	0.798	0.798	0.798
	Panel B: Market access approach			
CA-crop	-0.434***	-0.348***	-0.326***	-0.249*
	(0.146)	(0.121)	(0.116)	(0.136)
LogMarketAccess	0.023	0.021	0.018	0.019
	(0.027)	(0.027)	(0.028)	(0.029)
CA-crop*LogMarketAccess	0.041***	0.032***	0.028**	0.020
	(0.015)	(0.012)	(0.012)	(0.014)
N	35154	35154	35154	35154
adj. R^2	0.798	0.798	0.798	0.798

Notes: Standard errors are clustered at village level. The dependent variable is log village crop price. The analysis includes 25 major non-tree and crops over 450 nationally representative rural villages. Each column represents different cutoff for definition of CA-crops. In column 1, CA-crops are crops in the top 20% of within village ranking of crops based on yield relative to national average. In column 2 CA-crops are top 30% in the ranking, and so forth. All regressions include village, crop-month and year fixed effects. Panel A uses binary treatment dummy. Panel B uses market access measure constructed from the entire road network before and after the program, and pre-program spatial population distribution. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.7: New road construction and prices of CA-crops: DID estimation

	Definition of CA-crops:			
	Top 20%	Top 30%	Top 40%	Top 50%
	Panel A: Binary treatment			
CA-crop	-0.040***	-0.042***	-0.048***	-0.039***
	(0.013)	(0.011)	(0.011)	(0.011)
Post*URRAP	0.025	0.022	0.023	0.041**
	(0.019)	(0.019)	(0.019)	(0.021)
Post*URRAP * CA-crop	0.041*	0.028	0.023	-0.009
	(0.021)	(0.018)	(0.017)	(0.018)
N	41594	41594	41594	41594
adj. R^2	0.799	0.786	0.799	0.799
	Panel B: Market access approach			
CA-crop	-0.352**	-0.282**	-0.298***	-0.196
	(0.140)	(0.114)	(0.104)	(0.121)
LogMarketAccess	0.015	0.014	0.010	0.012
	(0.026)	(0.026)	(0.027)	(0.027)
LogMarketAccess * CA-crop	0.033**	0.025**	0.026**	0.016
	(0.014)	(0.012)	(0.011)	(0.012)
N	41594	41594	41594	41594
adj. R^2	0.799	0.799	0.799	0.799

Notes: Standard errors are clustered at village level. The dependent variable is log village crop price. The analysis includes 25 major non-tree and crops over 450 nationally representative rural villages. Each column represents different cutoff for definition of CA-crops. In column 1, CA-crops are crops in the top 20% of within village ranking of crops based on yield relative to national average. In column 2 CA-crops are top 30% in the ranking, and so forth. All regressions include village, crop-month and year fixed effects. Panel A uses binary treatment dummy. Panel B uses market access measure constructed from the entire road network before and after the program, and pre-program spatial population distribution. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.8: New road construction and village price index: Matching-Based DID estimation

	LogPriceIndex	LogPriceIndex
	Panel A: Based on AgSS yield	
Post	0.237***	0.222***
	(0.025)	(0.025)
Post * URRAP	0.030**	
	(0.015)	
LogMarketAccess		0.074***
		(0.021)
N	8193	8193
adj. R^2	0.977	0.977
	Panel B: Based on GAEZ yield	
Post	0.237***	0.222***
	(0.026)	(0.026)
Post * URRAP	0.030**	
	(0.015)	
LogMarketAccess		0.074***
		(0.021)
N	8193	8193
adj. R^2	0.977	0.977

Notes: Robust standard errors in parenthesis. The regression includes 450 villages. The price index is computed from the 25 major non-tree crops in the main analysis and using elasticity of substitution estimated in section 1.5. All regressions include village and month-year fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.9: New road construction and village price index: DID estimation

	LogPriceIndex	LogPriceIndex
	Panel A: Based on AgSS yield	
Post	0.230***	0.228***
	(0.023)	(0.023)
Post * URRAP	0.028*	
	(0.014)	
LogMarketAccess		0.038*
		(0.020)
N	9680	9680
adj. R^2	0.978	0.978
	Panel B: Based on GAEZ yield	
Post	0.230***	0.228***
	(0.023)	(0.023)
Post * URRAP	0.028*	
	(0.014)	
LogMarketAccess		0.038*
		(0.020)
N	9680	9680
adj. R^2	0.978	0.978

Notes: Robust standard errors in parenthesis. The regression includes 450 villages. The price index is computed from the 25 major non-tree crops in the main analysis and using elasticity of substitution estimated in section 1.5. All regressions include village and month-year fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.10: The welfare gains from new rural roads: Matching-Based DID estimation

	ln$W$	lnΛ	ln$W$
	Panel A: Binary treatment		
Post*URRAP	0.125*	0.162*	0.050
	(0.067)	(0.096)	(0.050)
lnΛ			0.461***
			(0.016)
N	3268	3268	3268
R^2	0.863	0.814	0.923
	Panel B: Market access approach		
LogMarketAccess	0.174**	0.113	0.122**
	(0.080)	(0.112)	(0.059)
lnΛ			0.461***
			(0.015)
N	3268	3268	3268
R^2	0.863	0.814	0.923

Notes: Robust standard errors in parenthesis. All regressions include year and village fixed effects. The estimation is based on 25 non-tree crops for which full information on prices and quantities of production is available. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.11: The welfare gains from new rural roads: DID estimation

	ln$W$	lnΛ	ln$W$
	Panel A: Binary treatment		
Post*URRAP	0.169***	0.264***	0.050
	(0.061)	(0.088)	(0.046)
lnΛ			0.453***
			(0.014)
N	4240	4240	4240
adj. R^2	0.789	0.686	0.881
	Panel B: Market access approach		
LogMarketAccess	0.224***	0.211**	0.129**
	(0.073)	(0.103)	(0.055)
lnΛ			0.453***
			(0.014)
N	4240	4240	4240
adj. R^2	0.789	0.685	0.881

Notes: Robust standard errors in parenthesis. All regressions include year and village fixed effects. The estimation is based on 25 non-tree crops for which full information on prices and quantities of production is available. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.12: Welfare gain: cereal vs non-cereal villages

	Non-matched	Matched sample
	Panel A: Binary treatment	
Post*URRAP	0.165***	0.122*
	(0.061)	(0.067)
Post*URRAP*CerealShare	-0.093**	-0.052
	(0.046)	(0.049)
N	4180	3266
adj. R^2	0.784	0.725
	Panel B: Market access approach	
LogMarketAccess	0.239***	0.178**
	(0.073)	(0.080)
LogMarketAccess*CerealShare	-0.009***	-0.008***
	(0.003)	(0.003)
N	4180	3266
adj. R^2	0.785	0.727

Notes: Robust standard errors in parenthesis. The first column reports results without matching. The second column reports results based on matching. All regressions include year and village fixed effects. The estimation is based on 25 non-tree crops for which full information on prices and quantities of production is available. $CerealShare$ is the share of farmland allocated to cereal crops in a village. Cereal crops include: Barley, Wheat, Maize, Teff, Sorghum, Millet, and Enset. Non-cereal crops include all vegetables, legumes and cash-crops which are predominantly produced for market. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.13: Remoteness and welfare gain from new roads

	The dependent variable is log real revenue ln$W$		
	Distance to town	Distance to trunk roads	Distance to pre-URRAP roads
	Panel A: Binary treatment		
Post*URRAP	0.091	0.113*	0.113*
	(0.063)	(0.062)	(0.062)
Post*URRAP*Distance	-0.145***	-0.061	-0.026
	(0.056)	(0.050)	(0.051)
N	3989	3989	3989
adj. R^2	0.759	0.759	0.758
	Panel B: Market access approach		
LogMarketAccess	0.092	0.219***	0.193**
	(0.079)	(0.078)	(0.078)
LogMarketAccess * Distance	-0.218**	-0.205***	-0.128**
	(0.085)	(0.063)	(0.061)
N	3989	3989	3989
adj. R^2	0.760	0.760	0.759

Notes: Robust standard errors in parenthesis. Distance measure is standardized so that the coefficients are interpreted as the effect of one standard deviation increase in distance. The first column uses distance to the nearest town with above 20,000 population. The second column uses distance to the nearest trunk road (inter-state roads and roads connecting zone capitals to the center). The third column uses distance to pre-existing road network. All regressions include year and village fixed effects. The estimation is based on 25 non-tree crops for which full information on prices and quantities of production is available. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendices

A Construction of market access measure

I follow Donaldson and Hornbeck (2016) to construct market access measure for each village derived from general equilibrium trade models. I use spatial distribution of population before the road program, and pre- and post-program entire road network of the country to construct the market access measure.

$$MarketAccess_{ot} = \sum_{d} \tau_{odt}^{-\theta} Population_d \tag{11}$$

where $Population_d$ is destination village population from the 2007 census (before the onset of the URRAP program). Using pre-URRAP population distribution is necessary because population distribution is likely to respond to improvement in road infrastructure. θ is the inverse land heterogeneity parameter (which governs trade elasticity), which I estimate in section 1.5.

τ_{odt} is the freight costs of transporting one ton of cargo from origin village o to destination village d along the least cost path, before ($t = 0$) and after ($t = 1$) the construction of URRAP roads.[15] I use the following procedure to estimate τ_{odt} for each year. First, I construct a link from each village centroid to the nearest available road in year t. Next, I use data on costs of moving weight (in USD per ton-kilometer) for five different road quality levels: asphalt, major gravel, cobbled road, minor gravel, and earth road. Because there is no similar cost estimates along the link roads, I scale up the costs along earth road by the factor of $\frac{Cost\ along\ earth\ road}{Cost\ along\ minor\ gravel}$ to obtain estimate of cost along the links.[16] After assigning each road type (including the links) with the estimated costs in USD per ton-kilometer, I use ArcGIS tools to calculate the costs (in USD) of moving a ton of weight from origin o to destination d along the least cost path, in each year. I use these estimates as τ_{odt}. As can be seen

[15] Alternatively, I use travel time along the least (time) cost path, instead of freight costs. The market access measures are strongly correlated (correlation of 0.92).

[16] I show that the results are robust to using alternative scales that are half or twice of the baseline scale $\frac{Cost\ along\ earth\ road}{Cost\ along\ minor\ gravel}$.

in equation 14, a change to a village's market access comes only from changes in τ_{odt}.

B Derivation of the conditional distribution of productivity and rental rate

Because the distribution of rental rate of a plot depends on the distribution of productivity of land, we need to first derive the distribution of land productivity conditional on the land being used for crop k, i.e., $z_i^k(\omega)|\omega \in \Omega_i^k$, which I denote as $G_i^k(t)$. This derivation of conditional distribution of land quality is similar to Sotelo (2018):

$$
\begin{aligned}
G_i^k(t) &= \mathcal{P}\Big[z_i^k(\omega) < t | p_i^k z_i^k(\omega) = \max_l p_i^l z_i^l(\omega)\Big] \\
&= \frac{\mathcal{P}\Big[z_i^k(\omega) < t \wedge p_i^k z_i^k(\omega) = \max_l p_i^l z_i^l(\omega)\Big]}{\mathcal{P}\Big[p_i^k z_i^k(\omega) = \max_l p_i^l z_i^l(\omega)\Big]} \\
&= \frac{1}{\eta_i^k}\mathcal{P}\Big[z_i^k(\omega) < t \wedge p_i^l z_i^l(\omega) < p_i^k z_i^k(\omega), \quad \forall l\Big] \\
&= \frac{1}{\eta_i^k}\mathcal{P}\Big[\frac{p_i^l}{p_i^k} z_i^l(\omega) < z_i^k(\omega) < t, \quad \forall l\Big] \\
&= \frac{1}{\eta_i^k}\int_0^t \Pi_{l\neq k}\mathcal{P}\Big[\frac{p_i^l}{p_i^k} z_i^l(\omega) < v\Big] f_i^k(v)dv
\end{aligned}
$$

Using the distribution of $z_i^k(\omega)$:

$$
\begin{aligned}
G_i^k(t) &= \frac{1}{\eta_i^k}\int_0^t \Pi_{l\neq k}\exp\Big((-A_i^l)^\theta(\frac{p_i^k}{p_i^l}v)^{-\theta}\Big) \quad \theta(A_i^k)^\theta v^{-\theta-1}\exp(-(A_i^k)^\theta v^{-\theta})dv \\
&= \frac{1}{\eta_i^k}\int_0^t \exp\Big(-(p_i^k v)^{-\theta}\sum_{l\neq k}(A_i^l p_i^l)^\theta\Big)\exp(-(A_i^k)^\theta v^{-\theta} \quad \theta(A_i^k)^\theta v^{-\theta-1})dv \\
&= \frac{1}{\eta_i^k}\int_0^t \exp\Big(-(p_i^k v)^{-\theta}\sum_{l}(A_i^l p_i^l)^\theta\Big) \quad \theta(A_i^k)^\theta v^{-\theta-1}dv \\
&= \int_0^t \exp\Big(-(p_i^k v)^{-\theta}\Phi_i^\theta\Big) \quad \theta\Phi_i^\theta(p_i^k)^{-\theta}v^{-\theta-1}dv \\
&= \exp\Big(-\Big(\frac{\Phi_i}{p_i^k}\Big)^\theta t^{-\theta}\Big)
\end{aligned}
$$

Thus, the distribution of productivity of the set plots in village i's which are covered by crop k is a Fréchet with the parameters $\frac{\Phi_i}{p_i^k}$ and θ. Notice that the average productivity of land covered with a crop decreases with the crop price. Intuitively, more and more land is allocated to a crop with higher price which leads to a decrease in the average quality of land allocated to the crop.

Recall that the rental rate on plot ω, conditional on ω being used for crop k, is given by $r(\omega) = \max_k\{p_i^k z_i^k(\omega)\}$. Thus the conditional distribution of rental rate $r(\omega)|\omega \in \Omega_i^k$ is Fréchet with parameters Φ_i and θ. That is, the rental rate of plots covered with different crops have the same distribution regardless of which crops are planted. This result follows from the property of the Fréchet distribution and the fact that $r(\omega)$ is homogeneous of degree one in crop prices.

C Appendix Tables

Table A.1: Road construction and reallocation of farmland towards CA-crops (all crops) : Matching-based DID estimation

	Definition of CA-crops:			
	Top 20%	Top 30%	Top 40%	Top 50%
	Panel A: Binary treatment			
CA-crop	0.021***	0.027***	0.031***	0.031***
	(0.002)	(0.002)	(0.002)	(0.001)
Post*URRAP	-0.003***	-0.003***	-0.003**	-0.002
	(0.001)	(0.001)	(0.001)	(0.001)
CA-crop*Post*URRAP	0.009***	0.008***	0.005**	0.002
	(0.003)	(0.002)	(0.002)	(0.002)
N	45412	45412	45412	45412
adj. R^2	0.338	0.344	0.347	0.345
	Panel B: Market access approach			
CA-crop	-0.040*	-0.024	-0.015	0.012
	(0.022)	(0.020)	(0.017)	(0.015)
LogMarketAccess	-0.002	-0.002*	-0.003**	-0.002
	(0.001)	(0.001)	(0.001)	(0.001)
CA-crop*LogMarketAccess	0.006***	0.005***	0.005***	0.002
	(0.002)	(0.002)	(0.002)	(0.002)
N	45412	45412	45412	45412
adj. R^2	0.338	0.344	0.347	0.345
Mean land share (CA-crops) in 2012	0.097	0.101	0.101	0.096
Mean land share (all) in 2012	0.073	0.073	0.073	0.073

Notes: Standard errors are clustered at village level. The analysis includes 45 crops. Each column represents different cutoff for definition of CA-crops. In column 1, CA-crops are crops in the top 20% of within village ranking of crops based on yield relative to national average. In column 2 CA-crops are top 30% in the ranking, and so forth. All regressions include village, crop and year fixed effects. Panel A uses binary treatment dummy. Panel B uses market access measure constructed from the entire road network before and after the program, and pre-program spatial population distribution. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.2: Road construction and reallocation of farmland towards CA-crops based on GAEZ yield: DID estimation

	Definition of CA-crops:			
	Top 20%	Top 30%	Top 40%	Top 50%
	Panel A: Binary treatment			
CA-crop	0.011*	0.033***	0.045***	0.049***
	(0.006)	(0.006)	(0.005)	(0.005)
Post*URRAP	-0.005	-0.007	-0.009*	-0.015***
	(0.004)	(0.004)	(0.005)	(0.005)
CA-crop*Post*URRAP	0.032***	0.029***	0.028***	0.032***
	(0.010)	(0.009)	(0.008)	(0.008)
N	22482	22482	22482	22482
adj. R^2	0.041	0.045	0.051	0.053
	Panel B: Market access approach			
CA-crop	-0.117	-0.093	-0.020	0.034
	(0.072)	(0.067)	(0.064)	(0.060)
LogMarketAccess	-0.000	-0.001	0.001	0.003
	(0.005)	(0.005)	(0.005)	(0.005)
CA-crop*LogMarketAccess	0.014*	0.014**	0.007	0.002
	(0.007)	(0.007)	(0.007)	(0.006)
N	22482	22482	22482	22482
adj. R^2	0.040	0.045	0.050	0.052
Mean land share (CA-crops)	0.166	0.173	0.176	0.172
Mean land share (all)	0.144	0.144	0.144	0.144

Notes: Standard errors are clustered at village level. This regression analysis includes 19 crops for which GAEZ yield measure is available. Each column represents different cutoff for definition of CA-crops. In column 1, CA-crops are crops in the top 20% of within village ranking of crops based on yield relative to national average. In column 2 CA-crops are top 30% in the ranking, and so forth. All regressions include village, crop and year fixed effects. Panel A uses binary treatment dummy. Panel B uses market access measure constructed from the entire road network before and after the program, and pre-program spatial population distribution. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.3: Road construction and prices of CA-crops based on GAEZ yield: Matching-based DID estimation

	Definition of CA-crops:			
	Top 20%	Top 30%	Top 40%	Top 50%
	Panel A: Binary treatment			
CA-crop	-0.083***	-0.090***	-0.093***	-0.075***
	(0.019)	(0.018)	(0.018)	(0.017)
Post*URRAP	0.001	-0.021	-0.035	-0.041
	(0.023)	(0.023)	(0.025)	(0.027)
CA-crop*Post*URRAP	0.075***	0.106***	0.117***	0.113***
	(0.024)	(0.024)	(0.028)	(0.028)
N	18178	18178	18178	18178
adj. R^2	0.759	0.760	0.760	0.759
	Panel B: Market access approach			
CA-crop	-0.558***	-0.456**	-0.366*	-0.317
	(0.205)	(0.193)	(0.196)	(0.196)
LogMarketAccess	0.030	0.027	0.029	0.028
	(0.034)	(0.034)	(0.034)	(0.034)
LogMarketAccess * CA-crop	0.050**	0.040**	0.031	0.028
	(0.021)	(0.020)	(0.020)	(0.020)
N	18178	18178	18178	18178
adj. R^2	0.759	0.759	0.759	0.758

Notes: Standard errors are clustered at village level. The dependent variable is log village crop price. This regression analysis includes 19 crops for which GAEZ yield measure is available. Each column represents different cutoff for definition of CA-crops. Each column represents different cutoff for definition of CA-crops. In column 1, CA-crops are crops in the top 20% of within village ranking of crops based on yield relative to national average. In column 2 CA-crops are top 30% in the ranking, and so forth. All regressions include village, crop-month and year fixed effects. Panel A uses binary treatment dummy and $RoadAccess = Post \times Treatment$. Panel B uses market access measure constructed from the entire road network before and after the program, and pre-program spatial population distribution. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Bibliography

Adamopoulos, T. (2018). Spatial integration, agricultural productivity, and development: A quantitative analysis of ethiopia's road expansion program. ICG Working Paper F-32404-ETH-1.

Allen, T. and Atkin, D. (2018). Volatility and the gains from trade. *Econometrica*, Forthcoming.

Arkolakis, C., Costinot, A., and Rodríguez-Clare, A. (2012). New trade models, same old gains? *American Economic Review*, 102(1):94–130.

Asher, S. and Novosad, P. (2019). Rural roads and local economic development. Dartmouth College.

Behrman, J. R. and Deolalikar, A. B. (1989). Is Variety the Spice of Life? Implications for Calorie Intake. *The Review of Economics and Statistics*, 71(4):666–672.

Costinot, A. and Donaldson, D. (2016). How large are the gains from economic integration? theory and evidence from u.s. agriculture, 1880-1997. *Econometrica*, Forthcoming.

Donaldson, D. (2018). Railroads of the raj: Estimating the impact of transportation infrastructure. *American Economic Review*, 108(4-5):899–934.

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

Eaton, J. and Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, 70(5):1741–1779.

Ethiopian Road Authority, E. (2016). Road sector development program: 19 years assessment. page .

Funk, C., Peterson, P., Landsfeld, M., Pedreros, D., Verdin, J., Shukla, S., Husak, G., Rowland, J., Harrison, L., Hoell, A., and Michaelsen, J. (2015). The climate hazards infrared precipitation with stations—a new environmental record for monitoring extremes. *Scientific Data*.

Gebresilasse, M. (2018). Rural roads, agricultural extension, and productivity. *Job Market Paper*.

Hummels, D. (2007). Transportation costs and international trade in the second era of globalization. *Journal of Economic Perspectives*, 21:131–154.

Shamdasani, Y. (2018). Rural road infrastructure & agricultural production: Evidence from india. page .

Shrestha, S. A. (2018). Roads, participation in markets, and benefits to agricultural households: Evidence from the topography-based highway network in nepal. page .

Sotelo, S. (2018). Domestic trade frictions and agriculture. *Journal of Political Economy*, Forthcoming.

Chapter 2

Market Integration and Separability of Production and Consumption Decisions in Farm Households

2.1 Introduction

Agricultural sector in developing countries is dominated by smallholder farmers, who produce not only for the market but also for their own consumption needs; and use both purchased inputs such as fertilizers and non-purchased inputs such as family labor. Whether households' production decision is separable from their consumption preferences, and whether these two decisions can be analyzed separately/recursively has been a subject of debate among policy practitioners and academicians (Singh et al. 1986). There are a number of reasons why separability is an important subject.

First, separability determines how households respond to policy interventions, such as those that target increase in agricultural exports. For instance, governments' attempts to boost the production and export of cash crops is unlikely to materialize if households had to self-produce most of the crops they need for consumption – because that would make reallocation of land to cash crops more difficult. Second, separability has implications for efficient utilization of resources such as land and labor. For example, if separability holds, the household's crop production choices would be dictated by market prices and productivity of their land in the potential crops and the household would obtain higher return from its resources, compared to the situation where the household's consumption preferences dictate production decisions. Third, whether household production decisions can be analyzed independent of their consumption decisions is important in academic and policy research. Agricultural Household Models (AHMs) in which production and consumption decisions are made jointly lack tractability, which limits their widespread use in empirical research.

The vast majority of empirical studies invoke the assumption of complete markets to obtain a tractable AHMs (Singh et al. 1986, Taylor and Adelman 2003). Under this assumption, household decisions can be modeled as *recursive*. In the first stage, the household makes its production decisions independent of its consumption preferences, and in the second stage, the household makes its utility maximization/consumption decisions given its farm profits from the first stage. For instance, if a farmer can always buy a rice at a fixed market price and can hire labor at a fixed market wage, the quantity of rice produced by the farmer would be independent of how much rice it wants to consume and the amount of labor used to produce rice would be independent of the amount of family labor supply. There are at least two issues with the assumption of complete market. First, while the assumption significantly simplifies modeling of the farm household behaviour, it is unlikely to reflect the reality under which the farmers operate. For instance, land (rental) markets are very limited in many countries and functional labor markets too are far from reality.[1] Second, even in the case of more functioning markets, such as crop markets, high trade costs due to poor infrastructure makes these markets too thin for the farmers to rely on. As a consequence, farmers may choose to self-produce most of the crops they need for consumption, instead of specializing in few crops and sourcing their consumption

[1]Scattered settlements and lack of transport infrastructure, typical in rural areas of developing countries, would make rural labor markets too thin to be reliable source of labor input.

from markets. The bottom line is that the mere existence of rudimentary markets may not guarantee recursiveness in farm production and consumption decisions. Hence, separability has been a subject of empirical test. Existing empirical tests look at the link between on-farm labor demand and household demographic characteristics, where a positive correlation is considered as evidence for rejection of separability.

In this chapter, I develop a new approach to test recursiveness and investigate how it varies across households with varying access to markets. My approach relies on a simple model of household decisions on crop production and consumption in an environment where a household can engage in costly trade. On the consumption side a household maximizes utility by choosing how much of different crops to consume given its tastes for different crops, its income, and local prices. On the production side, the household decides how to allocate its limited land across potential crops given productivity of its land in the crops and local crop prices. If the household does not face significant trade barriers, its production decision is separable from its consumption preferences. Hence, the household's land allocation across crops should not be correlated with its tastes for these crops. Otherwise, the household land allocation across crops will be dictated by the household's tastes – and the extent to which tastes dictate crop production choices depends on the level of trade costs the household faces. A decrease in trade costs due to road construction would thus weaken the link between consumption preference and production choices by improving households' opportunities to trade.

I empirically implement this approach using a very rich panel data from Ethiopia on household production and consumption disaggregated by crops. I use a large rural road construction project as a source of variation to the household's market access/trade costs. I first estimate household crop tastes using a modified version of Almost Ideal Demand Systems (AIDS) (Deaton and Muellbauer, 1980). I then test the separability hypothesis by regressing land allocation across crops on the estimated crop tastes, and explore how this correlation varies across households of varying level of market access. I find that household crop tastes significantly affect the fraction of land allocated to the crop, which implies rejection of the separability hypothesis. Moreover, the effect of tastes on household land allocation is stronger for household that reside further from market centers and roads, and improvements in market access due to massive rural road construction project leads to significant decreases in the correlation between household land allocation and tastes.

The empirical approach suggested here has several advantages over the previous studies testing separability. First, previous studies test separability by using the correlation between household on-farm labor demand and the household demographic characteristics. However, in the context of farm households, who are predominantly self-employing, on-farm labor demand is likely to be poorly measured. On the contrary, in most agricultural surveys, including the survey used in this chapter, land area is measured using GPS tools by well trained enumerators. Second, the current approach makes the link between separability and market integration straightforward. Given information on physical location of households and their nearest market centers, one can obtain a measure of households' proximity to crop markets and explore how separability varies across households with varying level of proximity to markets. Such exercise is difficult to come by for labor markets because there is no physical location for labor markets.

This chapter is closely related to studies that empirically test separability. The seminal paper by Benjamin (1992) tests separability using the relationship between household on-farm labor demand and the household's demographic characteristics. The basic idea is as follows. If markets are complete and farm household's production decisions are independent of the household's preferences, household's on-farm labor demand should be independent of the household's demographic composition, such as the number of active age persons in the household. Using data from rural Indonesia, Benjamin (1992) runs a regression of on-farm labor demand on different demographic characteristics and fails to reject the assumption of separability. However, his crossection data did not allow him to address a number of confounding factors. Using (better) panel data from the same country, LaFave and Thomas (2016) conduct a followup study to Benjamin (1992) in which they run similar regressions to the latter but use panel data specification. They strongly reject separability, contrasting Benjamin (1992).

More recently, LaFave et al. (2020) suggest a new consumption based test for separability. The central idea of the test is that if household production and consumption decisions are recursive, input prices affect household demand for goods only through their effect on profits. This implies that the ratio of the effects of two inputs on demand for a good is equal across all goods. They implement this test using demand estimations and Wald tests of non-linear coefficient restriction. The downside of this approach is that the test lacks power to discern separability, particularly when the number of farm inputs are many. Moreover, in general equilibrium, input prices should not affect demand for goods once the goods' prices are controlled for (because the output prices are themselves determined by the input prices) unless the households are the owners of the inputs. This perhaps explains why most of the input coefficients in the demand estimation are either statistically insignificant or enter with signs that are inconsistent with their theory in LaFave et al. (2020).

This chapter is also related to the literature on the development impact of rural roads. Asher and Novosad (2019) exploit strict implementation rule of India's massive rural road expansion project called Pradhan Mantri Gram Sadak Yojana (Prime Minister's Village Road Program, or PMGSY) to identify the program's causal effect using fuzzy regression discontinuity design. They find that the roads' main effect is to facilitate the movement of people out of agriculture, with little or no effect on agricultural income and consumption. Shamdasani (2018) studies the effect of a large road-building program in India and finds that remote farmers who got access to road diversified their crop portfolio by starting to produce non-cereal hybrids, adopted complementary inputs and improved technologies, and hired more labor. Gebresilasse (2018) studies how rural roads complement with an agricultural extension program that trains farmers on how to use best agricultural practices and technology adoption in Ethiopia. Shrestha (2018) finds that a 1% decrease in distance to roads due to expansion of highways resulted in a 0.1–0.25% increase in the value of agricultural land in Nepal. I contribute to this literature by providing evidence on another potential channel through which rural roads affect resource allocation and welfare, which is increased separability of household production and consumption decisions.

The rest of the chapter is organized as follows. In section 2, I present the data and a series of descriptive evidences motivating the theoretical and empirical methods. Section 3 presents the theoretical model and section 4 discusses the empirical implementation of the theoretical model. Sections 5 presents the results. Section 6 concludes the chapter.

2.2 Data

2.2.1 Sources

Agricultural production and consumption data: I use two data sources for agricultural production. The first is Ethiopian Socioeconomic Survey (ESS), which is an exceptionally detailed panel data of about 4,000 nationally representative farm households for the years 2011, 2013 and 2015. The data includes farm household's production, consumption and market participation, disaggregated by crops. The main advantage of the ESS dataset is its richness as it includes both production and consumption information, land and labor utilization, and a number of household demographic and geographic information. That is, I observe a household's production of each crop as well as consumption of each crop disaggregated by source (whether is comes from own production or purchase).[2] I use this data to estimate household crop tastes and to test separability between production decisions and consumption preferences. A big caveat of this data set is that it covers households in only about 330 villages.

I also use the Agricultural Sample Survey (AgSS), which is the largest annual

[2]The consumption information is based on a seven-day recall of basic consumption items, which are predominantly crops. However, household's crop utilization information also gives how much of each crops produced is consumed within the household.

agricultural survey in the country covering over 40,000 farm households in about 2000 villages. While this dataset goes back as far as 1995, villages were resampled every year until 2010. Staring from 2010, Central Statistical Agency (CSA) kept the sample of villages fixed but took a random sample of about 20 farmers per village every year. This dataset includes detail production information: areas of land covered by each crop, application of fertilizer and other inputs, and quantities of harvest. Moreover, every three-year starting from the year 2009/2010, CSA also gathered crop utilization information, i.e., the fraction of crop production used for own consumption, the fraction sold, the fraction used to pay wages, the fraction used for seeds, etc, for all crops. However, this dataset does not include information about household consumption. I use this data mainly to estimate productivity of villages in some crops which are not covered in GAEZ data (see below).

Price data: The price data comes from three different sources. I construct village level prices of crops by combining two sources. The first is the Agricultural Producer Price Survey (AgPPS), which is a monthly survey of farm-gate prices at a detailed geography (villages) for almost all crops and many other agricultural produces.[3] In villages that are not covered by AgPPS, I use ESS's price survey. Unfortunately ESS's price survey is not exhaustive in its coverage of crops. I overcome this problem by using the sample of households who report a positive purchases/sales of crops to construct village level unit values of crops in the cases where AgPPS prices are missing.

I also use the Retail Price Survey (RPS), which is a monthly survey of prices of almost all crops and non-agricultural commodities in major urban centers throughout the country. RPS dataset covers over 100 urban centers across all administrative zones of the country. Both AgPPS and RPS are collected by CSA and go back to at least 1996. Importantly, the agricultural products covered in both datasets overlap almost fully. I use RPS, together with village prices constructed using the above

[3]CSA claims that the prices in this survey can be considered as *farm-gate* price because they are collected at the lowest market channel where the sellers are the producers themselves, i.e., no intermediaries involved.

procedure, to explore how rural road expansion affected urban-rural price gaps.

Rainfall and agro-climatic data: I use FAO/GAEZ agro-climatically attainable yield for low/intermediate input use to construct villages' crop suitability, which is used in the separability test and to test how road affects the relationship between local comparative advantage and local prices. Unfortunately the GAEZ data doesn't include some of the most widely grown crops in Ethiopia such as *Teff*. For such crops, I use the AgSS data to construct village level suitability of land to the crops from the average yield in the villages over the period 2010-2016. The high correlation between yield estimates provided by GAEZ and AgSS for the sample of crops that exist in both data ensures that this approach gives a remarkably credible estimate of land-suitability.

The rainfall data comes from Climate Hazards Group InfraRed Precipitation with Station data (CHIRPS), which provides rainfall dataset starting from 1981. CHIRPS incorporates 0.05 resolution satellite imagery with station data to create gridded rainfall time series for trend analysis and seasonal drought monitoring. It is widely used to monitor drought in East Africa (Funk et al., 2015).

Road data: I use administrative data on the entire road-network in the country. This data includes the attributes of the roads (such as surface type), the role of the road (trunk road, link road, etc), and ownership (federal government, regional government, etc). In this chapter I use the massive rural road expansion under URRAP as a source of variation to villages' access to road/market. Over the period 2011-2015, the Ethiopian government gave exclusive focus to the URRAP and constructed over 62,413kms of new all-weather roads connecting village centers to the nearest road or district capital, which ever is shorter. Figure 2.1 shows map of the road network before and after URRAP.

The main objective of this project was to improve villages' access to product and input markets. The program increased the overall road density per 1000 square-km from 44.4 in 2010 to 100.4 in 2015 (Ethiopian Road Authority, 2016). Though the

URRAP was launched in 2011, very few roads were commenced in the years 2011 and 2012, which are officially considered as capacity building years. Almost all the rural roads constructed under the first-round of this program were completed between 2013-2015. This enables me to use the 2011 and 2013 round of ESS to test the parallel trend assumption when evaluating the impact of URRAP.

2.2.2 Identification issues

One objective of this chapter is to explore the link between market access and separability using the massive rural road expansion project, URRAP. There are three challenges to identify the causal effects of URRAP on the separability – selection bias, heterogeneity in treatment intensity, and spillover effects of road connectivity. Selection bias is a concern because villages are selected for the URRAP based on some demographic, geographic, social, and economic factors.[4] Villages that get connected to a dense network may gain more from the road than those that get connected to sparse network, implying heterogeneity in treatment intensity. Spillover effects is a concern because when a village is connected to the preexisting road network or to the nearest urban center, all its neighbors which did not get direct connection would also have improved access to market via the connected village. This would lead to underestimation of the causal effect of URRAP on separability.

I address the potential selection bias by using a Matching-Based Difference-in-Differences (MB-DID) strategy where I first obtain a matched sample of treated and non-treated villages based on their observable characteristics that might be relevant for selection of villages for URRAP and then conduct DID estimation based on these matched sample of treated and non-treated villages. Combining matching with DID strategy is a powerful approach to address the selection problem. The matching step enables me to compare treated villages with non-treated villages that have similar

[4]Unfortunately there was no official guideline as to which villages should be selected for the URRAP in a given year. Even though the project was fully funded by the federal government, implementation of URRAP was completely decentralized to regional governments. Within each regional government, districts propose list of villages that should get a road during a particular year and the regional governments approve villages based the available regional budget.

observed characteristics and hence similar treatment probability. The DID strategy on these matched samples helps me to washout unobserved time-invariant village characteristics that may confound the treatment effect. I identify the following village characteristics for matching treated and non-treated villages in consultation with officials at Ethiopian Roads Authority (ERA): distance to nearest town, distance to preexisting road network, population size, average slope of land in the village, average elevation in the village, and average rainfall over 1990-2010 period. I use Digital Elevation Model (DEM) data and ArcGIS tools to calculate average slope and elevation of each village.

To address the heterogeneity in treatment intensity and spillover effects, I use market access approach (Donaldson and Hornbeck, 2016) which excellently captures treatment benefits from both direct and indirect connectivity, and properly accounts for the density of the network to which a village is connected. The market access measure is derived from general equilibrium trade model and calculated using the entire road network and the distribution of population across Ethiopian villages. See Appendix A for details in the construction of market access measure. The constructed market access measure increases both for villages that are directly connected and those that are not by 47%, on average, but it increases more for the directly connected villages by about 40%.

2.2.3 Descriptive statistics

In this section I present some descriptive statistics about farm households in rural Ethiopia to guide the theoretical framework and empirical analysis.

Large barriers to trade: These barriers to trade are both physical and pecuniary. Table 2.1 shows the modes of transport used by farmers to get to market to sell their produce. The most frequently used mode of transport are *on foot* and *pack animals*, together accounting for more than 85% of transaction cases. Vehicle transport accounts for just 2.34% in 2011, and increases to 5.69% in 2015. Though vehicle transport is the least frequently used, it accounts for about one-third of the volume of transaction by value and quantity. The ad valorem trade cost (transport cost per value of transaction) on vehicle is very high (the median is 6.49 % and the mean is about 11%). The size of this cost is comparable to international trade costs estimated by Hummels (2007) for US and New Zealand import, although here the distance traveled is just few kilometers. Perhaps the low share of vehicle transport is attributed to farmers choosing not to use this option due to its higher pecuniary cost. The last row of table 2.1 shows inflation adjusted median transport fare from a village to district capital decreases from 0.7 Birr/km to 0.523Birr/km between 2011 and 2015.[5]

Households are less likely to consume a crop that they do not produce: Table 2.2 shows the fraction of households who have reported a positive amount of consumption of a crop and the fraction who consumed a positive amount of a crop but did not produce the crop (consumed from purchase).[6] The first two columns report the statistics for a sub-sample of households from small towns (a population of less than 10,000) while the last two columns are for rural households. There is a clear distinction between small town and rural households: (1) households in small towns are more likely to consume vegetables and relatively more expensive cereals such as *Teff* compared to their counterparts in the rural areas (on the contrary, rural households are more likely to consume cheaper cereals such maize, sorghum and millet compared to their urban counterparts), and (2) households in small towns are more likely to consume a crop that they did not produce compared to rural households. For example, about 59% of rural households report consumption of maize while only 23% consumed from purchase (in other words only 40% (23/59) of the households who consumed maize purchased the maize, the rest consumed from own production). On the contrary, in small towns, most of those who consumed a crop did not produce the crop.

[5]Ethiopia's currency is called Birr. One USD is sold for about 17 Birr in 2011.

[6]ESS asks households how much of each crops they consumed over the seven days before the interview day, disaggregated into from purchase, and from own production.

While the difference between households in small towns and those in rural areas could in part be driven by income gaps and by the fact that households in small town are more likely to engage in non-farm activities (though over 75% of the sample households in small town and 94% of household in rural villages did not have any non-farm income), a significant part might be attributed to better access to markets in small towns. In small towns there are more frequent markets, there are shops, the markets are larger since most of the surrounding villages transact with them and more importantly, the towns are connected to the rest of the country via all-weather roads.

Most of crop production is consumed within the household: Table 2.3 reports crop utilization within a household. On average, about 71% of all crop production is consumed within the household and only 13% is marketed. However, there is significant variation across crops.

Positive correlation between land and expenditure shares of crops: I use ESS data and focus on 20 crops for which complete information is available on both production and consumption. I calculate the share of each crop in household consumption expenditure[7] and the share of land allocated to the production of each crop. I run the following regression:

$$\eta^k_{hvt} = \beta_0 + \beta_1 s^k_{hvt} + \beta_2 p^k_{vt} + \beta_3 y^k_v + \gamma^k_t + \gamma_h + \varepsilon^k_{hvt} \tag{2.1}$$

where η and s are the land and expenditure share of crop, p is price, y is the GAEZ yield/productivity estimate which measures agro-climatic suitability of a village in each crop , h is household, v is a village, k is crop, and t is year. It is crucial to control for prices and yield in this regression because both production and consumption decisions are functions of these variables, directly or indirectly. Any significant positive correlation between the land and expenditure share of crop with

[7] Household consumption from own production is valued at village level prices.

in a household is suggestive evidence against separability. Under autarky near perfect correlation between household land and expenditure shares of crops is expected. I run the regressions for each of the survey rounds separately to show how the estimated correlation changed over time.

Table 2.4 reports the results. Panel A reports the estimated correlations between the land and expenditure share for the three rounds of survey. The estimated correlation is 0.47 for the year 2011, which slightly increases to 0.53 for the year 2013 before it decreases significantly to 0.20 for the year 2015. We fail to reject the hypothesis that the correlations for the years 2011 and 2013 are equal, but similar hypotheses between 2015 and 2011 or 2013 are rejected at 1% significance level. In panel B, I run analogous regressions where I use data on *plot level* labor use (both planting and harvesting hours of labor) which I convert to *crop level* labor use given the information about which crops covered a plot during a given year. Given this, I calculate the labor share of crop in exactly analogous way to the land share of crop. I then redo all the above regressions using the labor share of a crop as the dependent variable. The results looks very similar to the one we obtained using the land allocation. The correlation between the labor and expenditure share of crops slightly increase from 0.45 to 0.51 between 2011 and 2013 before it significantly decreases to 0.19 for the year 2015. To sum up, these correlation strongly suggest that household resource allocation is at least partially dictated by their consumption preferences.

New roads decrease the correlation land and expenditure share of crops: Before I explore the effects of URRAP roads on the correlation between the land and expenditure shares of crops, I provide evidence on whether the URRAP roads indeed decreased trade costs and improved market integration. In appendix B, I show that URRAP roads significantly decreased the urban-rural price gaps for crops, particularly for perishable crops such as vegetables. I also show that, the construction of roads significantly weakened the inverse relationship between local prices and local productivity of crops. These evidences imply that the URRAP roads have indeed improved market integration of rural areas.

Next, I estimate the effects of road construction under URRAP on the correlation between land/labor and expenditure share of crops. Even though URRAP was launched in 2011, very few roads were completed before 2013. Almost all the roads were completed between 2013 and 2015. Thus I use these two years as pre- and post-program periods. Table 2.4 shows that the correlation between land and expenditure shares decreases significantly between 2013 and 2015. I estimate how much of this decline is attributed to the URRAP roads in the DID framework:

$$\begin{aligned} \eta^k_{hvt} &= \beta_0 + \beta_1 s^k_{hvt} + \beta_2(Post_t * URRAP_v) + \beta_3(s^k_{hvt} * Post_t * URRAP_v) \\ &\quad + \delta Z_{vt} + \gamma^k_t + \gamma_v + \varepsilon^k_{hvt} \end{aligned} \tag{2.2}$$

where $Post_t * URRAP_v$ is a dummy variable indicating whether village v got new road connectivity under URRAP, which equals zero in 2013 and equals one in 2015 for villages that get new roads, and Z_{vt} includes the vector of control variables in equation 2.1. The main parameter of interest is β_3, which captures the causal effect of road connectivity under the assumption that assignment of roads is not endogenous to the the correlation between land and budget shares of crops.

Table 2.5 presents the results. Note that the first two columns do not include village fixed effects but instead include population density and village distance to the baseline network to account for village selection for the road program based on these observed characteristics. The last two columns include village fixed effects. Columns 1 and 3 are use the land share of a crop as dependent variable while column 2 and 4 use labor share of a crop. The results clearly show that road construction under URRAP caused a significant decline in the correlation between land/labor and expenditure shares of crops. Households in villages that got road connection between 2013-2015 have seen a decrease in the correlation between land and expenditure shares by about 0.16, compared to households in villages that were not directly exposed to the program. This is a large effect, roughly about 25% of the baseline correlation in 2011.

Overall, the above results suggest that household production and consumption decisions are likely made jointly. Moreover, the link between production and consumption decisions seems to be significantly influenced by the level of underlying market integration. Below, I build on these evidences to suggest a formal framework to test whether household production decision is dictated by the household's consumption preferences and how improvements in market access affects the link between the two.

2.3 Theoretical framework

Informed by the above facts, in this section I develop a theoretical framework to test the separability hypothesis. In doing so, I borrow tools from the Ricardian trade models. Particularly, I build on Eaton and Kortum (2002), Donaldson (2018), and Sotelo (2018).

Consider an economy constituting villages $v = 1, ..., V$. Each village is populated by I households indexed by $i = 1, ..., I_v$. The household derives utility from consumption of K homogeneous crops indexed by $k = 1, ..., K$ that can be potentially produced or purchased.

Preferences: A farm household spends all its income on crops and its preferences over different crops is given by

$$U_{ivt} = f\left(\mu_i^k; q_{it}^k\right)$$

where $f(.)$ is a common utility function across households in the country, q_i^k is the quantity of crop k consumed by household i, and μ_i^k is the household taste for crop k, which is assumed to be fixed over the short to medium period. The household crop tastes act as pure demand shifters. The household maximizes this utility subject to the following budget constraint:

$$\sum_k p_{vt}^k q_{it}^k \leq \Pi_{it} \tag{2.3}$$

where p^k_{vt} is village level crop price, and Π_{it} is household farm income/profit.

Production: I follow Sotelo (2018) to describe the farmer's production problem. Each farmer owns L_i amount of land, which is divided into a continuum of plots of size one indexed by $\omega \in \Omega_i$, where Ω_i is the set of plots owned by farmer i such that $\int_{\Omega_i} \omega d\omega = L_i$. Each of the plot is different in how well it is suited to growing different crops, which I denote as $z^k_i(\omega)$. Assuming that a given plot can only be used to grow one crop at a time (plots cannot be divided), the production function is given as:

$$y^k_i(\omega) = g\Big(z^k_i(\omega), \mathbf{x}_i(\omega); \boldsymbol{\alpha}^k\Big)$$

where $y^k_i(\omega)$ is the quantity of crop, $\mathbf{x}_i(\omega)$ is the amounts of vector of variable inputs (such as labor and fertilizer) used on the plot, and $\boldsymbol{\alpha}^k$ denotes parameters.

The farmer draws $z^k_i(\omega)$ independently for each plot-crop from a Fréchet distribution with the following cumulative distribution function:

$$F^k_i(z) = Pr(Z^k_i < z) = \exp(-(A^k_i)^\theta z^{-\theta})$$

where A^k_i is the location parameter for the distribution of crop-suitability of land across the set of plots owned by a farmer, Ω_i. It can be interpreted as the average productivity of farmer i's land in crop k, as determined by agro-climatic conditions of the village, and soil, slope, and other characteristics of the farmer's plots. In villages that have agro-climatic conditions impossible to produce crop k, A^k_i is set to zero for all farmers in the village. θ is the degree of homogeneity in the set of plots owned by a farmer, and it is constant across crops.

Farmers are geographically separated and there is an iceberg trade cost of $\tau^k_{vv'} \geq 1$ between farmers in villages v and v' in crop k.[8] Motivated by the result in appendix B, which shows that spatial price variation differs across crops, trade costs are

[8]For simplicity, I assume that within village trade costs between farmers are negligible. The median village has area of about $25km^2$. While distance is not a big impediment to trade within village, the fact that farmers within a village share similar agro-climatic condition implies that there is less room for crop trade within a village.

assumed to vary across crops to reflect that some crops, such as vegetables, are more costly to trade (e.g., perishable) than others such as cereals. I assume that $\tau^k_{vv} = 1, \forall k$, and impose standard assumption of triangle inequality in trade costs, $\tau^k_{vv'} \times \tau^k_{v'v''} \geq \tau^k_{vv''}, \forall k$.

2.3.1 Two extreme cases

To motivate the separability test, it suffices to considering the farmers' problem under two extreme cases so that we can characterize which of the two cases closely matches the farmer's observed choices. The first is the case where farmers are allowed to trade with each other paying reasonable trade costs, and the second is the case where trade costs are too high for the farmers to engage in trade. I discuss how we can generalize from these two extreme cases and provide a general proof of the link between separability and trade costs in appendix C.

Case-I: Separability $\tau^k_{vv'} << \infty, \forall k$. Suppose trade costs are such that farmers can buy and sell any crop at a prevailing market price. Assuming perfect competition, no arbitrage condition implies that, for any two villages v and v', equilibrium crop prices satisfy $p^k_{vv'} = \tau^k_{vv'} p^k_{vv}$ where $p^k_{vv'}$ is price in village v' of crop k originating from village v, p^k_{vv} is price in village v of crop k originating from the same village v.

Under this case, the farmer takes local crop prices p^k_v and a vector of local input prices $\mathbf{r}_v$ as given, and allocates land across crops. The fraction of household land allocated to crop k is given by:

$$\eta^k_i = h\Big(p^k_v, \mathbf{r}_v, A^k_i; \theta, \boldsymbol{\alpha}^k\Big) \tag{2.4}$$

This implies that the quantity of crops produced and profit from each crop are given, respectively, by:

$$y^k_i = \mathcal{Y}\Big(p^k_v, \mathbf{r}_v, A^k_i, L_i; \theta, \boldsymbol{\alpha}^k\Big), \quad \text{and} \tag{2.5}$$

$$R^k_i = \mathcal{R}\Big(p^k_v, \mathbf{r}_v, A^k_i, L_i; \theta, \boldsymbol{\alpha}^k\Big) \tag{2.6}$$

Given farm profit $\Pi_i = \sum_k R_i^k - \mathbf{r}_v.\mathbf{x}_v$, the farmer then maximizes utility subject to the budget constraint. The optimal quantities of each crop are given by:

$$q_i^{*k} = \mathcal{C}\Big(\mu_i^k, p_v^k, \Pi_i\Big) \tag{2.7}$$

Equations 2.4-2.7 show that: (i) household land allocation across crops and quantities of crops produced are independent of crop tastes μ_i^k, (ii) tastes affect household demand for crops but not production decisions, and (iii) household production decisions affect household demand only through its effect on farm profits. These imply that household decision is recursive: the household first makes production decision to maximize its farm profits given local crop prices, inputs prices and productivity, and in the second stage the household chooses optimal quantities of crops to consume given local crop prices, tastes, and farm profit.

Case-II: Autarky $\tau_{vv'}^k \longrightarrow \infty, \forall k$. Under this case, there is no crop market and hence, no market prices which the farmer takes as given. Instead, the farmer's decision is based on shadow price which is a function of the household tastes and other household characteristics: $\tilde{p}_i^k = \mathcal{P}\Big(\mu_i^k, \mathbf{r}_v, A_i^k, L_i, U_i; \theta, \boldsymbol{\alpha}^k\Big)$.

The household makes production decision given the shadow prices and productivity distribution parameters. Plugging this in the land allocation we have the following land allocation rule under autarky.

$$\tilde{\eta}_i^k = h\Big(\tilde{p}_i^k, \tilde{\mathbf{r}}_v, A_v^k; \theta, \boldsymbol{\alpha}^k\Big) \tag{2.8}$$

In equation 2.8, the fraction of land allocated to crop k depends on the household taste for the crop via the shadow price. That is, household production decision is not independent of its consumption preferences. This is a key result used to design the separability test in this chapter.

2.3.2 Trade costs and separability

Here I describe the intuition for generalizing the link between trade costs and separability, postponing formal proof to appendix C. To make the generalization clear, consider the case where, due to lack of fast transport, some goods are non-tradable. Perishable vegetables are good examples in rural areas of developing countries. Because the households have to rely on self-production for these high trade cost crops, the separability assumption no longer holds. The fraction of land allocated to such crops would be dictated by the households' tastes for these crops. In general, the probability that a farmer is the cheapest supplier of any given crop to itself, compared to all other farmers in the country, increases with trade costs. On the other hand, the probability that a farmer is the cheapest supplier of a given crop to any other farmer decreases with trade costs. These two probabilities determine household land allocation rule as a function of trade costs, tastes, prices and productivity for any level of trade costs. Given this land allocation rule, one can obtain different comparative statics, including comparative statics of trade costs on the correlation between tastes and land allocation (see appendix C).

2.4 Empirical methodology

2.4.1 Estimating household crop tastes

I follow Atkin (2013) to estimate household tastes for crops. Suppose household preference for crops is represented by the following expenditure function corresponding to Almost Ideal Demand System (AIDS) (Deaton and Muellbauer, 1980) where the coefficients of the first-order price terms are allowed to vary across households to allow for taste variations:

$$\ln e(u, \mathbf{p}_{vt}; \Theta) = \mu_0 + \sum_k \mu_i^k \ln p_{vt}^k + \frac{1}{2} \sum_k \sum_{k'} \gamma^{*kk'} \ln p_{vt}^k \ln p_{vt}^{k'} + u\beta_o \prod_k {p_{vt}^k}^{\beta_k} \tag{2.9}$$

where k indexes crops, v indexes villages, and t represents years. Applying Shefard's Lemma and replacing u by indirect utility function gives the following expression for the expenditure share of crops:

$$s_{it}^{k} = \mu_{i}^{k} + \sum_{k'} \gamma^{kk'} \mathrm{ln} p_{vt}^{k'} + \beta_{k} \mathrm{ln} \frac{m_{it}}{P_{vt}} \tag{2.10}$$

where i indexes household, $\gamma^{kk'} = \frac{1}{2}(\gamma^{kk'*} + \gamma^{k'k*})$, m_i is household nominal expenditure on food, and $\frac{m_i}{P_v}$ is real expenditure. Following Deaton and Muellbauer (1980) and Atkin (2013), I use Stone index for for village price index P_v, $\mathrm{ln}P_v = \sum_k \bar{s}_v^k \mathrm{ln} p_v^k$ where $\bar{s}_v^k$ is the average expenditure share of crop k in village v.

Household crop tastes μ_i^k are thus demand shifters, conditional on prices and real income of the household. The key assumption here is that tastes for crops do not change over short period of time. Atkin (2013) shows that regional tastes are indeed stable over time due to habit formation.

Atkin (2013) discusses two necessary conditions for identification of tastes in a similar equation to 2.10, which are satisfied in my setup. The first is the existence of temporary and supply driven price variation within village. In my setup this condition is satisfied by price variation due to rainfall fluctuations. See table A.4 in appendix E for evidence of price volatility in response to rainfall fluctuations. The second condition, which is assumed to hold, is the existence of a common preference structure across rural households in Ethiopia, conditional on taste differences, and that AIDS function approximates this preferences reasonably well.

I estimate the following equation to identify household crop tastes:

$$s_{it}^{k} = \mu_{i}^{k} + \sum_{k'} \gamma^{kk'} \mathrm{ln} p_{vt}^{k'} + \beta^{k} \mathrm{ln} \frac{m_{it}}{P_{vt}} + N_{it} + \delta_t + \varepsilon_{it}^{k} \tag{2.11}$$

where N denotes household size and other demographic characteristics, δ_t is year fixed effects and ε_{it}^k is the error term. One concern in estimating equation 2.11 using OLS is that unobserved factors correlated with both village prices and household idiosyncratic tastes could bias the estimated price coefficients and the taste parameters. Similar to

Atkin (2013), I address this concern by instrumenting prices by prices in the nearest villages (Hausman, 1994).

As a robustness check, I also estimate a specification where households within a village share the same crop tastes, i.e., $\mu_i^k = \mu_v^k, \forall i \in v$. The motivation for this is that village sizes are small (a median village in my sample has an area of about $25km^2$) and village population share the same culture including ethno-linguistic culture, and perhaps also the same food culture. As is shown below, the empirical result also strongly supports this conjecture – about 81% of variation in crop tastes comes from across village variation. The taste parameters estimated following the two approaches have a correlation of about 0.90.

2.4.2 Testing separability

Once I obtain the household crop tastes, I test the separability hypothesis by looking at whether household land allocation across different crops is independent from the household crop tastes, conditional on village crop prices and village crop yields (agro-climatic suitability of village for each crop). I estimate the following regression:

$$\eta_{ivt}^k = \beta_0 + \beta_1 \mu_i^k + \beta_2 \ln p_{vt}^k + \beta_3 \ln Y_v^k + \beta_4^k \ln \text{Rainfall}_{vt} + \gamma_t^k + \gamma_v + \epsilon_{ivt}^k \tag{2.12}$$

where η_{ivt}^k is the share of land allocated to crop k. Recursiveness requires that $\beta_1 = 0$, that is, there is no significant correlation between household land allocation across crops and the household crop tastes. On the other hand, a positive and statistically significant β_1 is evidence against recursiveness. The higher β_1, the closer the village economy is to an autarky.

The role of market access: Next, I explore how infrastructure and market integration affect the link between household production and consumption choices. The theoretical model implies that decreases in trade costs lead to a decrease in the correlation between the land share of crops and crop tastes (see appendix C for a formal proof). I run the following regression:

$$\eta_{ivt}^k = \beta_0 + \beta_1\mu_i^k + \beta_2\text{MA}_{vt} + \beta_3(\mu_i^k \times \text{MA}_{vt}) + \beta_4\text{ln}p_{vt}^k + \beta_5\text{ln}Y_v^k \qquad (2.13)$$
$$+ \beta_6^k\text{lnRainfall}_{vt} + \gamma_t^k + \gamma_v + \epsilon_{ivt}^k$$

where MA is a measure of village market access derived from general equilibrium trade models (Donaldson and Hornbeck, 2016). The market access (MA) is calculated using data on (i) the entire road network in Ethiopia, (ii) the spatial distribution of population across the country and (iii) the freight costs of transporting one ton of cargo from origin village to destination village along the least cost path, before and after the construction of URRAP roads, and trade elasticity parameter. The massive rural road expansion between 2013 and 2015 led to significant decreases in freight costs, increasing MA for all villages, particularly those villages that got direct road connectivity under the program (see appendix A for the detailed procedure followed in constructing MA measures). In equation 2.13 a negative and statistically significant β_3 would imply that market integration plays important role in weakening the link between household production and consumption choices.

2.5 Results

2.5.1 Estimating tastes and the separability test

The taste estimates: It is worth mentioning some points about the estimated taste parameters. First, both OLS and IV estimation of equation 2.11 give very similar result. The correlation between the taste estimates obtained from these approaches is about 0.96. Second, the estimated tastes for crops show significant variation across households. However, most of the variation comes from across village variations – on average, 81% of the variation in tastes comes from across villages. Third, because of small within village variation in tastes for crops, estimating tastes at village level gives very similar result to household level taste estimates when both

OLS and IV estimation is used. Overall, the IV passes the under-identification and weak identification tests remarkably and borderline passes the weak instrument test with average first-stage F-statistics of about 10.

Testing separability: Next, I explore how the estimated taste parameters correlate with household land allocation. The theoretical results in section 2.3 suggest that, if household production decisions are independent of their consumption preferences, a household taste for a crop should not affect the fraction of land the household allocates to the crop. Table 2.6 reports the results for estimation of equation 2.12. I allow the coefficient of taste to vary across years in order to see whether the estimated coefficient changes over time. The first columns uses OLS taste estimates while the second column uses IV taste estimates. Across all rows and columns, we observe that tastes significantly affect household land allocation, implying rejection of the separability/recursiveness hypothesis. The coefficients of taste slightly increase from 0.68 to 0.70 between 2011 and 2013, before it declines to about 0.66 in 2015. However, these coefficients are not statistically different from each other.

2.5.2 Separability and proximity to market

Before turning to the effect of road expansion under URRAP, I first explore how the correlation between land allocation and tastes varies across households with varying proximity to population centers (towns with above 20,000 population) and to all-weather roads. Towns serve as hubs and market centers for the surrounding villages. Also, most villages access the rest of the country via the nearest towns. Hence, proximity to towns is important for market access. Similarly, proximity to all-weather roads improve the village's access to the rest of the country. I use distances to nearest population centers and nearest towns to measure proximity. The first is time invariant while the latter decreases for households residing in villages that obtained new roads under URRAP.

If lack of access to market is a driving factor for the observed correlation between land allocation and tastes, one would expect that the correlation would be stronger for household that live further from towns or roads. Table 2.7 reports the result for the effect of proximity to towns and roads. Panel A shows that the correlation between land allocation and the household taste significantly increases with distance to nearest population center. Using the result in the first column and range of log distance to population center of about 6, the correlation between land allocation and taste ranges from about 0.46 for the nearest to 0.81 for the furthest household to population center. Panel B reports similar result using distance to nearest road. The correlation between land allocation and tastes increases significantly with distance from road, even though distance to road has weaker effect compared to distance to distance to towns.[9]

2.5.3 The effect of URRAP on separability

Finally, I explore the effect of road expansion under URRAP on the correlation between land allocation and tastes. As mentioned in section 2.2, I use a matching-based DID estimation strategy to minimize selection bias. That is, I first obtain a matched sample of treated and non-treated villages based on a set of village characteristics before conducting DID estimation. Figure 2.3 shows the histogram of propensity score by treatment status and table 2.8 reports the balancing of the matching variables. I also report DID estimation results without matching for comparison, but my discussions will be based on the results from Matching-Based DID estimation.

Table 2.9 reports the matching-based DID estimation results for equation 2.13. The first two columns use binary treatment dummy while the last two columns use a continuous market access measure. To facilitate interpretation, market access measure is standardized. Across all columns, we see that road connections under URRAP led to significant decreases in the correlation between land allocation and tastes. The first two columns show that the correlation between the land share of

[9]This is partly because there is less variation in household distance from road compared to distance from population center. The standardized coefficients similar across the two variables.

crops and the household tastes for the crops decreases by 0.054-0.056 for villages that got direct road connection under URRAP compared to the control villages. Columns 3 and 4 show that one standard deviation increase in market access leads to 0.071-0.077 decrease in the correlation between the land share of crop and the household crop tastes.

Overall, the results in table 2.9 clearly show that improvement in access to market due to URRAP has led to decreases in the correlation between land allocation and tastes. That is, road connection under URRAP has led to more separability between household production decision and consumption preferences. Moreover, the estimated decrease in correlation between land allocation and tastes is significant considering the fact that the time span after the roads were completed is is too short for the village economy to adjust fully to the expansion of infrastructure. One would expect that the correlation would decrease more in the long term because the infrastructure expansion would lead to over time improvement in transport options and the thickness of local crop markets, which would alter household land allocation rule.

Table 2.10 reports the result for DID estimation without matching for comparison. The results in this table look similar to those in table 2.9, except that the estimated effects of URRAP are smaller and statistically insignificant in the first two columns.

2.6 Conclusions

In this chapter I suggest a new approach to test the separability/recursiveness hypothesis in agricultural household production and consumption decision and explore how it is related to market integration due to massive rural road expansion. My approach relies on a simple theoretical insight that if household production decision is independent of its consumption preferences, the household's tastes for different crops would not affect household land allocation across crops. The theoretical model also suggests that the extent to which tastes affect household land allocation across crops depends on the level trade costs the households are facing. In particular, if the household faces autarky for at least one crop, the household land allocation would be affected by the household tastes.

I implement this test using a very rich household panel data from Ethiopia. The dataset includes household production and consumption information disaggregated by crops and coincides with period of massive rural road expansion. I first estimate household level crop tastes suing a version of Almost Ideal Demand Systems (AIDS) where household taste for crop is inferred from shift in expenditure share of a crop conditional on prices of all the crops, household real total expenditure, and household demographic characteristics. Next, I conduct the separability test by regressing the land share of crop on the estimated household crop tastes and find that the separability hypothesis is strongly rejected. Finally, I explore how the correlation between household tastes and land allocation vary with the household's proximity to markets and road. I find that: (i) proximity to market centers play significant role in the link between land allocation and tastes, and (ii) rural road expansion led to significant decrease in the effect of tastes on household land allocation by improving the access to market.

Figure 2.1: Rural road expansion under URRAP

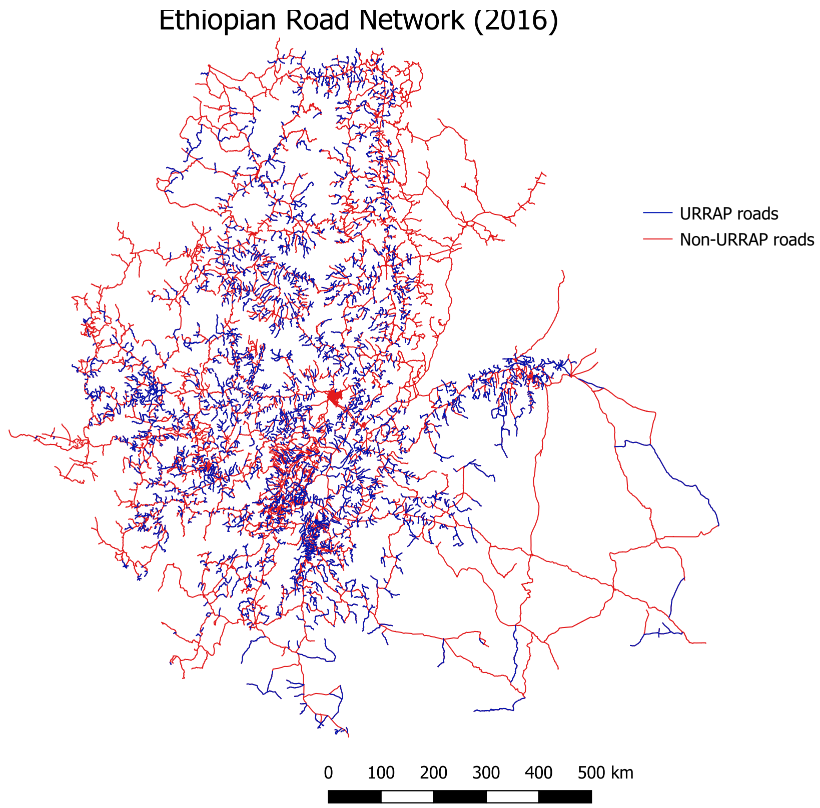

Figure 2.2: Completed URRAP roads (pictures are taken from Oromia Roads Authority).

Figure 2.3: Common support of propensity score matching

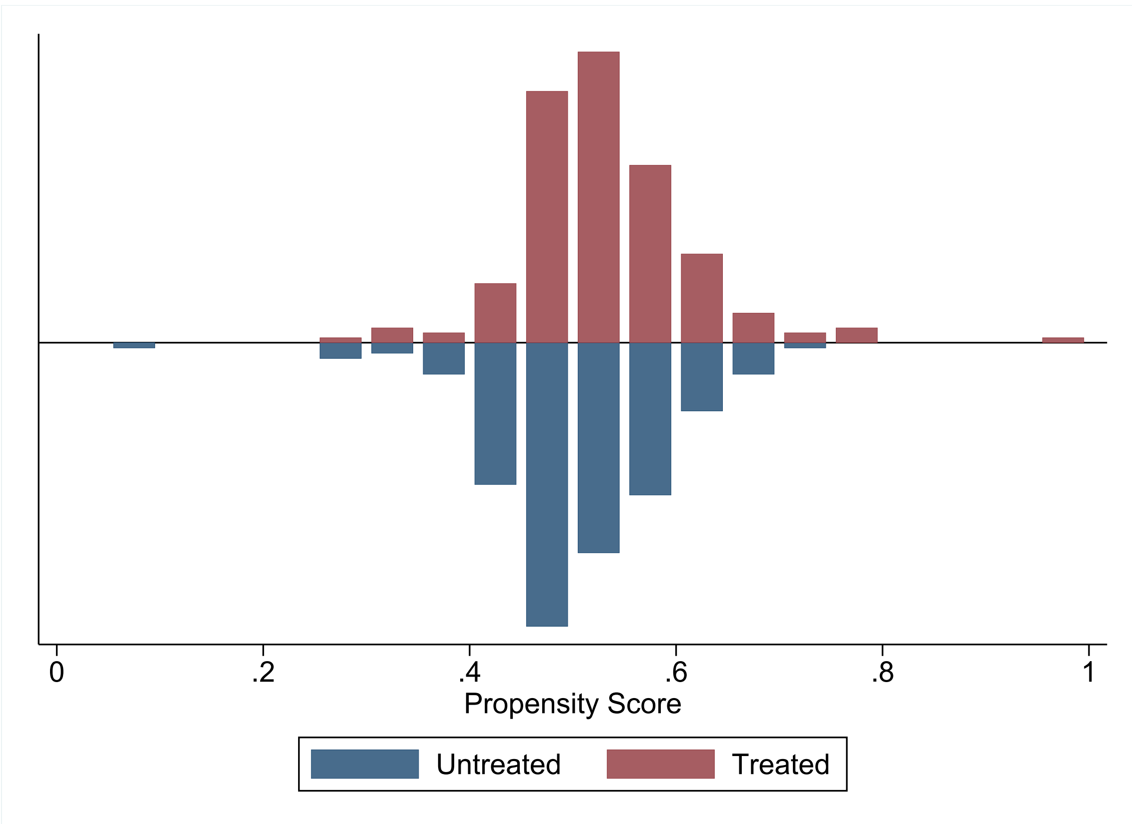

Table 2.1: Transport modes to market

Mode of transport	2011	2013	2015
On Foot	43.6	49.9	41.9
Pack Animals	45.8	41.80	43.9
Own Bicycle or Oxcart	6.74	1.76	4.78
Vehicle	2.34	2.53	5.69
Others	1.43	3.9	3.62
Ad valorem trade cost vehicle (mean)	11.37	6	6.4
Ad valorem trade cost vehicle (median)	6.49	4.3	3
distance to all-weather road (median KM)	10	10	8.5
distance to population centers (median KM)	30	30	30
distance to district (woreda) town (median KM)	17	17	17
distance to nearest weekly market (median KM)	12	8	8
Median transport fare to district capital (real Birr/KM)	0.7	0.597	0.523

Notes: This table is based on ESS data.

Table 2.2: Fraction of households who consume a positive amount of a crop, and those who consume and do not produce

	Small towns		Rural villages	
	Cons.	Cons.& Not prod.	Cons.	Cons. &Not prod.
Teff	0.719	0.640	0.349	0.114
Maize	0.438	0.382	0.593	0.232
Wheat	0.442	0.390	0.401	0.202
Enset	0.145	0.092	0.184	0.057
Barley	0.177	0.145	0.198	0.049
Sorghum	0.326	0.276	0.462	0.127
Millet	0.049	0.042	0.112	0.023
Field pea	0.432	0.399	0.232	0.151
Lentils	0.356	0.351	0.134	0.110
Linseed	0.044	0.042	0.074	0.043
Haricot beans	0.095	0.084	0.179	0.079
Horse beans	0.466	0.433	0.401	0.242
Onions	0.878	0.872	0.710	0.683
Potatoes	0.586	0.573	0.285	0.231
Tomatoes	0.660	0.656	0.350	0.333
Banana	0.273	0.259	0.161	0.100
Coffee	0.773	0.736	0.709	0.557
Total	0.560	0.536	0.455	0.366

Notes: This table shows fraction of households consuming a given crop and the source (own production or purchase) of the consumption. I present the statistics for rural areas and small towns separately to emphasize the potential role of access to market. Small towns are towns with a population of below 10,000. For each location groups, the table reports the fraction of households who consumed a specific crop and the fraction that consumed the crop and not produced it (i.e., the fraction who consumed a crop from purchase). The statistics is an average across the years 2011, 2013 and 2015.

Table 2.3: Crop utilization by farm households

	Consumed	Kept for seed	marketed
Barley	68.18	19.07	7.58
Maize	80.46	7.11	8.56
Millet	78.29	10.17	5.61
Oats	66.72	19.14	9.83
Rice	81.64	14.07	4.29
Sorghum	80.44	8.81	6.50
Teff	58.66	13.34	22.46
Wheat	62.35	17.76	14.76
Mung bean	20.84	12.11	62.76
Cassava	50.00	35.00	15.00
Chick pea	69.82	14.90	12.01
Haricot beans	85.28	7.97	5.76
Horse beans	71.07	14.02	11.48
Lentils	37.98	20.05	40.65
Field pea	63.88	18.11	13.97
Vetch	60.28	16.99	18.88
Gibto	29.23	26.31	43.69
Soya beans	14.59	13.54	69.20
Red kidney beans	75.43	8.78	14.19
Total	70.80	12.20	12.74

Notes: This table shows crop utilization by households. The first column shows the percent of production consumed within the household. Column 2 shows the percent kept for seed (input for next planting season), and column 3 shows the percent sold.

Table 2.4: Correlation between household production and consumption decisions

	2011	2013	2015
	Panel A: Land share of crops		
Expenditure Share	0.468***	0.530***	0.200***
	(0.025)	(0.028)	(0.021)
N	71067	71130	69932
R^2	0.284	0.331	0.177
	Panel B: Labor share of crops		
Expenditure Share	0.453***	0.513***	0.193***
	(0.025)	(0.030)	(0.021)
N	71067	71130	69932
R^2	0.292	0.328	0.182

Notes: Standard errors are clustered at village level. In panel A the dependent variable is the share of household land allocated to each crop while in panel B it is the share of labor allocated to each crop. All regressions include the control variables of village crop prices and yields, household fixed effects, and crop fixed effects. Observations are weighted by the household sampling weight. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.5: URRAP roads and the correlation between production and consumption decisions

	No village FE		With Village FE	
	Land	Labor	Land	Labor
Expenditure Share	0.372***	0.361***	0.370***	0.359***
	(0.026)	(0.027)	(0.026)	(0.027)
Post*URRAP	0.008***	0.008***	0.008***	0.008***
	(0.002)	(0.002)	(0.002)	(0.002)
Expenditure Share*Post*URRAP	-0.165***	-0.162***	-0.164***	-0.161***
	(0.032)	(0.032)	(0.032)	(0.032)
N	133372	133370	135461	135458
R^2	0.214	0.217	0.217	0.220

Note: Standard errors are clustered at village level. In columns 1 and 2 I include log distance to population centers and log distance to roads in 2011 (before the onset of URRAP). In columns 3 and 4, I include village fixed effects. All regressions include crop and year fixed effect. Observations are weighted by the household sampling weight.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.6: The separability test

	OLS Taste	IV Taste
Taste*2011	0.666***	0.677***
	(0.030)	(0.030)
Taste*2013	0.693***	0.702***
	(0.032)	(0.032)
Taste*2015	0.659***	0.655***
	(0.034)	(0.034)
N	153293	153293
R^2	0.327	0.330

Note: Standard errors are clustered at village level. All regressions include the following control variables: log village prices, log village yields, and log rainfall with crop specific coefficients. All regressions include village, crop, and year fixed effect. Observations are weighted by the household sampling weight.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.7: Separability and proximity to market

	Panel A: Distance to Population center	
	OLS Taste	IV Taste
Taste	0.454***	0.504***
	(0.112)	(0.104)
Log Dist. to Pop. Center	-0.000	-0.002**
	(0.000)	(0.001)
Taste*Log Dist. to Pop. Center	0.059**	0.047*
	(0.028)	(0.026)
N	153293	153293
R^2	0.322	0.324
	Panel B: Distance to Road	
Taste	0.613***	0.616***
	(0.052)	(0.051)
Log Dist. to Road	-0.000	-0.001*
	(0.000)	(0.001)
Taste*Log Dist. to Road	0.029*	0.029*
	(0.016)	(0.015)
N	153293	153293
R^2	0.321	0.324

Note: Standard errors are clustered at village level. All regressions include the following control variables: log village prices, log village yields, and log rainfall with crop specific coefficients. All regressions include village, crop, and year fixed effect. Observations are weighted by the household sampling weight.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.8: Balancing of variables for Average Treatment Effect on Treated (ATT)

	Treated	Control	% bias	t-stat	p-value
Population	5992.6	5983.3	0.2	0.23	0.822
Distance to nearest asphalt road	39.284	39.333	-0.1	-0.12	0.906
Distance to Woreda town	17.017	17.017	0.0	0.00	1.000
Distance to nearest major town	63.647	63.626	0.0	0.05	0.962
Distance to the nearest weekly market	7.3438	7.3438	0.0	0.00	1.000
Land slope	2.6623	2.6639	-0.1	-0.12	0.905
Fraction of land covered by forest	14.342	14.429	-0.6	-0.65	0.516
Average rainfall (1990-2010)	1149.8	1150.5	-0.1	-0.16	0.870

Notes: Population and rainfall correspond to the period before URRAP. Land slope is categorical variable with Flat=1, Slightly Slopeing=2, Moderately Sloping=3, Seeply sloping=4, and Hilly=5.

Table 2.9: The effects of URRAP on separability: Matching-based DID estimation

	Binary treatment		Market access approach	
	OLS Taste	IV Taste	OLS Taste	IV Taste
Taste	0.735***	0.738***	0.728***	0.730***
	(0.035)	(0.035)	(0.034)	(0.034)
Post*URRAP	-0.002**	0.000		
	(0.001)	(0.001)		
Taste*Post*URRAP	-0.056*	-0.054*		
	(0.029)	(0.029)		
Market Access			-0.000	0.002*
			(0.001)	(0.001)
Taste*Market Access			-0.077***	-0.071***
			(0.026)	(0.023)
N	77872	77712	77109	76949
R^2	0.343	0.345	0.345	0.347

Note: Standard errors are clustered at village level. Market access measure is standardised so that the coefficient can be interpreted as the effect of one standard deviation increase in market access. All regressions include the following control variables: log village prices, log village yields, and log rainfall with crop specific coefficients. All regressions include village, crop, and year fixed effect. Observations are weighted by the household sampling weight.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.10: The effects of URRAP on separability

	Binary treatment		Market access approach	
	OLS Taste	IV Taste	OLS Taste	IV Taste
Taste	0.680***	0.684***	0.685***	0.688***
	(0.035)	(0.035)	(0.034)	(0.034)
Post*URRAP	-0.001*	0.000		
	(0.001)	(0.001)		
Taste*Post*URRAP	-0.043	-0.038		
	(0.030)	(0.029)		
Market Access			-0.000	0.002
			(0.001)	(0.001)
Taste*Market Access			-0.068**	-0.063**
			(0.028)	(0.025)
N	102552	102552	98029	98029
R^2	0.366	0.368	0.370	0.372

Note: Standard errors are clustered at village level. Market access measure is standardised so that the coefficient can be interpreted as the effect of one standard deviation increase in market access. All regressions include the following control variables: log village prices, log village yields, and log rainfall with crop specific coefficients. All regressions include village, crop, and year fixed effect. Observations are weighted by the household sampling weight.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Appendices

A Construction of market access measure

The major concerns in identifying the effects of road connectivity based on a binary treatment dummy include: (i) heterogeneity in treatment intensity across villages that get connected to sparse network and those that get connected to dense network, and (ii) the potential spillover effects of the roads to villages that are not directly connected. When a given village is connected to the pre-existing road network or to the nearest urban center, all its neighboring villages also have improved access to market via the connected village. As a result, non-connected villages may not serve as control groups in identification of the effects of road connection. Both these concerns can be addressed by using a treatment measure that takes into account change *market access* from both direct and indirect connectivity, and the density of the network to which a village gets connected. I use market access measure derived from general equilibrium trade models (see Donaldson and Hornbeck (2016)) that are calculated using the entire road network and the distribution of population across Ethiopian villages:

$$MarketAccess_{ot} = \sum_{d} \tau_{odt}^{-\theta} Population_d \tag{14}$$

where $Population_d$ is destination village population from the 2007 census (before the onset of the URRAP program). Using pre-URRAP population distribution is necessary because population distribution is likely to respond to improvement in road infrastructure. θ is trade elasticity parameter which I estimate as discussed below.

τ_{odt} is the freight costs of transporting one ton of cargo from origin village o to destination village d along the least cost path, before $(t = 0)$ and after $(t = 1)$ the construction of URRAP roads. I use the following procedure to estimate τ_{odt} for each year. First, I construct a link from each village centroid to the nearest available road in year t. Next, I use data on costs of moving weight (in USD per ton-kilometer) for five different road quality levels: asphalt, major gravel, cobbled road, minor gravel, and earth road. Because there is no similar cost estimates along the link roads, I scale up the costs along earth road by the factor of $\frac{\textit{Cost along earth road}}{\textit{Cost along minor gravel}}$ to obtain estimate of cost along the links.[10] After assigning each road type (including the links) with the estimated costs in USD per ton-kilometer, I use ArcGIS tools to calculate the costs (in USD) of moving a ton of weight from origin o to destination d along the least cost path, in each year. I use these estimates as τ_{odt}. As can be seen in equation 14, a change to a village's market access comes only from changes in τ_{odt}.

Estimation of θ: I follow Sotelo (2018) in the estimation of θ. I assume that the average productivity of a farmer's land is related to the GAEZ village yield measure in the following equation:[11]

$$A_i^k = \delta^k Y_v^k \exp(-u_i^k)$$

where A_i^k is the average productivity of a farmer's land in crop k, Y_v^k is the village yield, $\exp(-u_i^k)$ is a random noise, and δ^k is crop-specific constant. Plugging this for A_i^k in the land share equation[12] $\eta_i^k = \frac{(p_i^k A_i^k)^\theta}{(\Phi_i)^\theta}$, where $\Phi_i = \left(\sum_{l=1}^K (p_i^l A_i^l)^\theta\right)^{\frac{1}{\theta}}$, and taking logs gives:

$$\ln\left(P_i^k Y_i^k\right) = \frac{1}{\theta}\eta_i^k + \ln\Phi_i - \ln\delta^k + u_i^k$$

[10] I show that the results are robust to using alternative scales that are half or twice of the baseline scale $\frac{\textit{Cost along earth road}}{\textit{Cost along minor gravel}}$.

[11] Some of the crops in my sample do not have GAEZ productivity estimate. For these crops, I rely on AgSS village level crop yield estimate that is constructed based on a random sample of crop cut. To purge out the noise in yield estimate and fluctuations due to whether conditions, I take the average across four years (2012-2016) to obtain a time invariant measure of yield for a crop in a village. I make sure that this approach gives reasonable village productivity estimate from comparison of GAEZ and AgSS village productivity measures for those crops that are covered in both datasets.

[12] I drop the variable input prices from the land share equation, as they do not affect the estimation equation to obtain θ (they are subsumed in village fixed effects).

The empirical counterpart of this is:

$$\ln\left(P_v^k Y_v^k\right) = \frac{1}{\theta}\eta_v^k + \gamma_v + \gamma^k + u_v^k$$

where γ_v and γ^k are village and crop fixed effects respectively. Notice that because the left hand side varies at village level (because both price and yield vary at village level), I aggregate the land share of a crop at village level as well. Thus the estimated θ is essentially a measure of within village land homogeneity.

I obtain a value of $\hat{\theta} = 2.7$ for productivity heterogeneity, which is larger than the estimate of Sotelo (2018) around 1.7, but smaller than that of Donaldson (2018) who reports a mean of about 7.5 across the 17 crops in his data.

Once I obtain the estimate for θ, I plug in equation 14 along with data on population and freight costs to obtain village level market access measures, before and after the URRAP program. On average, market access increased by about 50% among all villages due to URRAP roads. Market access increased significantly more, by about 38%, for villages that got direct connection relative to villages that did not get direct connection.

B URRAP roads and market integration

I use two measures of market integration to provide evidence on the effect of URRAP on market integration. The first measure is urban-rural price gap while the second measure is correlation between local prices and local yields for crops.

URRAP decreased trade costs: The main objective of URRAP roads was to integrate rural villages to market centers (Ethiopian Road Authority, 2016). If URRAP roads really integrated rural villages to local market centers, we would see the price gap between the rural villages and the market centers decreasing for villages that got road connection relative to villages that did not get roads. I test whether this was achieved by looking at the difference in crop prices between zone capitals and the villages within the zones using the two rich price surveys, AgPPS and RPS.

I run the following regression:

$$lnP^k_{zmt} - lnP^k_{zvmt} = \alpha_1 Post_t + \alpha_2(Post_t * URRAP_v) + \gamma_v + \gamma^k_m + \gamma_t + \varepsilon^k_{zvmt}$$

where k denotes crop, v is village, z is zone capital, m is month, t is year, $Post$ equals zero for all month-years before URRAP and one for all month-years after URRAP; $Road_v$ is a dummy variable representing whether a village got URRAP road between 2013 and 2015; and γ^k_m is crop-month fixed effect which captures possible seasonality of crop prices.

The result is reported in Table A.1. It shows that road connection significantly decreased the urban-rural price gap. The first column pools all 56 crops for which data is available on both urban and rural prices. It shows that trade cost, as proxied by the ratio of urban to rural prices, decreased by about 3% for villages that got road connection, relative to villages that did not get road connection. In column 2, the estimation is restricted to perishable products, vegetables and fruits. The estimated decrease in trade cost for these products is more than twice the estimate for all crops: trade cost for vegetables and fruits decreased by about 8%. This is not surprising because trading such products is difficult when there is no road passable by vehicle connecting a village to the urban center due to their perishability. In the last column, the sample is restricted to observations in which urban prices are higher than rural prices, which is what one would expect if villages are net exporters of crops to urban centers.[13] The gap between these two prices are plausibly capturing trade costs, which decrease by about 2.4%.

URRAP decreases the correlation between local prices and yields: One key indicator of an integrated market is that local prices are less sensitive to local supply. Under autarky, prices are relatively lower (higher) for the goods in which a region has a comparative advantage (disadvantage). Market integration weakens this inverse relationship between local prices and local comparative advantage. I run the

[13] Note that about 80% of observations (67,147 out of 82,944) conform with this expectation.

Table A.1: URRAP road access and trade costs

	Dependent Variable: log(Price in Zone Capital/Price in village)		
	All crops	Vegetables and Fruits	Cases where dep. var >0
$Post_t * URRAP_v$	-0.031**	-0.079*	-0.024*
	(0.016)	(0.044)	(0.013)
N	82944	24468	67147
R^2	0.378	0.360	0.493

Notes: Standard errors are clustered at village level. This table is based on AgPPS and RPS datasets. The regression includes 422 villages, 57 urban centers, and 56 crops. All regressions include village, crop-month, and year fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

following generalized difference-in-differences regression to investigate this:

$$lnP_{vt}^k = \alpha_1 lnA_v^k + \alpha_2(Post_t * URRAP_v) + \alpha_3(lnA_v^k * Post_t * URRAP_v) \\ + \gamma_v + \gamma_k + \gamma_t + \varepsilon_{vt}^k$$

where P_{vt}^k is price of crop k in village v, A_v^k is a village's productivity in crop k which is proxied by GAEZ potential yield for the crop.

The result is presented in Table A.2. We see that there is a negative relationship between local prices of a crop and local comparative advantage, and that this negative relationship is significantly weakened when a village gets road connection. The elasticity of village price to village yield is 2.7% for a village with no road connection and a road connection decreases this estimate to 1.7%.[14]Panel B of table A.2 reports the corresponding estimation result using market access measure instead of binary treatment dummy. The result clearly shows that in villages that see an increase in their market access, the negative correlation between crop price and yield becomes significantly weaker.

[14]Alternatively, a positive α_3 would imply that road connectivity increases the prices of crops in which a village has a comparative advantage.

Table A.2: Rural roads and the link between local prices and local yield: the dependent variable is crop-village level prices in 2012 and 2015.

	(1)	(2)	(3)
	Panel A: Binary Treatment		
LogYield	-0.036***		-0.027***
	(0.003)		(0.003)
Post*URRAP		0.017	-0.083***
		(0.023)	(0.023)
LogYield*Post*URRAP			0.009***
			(0.003)
N	59270	59270	59270
R^2	0.752	0.739	0.776
	Panel B: Market access approach		
LogYield	-0.036***		-0.099***
	(0.003)		(0.026)
LogMarketAccess		-0.026**	-0.043**
		(0.012)	(0.020)
LogYield*LogMarketAccess			0.006**
			(0.003)
N	59270	59270	59270
R^2	0.790	0.780	0.795

Notes: Standard errors are clustered at village level. The regression includes 277 villages, and 20 crops. All regressions include crop and year fixed effects, and log rainfall as a control. The last column includes village fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

C Proof of the effect of trade costs on correlation between land share and taste

Proof. Let r_v is the rental rate of a plot of land in village v, which is determined in equilibrium (see Sotelo 2018). The unit cost of production $c_i^k = \frac{r_v}{Z_i^k}$ is stochastic because it is a function of stochastic productivity Z_i^k. As a result, the price at which farmer i supplies crop k to farmer j, $P_{ij}^k = \frac{r_v}{Z_i^k}\tau_{ij}^k$, is stochastic.

Using the distribution of Z_i^k, we obtain the following distribution of the prices of crop k that farmer j is offered by another farmer i:

$$G_{ij}^k(p) = 1 - \exp\Big(- (A_i^k)^\theta (r_v \tau_{ij}^k)^{-\theta} p^\theta \Big)$$

Because crops supplied by different farmers are homogeneous, farmer j buys each crop k from any farmer that supplies the crop at the lowest price. Thus, the distribution of the price of crop k that is actually paid by farmer j is the distribution of the lowest prices across all other farmers and is given by:

$$\begin{aligned} G_j^k(p) &= 1 - \prod_{i=1}^{I}(1 - G_{ij}^k(p)) \\ &= 1 - \exp\Big(- p^\theta \sum_{i=1}^{I} (A_i^k)^\theta (r_v \tau_{ij}^k)^{-\theta} \Big) \end{aligned} \tag{15}$$

Given this probability distribution, we can derive the probability that farmer i is the cheapest supplier of crop k to farmer j (the probability that farmer i's productivity draw adjusted for trade costs and rental rates is the highest compared to all other potential farmers trading with farmer j) as:

$$\begin{aligned} \pi_{ij}^k &= \Pr\Big[P_{ij}^k \leq \min_n \{P_{nj}^k\} \Big] \\ &= \frac{(A_i^k)^\theta (r_i \tau_{ij}^k)^{-\theta}}{\sum_i (A_i^k)^\theta (r_n \tau_{ij}^k)^{-\theta}} \end{aligned}$$

which is increasing in the average productivity of farmer $i's$ plots in crop k, A_i^k and decreasing in the trade cost, τ_{ij}^k and the rental rate of farmer i's plot r_i relative to other farmers.[15]

The probability that a farmer will be the cheapest supplier of a crop to *itself* is

$$\begin{aligned}\pi_{ii}^k &= \Pr\Big[P_{ii}^k \leq \min_{\{n\neq i\}}\{P_{ni}^k\}\Big] \\ &= \frac{(A_i^k)^\theta r_i^{-\theta}}{\sum_n (A_n^k)^\theta (r_n \tau_{ni}^k)^{-\theta}}\end{aligned}$$

Farmer i is more likely to self-produce crop k if the farmer is more productive in the crop relative to other farmers and/or the higher the trade cost between farmer i and other farmers.

Note that $\pi_{ij}^k \leq \pi_{ii}^k$ because of trade costs: a farmer is more likely to be the cheapest supplier of a crop to itself than the cheapest supplier to any other farmer. Figure A.1 illustrates this.

Figure A.1: This figure illustrates the probability of a farmer being a cheapest supplier of a crop to itself and to any other farmer.

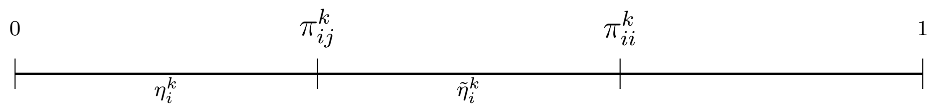

Let ψ_i^k is the fraction of land allocated to crop k by farmer i. ψ_i^k is given by:

$$\psi_i^k = \pi_{ij}^k \eta_i^k + (\pi_{ii}^k - \pi_{ij}^k)\tilde{\eta}_i^k$$

where η_i^k and $\tilde{\eta}_i^k$ are given in section 2.3. Taking derivative with respect to crop taste a_i^k gives

$$\frac{\partial \psi_i^k}{\partial \mu_i^k} = (\pi_{ii}^k - \pi_{ij}^k)\frac{\partial \tilde{\eta}_i^k}{\partial \mu_i^k}$$

[15]Note that, to save on notation I use r_i instead of r_v, even though rental rates are the same across farmers in the same village. This is without loss of generality because $r_i = r_j, \forall i, j \in v$.

which is positive given the expression for $\tilde{\eta}_i^k$. That is, more land is allocated to a crop for which the household has higher taste. Now to show that the effect of taste on land share is stronger if the household's trade cost for crop k with any other farmer j is higher, we take derivative of the above equation with respect to τ_{ij}^k:

$$\frac{\partial^2 \psi_i^k}{\partial \tau_{ij}^k \partial \mu_i^k} = \frac{\partial(\pi_{ii}^k - \pi_{ij}^k)}{\partial \tau_{ij}^k} \frac{\partial \tilde{\eta}_i^k}{\partial \mu_i^k}$$

Notice that the first term is positive because $\frac{\partial \pi_{ii}^k}{\partial \tau_{ij}^k} > 0$ (a farmer is more likely to be its own cheapest supplier the higher is the trade cost) and $\frac{\partial \pi_{ii}^k}{\partial \tau_{ij}^k} < 0$ (a farmer is less likely to be the cheapest supplier to any other farmer the higher is the trade costs). As a result, $\frac{\partial^2 \psi_i^k}{\partial \tau_{ij}^k \partial \mu_i^k} > 0$, which completes the proof that the correlation between the land share of a crop and the household taste becomes stronger with increase in trade costs. □

D An alternative test of separability

The next robustness check exploits the richness of the ESS data to test separability following the classic approach introduced by Benjamin (1992). This approach tests separability using the relationship between household on-farm labor demand and the household's demographic characteristics. The basic idea is as follows.[16] If labor market is complete and farm household's production decisions are independent of the household's preferences, household's on-farm labor demand should be independent of the household's demographic composition, such as the number of active age persons in the household.

The critical challenge in testing separability in this approach is that unobserved factors may affect both the household demographic composition and the household's farm labor demand. For example, household's land holding and/or the quality of the land may affect both household labor demand and household size (which is likely

[16] I refer interested readers to Benjamin (1992) and LaFave and Thomas (2016) for detailed discussions on the theoretical frameworks underlying this approach.

to be endogenously chosen based on wealth/land holding). While household land holding is reported in many surveys, accounting for land quality is often quite difficult. Another example includes shocks (such as weather shock) that effect both farm labor demand and household size through migration of family members. Drought decreases farm labor demand and may also lead some of the household members to migrate to cities for non-farm employment. Household specific shocks such as death and giving birth affect both labor demand and household demography.

Equipped with a panel data and a significant geographic variation in my sample households, I mitigate most of these problems using fixed effects. Time invariant household characteristics such as land size/quality are subsumed into household fixed effects. Shocks that uniformly affect households at a given location are accounted for by location-year fixed effects. The effect of household specific shocks that are likely to be correlated with household labor demand and demographic characteristics are addressed by restricting estimation to sub-samples with constant household size across the sample period.

I run similar specifications as Benjamin (1992) and LaFave and Thomas (2016) to compare my results with theirs. In my data, labor is measured in hours of work, and I observe hours spent on *planting* and *harvesting* separately. I report results for *total* labor demand (harvesting *plus* planting hours), and separately for planting and harvesting labor. Table A.3 reports the estimation results.

Table A.3 reports the estimation results. In my data labor is measured in hours of work, and I observe hours spent on *planting* and *harvesting* separately. I report results for *total* labor demand (harvesting *plus* planting hours), and separately for planting and harvesting labor. The result shows an unambiguous rejection of separability – household demographic composition significantly affects household labor demand. This result is robust across specifications that include household fixed effects and those that do not, and across planting and harvesting labor. Panel A includes the effects of the number of *males* of different age groups. Higher number of males of any age group is positively associated with on-farm labor demand throughout the specifications, with the effect peaking at the age group 35-49 for the preferred specification (those with household fixed effects). Panel B reports the effect of number of females of different age groups on labor demand. Clearly the number of female members of a household is not significantly associated with farm labor demand regardless of their age groups. This is less of a surprise for those who are familiar with agriculture in least developed countries such as Ethiopia. Farming in these part of the world is extremely physical, and women participation is limited to less physical activities such as weeding. Also important is the traditional division of labor where men work in the fields and women stay at home taking care of children and household activities such as cooking and cleaning.

Panel C reports the joint significance test of the coefficients for different age and sex groups. Both the F-statistics and the p-values are reported. Consistent with the statistical significance of the individual coefficients we observe that the coefficients for male members of different age groups are jointly statistically significant across all the specifications while the coefficients for females is jointly statistically significant only in the specifications without the household fixed effects and in the labor demand for planting (women are more likely to take part in planting activities such as weeding). Overall, the demographic variables are jointly statistically significant as shown by the F-statistics and the p-values of *all* age and sex groups, and in particular the joint significance of the *prime-age* groups (ages 15-64).

The result implies an unambiguous rejection of separability – household demographic composition significantly affects household labor demand. This is consistent with the new test suggested in this chapter. However, there are important differences in the two approaches. While any market incompleteness can lead to rejection of separability in the Benjamin's test, the test can be considered as a direct test of *missing* or *thin* labor markets. On the other hand, the approach suggested in the current chapter can be considered as a direct test of *missing* or *thin* crop markets. In this sense, the two approached also complement each other. Also important is that the Benjamin's approach relies on recall based data on labor input. Given the fact that most of the farm households are self-employing, the reliability of such data is questionable. The method suggested in the current chapter is *less* prone to such problem since land area is measured by trained enumerators using GPS tools.

Table A.3: The effect of household composition on farm labor demand: labor demand is measured as log-hours

	Pooled		Household Fixed effect		
	Total	Total	Total	Harvesting	Planting
	(1)	(2)	(3)	(4)	(5)
A. Number of Males					
age0_14	0.349***	-	0.136***	0.064	0.144***
	(0.025)		(0.036)	(0.043)	(0.044)
age15_19	0.275***	0.483*	0.200***	0.176**	0.253***
	(0.048)	(0.248)	(0.059)	(0.069)	(0.066)
age20_34	0.561***	0.999***	0.300***	0.244***	0.338***
	(0.050)	(0.22)	(0.058)	(0.068)	(0.063)
age35_49	0.691***	1.205***	0.295***	0.261**	0.338***
	(0.075)	(0.311)	(0.095)	(0.103)	(0.102)
age50_64	0.840***	2.305***	0.182	0.242*	0.156
	(0.084)	(0.327)	(0.111)	(0.127)	(0.122)
age65_above	0.413***	0.977***	0.087*	0.087	0.098*
	(0.038)	(0.140)	(0.047)	(0.054)	(0.053)
B. Number of females					
age0_14	0.286***	-0.280	0.054	0.020	0.088**
	(0.027)	(0.173)	(0.038)	(0.044)	(0.043)
age15_19	0.118**	-0.218	0.014	0.011	0.043
	(0.052)	(0.249)	(0.054)	(0.058)	(0.061)
age20_34	0.033	-0.478*	0.017	0.032	0.041
	(0.061)	(0.266)	(0.065)	(0.070)	(0.072)
age35_49	0.188**	0.198	0.120	0.065	0.151
	(0.081)	(0.297)	(0.090)	(0.100)	(0.097)
age50_64	0.622***	0.984***	0.059	0.189	0.023
	(0.086)	(0.265)	(0.117)	(0.133)	(0.123)
age65_above	0.032	-0.168	-0.052	-0.031	-0.046
	(0.042)	(0.152)	(0.046)	(0.052)	(0.052)
Log household size		1.835***			
		(0.074)			
C. Joint tests of significance					
All groups	62.46***	14.54***	3.81***	1.78**	4.31***
	(0.000)	(0.000)	(0.000)	(0.046)	(0.000)
Males	82.37***	21.31***	5.80***	2.84***	6.31***
	(0.000)	(0.000)	(0.000)	(0.001)	(0.000)
Females	26.55***	4.85***	1.64	0.57	2.23**
	(0.000)	(0.000)	(0.132)	(0.753)	(0.037)
Prime age	50.13***	12.42***	4.01***	2.45**	4.88***
	(0.000)	(0.000)	(0.000)	(0.012)	(0.000)
N	10353	10349	10264	10264	10264
R^2	0.354	0.380	0.864	0.820	0.830

Standard errors are clustered at household level. All regressions include Zone-Year fixed effects. The first three columns use the sum of planting and harvesting labor as dependent variable. Column 2 uses household size and shares of age groups in the household as regressors (see Benjamin (1992), and LaFave and Thomas (2016)). Prime age is defined as ages 15-64.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

E Appendix Tables

Table A.4: The effect of rainfall on village prices

	(1)	(2)	(3)
Log Rainfall	-0.087***	-0.087***	-0.087***
	(0.028)	(0.028)	(0.029)
Log GAEZ Yield		-0.025***	
		(0.003)	
Crop×YearFE	Yes	Yes	Yes
Village FE	Yes	Yes	.
Village ×CropFE	No	No	Yes
N	208324	208324	208324
R^2	0.809	0.813	0.920

Notes: Standard errors are clustered at village level. The regression includes 333 villages, and 20 crops. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Bibliography

Asher, S. and Novosad, P. (2019). Rural roads and local economic development. Dartmouth College.

Atkin, D. (2013). Trade, tastes, and nutrition in india. *American Economic Review*, 103:1629–63.

Benjamin, D. (1992). Household composition, labor markets, and labor demand: Testing for separation in agricultural household models. *Econometrica*, 60:287–322.

Deaton, A. S. and Muellbauer, J. N. J. (1980). American economic association an almost ideal demand system. *The American Economic Review*, 70:312 – 326.

Donaldson, D. (2018). Railroads of the raj: Estimating the impact of transportation infrastructure. *American Economic Review*, 108(4-5):899–934.

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

Eaton, J. and Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, 70(5):1741–1779.

Ethiopian Road Authority, E. (2016). Road sector development program: 19 years assessment. page .

Funk, C., Peterson, P., Landsfeld, M., Pedreros, D., Verdin, J., Shukla, S., Husak, G., Rowland, J., Harrison, L., Hoell, A., and Michaelsen, J. (2015). The climate hazards infrared precipitation with stations—a new environmental record for monitoring extremes. *Scientific Data*.

Gebresilasse, M. (2018). Rural roads, agricultural extension, and productivity. *Job Market Paper*.

Hausman, J. A. (1994). Valuation of new goods under perfect and imperfect competition. Working Paper 4970, National Bureau of Economic Research.

Hummels, D. (2007). Transportation costs and international trade in the second era of globalization. *Journal of Economic Perspectives*, 21:131–154.

LaFave, D. and Thomas, D. (2016). Farms, families, and markets:new evidence on completeness of markets in agricultural settings. *Econometrica*, 84:1917–1960.

LaFave, D. R., Peet, E. D., and Thomas, D. (2020). Farm profits, prices and household behavior. (26636).

Shamdasani, Y. (2018). Rural road infrastructure & agricultural production: Evidence from india. page .

Shrestha, S. A. (2018). Roads, participation in markets, and benefits to agricultural households: Evidence from the topography-based highway network in nepal. page .

Singh, I., Squire, L., and Strauss, J. (1986). *Agricultural Household Models: Extensions, Applications, and Policy.*

Sotelo, S. (2018). Domestic trade frictions and agriculture. *Journal of Political Economy*, Forthcoming.

Taylor, J. E. and Adelman, I. (2003). Agricultural household models: Genesis, evolution, and extensions. *Review of Economics of the Household*, 1(1):33–58.

Chapter 3

Commodity Price Shocks, Consumption Smoothing, and Child Malnutrition: Evidence from Ethiopia

3.1 Introduction

What is the effect of frequent booms and busts in international commodity prices on rural households in poor countries who rely on production of these commodities for livelihood? A number of studies have looked at the effect of such shocks on different outcomes such as child labor and schooling (Kruger, 2007; Cogneau and Jedwab, 2012; Carrillo, 2019); child mortality (Miller and Urdinola, 2010; Haaker, 2018); and long term adult mental health (Adhvaryu et al., 2019). In this chapter, I estimate the effect of coffee price shocks on household consumption and child health in Ethiopia. Ethiopia provides an ideal setting to answer these questions compared to most of the settings studied in the literature. Over 30% of the households live below poverty line and child malnourishment is among the highest in the world. Moreover, coffee production is the main source of livelihood for about 20% of the population, who are geographically concentrated due to uneven spatial patterns in agro-climatic suitability for coffee.

This chapter addresses two main gaps in the literature. First, the existing literature presumes that the effect of commodity price shocks on child health is mediated through changes in household consumption. However, empirical evidence on how household consumption responds to commodity price shocks is scant. In fact, a number of studies show that informal risk sharing arrangements among members of relatives or communities insure (perhaps imperfectly) household consumption against income shocks.[1] Second, existing literature mostly rely on households' geographic location and spatial variation in crop productivity to measure the households' exposure to commodity price shocks. This approach masks massive variation in households' choice of crop portfolio within a narrowly defined geographic unit. In Ethiopia's top coffee producing district, for instance, the fraction of farmland allocated to coffee ranges from less than 10% for some households to 100% for others, averaging about 60%.[2] I explore variation in exposure to coffee price shocks among households within a narrowly defined geographic unit. This enables me to evaluate whether studies using geographically based measures of exposure can correctly predict the effect of price shocks on average household.[3]

In this chapter, I first use nationally representative panel data on household consumption and land utilization to estimate the effects of coffee price shocks on household consumption. Households have varying levels of exposure to coffee price shocks depending on the fraction of their farmland allocated to coffee. I show that,

[1]See, for instance, Townsend, 1995; Jacoby and Skoufias, 1997; Fafchamps and Lund, 2003; and Gertler and Gruber, 2002. However, it is important to note that these studies look at idiosyncratic income shocks, such health shock to household head. Informal risk sharing is likely to be less effective against commodity price shocks or weather shocks since these shocks would affect most members of the community at the same time.

[2]Similar pattern exists in other major coffee districts of the country and looking at more narrowly defined geography such as *Kebele*, which is the lowest administrative unit, gives similar conclusion.

[3]This also applies to studies that estimate the effects of rainfall shocks. These studies also use spatial variation in exposure to rainfall shocks. However, exposure to rainfall shocks is likely to vary among households within a narrowly defined geography because households would, for instance, diversify their income sources, use irrigation or plant drought resistant varieties to protect their consumption during adverse rainfall shocks. Hence, it is not clear apriori whether using geographic based measure of exposure to rainfall shocks gives a reasonable prediction of the average effect in a given location.

for a household that allocates all its land to coffee production, a 10% decrease in coffee prices leads to 10% decrease in adult-equivalent per capita consumption. In a top coffee district, an average household allocates about 60% of farmland to coffee, which implies 6% decrease in consumption following 10% decrease in coffee prices. This is economically significant, considering the fact that international coffee prices increased by more than 100% between 2009 and 2011 and dropped by more than 50% between 2011 and 2013.

I also show that using district level measure of exposure to coffee price shocks gives a reasonably close prediction of the average effect of coffee price shocks on household consumption obtained based on household level measure of exposure. This is a crucial result because data on household level measure of exposure to shocks is rarely available, and it is not clear a priori whether the two measures give comparable prediction of the average effect of shocks.

Next, I estimate the effect of coffee price shocks on child health. My empirical strategy relies on comparing children born in different rural districts with varying levels of agro-climatic suitability to grow coffee, and cohorts that are born during periods of high and low international coffee prices between 1995 and 2016. This source of variation addresses a number of identification challenges facing similar studies that use rainfall variability as a proxy to income shocks. First, rainfall shocks in a given location might have large spillovers to neighboring locations in the form of higher food prices. Coffee price shocks affect income of households in coffee districts with plausibly little or no spillover to food prices in non-coffee districts.[4] Second, a rainfall shock in one period usually has sustained effect on future income through its effect on household's productive assets such as loss of oxen (or sales of oxen to smooth out consumption).[5] Thus studies that are based on comparison of locations

[4]However, one cannot rule out spillover effects through demand and labor market channels. Income shocks to coffee districts might still affect food prices in neighboring districts through increased demand. It may also lead to temporary labor flow to the coffee districts, suppressing wages in coffee districts and increasing it in non-coffee districts.

[5]One may suppose that sustained high coffee prices could enable the households build wealth which would then reduce their future exposure to price drops. While this is plausible, the fact that these households live on very low income and the apparent vulnerability of the households to recurrent price shocks refute this hypothesis.

that have seen different levels of rainfall shocks and cohorts born during and after the rainfall shocks are likely to be biased.

I consider a number of standard child health measures such as Weight-to-Age Z-score (WAZ), Height-to-Age Z-score (HAZ), Weight-to-Height Z-score (WHZ) and anemia status. I also look at extreme health conditions such as whether a child is underweight, wasted, and stunted — government policies and aid organizations usually target reducing these extreme health outcomes. I find that children born in coffee producing districts during high coffee price periods have significantly better health outcomes compared to children born in the non-coffee districts during the same periods. Their WAZ, WHZ, and HAZ are higher, respectively, by 0.275, 0.156, and 0.251, compared to their peers born in non-coffee districts. They are also less likely to be underweight, wasted, stunted, and anemic, respectively, by 6.8%, 2.7%, 4.8% and 5.2%, compared to their peers born in non-coffee districts. I also find that their under-five mortality rate is lower by 1.8 percentage points.

One of the striking results is that children in the intermediate age group, roughly from 15-45 months, are those more affected by income shocks compared to their younger and older counterparts. A plausible explanation is related to children transiting between different food regimes – from breast-feeding to baby-foods to adult-foods. When a child is very young, mother's breast-milk might shield the child from adverse income shocks, and when the child has transited from baby-foods to adult-foods, the effect of income shocks would be lessened because adult-foods are relatively cheaper than baby-foods. Another surprising result is that the effects of income shocks are not significantly different for children from *relatively* wealthier households. While children from wealthier households have, on average, better health outcomes compared to their peers from poorer households, they are equally affected by adverse income shocks as those from poorer households. This result is in contrast to the presumption that wealthier households can self-insure against income shocks. But, it could be explained by the possibility that even households in the *richest* category are poor in absolute terms, with little or no savings to rely on to smooth out consumption during income shocks. An alternative explanation is that these households own non-liquid assets that cannot be readily converted into cash to smooth out consumption.[6]

This chapter is closely related to studies that estimate the effects of commodity price fluctuations on child health. Miller and Urdinola (2010) show that periods of high coffee prices are associated with higher child mortality in Colombia. Haaker (2018) find that a drop in coca price increased child mortality in Peru. This chapter finds that decreases in coffee prices lead to higher under-five mortality in coffee dependent Ethiopian districts. Perhaps the most closely related to this chapter is Cogneau and Jedwab (2012) who study the effect of a cut in government administered producer price for cocoa on child schooling and health in Côte d'Ivoire. They measure exposure to the the price cut using a dummy variable indicating whether a household reported a positive cocoa production. However, they do not have household level panel that covers pre- and post-shock periods to tightly identify the effect of the shock on household consumption. To the best of my knowledge, this chapter is the first to use a panel data and household level measure of exposure to commodity price shocks to estimate the effect of such shocks on household consumption.

A number of studies estimate the effect on child health of income volatility due to macroeconomic business cycles or commodity price movements, typically in high or middle income settings including: Dehejia and Lleras-Muney, 2004; Neumayer, 2004; Paxson and Schady, 2005; Ferreira and Schady, 2008; Miller and Urdinola, 2010; and Page et al., 2016. These studies mainly seek to address the theoretical ambiguity on the effect of aggregate income shocks on child health. On one hand, a positive aggregate income shock, e.g., macro economic boom, implies higher *income* for families, which improves child outcomes through investments in food and health of children. On the other hand, periods of booms also imply higher *opportunity costs*

[6]DHS classifies households into five wealth groups: –*Poorest*, *Poor*, *Middle*, *Rich*, and *Richest* – based on survey of ownership of a number of assets. This classification is purely *relative* among rural households in Ethiopia, and conveys no information about the *absolute* level of wealth or poverty situation of the households. See `https://www.dhsprogram.com/topics/wealth-index/Index.cfm` for details on how this index is created.

of mother's time, which would adversely affect child outcome if mothers spend less time on taking care of the children. These studies find that the *opportunity cost* channel dominates the *income* channel, i.e., child health measures are counter-cyclical in high income countries. This chapter finds the exact opposite in a low-income country setting. Because most households in poor countries live on a subsistence income, shocks to income have a first-order effect on spending on child nutrition and health.[7]

Recently, few papers look at the long term effects of commodity price shocks. Carrillo (2019) studies the effect of coffee price shocks during school-going age on long-run human capital accumulation and earnings in Colombia and Adhvaryu et al. (2019) use fluctuations in cocoa prices to study the effect of income shocks during childhood on adult mental health. Kruger (2007) looks at the contemporaneous effect of coffee price on child labor and schooling in Brazil. These studies measure exposure to commodity price shocks based on households' geographic location. This chapter complements these studies by showing that using measures of exposure at geographic unit gives reasonably close approximation of the average effect of shocks obtained based on household level measure of exposure.

This chapter is also related to studies that use rainfall shocks as a source of variation to address various questions related to household consumption, child health, and human capital accumulation. Earlier works include: Paxson, 1992; Townsend, 1995; Morduch, 1995; Jacoby and Skoufias, 1997; and Jensen, 2001. More recent ones include: Tiwari et al., 2013; Mendiratta, 2015; and Shah and Steinberg, 2017.

The rest of the chapter is organized as follows. Section 2 briefly describes the datasets. Section 3 gives an overview of the state of child malnutrition and a brief summary of coffee production and marketing in Ethiopia. Section 4 describes the empirical strategy, while Section 5 presents the main results. In Section 6, I present robustness checks, and Section 7 concludes the chapter.

[7] According to UNICEF (2015) almost half of under-5 mortality is attributable to under-nutrition in poor countries.

3.2 Data

I use two main datasets. The first is the Ethiopian Socioeconomic Survey (ESS) data collected by Ethiopian Central Statistical Agency (CSA) in collaboration with the World Bank.This is an exceptionally rich nationally representative panel data of about 4,000 (rural sample) households. There are three rounds of this survey for the years 2011, 2013, and 2015. It includes information on the production and consumption, access to services and facilities, and others. Importantly, the data is geographically representative of the country as it covers coffee producing and non-coffee producing districts. I use this dataset mainly to establish the effect of coffee price shocks on household consumption, and to show how results based on household level measure of exposure compare with those based on measures of exposure at geographic unit. This is crucial because most studies rely on measures of exposure at some geographic unit, and it is not clear apriori whether this approach gives similar conclusion to using household level measure of exposure.

The dataset on child health comes from the Ethiopian Demographic and Health Survey (DHS). All the four rounds 2000, 2005, 2011, and 2016 are utilized. Each round covers more than 16,500 nationally representative households. The DHS data covers children born since 1995 and the sampling covers almost all districts in the country (see Figure 3.1). DHS data is the richest and the most widely used to monitor child health for developing countries. Due to its exclusive focus on children and maternal health, the data gives detail information about the health status of children and mothers, and household and environmental factors that might influence them. For Ethiopia, all the survey rounds give anthropometeric data for children under 5 years of age. With the exception of 2000 round, the survey also includes hemoglobin measures. Hence all the analysis on hemoglobin measure and anemia status is based on the last three rounds of the survey.

I use FAO/GAEZ agro-climatically attainable yield for low/intermediate input use to construct districts' suitability for coffee production. This data gives estimates of coffee yield in quintals per hectare under different scenarios of intensity of input use. I use *intermediate input* scenario in my analysis, which is more likely to reflect the actual coffee production practice based on results from survey data.

Finally, my international coffee price data comes from *International Coffee Organization* which maintains historical statistics on international coffee price and trade.

3.3 Background

3.3.1 The state of child malnutrition in Ethiopia

Child malnourishment in Ethiopia has been among the worst in the world. Ethiopia performs worse than Sub-Saharan African (SSA) countries' average in the proportion of under-5 children stunted, underweight, and wasted. Ethiopia has the highest percent (76%) of children who have not received any of the eight EPI immunizations in SSA countries. See Kanamori and Pullum (2013) for a detailed comparison of child health outcomes across 30 SSA countries using DHS data.

However, there is moderate improvement in the last two decades though the substantial gap between rural and urban children remained steady. Table 3.1 shows trends in key child health outcomes for both rural and urban samples of the DHS data. The proportion of under-5 children who are stunted decreased from 57% to 36% in *rural* areas and from 38% to 21% in *urban* areas between the 2000 and 2016 rounds of survey. Under-5 mortality decreased from 12% to 7% in rural areas and from 10% to 3% in urban areas over the same period. The proportion of children who are moderately or severely anemic increased in both rural and urban areas between the 2005 and 2016 rounds.[8]

Gender gap in child health is not significant, even though it is slightly worse for boys.

[8]Perhaps this could be due to selection issue. In the 2005 survey a significant proportion of eligible children refused to take the hemoglobin measure. In contrast, in the 2016 survey almost all eligible children were measured.

3.3.2 Coffee production and marketing in Ethiopia

Ethiopia is a top coffee producer in Africa and ranks 5th in the world after Brazil, Vietnam, Colombia, and Indonesia in the year 2015/16. Ethiopia produced about 7 million 60-kg bags of coffee in 2015/16 , which is about 9% of the world coffee production. About 95% of coffee is produced by smallholder farmers with about a median land size of one hectare. The remaining 5% is produced by government owned farms and large scale private farms. Coffee production in Ethiopia is largely concentrated in the southern and southwestern parts of the country, where there is high rainfall and forest cover, the two essential ingredients of *Arabica* coffee production in Ethiopia.

After harvesting, coffee cherries are processed in two different ways. The dominant approach is dry-processing, meaning coffee cherries are dried in the sun on mat or cement floors and the outer layer of the cherries are removed to obtain the green beans. The wet processing, predominantly used in the coffee producers in the northeastern Oromia zones (West and East Hararge), involves coffee cherries pulped, fermented in tanks and then washed in clean waters.

About 60% of coffee production is exported, and coffee export accounts for about a quarter of the country's export revenue. Coffee export passes through a series of market chains. Because of the importance of coffee as source foreign exchange, the Ethiopian government tightly monitors the local and export trade through licensing and regulations. The Derg military regime(1975-1991) required farmers to supply to the government a quota of coffee production at a specified price. After the downfall of Derg, there has been a gradual liberalization of coffee trade but the government still puts a heavy hand in allocating licenses and strictly monitoring to ensure that export quality coffee is not being sold in local markets. Coffee is exported by Cooperatives Unions (which are essentially parastatals) or by licensed private exporters. Cooperative Unions source their export coffee from (licensed) intermediate suppliers or hulling firms while private exporters source it from coffee auction markets in Addis Ababa or Dire Dawa.

3.4 Empirical Strategy

3.4.1 Coffee price shocks and consumption smoothing

First, I use ESS data to establish whether there is a link between coffee price shocks and household consumption. ESS data allows me to measure household level exposure to coffee price shocks using the share of household farmland allocated to coffee. The estimation equation is as follows:

$$\mathrm{LogC}_{it} = \beta_0 + \beta_1(\mathrm{LandShare}_i \times \mathrm{LogPrice}_t) + \gamma_i + \gamma_t + \varepsilon_{it} \tag{3.1}$$

where i indexes household. t is year. C denotes consumption. ESS data reports household level adult-equivalent per capita consumption expenditures disaggregated into food and non-food. Significant number of households report zero non-food expenditure. Hence, I focus on total expenditure and food expenditure, even though the two are highly correlated. LandShare is the share of household land allocated to coffee in the beginning year (2011). *Price* is the average international coffee price in the quarter immediately before the survey months. In this equation, the effect of a percentage change in coffee price is given by $\beta_1 \times$ LandShare which varies across households depending on the fraction of household farmland allocated to coffee. For a household that allocates all its farmland to coffee, LandShare $= 1$, the effect of a percentage change in coffee price is given by β_1. For a household that does not produce any coffee, the effect of the price change is zero.

As a robustness check, and in order to facilitate comparison with the next sections where data on household level share of land allocated to coffee is not available, I also run similar regression where a household's exposure to coffee prices is measured by its district's coffee suitability:

$$\mathrm{LogC}_{idt} = \beta_0 + \beta_1(\mathrm{CoffeeDistrict}_d \times \mathrm{LogPrice}_t) + \gamma_i + \gamma_t + \varepsilon_{idt} \tag{3.2}$$

where d indexes districts.[9] Coffee District is a district with above 8 quintals per hectare of coffee yield (districts above national average yield — the ratio of national production in quintals to total area of land allocated to coffee nationally in hectares) based on GAEZ potential yield estimate. This approach assumes that all households that reside in a coffee producing area are equally exposed to coffee price shocks. An alternative interpretation is that β_1 gives the average effect of coffee price shocks across households in coffee districts.

I will explain how the results across the above two specifications compare, as it has crucial implications for the remaining analysis in this chapter as well as for related studies that rely on measures of exposure at some geographic unit.

3.4.2 Coffee price shocks and child malnutrition

Next, I explore the effect of coffee price shocks on child health using DHS data. My main identification strategy relies on comparing cohorts based on their birth locations (districts) and the average price of coffee between the time when they are born to the time when their health outcome is measured in the survey. The identification comes from large swing in the international coffee price between the period January 1995 and December 2016 (see Figure 3.2) and the significant variation in the coffee suitability across the Ethiopian districts (see Figure 3.3). DHS data does not include information about production, such as the fraction of household farmland allocated to coffee production. As a result, a household's exposure to coffee price changes is measured using the potential coffee yield of the household's location. However, the results from the previous subsection would give us a crucial insight on to what extent this approach gives reasonably close prediction to the average effect of the shocks obtained based on household level measure of exposure.

To facilitate ease of interpretation, I run a specification that uses dummy variables

[9] In fact, the geographic unit in this estimation equation is sub-village or Enumeration Area (EA). But to save notation and be consistent with the next subsections I refer to them as districts. In almost all cases, there is one EA sample per district in ESS data.

for coffee vs. non-coffee district and for periods of high vs. low coffee prices.

$$H_{idts} = \beta_0 + \beta_1(\text{CoffeeDistrict}_d \times \text{HighPrice}_{ts}) + \beta_2\text{HighPrice}_{ts} + X_i\delta + FEs + \gamma_d + \varepsilon_{idts} \tag{3.3}$$

where H_{idts} is a measure of health outcome for child i in district d born in month t surveyed in month s. Coffee districts are defined under equation 3.2. High price dummy assumes a value of one if $Price_{ts}$ for a child is higher than the median $Price_{ts}$, or zero otherwise. $Price_{ts}$ is the average coffee price between the birth month of a child and the month when the child's health is measured. X includes a vector of child and household characteristics including: child's age and gender, succeeding birth interval, mother's education and whether source of drinking water is tab water. FEs includes a vector of fixed effects including: year of birth fixed effects, month of birth fixed effects, dummy variables for survey round and birth order dummies. γ_d is district fixed effects, and ε_{idts} is the error term. The parameter of interest is β_1.

To exploit full variation in the data, I also run a generalized DID specification:

$$H_{idts} = \beta_0 + \beta_1(\text{LogYield}_d \times \text{LogPrice}_{ts}) + \beta_2\text{LogPrice}_{ts} + X_i\delta + FEs + \gamma_d + \varepsilon_{icd} \tag{3.4}$$

where $Yield$ is FAO/GAEZ estimation of district's potential coffee yield in quintals per hectare.

I also run a specification in which the dummy for coffee district is interacted with LogPrice_{ts} to check robustness to specifications.

Outcome variables: I consider a number of standard measures of child health including Weight-for-Age Z-score (WAZ), Height-for-Age Z-score (HAZ), Weight-for-Height Z-score (WHZ), and anemia status. I also consider extreme health outcomes including whether the child is underweight, wasted, stunted, and whether the child is alive. Considering extreme health outcomes is crucial. First, it facilitates sharp interpretation. Second, the main target of child health programs is reducing these extreme health outcomes. For instance, *wasting*, substantial weight loss due to starvation and/or disease, is strongly associated with mortality and is often used to assess severity of emergencies such as drought or other natural or man made calamities. Similarly, death of about 4 million under-5 children is associated with the underweight status of the children themselves or their mother (World Food Program (2013)). *Stunting* is an indicator of chronic malnutrition.

3.5 Results

3.5.1 Income shocks and consumption smoothing

Table 3.2 reports the estimation result for the effect of coffee price shocks on household consumption, using the share of household farmland allocated to coffee as measure of household exposure. The first two columns of the table show that coffee price shocks have a significant effect on household total and food consumptions. A 10% decrease in international coffee price decreases household consumption by about 10% for a household that allocates all its farmland to coffee production. For a household that allocates just half of its farmland to coffee, the effect is about 5% decrease in consumption. This is sizable given that coffee prices often change significantly. For instance, coffee prices increased by more than 100% between 2009 and 2011, and decreased by more than 50% between 2011 and 2013.

The consumption movement in response to coffee price shocks implies complete pass-through of the coffee price volatility to household consumption for households that produce only coffee. Fortunately, most households, even those in the district with the highest coffee yield in the country, do not allocate all their farmland to coffee production. Households in the top coffee district based on GAEZ potential yield ranking allocate only 60% of their farmland to coffee production, on average. This figure decreases to just 18% for households in the coffee districts. That is, a 10% decrease in coffee price decreases the consumption of an average household by 6% in the top coffee producing district, and by 1.8% in an average coffee producing district. This suggests that crop diversification is an effective strategy to hedge against coffee price volatility.

Columns 3-4 investigate whether borrowing helps households to smooth consumption against coffee price shocks. The interaction term $\text{LandShare}_d \times \text{LogPrice}_t \times Credit$ is negative but statistically insignificant. This might be related to limited fungibility of the loans – about 60% of borrowing households report that the loans are received for purchase of farm inputs. In columns 5-6, I include analogous interaction term with with aid, $\text{LandShare}_d \times \text{LogPrice}_t \times Aid$.[10] While this interaction term enters with expected negative sign, it is statistically insignificant.

Table 3.3 reports results based on district level measure of exposure to coffee price shocks. The estimates in the first two columns suggest that a 10% decrease in coffee price decreases consumption by about 2% for an average household that lives in coffee districts.

Comparison of tables 3.2 and 3.3 shows that the two approaches give strikingly close estimates on the effect of coffee price shocks. Table 3.2 predicts that a 10% decrease in coffee price decreases household consumption by about 1.8%, on average, for households in coffee districts (districts with above 8 quintals per hectare of GAEZ potential yield). This is reasonably close to the 2% decrease in consumption predicted in table 3.3. This is a crucial result as it boosts our confidence in using district level measure of exposure to coffee price shocks in the next subsections. It also complements previous related studies that used a measure of exposure at a level of geographic unit due to lack of data on household level measure of exposure to shocks.

Overall, this subsection establishes that income shocks have significant effect on household consumption, implying that household consumption is not fully insured even when they have access to credit and assistance from government and non-

[10] $Credit = 1$ if the household received loan from any potential sources including: relatives, government and non-government agencies, microcredit institutions, etc. $Aid = 1$ if a household received aid from relatives, friends, government, and non-governmental sources.

governmental agencies. It also establishes that using district and household level measures of exposure give us reasonably comparable predictions of the effect of coffee price shocks. In the next subsections, we will explore the effects of such shocks on children – who are arguably the most exposed to shocks.

3.5.2 Income shocks and child malnutrition

The previous subsection has shown that coffee price shocks have significant effect on household consumption. What is the consequence of this on child health? Table 3.4 reports the main results for the effect of coffee price shocks on child health. In the first three columns, continuous measures of child health – WAZ, WHZ, and HAZ – are used as outcome variables while in the last three columns I use extreme health outcomes: underweight, wasted and stunted, which correspond to, respectively, WAZ<-2, WHZ<-2 and HAZ<-2. Overall, the results show that both continuous measures and extreme health outcomes are strongly pro-cyclical. Column 1 shows that children born in coffee districts during high coffee price periods have higher WAZ by 0.275, compared to children in non-coffee districts. This estimate is economically significant – about 20% of the standard deviation of WAZ. Column 2 shows that these children also have higher WHZ of 0.156 (12% of the standard deviation WHZ) and column 3 shows that their HAZ is higher by 0.252 (15% of the standard deviation for HAZ).

The last three columns show that children born in coffee district during high coffee price periods are less likely to be underweight, wasted, and stunt, respectively by 6.8%, 2.7% and 4.8%. This estimates are large, though the proportion of children who are underweight, wasted and stunt is also large (see Table 3.1).

As an alternative specification, I run a regression in which coffee district dummy is interacted with $\log Price_{ts}$. The coefficient in this regression can be interpreted as the effect of one log-unit increase in coffee price on children in coffee districts relative to children non-coffee districts. The result is reported in table A.1. The results in the first three columns show that one log-unit higher average coffee price between when a child is born and when the child health is measured increases the child's WAZ, WHZ, and HAZ by 0.331, 0.233, and 0.312, respectively, for a child born in a coffee district compared to a child born in non-coffee district. The last three columns show that the same increase in price decreases the probability that a child is underweight, wasted and stunted by 6.9%, 5.4%, and 8.4%, respectively. The mean, standard deviation and range of $\log Price_{ts}$ are 4.76, 0.31 and 1.6, respectively. Thus, for instance, a child born in a coffee district and exposed to the highest average coffee price is less likely to be underweight, wasted, and stunted by 11%, 8.6%, and 13.4%, compared to a child born during the lowest coffee prices.

Table A.2 presents results for the generalized DID model. The first three columns show that one unit increase in $\text{LogYield}_d \times \text{LogPrice}_{ts}$ leads to 0.149 higher WAZ, 0.097 higher WHZ, and 0.148 higher HAZ. This is large given the mean, standard deviation, and range of $\text{LogYield}_d \times \text{LogPrice}_{ts}$ of 4.7, 5.3, and 15, respectively. The last three columns show one unit increase in $\text{LogYield}_d \times \text{LogPrice}_{ts}$ reduces the probability of a child being underweight, wasted, and stunt, respectively by 3.2%, 2.3% and 3.7%.

Overall, these estimates are large but not implausible given that these households mostly live close to the poverty line and that children are the most vulnerable part of the population. A 10% lower price is equivalent to a 10% lower income and consumption for a household that allocates all its land to coffee production (since production of coffee is largely inelastic to price at least in the short and medium run). For a household that lives a subsistence life, a 10% loss of consumption is substantial and the effect on children in particular is likely to be large.

Table 3.5 reports results using a different measure of child malnourishment – whether a child is (severely) anemic. Low hemoglobin level in red blood cells are shown to be an important indicator of micronutrient deficiency such as iron, zinc, vitamin A and folate. In column 1, I use whether a child is moderately anemic or worse as an outcome variable. The estimate shows that a child born in a coffee district and during higher than median coffee price period is less likely to be anemic by 6.6%, compared to children in non-coffee districts. Column 2 reports the corresponding generalized DID estimation and shows that one unit higher $\text{LogYield}_d \times \text{LogPrice}_{ts}$ is associated with 4% lower probability of being anemic.

3.5.3 Income shock and child mortality

Next, I explore the effect of income shock on child mortality. According to UNICEF 2005, about half of under-5 mortality in low income countries is attributable to malnutrition. I use the same setup as above to see whether income shocks is associated with the probability that a child is alive at the time of survey. However, because most of the child death occurs in the first few weeks or months after birth, I use *in utero* and *month of birth* exposure to price shocks, as explanatory variables.

$$\begin{aligned}\text{Alive}_{idts} = \beta_0 + \beta_1(\text{LogYield}_d \times \text{LogPriceUtero}_t) + \beta_2\text{LogPriceUtero}_t \\ + X_i\delta + FEs + \gamma_d + \varepsilon_{idts}\end{aligned} \tag{3.5}$$

where $\text{Alive}_{idts} = 1$ if child i born in district d in month t is alive in survey month s, and PriceUtero_t is the average coffee price when the child was in utero. I also run the same specification where PriceUtero is replaced by price in month of birth PriceBirthMonth_t.

Table 3.6 reports the results. The first column shows that one unit increase in $\text{LogYield}_d \times \text{LogPriceUtero}_t$ increases the probability that a child survives up to the survey time by one percentage point. The second column shows that one unit higher $\text{LogYield}_d \times \text{LogPriceBirthMonth}_t$ increases the probability that the child survives up to the survey time by 0.008. These estimates imply that a child that is born in an average coffee district during average coffee price period is 1.8 percentage point more likely to be alive by the survey time. For a child that is born in the top coffee district during the highest coffee price period, the probability of being alive at the survey time is higher by 5 percentage points. In the rural sample of DHS data for Ethiopia (across all the four rounds), on which all the analyses in this chapter are based, 9% of the children die before the age of 5 years.

3.5.4 In-utero exposure

The fetal origins hypothesis stipulates that *in utero* exposure to shocks have a permanent effect on child health. Previous studies testing this hypothesis considered mother's exposure to extreme health shocks (such as the influenza pandemic), extreme hunger or Ramadan fasting during pregnancy. It is not obvious whether the fetal origins hypothesis also applies to more common contexts where the mother is exposed to significant, but not extreme, income and nutrition shocks during pregnancy. I test this using a similar regression to equation 3.5:

$$
\begin{aligned}
H_{idt} = \beta_0 &+ \beta_1(\text{LogYield}_d \times \text{LogPriceUtero}_t) + \beta_2\text{LogPriceUtero}_t \\
&+ X_i\delta + FEs + \gamma_d + \varepsilon_{idt}
\end{aligned}
\tag{3.6}
$$

where PriceUtero is the average coffee price during the months when the child is in utero.

Tables 3.7 report results for in utero exposure to income shocks. Unlike exposure to income shocks *after* birth, in utero exposures have weaker effect. The estimated effects are not only smaller in magnitude (compared to the results reported in table A.2) but also are less precisely estimated. This is in contrast to the literature on the fetal origins hypothesis which usually report a significant effect of in utero exposure. One potential explanation why the results in this chapter are not as strong is that most of the previous studies rely on extreme shocks such as the influenza pandemic or extreme famine to identify the causal effect of in utero exposure while here we are considering the effect of significant, but not extreme, income shocks.

3.5.5 Heterogeneous effects

Age at exposure: The above setup assumes that income shocks have similar effect regardless of at what age a child is exposed to the shock, which is a strong assumption. First, a child is at different stages of physical development at different ages which implies different level of exposure to malnutrition. Second, a child's food type and source varies across ages. For instance, income shocks during the first six months to one year (when the child is mostly reliant on breastfeeding) might not have the same effect as income shocks that occur immediately after the child switches to baby-foods. I test whether exposure to price shocks at different ages have different effects in the following series of regressions:

$$H_{idtr} = \beta_0 + \beta_1^a \text{LogPrice}_a + \beta_2^a \text{LogPrice}_a \times \text{LogYield}_d) + X_i\delta \qquad (3.7)$$
$$+ FEs + \gamma_d + \varepsilon_{idtr}$$

where H_{idtr} is a measure of health for child i in district d born in month t and measured in month r. $Price_a$ is the price of coffee when the child is a months old $0 \leq a \leq r - t$. Thus, this regression pools together all the children who are age a or older and runs a series of regressions on prices $\text{Price}_s, \{s = 1, ..., a\}$, one at a time. Note that as a increases, the sample size decreases. The estimated β_2^as and their confidence intervals are then plotted against a.

Figures 3.4-3.7 show the results. The horizontal axis measures age (in months) at which a child is exposed to income shocks and the vertical axis measures the effects of the shocks (the point estimates together with the 95% confidence interval). Figure 3.4 shows the effect on HAZ. While the effects of the shocks when a child is below 45 months are mostly significant, clearly the effect is stronger when the child experiences such a shock between ages 15-45. Even though the point estimates are statistically significant at least at 10% starting from age zero, the estimated effect increases both in magnitude and precision until age 45 months. Figure 3.5 shows similar trend using WAZ, the effects of the income shocks are larger and more precise for middle-aged children between 15-45 months. Figure 3.7 shows the result for the probability that a child is moderately anemic or worse. The effect of income shocks is again stronger for middle-aged children. A somewhat surprising result is Figure 3.6 which shows no significant effect of income shocks at each ages, even though tables 3.4, A.1, and A.2 show a significant effect of average coffee price between a child's month of birth and the measurement month.

Studies on child development in developing countries consider this intermediate age as *critical age*. In Ethiopia, the proportion of stunted and underweight children sharply increases between ages 6 - 25 months, and stays higher at least until 59 months. Some studies find similar result elsewhere. For instance, Glewwe and King (2001) find that malnutrition during the second year of life has a larger negative impact on child cognitive development than malnutrition in the first year using longitudinal data from Philippines. See also Tiwari et al. (2013), Hoddinott and Kinsey (2001), and Waber et al. (1981) who find similar results.

Overall, these results are consistent with children transiting from breastfeeding to baby food and then to adult food. Younger children who are on breastfeeding are relatively less exposed to the same income shock compared to their older counterparts who have just transited from breastfeeding to baby foods. This is particularly the case if baby food is relatively more expensive and has more income elasticity than adult food, so that household to substitute adult food for baby food during adverse income shock – that is, they start feeding the child adult food during income shock. Similarly, older children who have already transited from baby food to adult food are also less susceptible, compared to their younger counterparts who have not yet transited from baby food to adult food, to the same income shock.

Gender of the child: There are several descriptive evidences that households tend to favor boys over girls in resource allocation during financial stress, particularly in South Asian countries. In this section, I explore whether income shock affects boys and girls differently. Table 3.8 reports the result. The table shows that girls have better health outcomes than boys, perhaps due to their natural advantages at early childhood level. However, there is no evidence that parents favor a particular gender during income shocks; the interaction term of income shocks with gender of the child is not statistically and economically significant. Similar results are reported by Jensen (2001) and Cogneau and Jedwab (2012) in Côte d'Ivoire.

Heterogeneity in wealth status: The significant adverse effect of coffee price shocks on household consumption and child health reported in the previous subsections imply that it is difficult to insure against covariate income shocks through informal risk sharing mechanisms. One may also ask if households could self-insure against such shocks, that is, if exposure to such shocks varies across households of different wealth groups. DHS data includes household wealth index that classifies households into five wealth groups: –*Poorest*, *Poor*, *Middle*, *Rich*, and *Richest* – based on survey of ownership of a number of assets. Note that this classification is purely *relative* among rural households in Ethiopia, and conveys no information about the *absolute* level of wealth or poverty situation of the households. I use this measure to see if exposure to income shocks vary across wealth groups by interacting this wealth indicator with $\text{LogPrice}_{ts} \times \text{LogYield}_{d}$ in equation 3.4.

The result is reported in table 3.9. The interaction term involving the *Poorest* group is omitted so that interaction terms for other wealth groups are defined relative to this group. The table clearly shows that children from wealthier households have significantly better health outcomes and are significantly less likely to be stunted, underweight, or wasted. However, the interaction terms for wealth groups with $\text{LogPrice}_{ts} \times \text{LogYield}_{d}$ are all statistically insignificant, suggesting that children from different wealth groups are not deferentially exposed to income shocks. One potential explanation is that the wealthier households are themselves in a subsistence livelihood with little or no savings to reallocate consumption inter-temporally.

3.6 Robustness

3.6.1 Using district land allocation to measure district coffee intensity

In the analysis so far, a district's coffee intensity is measured using FAO/GAEZ's measure of potential coffee yield. However, actual land usage in a district may not follow potential yield measured from agro-climatic variables. For instance, limited inter-district trade opportunity, due to poor market integration, may force districts with high potential coffee yield to allocate significant fraction of land to other necessary crops, such as cereals, for consumption needs.

An obvious way to address this is to use actual fraction of district land allocated to coffee as a measure of a district's coffee intensity. Unfortunately there is no representative agricultural survey at a district level to construct the share of land allocated to coffee in each district. The Agricultural Sample Survey (AgSS) data collected by Central Statistical Agency (CSA) of Ethiopia is representative only at Zone and Village levels. However, one can still construct district level measure of land share of coffee from the corresponding statistics for the villages that lie within the district boarder. In this section I follow this procedure. Figure A.1 depicts the variation in land share of coffee across districts.

Due to concern of measurement error in the fraction of land allocated to coffee from using non-representative data, I use instrumental variables (IV) estimation where FAO/GAEZ measure of potential yield is used as an IV for the actual fraction of land allocated to coffee. The correlation between the two is fairly high, about 0.5. The First-stage F-statistics is 39.7, which implies that the instrumental variable is strong. Tables A.3 reports the IV estimation results. The results are qualitatively similar to results from similar specifications using FAO/GAEZ potential yield as a measure of district coffee intensity reported in table 3.4. However, the estimated coefficients are larger in table A.3 than in table 3.4. Though large, these results are not unreasonable.

3.7 Conclusions

In this chapter, I provide strong evidence that exogenous shocks to international coffee prices have significant effect on household consumption in rural Ethiopia using household level measure of exposure to coffee price shocks. Moreover, I show that using district level measure of exposure to coffee price shocks gives reasonably close approximation of the effect of the shocks on consumption of the average household in coffee districts. This is crucial result because data on household level measure of exposure to shocks is rarely available.

I then explore the consequences of these consumption fluctuations on child health, and show that children who are born in coffee districts during low coffee price periods are significantly more likely to be underweight, stunt, wasted, anemic and less likely to be alive, compared to their peers in non-coffee districts. In view of evidences that childhood malnourishment and poor health have lasting effects on the children, this result implies significant welfare cost of price fluctuations.

While children from wealthier households have better health outcomes compared to their peers from poorer households, they are equally exposed to income shocks, perhaps because even the households in the wealthiest category in relative terms are poor in absolute terms. Access to credit and aid also do not have significant consumption smoothing effects.

Future studies that seek experimental evidence on whether innovative insurance schemes, such as commodity price indexed insurances (similar to weather indexed insurance schemes), have a positive effect on welfare would be a step forward towards solution to this problem.

Figure 3.1: DHS sample locations

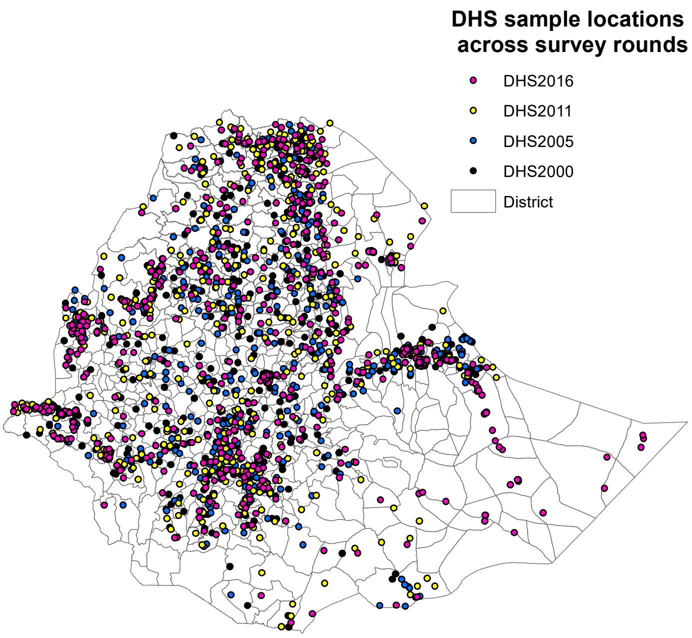

This figure shows the DHS sample locations for each survey round.

Figure 3.2: Monthly international coffee price index

This figure plots coffee price for the variety of coffee known as *Brazilian Naturals*, which is Arabica species that have been dried inside the fruit rather than after the fruit has been removed. It is one of the most widely traded coffee type, with Ethiopia as one of the major suppliers. *Data Source: International Coffee Organization.*

Figure 3.3: GAEZ yield (kg/hectare)

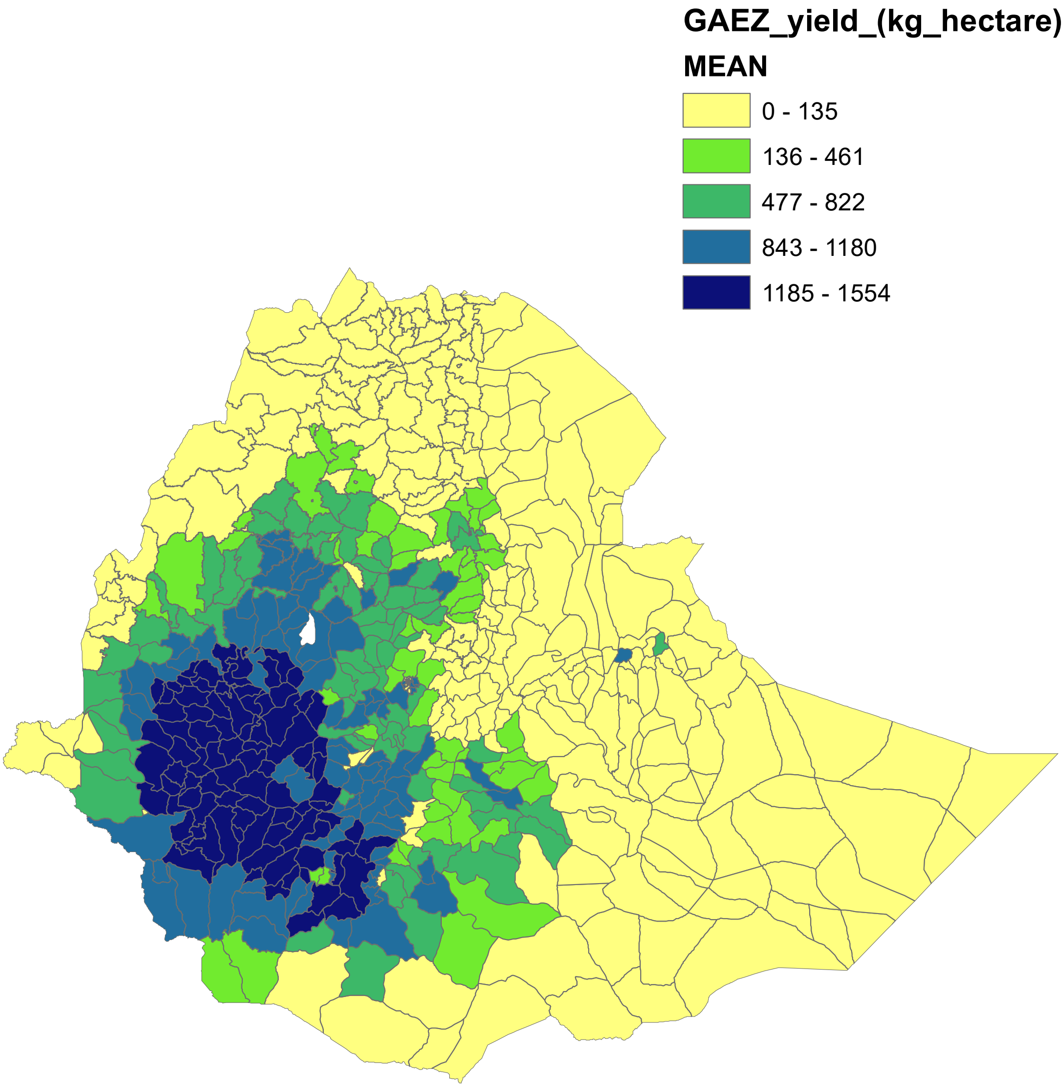

Notes: This figure shows GAEZ estimated yield for coffee based on rain-fed and intermediate input usage farming techniques. *Data Source: FAO/GAEZ.*

Figure 3.4: The effect on HAZ of exposure to price shock at different ages

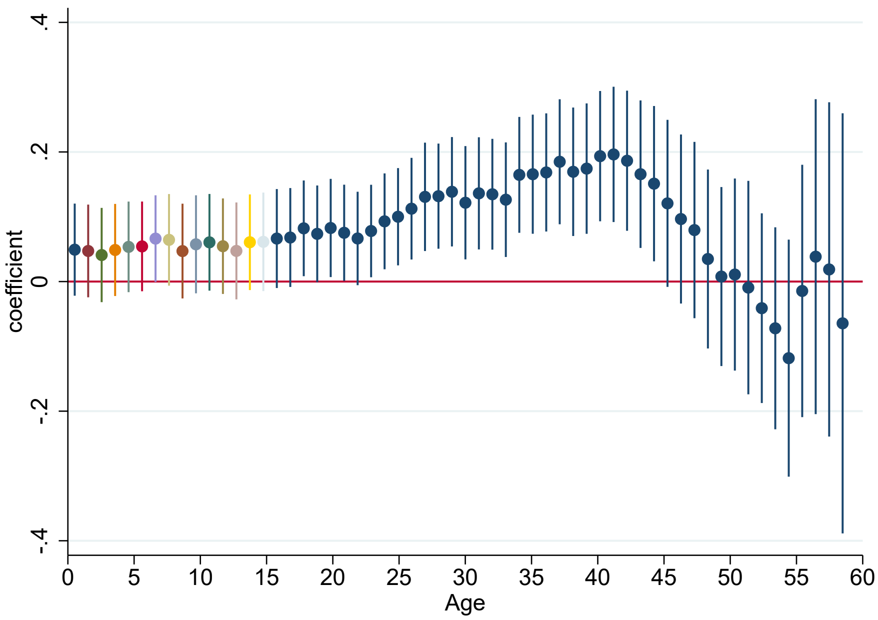

This figure plots the coefficients from regression of HAZ on shock at each age between the month when the child is born and the month when the child was surveyed (the child's age) . For example, a child whose outcome is measured when she was at age 30 months has been exposed to thirty different price shocks. Thus as we move to the right on the horizontal axis, the number of children included in the estimation decreases because, for instance, children younger than 40 months at the time of survey were not exposed to shocks above 40. The dots show point estimates and the vertical bars show 95% confidence interval.

Figure 3.5: The effect on WAZ of exposure to price shock at different ages

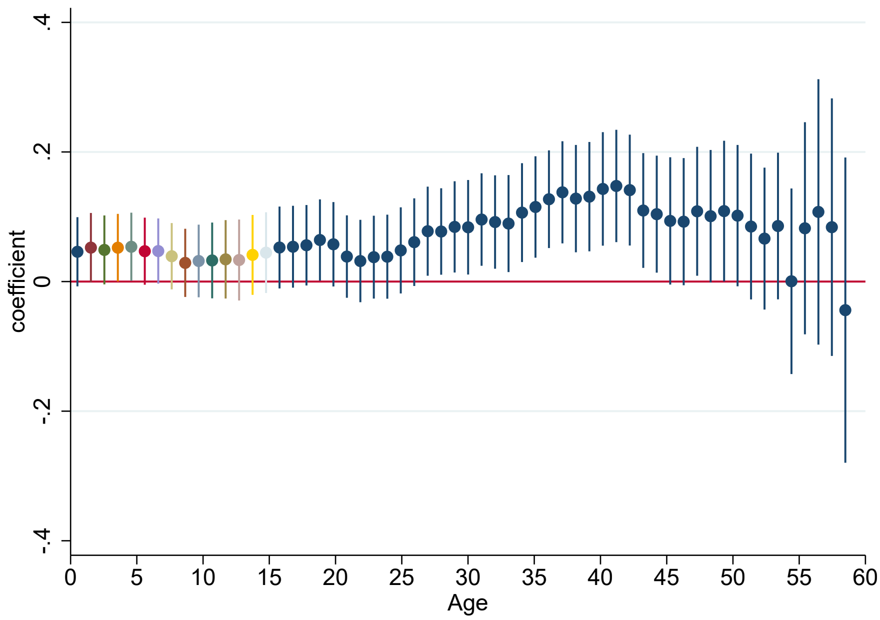

This figure plots the coefficients from regression of WAZ on shock at each age between the month when the child is born and the month when the child was surveyed (the child's age) . For example, a child whose outcome is measured when she was at age 30 months has been exposed to thirty different price shocks. Thus as we move to the right on the horizontal axis, the number of children included in the estimation decreases because, for instance, children younger than 40 months at the time of survey were not exposed to shocks above 40. The dots show point estimates and the vertical bars show 95% confidence interval.

Figure 3.6: The effect on WHZ of exposure to price shock at different ages

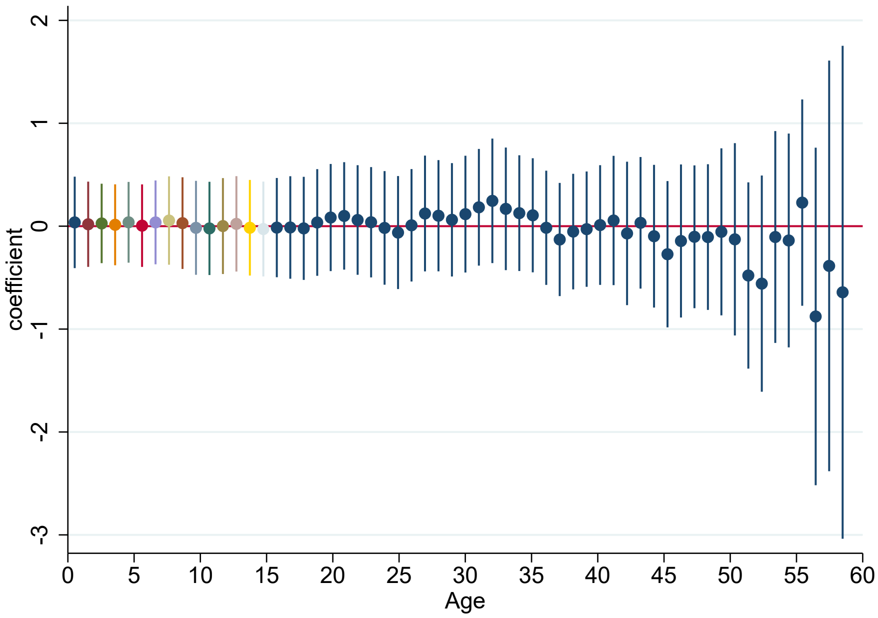

This figure plots the coefficients from regression of WHZ on shock at each age between the month when the child is born and the month when the child was surveyed (the child's age) . For example, a child whose outcome is measured when she was at age 30 months has been exposed to thirty different price shocks. Thus as we move to the right on the horizontal axis, the number of children included in the estimation decreases because, for instance, children younger than 40 months at the time of survey were not exposed to shocks above 40. The dots show point estimates and the vertical bars show 95% confidence interval.

Figure 3.7: The effect on hemoglobin level of exposure to price shock at different ages

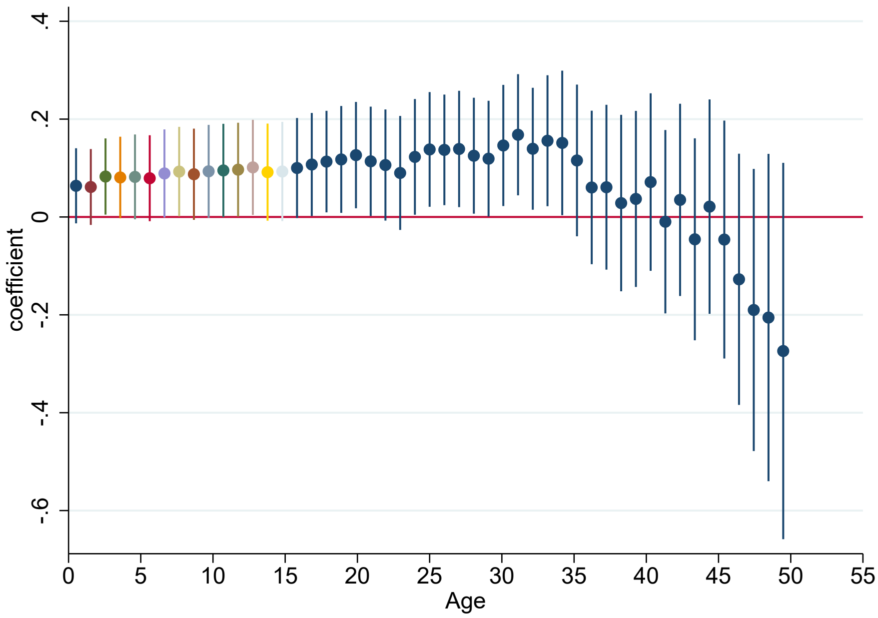

This figure plots the coefficients from regression of hemoglobin level on shocks at each age between the month when the child is born and the month when the child was surveyed (the child's age) . For example, a child whose outcome is measured when she was at age 30 months has been exposed to 25 different price shocks (from age 6 months when the child is eligible to hemoglobin measure to age 30). Thus as we move to the right on the horizontal axis, the number of children included in the estimation decreases because, for instance, children younger than 40 months at the time of survey were not exposed to shocks above 40. Only children above 6 months of age at the time of survey have their hemoglobin level measured, hence zero on the horizontal axis means age of six months. The dots show point estimates and the vertical bars show 95% confidence interval.

Tables

Table 3.1: Trends in child malnutrition and health in Ethiopia

Survey Round	Stunted	Underweight	Wasted	Child is alive	Moderately or severely anemic
			Urban Samples		
2000	0.38	0.20	0.09	0.90	.
2005	0.30	0.15	0.09	0.94	0.17
2010	0.26	0.16	0.09	0.94	0.17
2016	0.21	0.13	0.10	0.97	0.30
			Rural Samples		
2000	0.57	0.43	0.16	0.88	.
2005	0.52	0.35	0.15	0.91	0.27
2010	0.45	0.33	0.13	0.93	0.24
2016	0.36	0.26	0.13	0.93	0.38
			Rural Samples: coffee districts		
2000	0.57	0.47	0.17	0.87	.
2005	0.53	0.31	0.12	0.90	0.23
2010	0.43	0.28	0.09	0.92	0.16
2016	0.35	0.22	0.09	0.94	0.26
			Rural Samples: non-coffee districts		
2000	0.57	0.42	0.15	0.88	.
2005	0.52	0.36	0.16	0.91	0.28
2010	0.46	0.35	0.15	0.93	0.27
2016	0.37	0.27	0.14	0.93	0.41

Notes: Coffee districts are districts with GAEZ yield of above 8 quintals per hectare. A child is underweight if weight for age z-score (WAZ) is less than -2. A child is wasted if weight for height z-score (WHZ) is less than -2. A child is stunt if height for age z-score (HAZ) is less than -2.

Table 3.2: Income shocks and consumption smoothing: household level exposure

	Main		Credit Access		Aid	
	Total	Food	Total	Food	Total	Food
Land Share×LogPrice	0.997***	1.019***	1.003***	1.027***	1.004***	1.019***
	(0.203)	(0.219)	(0.203)	(0.219)	(0.209)	(0.226)
Land Share×LogPrice ×Credit			-0.024	-0.037		
			(0.027)	(0.031)		
Land Share×LogPrice ×Aid					-0.022	-0.012
					(0.042)	(0.051)
Credit			-0.016	-0.024		
			(0.021)	(0.024)		
Aid					0.028	0.045
					(0.027)	(0.030)
N	9648	9648	9648	9648	9648	9648
R^2	0.613	0.566	0.613	0.566	0.613	0.566

Notes: Robust standard errors in parenthesis. All regressions include household and year fixed effects. *Total* is log household adult-equivalent total consumption. *Food* is log household adult-equivalent food consumption. *Land Share* is the share of household land allocated to coffee. LogPrice is the average coffee price in the quarter before survey time. *Credit Access* = 1 if the household has received credit during the year from any source including: relatives, money lenders, microfinance, NGOs, banks, etc. *Aid* = 1 if the household has received assistance in-cash or in-kind from government or other sources including: productive safety net, food for work program, NGOs, relatives, etc.

Table 3.3: Income shocks and consumption smoothing: village level exposure

	Main		Credit Access		Aid	
	Total	Food	Total	Food	Total	Food
Coffee District×LogPrice	0.196*	0.220*	0.202*	0.226*	0.198*	0.223*
	(0.113)	(0.122)	(0.113)	(0.122)	(0.113)	(0.122)
Coffee District×LogPrice ×Credit			-0.025**	-0.026**		
			(0.010)	(0.012)		
Coffee District×LogPrice ×Aid					-0.008	-0.013
					(0.020)	(0.022)
Credit			0.030	0.019		
			(0.032)	(0.037)		
Aid					0.042	0.068
					(0.053)	(0.058)
N	9648	9648	9648	9648	9648	9648
R^2	0.609	0.563	0.611	0.564	0.610	0.563

Notes: Standard errors are clustered at village level (290 villages). All regressions include household and year fixed effects. *Total* is log household adult-equivalent total consumption. *Food* is log household adult-equivalent food consumption. Coffee districts are districts with GAEZ yield of above 8 quintals per hectare. LogPrice is the average coffee price in the quarter before survey time. *Credit Access* $= 1$ if the household has received credit during the year from any source including: relatives, money lenders, microfinance, NGOs, banks, etc. *Aid* $= 1$ if the household has received assistance in-cash or in-kind from government or other sources including: productive safety net, food for work program, NGOs, etc.

Table 3.4: Income shock and child malnourishment: basic result

	Continuous measures			Extreme health outcomes		
	WAZ	WHZ	HAZ	Underweight	Wasted	Stunted
High Price × Coffee District	0.275***	0.156***	0.251***	-0.068***	-0.027**	-0.048**
	(0.058)	(0.052)	(0.078)	(0.019)	(0.012)	(0.022)
High Price	-0.376***	-0.008	-0.540***	0.034	0.015	0.082***
	(0.071)	(0.064)	(0.090)	(0.022)	(0.016)	(0.022)
Succeeding Birth Interval	-0.001**	0.000	-0.001	0.000	0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)
Girl	0.074***	0.092***	0.084***	-0.026***	-0.022***	-0.028***
	(0.016)	(0.016)	(0.021)	(0.007)	(0.004)	(0.006)
Age in Months	-0.019***	-0.004*	-0.026***	0.004***	-0.001	0.005***
	(0.003)	(0.003)	(0.003)	(0.001)	(0.001)	(0.001)
N	23216	22757	21990	23216	22757	21990
R^2	0.184	0.075	0.233	0.100	0.061	0.148

Notes: Standard errors are clustered at district level. All regressions include fixed effects for district, twins, birth-order, month of birth, survey round, cohort (year of birth), Mother's education, and if source of drinking water is piped. Coffee districts are districts with GAEZ yield of above 8 quintals per hectare. High Price equals 1 if Price_{ts} for a child is higher than the median, where Price_{ts} is the average coffee price between the birth month of a child and the month when the child health is measured.

WAZ is weight for age z-score. A child is underweight if WAZ <-2.

WHZ is weight for height z-score. A child is wasted if WHZ<-2.

HAZ is height for age z-score. A child is stunt if HAZ is <-2.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.5: Income shock and child malnourishment: Anemia

	Moderately/severely anemic	
	(1)	(2)
High Price*Coffee District	-0.066**	
	(0.028)	
High Price	0.124**	
	(0.052)	
LogPrice*LogYield		-0.040***
		(0.015)
LogPrice		-0.086
		(0.116)
Succeeding Birth Interv	-0.000*	-0.000*
	(0.000)	(0.000)
Girl	-0.011	-0.011
	(0.008)	(0.008)
Age in Months	-0.016***	-0.013**
	(0.006)	(0.006)
N	14413	14413
R^2	0.182	0.182

Notes: Standard errors are clustered at district level. Columns 1 uses dummy variables for coffee district and average price. Coffee districts are districts with GAEZ yield of above 8 quintals per hectare. High Price equals 1 if Price_{ts} for a child is higher than the median, where Price_{ts} is the average coffee price between the birth month of a child and the month when the child is surveyed. $Yield$ is FAO/GAEZ's measure of a district's coffee suitability measured as potential coffee yield in quintals. All regressions include fixed effects for district, twins, birth-order, month of birth, survey round, cohort (year of birth), mother's education, and if source of drinking water is piped.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.6: Income shock and child mortality

	Dependent Variable: Child is alive (=1)	
	(1)	(2)
LogPriceUtero*LogYield	0.010**	
	(0.004)	
LogPriceUtero	-0.045***	
	(0.017)	
LogPriceBirthMonth*LogYield		0.008*
		(0.004)
LogPriceBirthMonth		-0.018
		(0.018)
Preceeding Birth Interval	0.001***	0.001***
	(0.000)	(0.000)
Girl	0.014***	0.014***
	(0.004)	(0.004)
N	22907	22907
R^2	0.041	0.041

Notes: Standard errors are clustered at district level. All regressions include fixed effects for district, twins, birth-order, month of birth, survey round, cohort (year of birth), Mother's education, and if source of drinking water is piped. Price Utero_t is the average coffee price when the child was in utero. $Yield$ is FAO/GAEZ's measure of a district's coffee suitability measured as potential coffee yield in quintals.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.7: Income shock and child malnourishment: in-utero exposure

	Continuous measures			Extreme health outcomes		
	WAZ	WHZ	HAZ	Underweight	Wasted	Stunted
LogPriceUtero	0.032	-0.023	0.066*	-0.001	-0.001	-0.014
×LogYield	(0.027)	(0.025)	(0.035)	(0.010)	(0.006)	(0.010)
LogPriceUtero	-0.079	-0.118	0.036	0.041	0.029	-0.029
	(0.088)	(0.089)	(0.118)	(0.032)	(0.024)	(0.039)
Succ. Birth Interv	-0.001*	0.000	-0.001	0.000	0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)
Girl	0.076***	0.092***	0.086***	-0.026***	-0.022***	-0.028***
	(0.017)	(0.016)	(0.021)	(0.007)	(0.004)	(0.006)
Age in Months	-0.026***	-0.004***	-0.035***	0.005***	-0.001***	0.007***
	(0.001)	(0.001)	(0.002)	(0.000)	(0.000)	(0.000)
N	23216	22757	21990	23216	22757	21990
R^2	0.181	0.075	0.231	0.099	0.061	0.148

Notes: Standard errors are clustered at district level. All regressions include fixed effects for district, twins, birth-order, month of birth, survey round, cohort (year of birth), Mother's education, and if source of drinking water is piped. Price Utero$_t$ is the average price when the child was in utero. $Yield$ is FAO/GAEZ's measure of a district's coffee suitability measured as potential coffee yield in quintals.
WAZ is weight for age z-score. A child is underweight if WAZ <-2.
WHZ is weight for height z-score. A child is wasted if WHZ<-2.
HAZ is height for age z-score. A child is stunt if HAZ is <-2.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.8: Income shock and child malnourishment: Gender difference?

	Continuous measures			Extreme health outcomes		
	WAZ	WHZ	HAZ	Underweight	Wasted	Stunted
LogPrice×LogYield	0.149***	0.096**	0.150***	-0.032**	-0.023**	-0.035**
	(0.040)	(0.039)	(0.054)	(0.013)	(0.009)	(0.014)
LogPrice×LogYield ×Girls	-0.002	-0.004	0.001	-0.000	0.001	-0.000
	(0.003)	(0.003)	(0.004)	(0.001)	(0.001)	(0.001)
LogPrice	0.978***	0.695***	1.486***	-0.257***	-0.075	-0.276***
	(0.206)	(0.195)	(0.267)	(0.070)	(0.052)	(0.070)
Succ. Birth Interv	-0.001**	0.000	-0.001	0.000	0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)
Girl	0.083***	0.111***	0.083***	-0.026***	-0.025***	-0.026***
	(0.021)	(0.021)	(0.028)	(0.010)	(0.006)	(0.008)
Age in Months	-0.019***	0.000	-0.024***	0.003***	-0.002***	0.005***
	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)
N	23216	22757	21990	23216	22757	21990
R^2	0.184	0.076	0.233	0.100	0.061	0.148

Notes: Standard errors are clustered at district level. All regressions include fixed effects for district, twins, birth-order, month of birth, survey round, cohort (year of birth), Mother's education, and if source of drinking water is piped. Price_{ts} is the average coffee price between the birth month of a child and the month when the child is surveyed. $Yield$ is FAO/GAEZ's measure of a district's coffee suitability measured as potential coffee yield in quintals.
WAZ is weight for age z-score. A child is underweight if WAZ <-2.
WHZ is weight for height z-score. A child is wasted if WHZ<-2.
HAZ is height for age z-score. A child is stunt if HAZ is <-2.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.9: Income shock and child malnourishment: Heterogeneity across wealth groups?

	Continuous measures			Extreme health outcomes		
	WAZ	WHZ	HAZ	Underweight	Wasted	Stunted
LogPrice*LogYield	0.132***	0.004	0.187***	-0.027**	-0.011	-0.047***
	(0.045)	(0.041)	(0.060)	(0.014)	(0.009)	(0.016)
LogPrice*LogYield *Poorer	-0.002	-0.005	0.001	0.000	0.001	-0.002
	(0.006)	(0.006)	(0.008)	(0.002)	(0.001)	(0.002)
LogPrice*LogYield *Middle	-0.009	-0.004	-0.008	0.004*	0.001	0.002
	(0.007)	(0.006)	(0.009)	(0.002)	(0.002)	(0.002)
LogPrice*LogYield *Richer	-0.004	-0.005	-0.000	0.004*	0.000	0.001
	(0.007)	(0.006)	(0.009)	(0.002)	(0.002)	(0.002)
LogPrice*LogYield *Richest	0.002	-0.006	0.011	0.001	0.003	-0.004
	(0.011)	(0.009)	(0.013)	(0.003)	(0.002)	(0.003)
Poorer	0.063*	0.041	0.026	-0.026*	-0.011	-0.003
	(0.036)	(0.041)	(0.059)	(0.013)	(0.010)	(0.015)
Middle	0.212***	0.082*	0.220***	-0.077***	-0.018	-0.060***
	(0.047)	(0.046)	(0.056)	(0.016)	(0.013)	(0.015)
Richer	0.291***	0.210***	0.221***	-0.127***	-0.037***	-0.068***
	(0.051)	(0.050)	(0.061)	(0.018)	(0.013)	(0.018)
Richest	0.391***	0.276***	0.274***	-0.115***	-0.056***	-0.098***
	(0.086)	(0.065)	(0.103)	(0.027)	(0.015)	(0.023)
LogPrice	0.434*	0.252	1.252***	-0.135	-0.054	-0.219**
	(0.255)	(0.261)	(0.352)	(0.087)	(0.069)	(0.089)
N	16572	16217	15721	16572	16217	15721
R^2	0.192	0.090	0.233	0.108	0.068	0.149

Notes: Standard errors are clustered at district level. All regressions include fixed effects for district, twins, birth-order, month of birth, survey round, cohort (year of birth), Mother's education, if source of drinking water is piped, and the control variables in table 3.4. All regressions include fixed effects for five wealth groups (ranging from "poorest" to "richest"). Price is the average coffee price between the birth month of a child and the month when the child is surveyed. *Yield* is FAO/GAEZ's measure of a district's coffee suitability measured as potential coffee yield in quintals. The number of observations is smaller compared to table A.2 because wealth measure is missing for significant number of households.
WAZ is weight for age z-score. A child is underweight if WAZ <-2.
WHZ is weight for height z-score. A child is wasted if WHZ<-2.
HAZ is height for age z-score. A child is stunt if HAZ is <-2.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Appendices

A Appendix Figures

Figure A.1: Percent of district land allocated to coffee

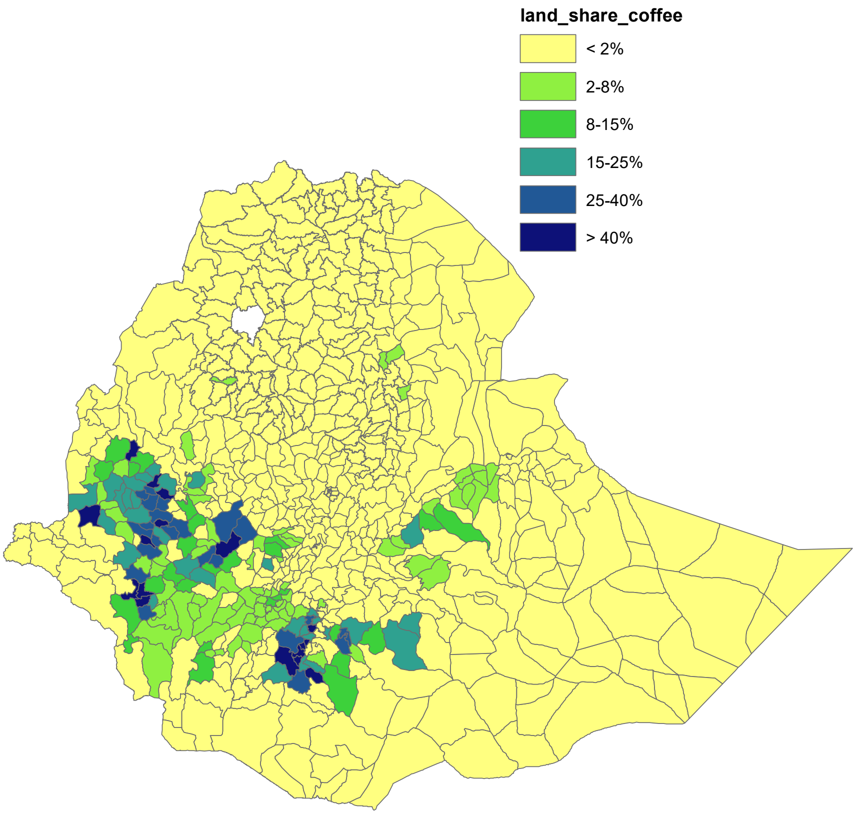

Notes: This figure shows the share of agricultural land allocated to coffee. Each district's statistics is average across villages in the district. *Data Source*: Central Statistical Agency (CSA)- Agricultural Sample Survey (AgSS) 2004-2016.

B Appendix Tables

Table A.1: Income shock and child malnourishment: alternative specification

	Continuous measures			Extreme health outcomes		
	WAZ	WHZ	HAZ	Underweight	Wasted	Stunted
Coffee District *LogPrice	0.331***	0.233**	0.312**	-0.069**	-0.054**	-0.084**
	(0.095)	(0.093)	(0.127)	(0.031)	(0.021)	(0.034)
LogPrice	0.996***	0.884***	1.620***	-0.290***	-0.130**	-0.310***
	(0.222)	(0.221)	(0.294)	(0.075)	(0.055)	(0.074)
Succ. Birth Interv	-0.001**	-0.000	-0.001	0.000	0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)
Girl	0.076***	0.094***	0.085***	-0.027***	-0.022***	-0.028***
	(0.017)	(0.016)	(0.021)	(0.007)	(0.004)	(0.007)
Age in Months	-0.020***	-0.003	-0.026***	0.004***	-0.001	0.005***
	(0.003)	(0.003)	(0.003)	(0.001)	(0.001)	(0.001)
N	23216	22757	21990	23216	22757	21990
R^2	0.184	0.076	0.233	0.100	0.061	0.149

Notes: Standard errors are clustered at district level. All regressions include fixed effects for district, twins, birth-order, month of birth, survey round, cohort (year of birth), Mother's education, and if source of drinking water is piped. Price_{ts} is the average coffee price between the birth month of a child and the month when the child is surveyed. $Yield$ is FAO/GAEZ's measure of a district's coffee suitability measured as potential coffee yield in quintals.
WAZ is weight for age z-score. A child is underweight if WAZ <-2.
WHZ is weight for height z-score. A child is wasted if WHZ<-2.
HAZ is height for age z-score. A child is stunt if HAZ is <-2.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.2: Income shock and child malnourishment: generalized DID estimation

	Continuous measures			Extreme health outcomes		
	WAZ	WHZ	HAZ	Underweight	Wasted	Stunted
LogPrice*LogYield	0.147***	0.090**	0.152***	-0.032**	-0.022**	-0.036**
	(0.041)	(0.039)	(0.055)	(0.013)	(0.009)	(0.014)
LogPrice	0.991***	0.887***	1.610***	-0.288***	-0.131**	-0.310***
	(0.222)	(0.222)	(0.295)	(0.075)	(0.055)	(0.074)
Succe. Birth Interv	-0.001**	-0.000	-0.001	0.000	0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)
Girl	0.076***	0.094***	0.086***	-0.027***	-0.023***	-0.028***
	(0.017)	(0.016)	(0.021)	(0.007)	(0.004)	(0.007)
Age in Months	-0.020***	-0.003	-0.026***	0.004***	-0.001	0.005***
	(0.003)	(0.003)	(0.003)	(0.001)	(0.001)	(0.001)
N	23216	22757	21990	23216	22757	21990
R^2	0.184	0.076	0.233	0.100	0.061	0.149

Notes: Standard errors are clustered at district level. All regressions include fixed effects for district, twins, birth-order, month of birth, survey round, cohort (year of birth), Mother's education, and if source of drinking water is piped. Price is the average coffee price between the birth month of a child and the month when the child is surveyed. $Yield$ is FAO/GAEZ's measure of a district's coffee suitability measured as potential coffee yield in quintals.
WAZ is weight for age z-score. A child is underweight if WAZ <-2.
WHZ is weight for height z-score. A child is wasted if WHZ<-2.
HAZ is height for age z-score. A child is stunt if HAZ is <-2.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.3: Using share of district land allocated to coffee as a measure of district coffee intensity: IV estimation

	Continuous measures			Extreme health outcomes		
	WAZ	WHZ	HAZ	Underweight	Wasted	Stunted
Coffee District	0.783***	0.410**	0.769***	-0.162**	-0.050	-0.117*
*High Price	(0.220)	(0.181)	(0.262)	(0.063)	(0.043)	(0.066)
High Price	-0.415***	-0.030	-0.589***	0.043**	0.014	0.088***
	(0.073)	(0.060)	(0.088)	(0.022)	(0.017)	(0.023)
Succ. Birth Interv	-0.001*	0.000	-0.001	0.000	0.000	0.000
	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)
Girl	0.078***	0.096***	0.084***	-0.028***	-0.022***	-0.028***
	(0.017)	(0.016)	(0.021)	(0.007)	(0.004)	(0.007)
Age in Months	-0.020***	-0.006**	-0.027***	0.004***	-0.000	0.005***
	(0.003)	(0.003)	(0.003)	(0.001)	(0.001)	(0.001)
N	23590	23117	22345	23590	23117	22345
R^2	0.185	0.079	0.235	0.105	0.065	0.151

Notes: Standard errors are clustered at district level. All regressions include fixed effects for district, twins, birth-order, month of birth, survey round, cohort (year of birth), Mother's education, and if source of drinking water is piped. A child is exposed to $HighPrice$ if the average coffee price between the birth month of a child and the month when the child health is measured is above median average prices similarly calculated for the all children. A district is defined as a *Coffee District* if it belongs to the top quartile in the share of agricultural land allocated to coffee. FAO/GAEZ potential coffee yield is used as an instrument for *Coffee District*.
WAZ is weight for age z-score. A child is underweight if WAZ <-2.
WHZ is weight for height z-score. A child is wasted if WHZ<-2.
HAZ is height for age z-score. A child is stunt if HAZ is <-2.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Bibliography

Adhvaryu, A., Fenske, J., and Nyshadham, A. (2019). Early Life Circumstance and Adult Mental Health. *Journal of Political Economy*, Forthcoming.

Carrillo, B. (2019). Present Bias and Underinvestment in Education? Long-run Effects of Childhood Exposure to Booms in Colombia. *Journal of Labor Economics*, page Forthcoming.

Cogneau, D. and Jedwab, R. (2012). Commodity Price Shocks and Child Outcomes: The 1990 Cocoa Crisis in Côte d'Ivoire. *Economic Development and Cultural Change*, 60:507–534.

Dehejia, R. and Lleras-Muney, A. (2004). Booms, Busts, and Babies' Health. *Quarterly Journal of Economics*, 119:1091–1130.

Fafchamps, M. and Lund, S. (2003). Risk-sharing networks in rural Philippines. *Journal of Development Economics*, 71(2):261 – 287.

Ferreira, F. H. G. and Schady, N. (2008). Aggregate Economic Shocks, Child Schooling and Child Health. *World Bank Development Group, Policy Research Working Paper*, 4701:.

Gertler, P. and Gruber, J. (2002). Insuring Consumption Against Illness. *American Economic Review*, 92:51–70.

Glewwe, P. and King, E. M. (2001). The Impact of Early Childhood Nutritional Status on Cognitive Development: Does the Timing of Malnutrition Matter? *The World Bank Economic Review*, 15:81–113.

Haaker, J. R. (2018). Price Shocks and Child Mortality: Evidence from Anti-Drug Policies in Peru. *Duke University, JMP*, page .

Hoddinott, J. and Kinsey, B. (2001). Child growth in the time of drought. *Oxford Bulletin of Economics and Statistics*, 63:0305–9049.

Jacoby, H. G. and Skoufias, E. (1997). Risk, Financial Markets, and Human Capital in a Developing Country. *The Review of Economic Studies*, 64(3):311–335.

Jensen, R. (2001). Agricultural Volatility and Investment in Children. *American Economic Review: Papers and Proceedings*, 90:399–404.

Kanamori, M. J. and Pullum, T. (2013). Indicators of Child Deprivation in Sub-Saharan Africa: Levels and Trends from the Demographic and Health Surveys. *DHS Comparative Reports*, 32:.

Kruger, D. (2007). Coffee Production Effects on Child Labor and Schooling in Rural Brazil. *Journal of Development Economics*, 82:448–463.

Mendiratta, V. (2015). Impact of Rainfall Shocks on Child Health: Evidence from India. *Paris School of Economics, Working Paper*, 33:.

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

Morduch, J. (1995). Income Smoothing and Consumption Smoothing. *Journal of Economic Perispective*, 9:103–114.

Neumayer, E. (2004). Recessions lower (some) mortality rates: evidence from Germany. *Social Science and Medicine*, 58:1037–1047.

Page, M., Schaller, J., and Simon, D. (2016). The effects of aggregate and gender-specific labor demand shocks on child health. *NBER Working Paper*, 22394:.

Paxson, C. and Schady, N. (2005). Child Health and Economic Crisis in Peru. *The World Bank Economic Review*, 19:203–223.

Paxson, C. H. (1992). Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand. *American Economic Review*, 82(1):15–33.

Shah, M. and Steinberg, B. M. (2017). Drought of Opportunities: Contemporaneous and Long-Term Impacts of Rainfall Shocks on Human Capital. *Journal of Political Economy*, 125:527–561.

Tiwari, S., Jacoby, H. G., and Skoufias, E. (2013). Monsoon Babies: Rainfall Shocks and Child Nutrition in Nepal. *Economic Development and Cultural Change*, 65:167–188.

Townsend, R. M. (1995). Risk and Insurance in Village India. *Econometrica*, 62:539–591.

Waber, D. P., Vuori-Christiansen, L., Ortiz, N., Clement, J. R., Christiansen, N. E., Mora, J. O., Reed, R. B., and Herrera, M. G. (1981). Nutritional supplementation, maternal education, and cognitive development of infants at risk of malnutrition. *The American Journal of Clinical Nutrition*, 34:807–813.

World Food Program, W. (2013). A Hunger Glossary. .