Essays in Development Economics and Political Economy

Essays in Development Economics and Political Economy

A dissertation presented

by

Lorenzo Casaburi

to

The Department of Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

Harvard University

Cambridge, Massachusetts

April 2013

Dissertation Advisor:
Professor Michael Kremer

Author:
Lorenzo Casaburi

Essays in Development Economics and Political Economy

Abstract

Chapter 1 studies the electoral response to the Ghost Buildings program, a nationwide anti tax evasion policy in Italy which used innovative monitoring technologies to target buildings hidden from tax authorities. The difference-in-differences identification strategy exploits both variation across towns in the ex ante program scope to increase enforcement as well as administrative data on actual building registrations. Local incumbents experience an increase in their reelection likelihood as a consequence of the policy. In addition, these political returns are higher in areas with higher speed of public good provision and with lower tax evasion tolerance, implying complementarity among enforcement policies, government efficiency, and the underlying tax culture. Chapter 2 uses a road-level regression discontinuity design in Sierra Leone to study the impact of improvements in rural road infrastructure on agricultural markets. We show that the improved roads reduced the market prices of local crops. These price effects are stronger in markets that are further from major urban centers and in less productive areas. We also find that these price effects are reversed in areas with better cell phone penetration. We show that our empirical findings are consistent with a search cost framework à la Mortensen, but inconsistent with other models, such as Bertrand competition, bilateral bargaining, and Cournot oligopsony. Chapter 3 present results from a randomized controlled experiment designed to study the multiple margins through which value is passed from traders to agricultural producers in the presence of interlinked transactions. Consistent with other studies, we find limited price pass-through in response to an increase in the trader resale price. However, there is a large response in credit provision. We develop a model of interlinked transactions that highlights the substitutability of price and credit pass-through across markets, and verify its predictions empirically. Calibration suggests that to ignore margins of pass-through other than price has substantial implications for welfare analysis.

Contents

List of Tables

List of Figures

Acknowledgments

Attending graduate school offers the unique opportunity to interact with brilliant faculty and colleagues. Exposure to these people greatly enriched my experience as a doctoral student.

I am particularly grateful to my advisors, Michael Kremer, Pol Antrás, and Tavneet Suri. Since my early years in Kenya, Michael has provided an intellectual model to follow. His dedication and creativity never cease to surprise me. Pol's efforts to expand the boundaries of the trade field provided inspiration and hope for my research. Tavneet has extended her work day to thirty hours to accommodate my meeting requests. Over these years, she has been an almost daily presence, as a mentor and friend.

I also benefited from conversations with Alberto Alesina, Raj Chetty, Rachel Glennerster, Elhanan Helpman, Asim Khwaja, Marc Melitz, Ted Miguel, Sendhil Mullainathan, Nathan Nunn, Dina Pomeranz, and David Yanagizawa-Drott.

Several colleagues and friends have played an important role through conversations, feedback on research ideas, and fun moments together: Sam Asher, Shai Bernstein, Joshua Gottlieb, Jacob Leshno, Massimo Maoret, Eduardo Morales, Tristan Reed, Dorothee Rouzet, Frank Schilbach, Andrea Stella, and Ugo Troiano. I am also particularly indebted to my coauthors — Rachel, Tavneet, Tristan, and Ugo — for their ideas and effort in our joint work.

My parents, Salvatore and Gloria, have shaped the person I am now. Throughout these years, I have always felt them right next to me, in spite of the geographical distance.

Finally, I doubt I would be at this stage without Selena's love, support, and reminders that academia is not the only important thing in life. Thank you for all of this.

Chapter 1

Ghost-House Busters: The Electoral Response to a Large Anti Tax Evasion Program [1]

1.1 Introduction

Government ability to enforce tax collection efficiently is one of the fundamental components of state capacity and, in turn, an important driver of historical economic development. Tax evasion generates both large losses in government revenues and large distortions.[2] The literature (e.g., Slemrod (2007); Besley and Persson (2012)) describes three main determinants of tax compliance: enforcement technology, political incentives, and cultural norms. This paper illustrates the interaction among these three factors. We estimate the electoral returns — the change in reelection likelihood — that local policymakers obtain from a nationwide anti tax evasion policy in Italy, based on an innovation in tax-payers monitoring technology. In

[1]Co-authored with Ugo Troiano

[2]Slemrod (2007) estimates overall noncompliance in the United States at 14 percent. Estimates from other developed countries deliver similar figures (for Italy, Marino and Zizza (2008)). In developing countries, where the share of the informal economy is typically larger, the figures are much higher (Gordon and Li (2009); Schneider, Buehn, and Enste (2010)).

addition, we study how these electoral returns depend on underlying social preferences for tax compliance and on the local government efficiency in public good provision. The paper provides, to the best of our knowledge, the first empirical evidence on voters' responses to anti tax evasion policies.

Measures to reduce tax evasion generate a conflict between voters. They hurt tax evaders, typically a minority of voters, while the majority of the population is likely to benefit from additional government expenditures, lower tax rates, or even directly from the punishment of former shirkers.[3] However, the magnitude of the individual costs tax evaders incur is potentially higher than the individual benefits non evaders derive. Anti tax evasion policies are thus canonical examples of policies that have an asymmetry in the concentration of costs and benefits (Tullock (1959); Olson (1965)). Depending on which of the two types of voters is more likely to change its voting behavior in response to changes in enforcement, fighting tax evasion might either benefit or harm politicians who seek reelection. The sign of this impact is ex ante ambiguous and, therefore, an empirical question.

In 2007, the Italian government instituted a nationwide anti tax evasion policy, the *Ghost Buildings* program. The program identified ghost buildings — properties not included in the land registry and thus hidden from tax authorities[4] — by overlaying aerial photographs and digital land registry maps.[5] The intervention detected more than two million land registry parcels with ghost buildings.[6] Following the completion of the mapping exercise, the government commenced a large registration program that targeted the identified ghost buildings. While the central government began the program and coordinated registration activities, municipality administrations circulated information about the program, collab-

[3]For experimental evidence on this channel, see Carpenter et al. (2009); Casari and Luini (2009); Ouss and Peysakhovich (2012).

[4]The value of registered buildings enters the tax base for personal income tax and property taxes, among other taxes.

[5]Other countries, such as Greece and Rwanda, have recently implemented policies using similar technologies.

[6]The unit of the Italian land registry maps is the parcel (*parcella*), which is defined as a portion of land belonging to a given physical or legal person. In the case where the land is shared across several owners, the parcel is split into several sub-parcels (*subalterni*).

orated with follow-up inspections, and enforced payment of overdue local taxes. Media reports highlight both the importance of local administrations in the registration process[7] and the heterogeneity in their actions in response to the program.[8]

The policy induced a large shift in tax enforcement, the intensity of which varied significantly across towns. In towns with a higher prevalence of detected ghost buildings, the program had larger scope to affect the level of building registration. We use a measure of *Ghost Building Intensity* — the ratio of the number of land registry parcels with ghost buildings identified by the program to the total number of land registry parcels in the town — to proxy for the scope of the program. Using a difference-in-differences approach, we test the impact of the anti evasion policy on local incumbent reelection by exploiting variations across municipalities in this intensity.[9] This strategy, under plausible assumptions, isolates the causal effect on electoral outcomes of the policy scope to increase enforcement from other mayor or voter characteristics that might have affected the actual levels of ghost building registration in the town.[10]

In local elections occurring after the beginning of the program, an increase of one standard deviation in the ghost building intensity raises the likelihood of reelection of the local incumbent relative to pre-program elections by approximately 2.5 percentage points, about 5.5% of the average reelection rate. Higher town-level ghost building intensity also lowers several measures of competitiveness of local elections. In particular, it reduces the number of candidates running for election, increases the margin of victory for the winner, and reduces the likelihood of a runoff. Guiding our empirical models is a retrospective voting framework of political agency. Such a theoretical framework helps us predict how a change in tax enforcement can impact voter choices and which factors affect this response.

[7]For example, Dell'Oste and Trovati (2011).

[8]Among many others, Bernardini (2011) and Barca (2008) discuss the particular way in which the city of Montecatini and some cities in the Reggio Emilia province implemented the program, respectively.

[9]In a recent contribution, Mian and Sufi (2012) adopted a similar empirical approach to study the effects of the fiscal stimulus in the US.

[10]We verify that the assumptions required by the identification strategy hold (i.e., no contemporary differential changes and no differential pre-trends in the outcome variable by treatment intensity).

Additional analysis of the actual building registrations induced by the program complements the reduced form analysis described above. For a given town-level program scope, a higher registration rate of ghost buildings under the incumbent local administration (i.e., the share of ghost buildings that gets registered prior to the local election date) has a positive effect on the likelihood of reelection. The result is robust to the inclusion of mayors' characteristics as controls and to an instrumental variable approach, based on the time elapsed between the program start date and the town election date.

We provide evidence for two channels that could drive the observed electoral response. First, towns where the government is more efficient in delivering public goods show a larger electoral response to the program. We also verify that towns with higher ghost building intensity experienced a differential increase in local government expenditures following the program inception. Second, using survey data on the self-reported tolerance for tax evasion among voters, we show that the program's positive impact on incumbent reelection is significantly higher in areas with lower tolerance for tax evasion. Finally, the empirical findings are inconsistent with two potential alternative interpretations on the impact of the program on voter support for incumbents. In the first, the program changes voter behavior by providing information on the existing stock of ghost buildings. In the second, it gives an electoral rent to the incumbent by giving her the option to *not* register identified ghost buildings.

Our approach can potentially be applied in different settings to study the political feasibility of upgrading tax administrations around the world using new electronic data, cross-checking technologies, and other monitoring devices (Bird and Zolt (2008)). Additionally, our analysis points at complementarity between technological innovations in tax enforcement and political incentives. When exposed to a reduction in monitoring costs, politicians exploit the new technologies and experience political gains. These findings have a direct bearing on the political feasibility of upgrading tax administrations around the world using new electronic data, cross-checking technologies, and other monitoring devices (Bird and Zolt (2008)). In addition, our study provides evidence that the underlying tax culture shapes the political incentives for tax enforcement and the political returns to these innovations (Torgler (2007); Rothstein (2000)). We discuss several policy implications arising from these findings. Finally, access to town-level nationwide administrative data from the program allows us to provide evidence on two additional fronts. First, we study the correlates of tax evasion at the town level. We find that geographical features, such as town size, are important determinants of tax evasion, consistent with Saiz (2010), and that social capital is negatively correlated with tax evasion (Putnam (2001)). Second, we document that mayor characteristics, such as education, gender, and age, do affect the extent to which the Ghost Buildings program increased tax enforcement (consistent with Alesina (1988); Besley and Coate (1997); Besley, Montalvo and Reynal-Querol (2012)).

This paper relates to several strands of literature. First, a recent set of studies uses microdata to shed light on enforcement technologies such as third-party reporting (Slemrod, Blumenthal, and Christian (2001); Saez (2010); Kleven et al (2011); Chetty, Friedman, and Saez (2012)), paper trails (Pomeranz (2012); Kumler, Verhoogen, and Frías (2011)), cross-checking (Carrillo, Pomeranz, and Singhal (2012)), and targeted auditing strategies (Almunia and Lopez-Rodriguez (2012); Aparicio (2012)).[11] By studying how technology-driven enforcement policies affect policymakers, we bridge this work with the one estimating the political returns to fiscal policies (Brender and Drazen (2008); Alesina, Carloni, and Lecce (2011)). In addition, by delving into the relation between incentives of political agents and tax evasion, our paper is related to Artavanis, Morse, and Tsoutsoura (2012), who find that tax evasion is higher in industries supported by parliamentarians. Finally, our results provide support to the existing literature that highlights the role of culture and social norms as determinants of tax evasion, either via cross-country analysis (Torgler (2003); Slemrod (2003)) or lab experiments (Spicer and Becker (1980); Alm, Jackson, and McKee (1992)).

The remainder of the paper is organized as it follows. Section 1.2 describes the Ghost Buildings program. Section 1.3 presents a simple framework that guides our empirical analysis. Section 1.4 describes the data and presents descriptive evidence. Section 1.5

[11] For a review of the literature, see Andreoni, Erard, and Feinstein (1998), Slemrod and Yitzhakil (2002).

lays out our empirical strategy to estimate the electoral response to the policy. Section 1.6 presents the results. Section 1.7 concludes.

1.2 The Ghost Buildings Program

The value of the buildings registered in the land registry enters the tax base for several national and local taxes, including "ICI", the local property tax, "IRPEF", the personal income tax,[12] and the local waste management tax. Italian legislation[13] requires that owners register new buildings at the local office of the *Agenzia del Territorio*, the agency managing the land registry, within thirty days after their completion.[14]

In 2006, the national government approved new anti tax evasion legislation, the Ghost Buildings program,[15] aimed at detecting buildings not registered in the land registry maps.[16] The *Agenzia del Territorio*, the national agency managing the land registry, coordinated the effort. The *Agenzia del Territorio* first juxtaposed the land and building registry maps to obtain the "Official Building Map". It subsequently compiled high-resolution (50cm) aerial photographs of the entire country in order to identify the ghost buildings. Figures 1.1a-1.1c summarize the steps of the identification. First, the aerial photograph for a particular location was created (Figure 1.1a). Second, the pictures were matched with the official building map for the corresponding area (Figure 1.1b). Finally, the ghost buildings were identified (Figure 1.1c)[17]. Ghost buildings include commercial, industrial, and residential

[12]"IRPEF" includes the inferred opportunity cost of living in the house, with both a local and a national component.

[13]*Legge 9 Marzo 2006 n.80 - Art. 34-quinquies*.

[14]All buildings in Italy need a building permit before their construction starts. Obtaining a building permit makes the building part of the City Plan. The process of obtaining building permits is administered independently from the registration in the land registry maps. Buildings not in the City Plan are required to be demolished.

[15]*Legge 24 novembre 2006, n. 286* subsequently modified by *Legge 30 Luglio 2010, n. 122*.

[16]The exercise did not cover one of the semi-autonomous regions, Trentino Alto-Adige, because in that region land registry maps are autonomously administered. The region contains less than two percent of the total population of Italy.

[17]According to the Law *Decreto Ministero delle Finanze 2 gennaio 1998, n.28.Art. 3* the following buildings do

stand-alone buildings, and also substantial extensions of previously registered buildings that should have been reported to the land registry.

Through this process, the *Agenzia del Territorio* identified approximately two million land registry parcels with unregistered buildings. Beginning in August 2007, the *Agenzia del Territorio* started to publish parcel-level data on unregistered properties in the *Gazzetta Ufficiale*, the official bulletin promulgating Italian laws and decrees, in order to induce registrations of the ghost buildings. Within three years, it coded detailed information on the number of ghost buildings in the universe of Italian municipalities (with the exception of Trentino Alto-Adige). The order of publication relied on the availability of digitized land registry maps at the time when the program started. The *Agenzia del Territorio* had 60% of the land registry maps of the Italian territory in digitized form before the Ghost Buildings program was approved. After 2006, the *Agenzia del Territorio* began digitizing the remaining land registry maps, proceeding by province (i.e., they simultaneously coded municipalities in the same province). It completed the identification exercise by the end of 2010.[18]

According to the initial legislation, owners could register the detected ghost building with the land registry by April 30, 2011.[19] Widespread media campaigns and local administration efforts helped the program achieve high registration rates. In particular, local administrators a) disseminated information about targeted parcels; b) collaborated on follow-up building inspections; c) proceeded with the collection of overdue local taxes up to five years before the program began; and d) verified the conformity of ghost buildings to the city plan and local zoning restrictions. Local administrations received a large share of the additional tax revenues generated by the program. Owners of ghost buildings that registered prior to the April 2011 deadline were required to pay overdue taxes dating back to 2007, or to

---

not increase the tax base of their owners and thus are not subject to registration requirements: (i) buildings that are not completed (ii) buildings particularly degraded (iii) solar collectors (iv) greenhouses (v) henhouses or others reserved for animals.

[18]Publication on the *Gazzetta Ufficiale* occurred in the following waves: August 2007, October 2007, December 2007, December 2008, December 2009, December 2010.

[19]This was the result of two previous deadlines of ninety days and seven months since the publication on the *Gazzetta Ufficiale*.

Figure 1.1A: Aerial Picture

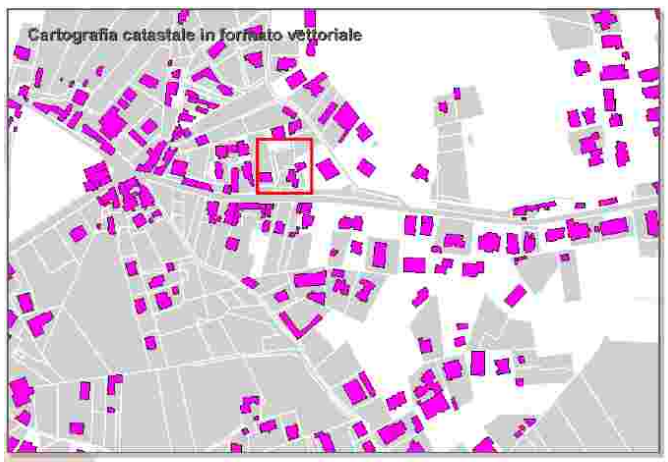

Figure 1.1B: Digital Land Registry Map

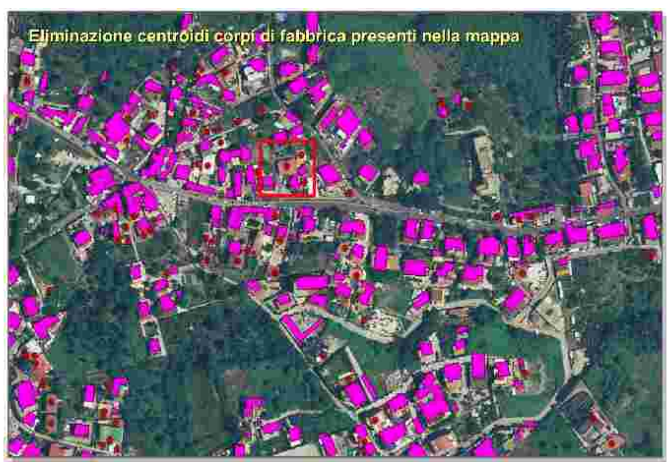

Figure 1.1C: Overlay

Source: Agenzia del Territorio

Figure 1.1: *The Ghost Building Identification Process*

the construction date for post-2007 cases, and to pay penalties for delayed payments.[20] After April 2011, the *Agenzia del Territorio*, with the support of local administrations and local contractors, proceeded with follow-up inspections[21] to impute the tax base for the remaining unregistered buildings.[22] Additional penalties and a fee for the extra inspection were assessed on owners of buildings for which the *Agenzia del Territorio* imputed the tax base after April 2011.

The *Agenzia del Territorio* published the detailed economic impact of the program for the year 2011. The program led to a substantial wave of registrations. According to the administrative data, roughly 40 percent of the ghost buildings were registered as of 30th of December 2011. According to the figures provided, the program increased total tax revenues by 472 million euros in that year.[23] We estimate that approximately 65 percent of those revenues come from local taxes. We run a "back-of-the-envelope" calculation using figures on the number of land parcels with ghost buildings, the registration rates, and the total additional tax revenues from the program. A one standard deviation increase in ghost buildings targeted by the program increased the tax revenues by around 3.5% of the median value. Using the same information, we find that, on average, the owner of a registered ghost building now faces an additional yearly tax burden of approximately 528 Euros.

1.3 Tax Enforcement and Retrospective Voting: A Conceptual Framework

This section outlines the impact of a change in tax enforcement on voters' electoral choices. We provide a simple framework based on modeling of retrospective voting (Barro

[20] Penalties were determined by *Legge 29 Dicembre 1990, n. 408*, subsequently modified by *Decreto Legislativo. 18 Dicembre 1997.*

[21] To increase incentives for the local administrations further, an additional bonus was introduced in 2011 for each registered ghost building.

[22] *Decreto Legge 79/2010, art. 10, 11.*

[23] This figure does not include payment for overdue taxes from previous years.

(1973); Ferejohn (1986)) and of tax evasion (Allingham and Sandmo (1972)). The main intuition of retrospective voting models is that citizens examine whether their welfare has increased under a politician's office, and vote accordingly.[24] While the specific application to tax evasion is novel, the discussion in this section heavily relies on the intuition from existing models.

The economy is populated by a unit mass of voters and by politicians. Voters are heterogeneous in their ability to evade. This ability could be a function of occupation type (employed vs. self-dependent) or evasion costs (economic and psychological). We consider a simple case with two fixed types of voters: evaders and non evaders. Evaders pay taxes only if enforcement occurs, while non-evaders always pay taxes.[25] The population share of evaders is λ. Enforcement of tax collection for each evader occurs with probability p. Enforcement draws are independent across evaders, and thus p is the share of evaders for which enforcement occurs. This is assumed to be a function of the politician type (a) and of an idiosyncratic component (v), whose distributions are $G(a)$ and $G(v)$, respectively. Voters do not observe the two components and are uncertain over the politician type (Banks and Sundaram (1998)). They use previous realizations to form expectations $\hat{a}$ and $\hat{p}$ (in the spirit of Holmstrom, 1982).

We assume an exogenous income level (normalized to 1) and tax rate (τ), constant across the population. Voters derive utility from disposable income and from the overall level of enforcement, for instance through increased public good provision and deficit reduction. This implies that enforcement has two effects on evaders' utility, which go in opposite directions. First, enforcement decreases the disposable income for evaders. However, cracking down on tax evasion increases the size of government, which benefits all

[24]This implies that incumbent's reelection is the main outcome to analyze when testing the empirical predictions of retrospective voting models. Retrospective voting models have received considerable empirical support in the context of government corruption or fiscal stabilization (Brender, 2003; Besley and Pratt, 2006; Ferraz and Finan, 2008, Nannicini, Stella, Tabellini and Troiano, 2012).

[25]For simplicity, we ignore the extra fines evaders pay when audited and, thus, the optimal individual evasion level they choose.

citizens, including evaders. The expected utility for evaders, V_E is defined as:

$$V_E(\hat{p}) = \hat{p}(U(1-\tau)) + (1-\hat{p})U(1) + \hat{p}W_E(\lambda, e), \tag{1.1}$$

where we highlight that V_E depends on the expected level of enforcement, $\hat{p}$. In Equation 1.1, $U(\cdot)$ is the monetary utility from disposable income and $W_E(\cdot)$ is the utility from tax collection enforcement.[26] W_E is increasing in λ, the share of evaders in the population, and e, the government efficiency in using tax revenues to produce public good.

We allow non-evaders to obtain an additional non-monetary benefit from enforcement. One example is the case where, because of fairness concerns, non-evaders derive direct utility from the enforcement of evaders' tax payments, independently from their monetary returns. Thus, the expected utility function for the non-evaders is:

$$V_N(\hat{p}) = \hat{p}W_N(\lambda, e, n) + U(1-\tau) \tag{1.2}$$

We notice that, in addition to λ and e, V_N also depends on n, a shifter that affects the non-monetary benefits from increases in enforcement. For instance, n captures the extent to which voters are averse to tax evasion ("tax culture"). In the model, we abstract from the utility arising from government services financed by the tax payments of the non-evaders since that does not depend on $\hat{p}$, the core variable of interest for our argument.

We now consider the voters' choice between an incumbent and a contender. We adopt a standard probabilistic voting approach (Lindbeck and Weibull, 1987). In the text below, $\hat{a}$ and $\hat{p}$ denote the voters' beliefs about incumbent type and enforcement, respectively. On the other hand $\overline{a}$ and $\overline{p}$ capture the expectations about the contender. In deciding whether to reelected the incumbent, the two groups of voters compare the utility under the expected incumbent's type with an average opponent. Voter i in group $j = \{E, N\}$ will reelect the incumbent if $V_j(\hat{p}) > V_j(\overline{p}) + \epsilon_{ij} + \delta$. The parameter ϵ_{ij} is an individual ideological bias

[26] In order to simplify the presentation, we assume that the utility from enforcement is proportional to the expected level of enforcement.

toward the contender, distributed uniformly over $[-\frac{1}{2\phi^j}, \frac{1}{2\phi^j}]$.[27] The parameter δ measures the average popularity of the contender in the population and is distributed uniformly over $U[-\frac{1}{2}, \frac{1}{2}]$.

Under the above assumptions, the ex-ante incumbent reelection probability (i.e., before the realization of δ) is:

$$\pi = (\hat{p} - \overline{p})\left[\lambda\phi_E(-U(1) + U(1-\tau) + W_E) + (1-\lambda)\phi_N W_N\right] \tag{1.3}$$

The following equation presents the electoral impact of an increase in the expected enforcement level under the incumbent, $\hat{p}$:

$$\frac{\partial \pi}{\partial \hat{p}} = \lambda\phi_E\big(-U(1) + U(1-\tau) + W_E\big) + (1-\lambda)\phi_N W_N \tag{1.4}$$

The first term of the outer sum represents the net electoral gains coming from evaders voting. These will be negative whenever the utility cost of the expected loss in disposable income, $U(1) - U(1-\tau)$, more than offsets the benefits from enforcement, W_E. The second term is the electoral gain from non-evaders (always positive). This duality is consistent with the discussion in Section 1.1: an increase in the perception of the enforcement type of the incumbent has ambiguous effects. The change generates a conflict across voters and the model parameters determine which channel prevails.

In addition, the model delivers intuitive comparative statics on the heterogeneity of the electoral impact arising from an increase in expected enforcement under the incumbent. Intuitively, both governmental efficiency in public good provision and the intensity of non-monetary benefits from the additional enforcement matter play a role. Specifically:

$$\frac{\partial^2 \pi}{\partial \hat{p} \partial e} = \lambda\phi_E \frac{\partial W_E}{\partial e} + (1-\lambda)\phi_N \frac{\partial W_N}{\partial e} \tag{1.5}$$

and

$$\frac{\partial^2 \pi}{\partial \hat{p} \partial n} = (1-\lambda)\phi_N \frac{\partial W_N}{\partial n}, \tag{1.6}$$

[27]The parameters ϕ^E and ϕ^N should be interpreted as proxies for the responsiveness of voters in each group to tax evasion enforcement. They might reflect for example the fact that the political power of a group can change depending on its size or ability to self-organize (Olson(1965)).

which are both positive.

To summarize, the simple model predicts that an increase in the expected level of enforcement under the incumbent:

i) Has an ambiguous impact on the incumbent reelection likelihood

ii) Is larger when government is more efficient in public good provision

iii) Is larger when the non-monetary returns from enforcement are larger

The Ghost Buildings program allows us to shed light on these predictions. The, program initiated by the central government, can be considered as a positive shock to enforcement. We argue that voters observe the increase in the building registration but have limited information about the specific "production function" of enforcement (i.e., information collected by the central government, local administration effort, and complementarity between the two sources) This in turn increases the belief voters hold about the local incumbent type, $\hat{a}$ (and thus on $\hat{p}$), and, according to the model, generates an ambiguous effect on the incumbent reelection probability.

Crucially, this result relies on the assumption that voters have limited information about the details of the Ghost Buildings program. They observe the change in enforcement and still attribute a part of it to the incumbent, thus extracting signal on her type. Models with *rational but poorly informed voters* have received growing attention in the literature. They can provide theoretical support for the empirical findings that voters' electoral choices respond to economic conditions (Wolfers (2009)), natural disasters (Cole, Healy, and Werker (2012)), corruption (Nannicini, Stella, Tabellini and Troiano, (2012)) and quasi-random targeted transfers (Manacorda, Miguel, and Vigorito (2011)). In addition, a recent wave of randomized experiments shows that information provision can significantly affect voter choices and political outcomes (for a review, see Pande (2011)). For the Ghost Buildings program, it is likely that inference about who exactly was causing the extra enforcement was difficult. Local administration efforts complemented the initial identification process. In addition, evidence from media reports and town bulletins suggest that at least some mayors claimed a role throughout the program, including the initial stages of building identification through aerial pictures (Cavallaro (2011), Corriere della Citta' (2012), Gazzetta del Mezzogiorno (2012)).

Finally, we notice that it is also possible to predict an impact on support for the incumbent in an alternative model where voters perfectly observe the nature of the Ghost Buildings program (while they are still uncertain about the type of mayor). In this alternative setting, the program provides an opportunity for voters to extract a more precise signal about the incumbent type, as in Bubb (2008). This can in turn either benefit or hurt the incumbent, depending for instance on voter risk preferences (Quattrone and Tversky (1988)) or on the skewness of the distribution of incumbent types (Caselli et al, (2013)). In the rest of the paper, we do not aim to differentiate the two classes of models in the data analysis. Rather, the insight that the net voter response to an enforcement policy is theoretically ambiguous, which is common to both models, motivates our empirical investigation.

1.4 Data and Descriptive Evidence

1.4.1 Data

The main database for the analysis includes information on the number of parcels containing Ghost Buildings in each town. The aerial photographs detected more than two million such parcels. We target the population of 7,720 of the 8,092 Italian towns (*Comuni*) for which we can define the measure of ghost building intensity. Additionally, we obtain data on registered ghost buildings up to the deadline of April 30, 2011. In order to analyze the electoral response to ghost building registration, we construct a measure of registration imputable to the incumbent administration. Specifically, we multiply the registration rate by the ratio between a) the time elapsed between program start date and election date and b) the time elapsed between program start date and April 2011.[28]

[28]In one of our robustness checks, we also compute a second measure of registration imputable to the incumbent under the assumption of a constant growth rate of 50% in the registration levels over years.

We complement this information with data from the Italian Deparment of the Interior (*Ministero degli Interni*) which contains outcomes for the universe of municipal elections from 1993 to 2011.[29] In Figure 1.2, we plot the number of elections per year. Towns vote in different years, according to predetermined waves. We distinguish between elections before and after the beginning of the Ghost Buildings program. There are almost 5,200 municipalities for which we have data on an election that occurred after program inception (about 67% of the total number of towns targeted by the program). It is also important to discuss two institutional reforms that occurred in the time span of our sample. First, in 1993, the starting year for our election sample, Italian municipal politics were overhauled: a new electoral law changed the mayoral electoral system from party to individual ballot. It also introduced a two-term limit. Second, in 2000, the length of the mayoral term was extended from four to five years.

In addition to the core data, we collect geographic and socio-economic data at the municipality level from the Italian National Statistical Office. Finally, we use two additional data sources to test the channels driving the electoral response: town-level government expenses (from the *Ministero degli Interni*) and a region-level standardized score to the question "Do you justify tax cheating?" from the *European Values Study*.

Table 1.1 presents summary statistics for the variables used in the paper. Panel A presents the main variables related to the Ghost Buildings program. Panels B and C include town-level geographical and socio-economic covariates, respectively. These are measured prior to the inception of the Ghost Buildings program, mostly in the 2001 Population Census. Panel D ("Mayor" variables") summarizes characteristics of the mayor in office at the time of the program inception in the town.[30] for each Italian city in our sample. In Panel E, we summarize the local election panel variables.[31] Tables A.1 and A.2 provide a detailed

[29]The Italian municipal government (*Comune*) is composed of a mayor (*Sindaco*), an executive committee (*Giunta*) appointed by the mayor, and an elected city council (*Consiglio Comunale*).

[30]Only about a half of mayors are matched to national parties. We therefore choose not to focus on this variable in our analysis. Including dummies for political alignment among the controls in the regressions does not affect the results we present later in the paper.

[31]Given that our main outcome of interest is the probability of reelection, Table 1.1 summarizes the variables

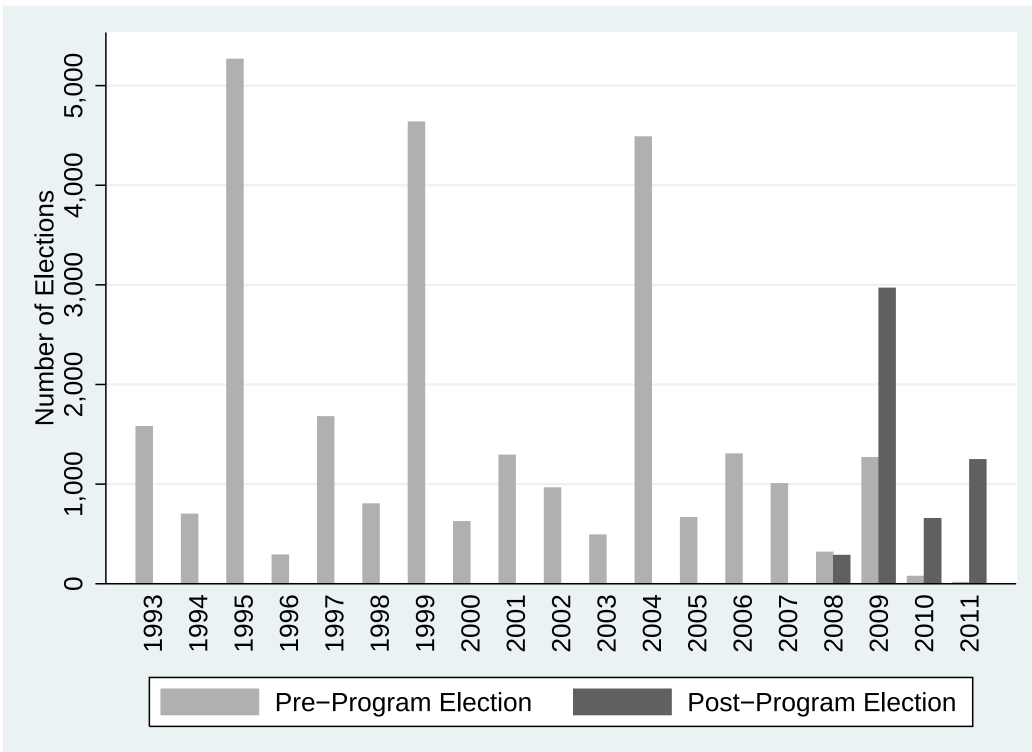

Notes: The figure shows, for each calendar year, the number of elections held before and after the inception of the Ghost Buildings program. The figure includes elections from 1993 to 2011.

Figure 1.2: *Number of Elections per Year*

description of data sources and variable definitions.

1.4.2 The Correlates of Tax Evasion

We use data from the Ghost Buildings program to study the correlates of tax evasion at the town level. Figure **??** presents our measure of ghost building intensity across Italian towns.[32] Tax evasion is more prevalent in Southern Italy, and it is less widespread in the North. Table 1.2 presents the correlates of ghost building intensity (per 1,000 of land registry parcels). In Column (1), we first study whether geographical factors (altitude, area size of the municipality, number of land registry parcels) are correlated with tax evasion. In Column (2), we add socio-economic controls (population, income per-capita, social capital, number of firms, urbanization rate). Finally, in Column (3), we show that our results are unaffected by the inclusion of regional fixed effects.[33]

We find that several geographic characteristics are strongly associated with tax evasion. In particular, controlling for other variables, tax evasion is higher in more widespread municipalities. Plausibly, in cities with wide geographical extension, the opportunities for unregistered buildings are higher as the enforcement of building registration is more difficult and resource-intensive. However, we cannot decisively interpret this evidence as causal. Previous literature has shown for example that borders are endogenously determined (see, among others, Alesina and Spolaore (1997); Alesina, Baqir and Hoxby (2004); Alesina, Easterly, and Matuszeski (2011)).

Finally, as expected, tax evasion is negatively associated with both social capital and income. In particular, the finding on social capital is consistent with Putnam (2001), who finds that the percentage of tax evasion, as measured by the Internal Revenue Service,

---

only for the elections in which the mayor does not face a binding a term limit.

[32]For presentation purposes, we choose to show the results of this section using a measure of ghost building intensity obtained by normalizing the number of parcels with ghost buildings (per 1,000 land registry parcels).

[33]For 3.5% of the towns in our sample we are missing at least one town-level control. In our regressions throughout the paper, for each control, we include a binary indicator which is equal to one if the control is missing. In addition, we replace missing values with an arbitrary unique value. The results of our paper are unchanged when we undo this and just drop observations with missing values for the control variables.

Table 1.1: *Summary Statistics*

Variable	Mean	Std. Dev.	Min.	Max.	N
Panel A: Ghost Building Town Variables					
Ghost Building Intensity	0.027	0.021	0	0.187	7720
Registered Ghost Building Intensity (Apr 2011)	0.006	0.006	0	0.051	7720
Ghost Building Registration Rate (Apr 2011)	0.243	0.181	0	1	7720
Panel B: Geographic Town Variables					
Town Area Size (sq km)	37.045	50.112	0.2	1307.71	7714
Altitude (mt)	510.535	461.655	0	3072.5	7714
Land Registry Parcels (1,000)	10.776	13.278	0.001	514.372	7720
Panel C: Socio-Economic Town Variables					
Population (1,000)	7.227	40.248	0.033	2546.804	7713
Disposable Income per capita (1,000 Euros)	13.487	3.05	5.013	44.949	7595
Urbanization Index	1.619	0.684	1	3	7713
Non-Profit Associations/1,000 pop	5.285	3.972	0.212	81.218	7461
Number of Firms per capita	0.077	0.027	0.018	0.344	7595
Panel D: Mayor Variables					
Mayor Age	49.03	9.68	21	83	7427
Mayor Education	5.36	0.78	3	9	7257
Mayor Born Same City (0/1)	0.47	0.5	0	1	7421
Mayor Term Number	1.3	0.46	1	2	7608
Mayor Woman (0/1)	0.1	0.3	0	1	7603
Panel E: Election Panel Variables					
Term Limit Indicator (0/1)	0.202	0.401	0	1	32501
Post Program Election (0/1)	0.145	0.352	0	1	25952
Years Elapsed since Program Inception ($= 0$ if ≤ 0)	0.306	0.816	0	4	25952
Incumbent Reelection (0/1)	0.453	0.498	0	1	25952
Incumbent Rerun (0/1)	0.572	0.495	0	1	25580
N. Candidates	2.763	1.303	1	17	24637
Victory Margin	25.953	26.938	0	100	23985
Runoff (0/1)	0.526	0.499	0	1	2297

Notes: **Socio-Economic Town Variables** are collected before the Ghost Buildings program inception. **Mayor Variables** refer to characteristics of the incumbent mayor at the time of program inception. Summary statistics for the **Election Panel Variables** are reported for the subsample of elections with no binding term limit, except for *Term Limit Indicator*. A detailed description and source of each variable is provided in the *Data Appendix*

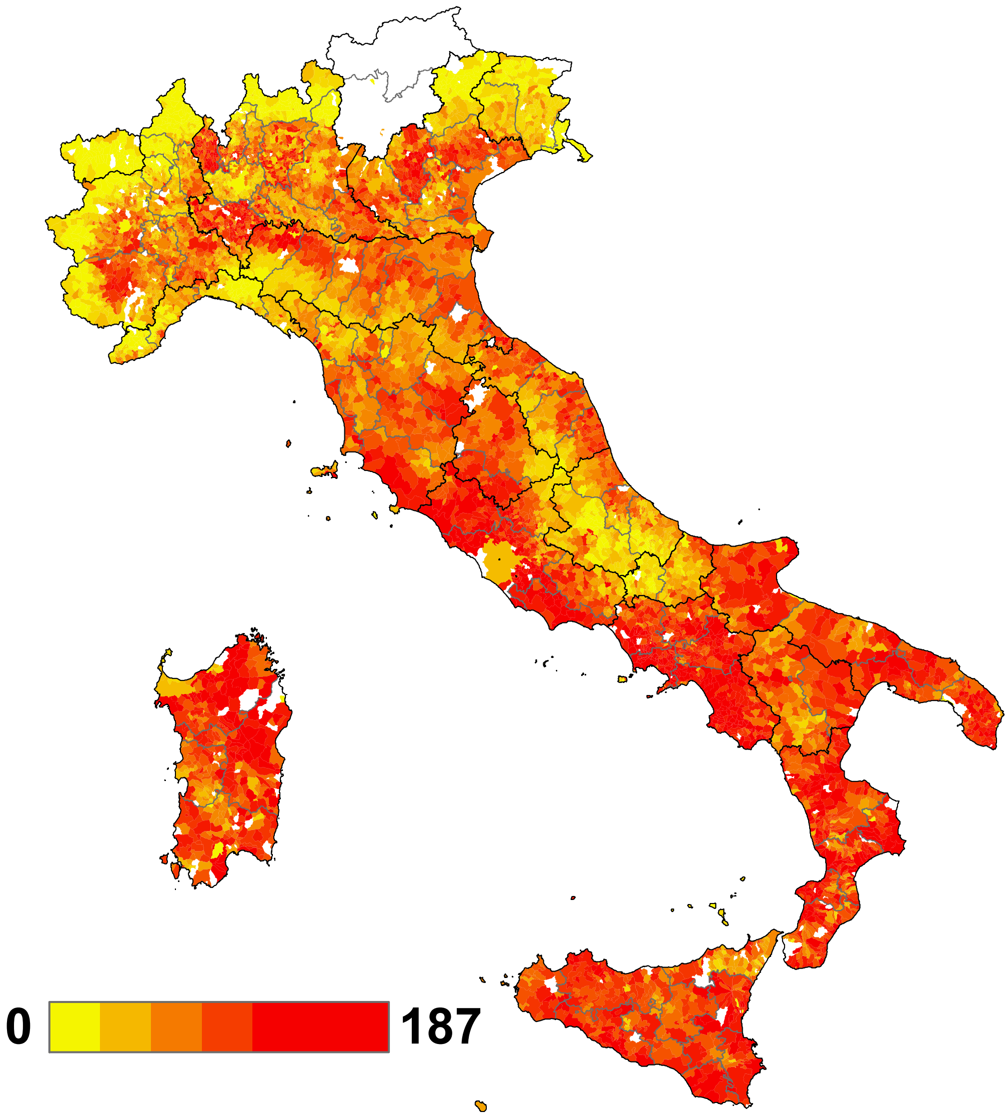

Notes: In this figure, Ghost Building Intensity is defined as the number of land registry parcels with ghost buildings per thousand of land registry parcels. White areas identify towns for which we do not have data.

Figure 1.3: *Ghost Building Intensity (per 1,000 land registry parcels)*

Table 1.2: *The Determinants of Ghost Building Intensity (per 1,000 land parcels)*

	(1)	(2)	(3)
Town Area Size (sq km)	0.102***	0.123***	0.098***
	(0.021)	(0.015)	(0.012)
Altitude (mt)	-0.015***	-0.009***	-0.011***
	(0.002)	(0.002)	(0.002)
Land Registry Parcels (1,000)	-0.236***	-0.328***	-0.270***
	(0.065)	(0.070)	(0.047)
Population (1,000)		0.020	0.004
		(0.016)	(0.011)
Disposable Income per capita (1,000 Euros)		-2.601***	-1.223***
		(0.363)	(0.291)
Urbanization Index		5.948***	4.407***
		(1.839)	(1.652)
Non-Profit Associations/1,000 pop		-0.458***	-0.211***
		(0.144)	(0.074)
Number of Firms per capita		56.100***	89.945***
		(20.040)	(17.842)
Region FE			X
Observations	7720	7720	7720

Notes: The dependent variable is the town-level ghost building intensity per thousand of parcels, defined as the ratio between the number of land registry parcels with ghost buildings and the total number of land registry parcels, multiplied by one thousand. Standard errors are clustered at provincial level. *p<0.1, **p<0.05, ***p<0.01.

is strongly related to differences in social capital at the state level. Additionally, when analyzing the DDB Chicago Lifestyles survey, he argues that by far the best predictor of tax evasion is the number of times in the course of the last year that respondents gave the "finger" to another driver.

1.4.3 The Political Determinants of the Ghost Buildings Registration

We now provide more details on the wave of registration of ghost buildings induced by the program. First, we show that the number of ghost buildings detected by the program is a good predictor of the number of ghost buildings that were registered in response to the policy. Figure **??** displays the relation between the number of land parcels with ghost buildings eventually registered by the April 2011 deadline ("registered ghost building intensity") and the number of parcels that were identified as containing ghost buildings ("ghost building intensity"), both as a share of the total number of land registry parcels. In the graph, the x-axis variable is partitioned into 25 quantiles. The scatter plot shows a clear increasing relation. In a linear regression analysis, an increase of one standard deviation in the detected intensity of ghost buildings raises the *registered* ghost building intensity at April 2011 by approximately 0.75 standard deviations ($p < 0.01$).[34] To summarize, the program scope at the town-level strongly predicts the actual impact of the program on tax enforcement. This premise motivates the strategy that we introduce in Section 1.5 to estimate the impact of the Ghost Buildings program on electoral outcomes.

Second, we analyze the ghost buildings registration rate, defined as the percentage of ghost building parcels that get registered by the April 2011 deadline. Figure **??** summarizes the ghost building registration rate and documents a substantial dispersion across towns. Table 1.3 documents the impact of characteristics of the mayor at the time of the program inception on this outcome. For a given level of the other covariates, the registration rate is higher when mayors are male, younger, more educated, or are born in the same city in which

[34] The relation is basically unchanged when adding town-level controls and regional fixed effects (results available on request).

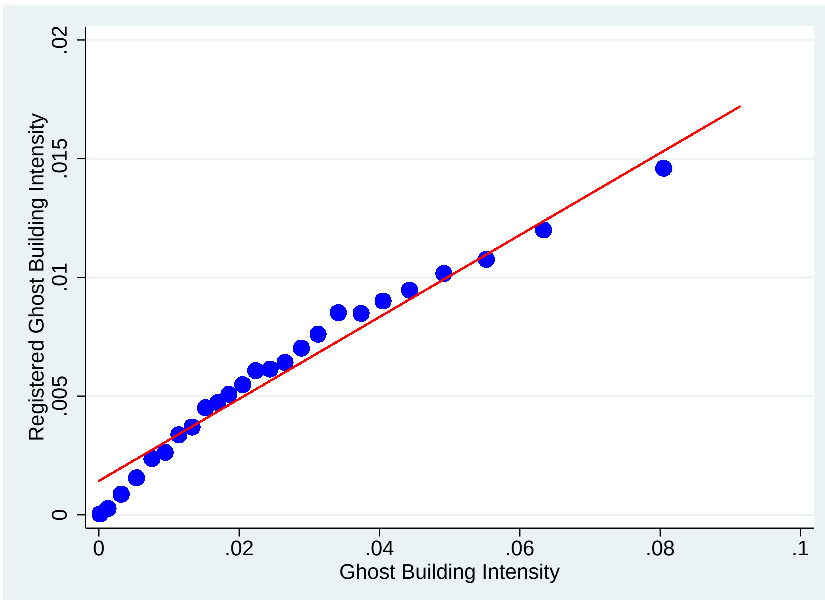

Notes: The scatter plots the relation between *Registered Ghost Building Intensity* (i.e., the fraction of land parcels with ghost buildings that get registered by April 2011) and *Ghost Building Intensity* (i.e., the fraction of land parcels with ghost buildings identified by the program). The x-axis is partitioned into 25 quantiles. The x-axis of each dot is the median value of the ghost building intensity in the quantile. The y-axis is the average value of the registered ghost building intensity in the quantile. We cut the top 1% of the x-axis values from the graph. The line plots the predicted values from a linear regression model.

Figure 1.4: *Registered Ghost Building Intensity*

they serve as mayor. The correlation between gender and policies in Italian municipality is potentially consistent with the results of Gagliarducci and Paserman (2012), who find that female policymakers usually face more difficulty in implementing policies when in office. To the extent education can be considered a proxy for politicians' quality (see, for example, Besley, Montalvo and Reynal-Querol (2011)), this set of results also supports the view that better policymakers fight tax evasion more. We highlight the correlation between the mayor's place of birth and the tax evasion enforcement. One possible explanation could be that mayors who are born in the same city have access to additional information that can facilitate tax evasion enforcement. We acknowledge that this evidence relies on cross-sectional correlation analysis and thus should be interpreted with caution. However, we also notice that the results are robust to the inclusion of geographical and socio-economic controls, in Columns (2) and (3), respectively. With these caveats in mind, the findings of this section suggest that mayor's characteristics did have a role in shaping registration activities across towns.

1.5 Empirical Strategy

1.5.1 The Electoral Response to the Ghost Buildings Program

In this section, we outline our approach to estimate the voter response to the Ghost Buildings program. We also aim to isolate the channels that drive this response. Our empirical strategy exploits variation across towns in the program scope to increase tax enforcement.[35] We implement a difference-in-differences approach based on the town-level ghost building intensity, which we defined above as the ratio between the number of land registry parcels with ghost buildings and total number of land registry parcels in each town.

In Section 1.4.3, we documented that mayor characteristics, such as age, education, and

[35]Importantly, the *Agenzia del Territorio* conducted the detection activities homogeneously throughout the country. Thus, heterogeneity in the number of detected unregistered buildings captures differences in of actual levels of non-registration at the time of the aerial photographing as opposed to differential intensity in the detection activity.

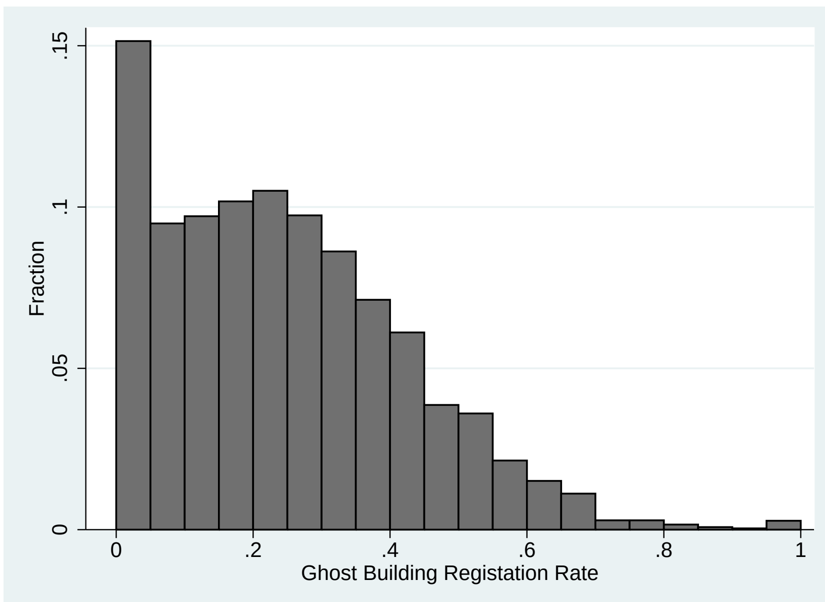

Notes: The histogram shows the distribution of the ghost building rate at April 30, 2011, defined as the ratio between the number of land registry parcels ghost buildings that get registered by April 2011 and the number of land registry parcels with ghost buildings identified by the program.

Figure 1.5: *Ghost Building Registration Rate*

Table 1.3: *The Determinants of the Ghost Building Registration Rate*

	(1)	(2)	(3)
Mayor Age	-0.067***	-0.068***	-0.057***
	(0.022)	(0.021)	(0.020)
Mayor Education	0.704***	0.630**	0.702***
	(0.265)	(0.262)	(0.249)
Mayor Born Same City (0/1)	1.051**	1.145***	0.968**
	(0.424)	(0.423)	(0.408)
Mayor Term Number	-0.174	-0.047	-0.041
	(0.362)	(0.353)	(0.364)
Mayor Woman (0/1)	-0.923	-1.246**	-1.202*
	(0.642)	(0.626)	(0.607)
Geograpic Controls		X	X
Socio-Economic Controls			X
Observations	7720	7720	7720

Notes: The dependent variable is the town-level ghost building registration rate, defined as the ratio between the number of land registry parcels with ghost buildings that get registered by April 2011 and the number of land registry parcels with ghost buildings identified at the beginning of the program. Refer to Table 1.1 for a description of the *Geographic* and *Socio-Economic* Controls. All the regressions include regional fixed effects and year-of-program-inception fixed effects. Standard errors are clustered at provincial level. *p<0.1, **p<0.05, ***p<0.01.

gender, predict the registration rate of the detected ghost buildings. In addition, the actual levels of registration might depend on voter preferences and responsiveness to the program. Thus, a naive analysis looking at the relationship between actual ghost building registrations and reelection outcomes will suffer from the standard omitted variable bias. This motivates our focus on *ex ante* program scope to increase enforcement.

The rationale for our identification approach is that program scope at the town level predicts the increase in enforcement induced by the Ghost Buildings program, as shown in Figure **??**. Towns with a higher share of parcels with *detected* ghost buildings also have, on average, a higher share of parcels with *registered* ghost buildings, as measured in April 2011.[36]

[36]Importantly, we argue that the intensity of ghost buildings is not a valid instrument for actual registration intensity. In principle, as we discussed in Section 1.1, the program could affect incumbent reelection probability through other channels besides registration. This would make the standard exclusion restriction required for an instrumental variable approach invalid.

Our baseline specification is therefore:

$$R_{imet} = \beta_0 + \beta_1 Post_{ie} \cdot Ghost\ Building\ Intensity_i + \eta_m \cdot Post_{ie} + \phi_i + \phi_t + \epsilon_{imet} \quad (1.7)$$

The dependent variable R_{iret} is a dummy that indicates whether the incumbent of municipality i in macro-area m is re-elected in election e in year t. Observations where the incumbent cannot be reelected because of a binding term limit are excluded from the regression sample. The dummy *Post* is equal to one when election e occurs after the beginning of the Ghost Buildings program in the town. The coefficients η_m capture post-program period fixed effects that are specific to the four Italia macro-areas[37] m where town i is located. We also include town fixed effects, ϕ_i, and election year fixed effects, ϕ_t. Finally, $Ghost\ Building\ Intensity_i$ is the intensity of ghost buildings in town i. The coefficient of interest, β_1, thus captures the differential impact of the Ghost Buildings program on incumbent reelection by ghost building intensity. Throughout the paper, we cluster standard errors at the provincial level to allow for spatial correlation in the error term.

We adopt a similar regression model to study the impact of the program on other electoral competitiveness outcomes. We focus on four variables: i) a binary indicator for whether the incumbent reruns for election; ii) the number of candidates running for the mayor office; iii) the difference in the percentage of votes between the first and the second candidate;[38] iv) a binary indicator equal to one if a runoff takes place, which occurs in towns with more than 15,000 inhabitants when none of the candidates obtain the absolute majority in the first-round.

One potential challenge to our identification strategy may arise from the town-specific timing of publications of the unauthorized building lists. On the one hand, if local administrators had influence over publication date, unpopular mayors in cities with high evasion rates might lobby to delay the publication. On the other hand, the central government might push to start the program earlier in those towns where mayors set a lower level of

[37] North, Center, South, Islands.

[38] For elections with a runoff, we use the first round results.

tax enforcement. In both these cases, our estimates of the impact of the Ghost Building program on reelection likelihood might just capture a selection effect. We handle this concern in several ways. First, as discussed in Section 1.2, we notice that the timing of the publication was primarily determined at the provincial level by the availability of digital land registry maps and was highly clustered by province. Figure 1.6 emphasizes the high level of provincial clustering in the publication years. Only about 7% of the post-program elections have values for the post-program indicator different from the one they would have had based on the modal date of publication in the province. To deal with these discrepancies we implement an instrumental variable approach. We code elections based on whether they occur before or after the modal date of publication of the unauthorized building lists in the province. We then instrument the actual *Post* dummy with this binary indicator at the provincial level. We adopt this strategy for our main specifications.[39] In addition, in Section 1.6.2, we present robustness checks using an alternative instrument for the post-program indicator using the *national* modal program inception year.

As is standard in difference-in-differences estimation, identification of the coefficient of interest relies on two assumptions. The first is the absence of contemporary events that differentially affected towns having higher ghost building intensity. We are not aware of other policies targeting this form of tax evasion happening concurrently with the Ghost Buildings program. However, it is still possible that other events, which differ in intensity by other variables correlated with ghost building intensity, occurred at the same time. We address the concern by presenting alternative specifications where we include interactions between a comprehensive set of geographical, socio-economic, and political controls, all measured before the beginning of the program, and the post-program binary indicator. The second assumption is the presence of parallel trends in the outcome variable. We assess this assumption using several tests and placebo exercises in the next section.

[39]The towns targeted by the program belonged to 101 provinces.

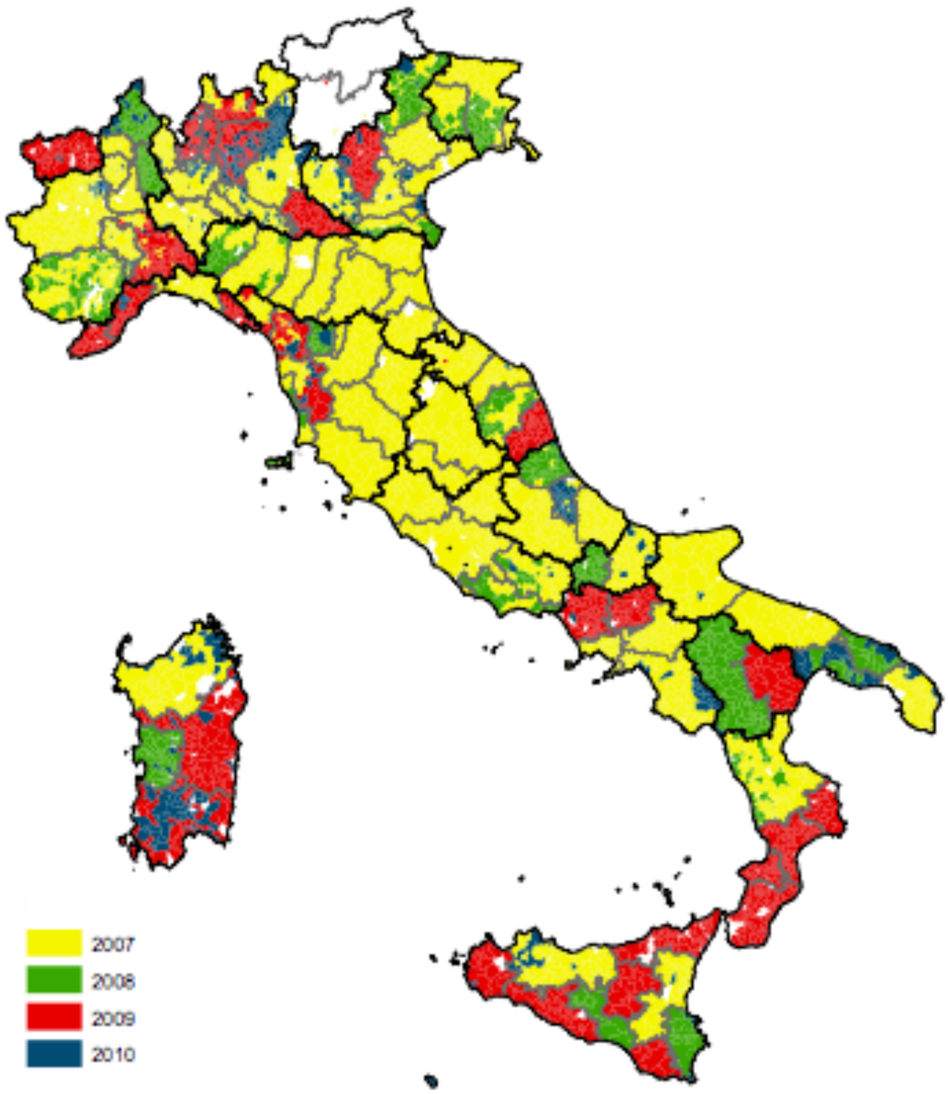

Notes: The figure shows the year of inception of the Ghost Building program (i.e., the year of publication of the list of ghost buildings) in each town. White areas identify towns for which we do not have data.

Figure 1.6: *Ghost Buildings Program Inception Year*

1.5.2 Channels

The reduced form approach presented thus far tests whether higher program scope to increase tax enforcement at the town level affects incumbent reelection likelihood in the post-program period. We complement this baseline regression with further analysis. First, we show that it is the registration induced by the program that drives the electoral response, as opposed to other potential interpretations. For this purpose, we use actual ghost building registration data. In Section 1.4, we emphasized several important measurement limitations of these data that warrant caution. With this caveat in mind, we test whether, for a given intensity of ghost buildings, a higher ghost building registration rate (*Registration Rate*) induced by the program has a positive effect on incumbent reelection likelihood:

$$\begin{aligned} R_{imet} &= \gamma_0 + \gamma_1 Post_{ie} \cdot Ghost\ Building\ Intensity_i \\ &\quad + \gamma_2 Post_{ie} \cdot Registration\ Rate_i + \zeta_m \cdot Post_{ie} + \mu_i + \mu_t + v_{imet} \end{aligned} \tag{1.8}$$

As we discussed above, an obvious threat to identification of γ_2 in Equation 1.8 arises from the fact that the registration effort is potentially correlated with many potential town-level confounders. We first check robustness of the results to the inclusion of mayor controls. In addition, the timing of the program provides a strategy that can alleviate this concern. Even if the program started in the same year in most of the towns, we can exploit the variation generated by the fact that Italian municipalities hold elections in different years. A longer time between the beginning of the program and the election date naturally leads to more registration activities. This generates variation across towns in the registration rate achieved prior to the local election date that is plausibly uncorrelated with mayor quality. We use this instrumental variable strategy to look at the impact of a change in registration rate on incumbent reelection likelihood.

Second, we shed light on the channels through which the program could affect voters' political preferences. Consistently with the theoretical framework, we investigate the interaction among the political returns to an increase in tax enforcement, local government efficiency in delivering public goods, and "tax culture", the stigma associated to evading taxes. We use the speed of public good provision as a proxy for the quality of the delivery at the *municipal* level. This indicator is measured as the ratio of paid outlays in the municipal financial report over the total outlays committed in the budget. The intuition is that the provision of public goods is more effective in places where the actual allocation delivered to citizens is closer to the amount allocated in the budget.[40] The following regression model tests whether the electoral response to the Ghost Building program varies by speed of public good provision:

$$R_{iret} = \delta_0 + \delta_1 Post_{ie} \cdot Ghost\ Building\ Intensity_i + \delta_2 Post_{ie} \cdot Speed\ Public\ Good\ Provision_i + \delta_3 \cdot Post_{ie} * GB_i \cdot Speed\ Public\ Good\ Provision_i + \xi_r \cdot Post_{ie} + \lambda_i + \lambda_t + \nu_{iret} \quad (1.9)$$

where δ_3 is the coefficient of interest.

We then use data from the *European Value Study* — the European component of *World Values Survey*— to study the role of tax culture. Specifically, we use the question: "Do you justify cheating on tax?". For these data, geographical identification of respondents is available only at the regional level. We thus compute and standardize region-level means. The regression model to capture heterogeneity by this variable is similar to the one presented in Equation 1.9.

Finally, we also assess the impact of the program on town-level public expenditures. To test whether the program scope to increase tax enforcement affected these expenditures, we adopt a specification similar to the one we presented in Equation 1.7, using the natural logarithm of the local government expenditures as the dependent variable.

---

[40] A similar proxy has been used by Grembi, Nannicini and Troiano (2013) and Gagliarducci and Nannicini (2013) to measure the quality of public good delivery. For our sample, we compute the speed of public good provision as the average across two pre-treatment years. The results are similar with alternative definitions. Results are available upon request.

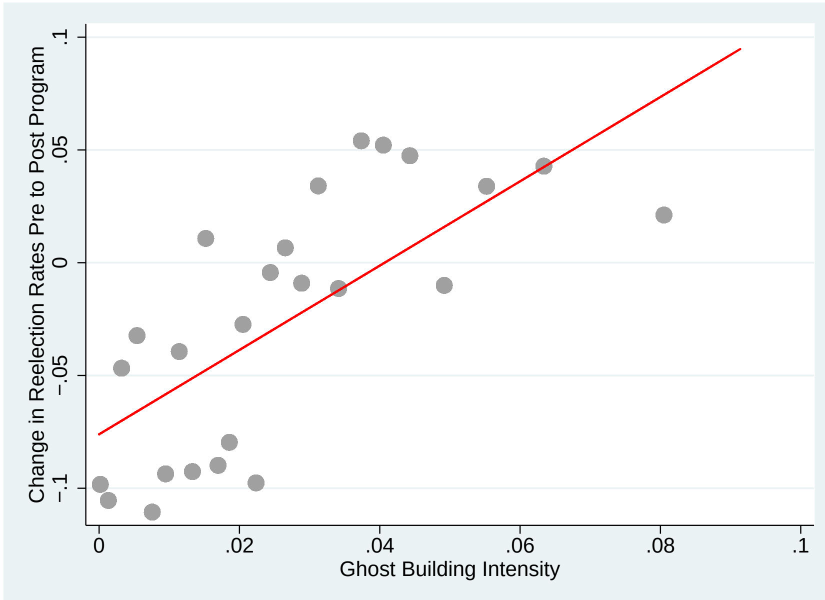

Notes: The scatter plots the relation between the change in the average (year-demeaned) reelection rate between the pre-program and the post-program periods and the *Ghost Building Intensity* (i.e., the ratio between the number of land registry parcels with ghost buildings and the total number of land registry parcels). The x-axis is partitioned into 25 quantiles. The x-axis of each dot is the median value of the ghost building intensity in the quantile. The y-axis is the average value of the registered ghost building intensity in the quantile. We cut the top 1% of the x-axis values from the graph. The sample includes elections with no binding term limit. The line plots the predicted values from a linear regression model.

Figure 1.7: *Difference in reelection rates pre- to post- Ghost Buildings program*

1.6 Results

1.6.1 Baseline Results

In this section, we investigate the electoral consequences of the Ghost Buildings program. Figure 1.7 provides a visual analysis of the relation between ghost building intensity and changes in the incumbent reelection likelihood — our main outcome variable — after the beginning of the program. On the x-axis, the ghost building intensity is partitioned into 25 quantiles. The scatter displays a clear increasing relation, with the quantile dots fairly closed aligned along the fitting line.

Table 1.4 formalizes the analysis above and presents the results of the difference-in-differences estimation discussed in Section 1.5. Column (1) reports the basic OLS specification ("Reduced Form") using the provincial post-program indicator. The coefficient remains stable with the inclusion of fixed effect (Column (2)) and of election year fixed effect (Column (3)). Starting in Column (4), we instrument the post-program indicator with the provincial post-program indicator. The coefficient is stable across the different specifications. The inclusion of year fixed effects and town fixed effects, in Columns (5) and (6) respectively, does not change our results. In Column (6) — the baseline specification for the rest of the analysis — the reported coefficient on the interaction between the ghost building intensity and the post-program indicator is 1.042, significant at 1%. This magnitude implies that a one standard deviation in the town-level program scope to increase enforcement, as measured by the ghost building intensity, raises the likelihood of reelection of the incumbent by approximately 2.5 percentage points in post-program elections, relative to pre-program ones (from a sample mean of 45.4). A back of the envelope calculation suggests that the effect of a one standard deviation increase in Ghost Buildings program scope on incumbent reelection probability is on the order of magnitude of i) 6% of the incumbency effect in U.S. House Elections (Lee (2008)) and ii) the effect of a 5% increase in town government spending in Brazil municipal elections (Litschig and Morrison (2012)).

We adopt an analogous regression strategy to study the impact of the program on other measures of election competitiveness as described in Section 1.5. For each of these variables, we report the specification used in Column (6) of Table 1.4. Table 1.5 presents a clear picture. An increase in ghost building intensity raises the likelihood that the incumbent re-runs, and decreases competitiveness of the elections. Specifically, a one standard deviation increase in ghost building intensity reduces the number of candidates by 0.03 standard deviations, increases the margin of victory by 0.05 standard deviations, though this last result is not statistically significant at conventional levels, and reduces the likelihood of a runoff by 15% of the mean value. This is consistent with the idea that potential entrants in the electoral competition correctly anticipate stronger incumbent advantage in response to the program.

Table 1.4: *Ghost Building Intensity and Incumbent Reelection: Baseline Results*

	Reduced Form			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
Ghost Building Intensity*Post				1.099***	1.114***	1.042***
				(0.357)	(0.373)	(0.378)
Ghost Building Intensity*Province Post	1.085***	1.061***	0.953***			
	(0.343)	(0.358)	(0.360)			
Town FE		X	X		X	X
Election Year FE			X			X
Observations	25893	25893	25893	25893	25893	25893

Notes: The dependent variable is a binary indicator equal to one if the incumbent mayor is reelected (mean 0.453). **Post** is a binary indicator equal to one if the election occurs after the Ghost Buildings program inception. **Province Post** is a binary indicator equal to one if the election occurs after the Ghost Buildings program modal inception year in the province. In the columns grouped under the header "2SLS", *Post* is instrumented with *Post Province*. **Ghost Building Intensity** is defined as the ratio between the number of land registry parcels with ghost buildings and the total number of land registry parcels. All the columns include an interaction between region fixed effects and either *Province Post* (Columns (1)-(3)) or *Post* (Columns (4)-(6)). Columns (1) and (4) include the Ghost Building Intensity level. The regression sample includes all the elections between 1993 and 2011 in which the incumbent does not face a binding term-limit. Standard errors are clustered at provincial level. *p<0.1, **p<0.05, ***p<0.01.

Table 1.5: *Ghost Building Intensity and Election Competitiveness*

	Incumbent Rerun	N. Candidates	Victory Margin	Runoff
	(1)	(2)	(3)	(4)
Ghost Building Intensity*Post Program	1.115**	-2.383**	42.063	-4.502***
	(0.457)	(1.057)	(25.984)	(1.383)
Dependent Variable Mean	0.572	2.761	25.999	0.525
Observations	25483	24441	23562	2216

Notes: **Incumbent Rerun** is a binary indicator equal to one when the current incumbent runs for reelection. **N. Candidates** is the number of candidates running for election. **Victory Margin** is the percentage point difference between the first and the second candidate in the elections (we use first-round percentages even for towns with a runoff). **Runoff** is a binary indicator, defined only for towns with more than 15,000 inhabitants, equal to one if the election requires a runoff. This occurs if the first candidate in the first round receives less than 50% of the votes. **Post** is a binary indicator equal to one if the election occurs after the Ghost Buildings program inception. In all the columns, *Post* is instrumented by *Province Post*, a binary indicator equal to one if the election occurs after the modal program inception year in the province. **Ghost Building Intensity** is defined as the ratio between the number of land registry parcels with ghost buildings and the total number of land registry parcels. All the regressions include town fixed effects, election-year fixed effects and an interaction between macro-areas fixed effects and *Post*. The regression sample includes all the elections between 1993 and 2011 in which the incumbent does not face a binding term-limit. The ghost building intensity is defined as the ratio between the number of land parcels with ghost buildings (as measured at the beginning of the program) and the total number of land parcels in the town. Standard errors are clustered at provincial level. *p<0.1, **p<0.05, ***p<0.01.

1.6.2 Identification Validity and Robustness Checks

In this section, we report several auxiliary results that support the validity of our identification strategy. We focus on our main outcome variable: the likelihood of incumbent reelection.[41] Figure **??** presents a placebo version of Figure 1.7. We look only at pre-program elections and define a placebo program year in the median year of this restricted sample. In this exercise, the relation between changes in incumbent reelection likelihood from the pre- to the post-program periods appears and ghost building intensity is very noisy and, if anything, negative.

Second, we test whether the inclusion of town-level controls affects our results. Specifically, we look at three sets of town-level variables that we present in Table 1.1: geographic features, socio-economic variables, and incumbent mayor characteristics at the time of the program commencement. We include the interaction between these variables and the post-program indicator in our regression model (the level is absorbed by the town fixed effect). If the findings in Table 1.4 were reflecting differential changes in the outcome along variables correlated to the ghost building intensity, we would expect the coefficient on the interaction between ghost building intensity and post-program indicator to be substantially affected by the inclusion of these extra controls. Table 1.6 presents the results of this analysis. We find that the baseline results are very robust to the inclusion of each set of controls both separately (Columns (2)-(4)) and together (Column (5)).

In Figure **??**, we check whether towns with different levels of evasion were on different trends before the treatment. We report point estimates and confidence intervals on ghost building intensity for each of the elections pre- and post-program. The figure shows that, before the Ghost Buildings program started, the probability of reelection of the incumbent was independent of tax evasion. None of the pre-program coefficients are either significantly different from zero (the normalized value of the coefficient in the ‚Äú-1‚Äù election, the omitted group) or significantly different from each other. However, after the beginning of the program there is a statistically and economically significant impact. Thus, the coefficient

[41] We obtain similar results for our auxiliary outcome variables. Results are available on request.

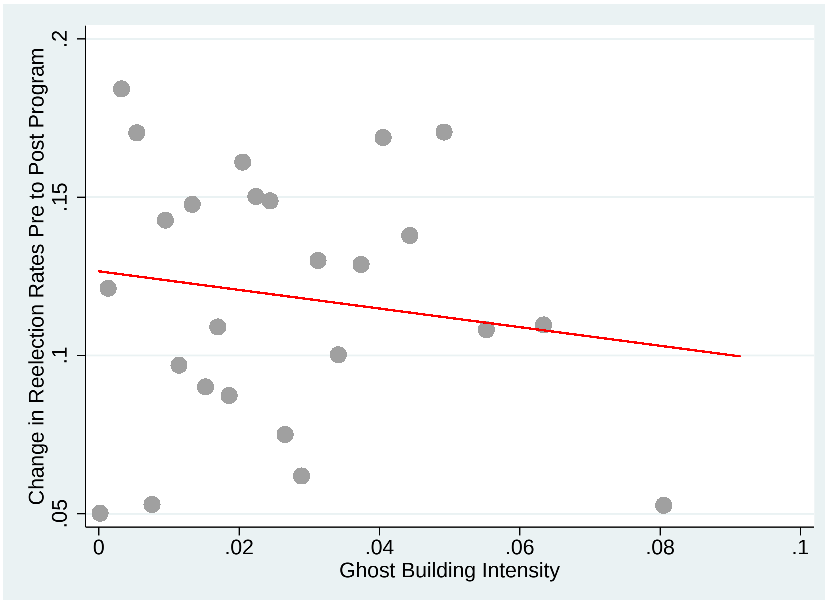

Notes: The scatter plots the relation between the change in the average (year-demeaned) reelection rate between the pre-program and the post-program periods and the *Ghost Building Intensity* (i.e., the ratio between the number of land registry parcels with ghost buildings and the total number of land registry parcels). The x-axis is partitioned into 25 quantiles. The x-axis of each dot is the median value of the ghost building intensity in the quantile. The y-axis is the average value of the registered ghost building intensity in the quantile. We cut the top 1% of the x-axis values from the graph. The sample includes elections with no binding term limit. The line plots the predicted values from a linear regression model. The **placebo subsample** of observations used for this graph only includes election that occurred before the actual program inception. In each town, the year of the placebo program start is defined as nine years before the actual publication. This roughly divides the graph sample in two equally sized groups of pre-placebo and post-placebo elections. The domain of the ghost building intensity is partitioned in 25 quantiles. The x-axis of each dot is the median value of the ghost building intensity in the quantile. The y-axis is the average value of the registered ghost building intensity in the quantile. We trim the top 1% of the x-axis values. The sample includes elections with no binding term limit. The line plots the predicted values from a linear regression model.

Figure 1.8: *Difference in reelection rates pre- to post- Placebo Ghost Buildings program*

Table 1.6: *Ghost Building Intensity and Incumbent Reelection: Additional Controls*

	(1)	(2)	(3)	(4)	(5)
Ghost Building Intensity*Post	1.042***	0.850**	1.053**	1.254***	1.096***
	(0.378)	(0.370)	(0.410)	(0.360)	(0.374)
Geographical Controls*Post		X			X
Socio-Economic Controls*Post			X		X
Mayor Controls*Post				X	X
Observations	25893	25893	25893	25893	25893

Notes: The dependent variable is a binary indicator equal to one if the incumbent mayor is reelected (mean 0.453). **Post** is a binary indicator equal to one if the election occurs after the Ghost Buildings program inception. In all the columns, *Post* is instrumented by *Province Post*, a binary indicator equal to one if the election occurs after the modal program inception year in the province. **Ghost Building Intensity** is defined as the ratio between the number of land registry parcels with ghost buildings and the total number of land registry parcels. Refer to Table 1.1 for a description of the *Geographic*, *Socio-Economic*, and *Mayor* Controls. All the regressions include town fixed effects, election-year fixed effects and an interaction between macro-areas fixed effects and *Post*. The regression sample includes all the elections between 1993 and 2011 in which the incumbent does not face a binding term-limit. Standard errors are clustered at provincial level. *p<0.1, **p<0.05, ***p<0.01.

pattern in Figure **??** suggests that the common trend assumption holds in our setting.[42]

Table 1.7 presents several additional robustness checks. Column (1) shows our baseline specification. In Column (2), we show that the results are robust to the inclusion between regional dummies and post-program indicators. In Column (3), we show that our results hold if we trim the ghost building intensity at the top 1 percent. If anything, the coefficient size grows. This suggests that outliers (i.e., cities with an abnormally large fraction of unregistered buildings) are not driving the results. In Column (4), we check robustness of the results to dropping small towns with a low number of land registry parcels (we drop the bottom 10%). One might be concerned that our results are driven by cities with a large number of non-residents who own a building. If this were the case, mayors would not pay the electoral cost of enforcing tax evasion. Therefore, we exclude in Column (5) cities that are classified as tourist destinations by *Ancitel*, the Italian Association of Cities.[43] The

[42]Appendix Figure A.1 shows that the parallel trend assumptions hold also for the other political outcome variables described above.

[43]Ancitel defines as tourist destinations those cities with a large percentage of touristic income over the total city income. As noted by the Italian Association of Real Estate Agents (*FIAIP*) in their yearly report cities with touristic activities are usually those with a large number of buildings not owned by residents.

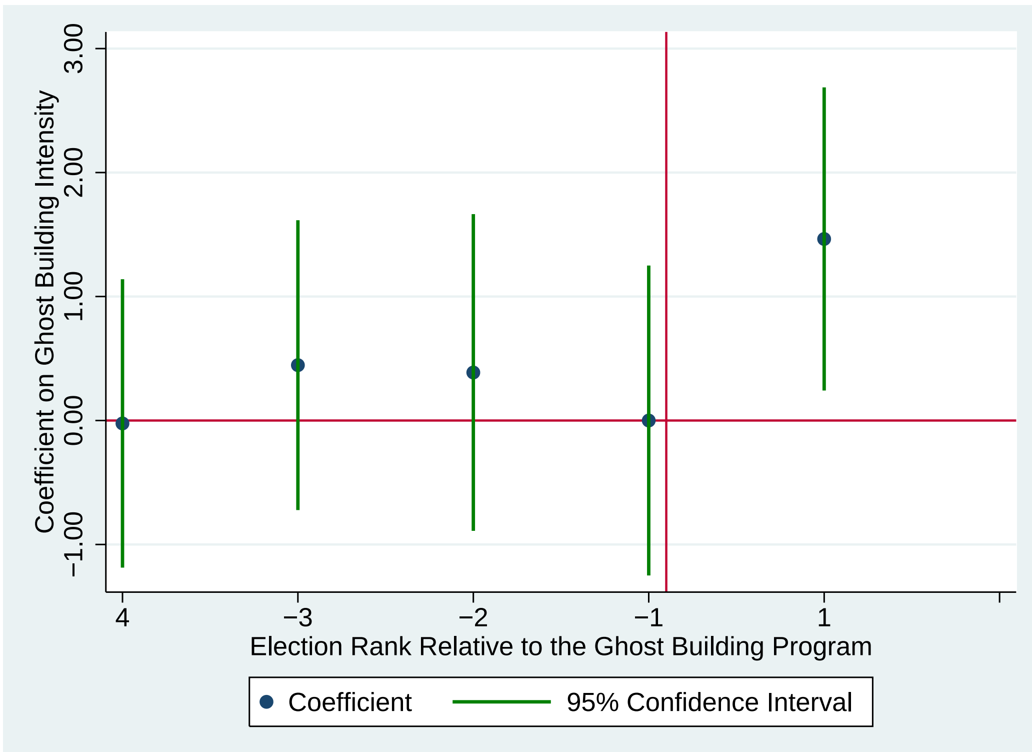

Notes: The graph reports the coefficients on the ghost building intensity for each election before and after the beginning of the Ghost Buildings program. On the x-axis, elections are ranked based on their occurrence relative to the program. The regression includes town and year fixed effects. For each election rank, we report the point estimate and the 95% confidence interval. The last election before the program ("-1") is the omitted category. The coefficient on ghost building intensity for this election is normalized to zero. Confidence interval width for this election is obtained as the mean of the confidence interval width in election -2 and election +1. The modal number of years between elections is five years between 1993 and 2001, and four afterwards.

Figure 1.9: *Ghost Building Intensity Coefficient by Election Pre/Post Program*

results are robust to this sample restriction. Column (6) reports an alternative definition of the post-program indicator. We instrument the town-level indicator with a binary indicator that takes value one for the years 2007-2011. This approach treats 2007, the first and modal program start year in the country, as the *intended* program start year for all the towns. It thus estimates the Local Average Treatment Effect for those towns that started the program in that year, which constitutes a different population of compliers relative to the main specification. The coefficient is about one standard-error larger than our baseline specification and still significant at 1%. Finally, we report an alternative normalization for our dependent variable. As the data reported by the *Agenzia del Territorio* included the number of parcels with unregistered buildings, we divided this outcome by the total number of land registry parcels in each municipality in our baseline specification. In order to show that this choice does not affect the result, in Column (7) we estimate our equation using the total number of buildings recorded in the town as a normalizing factor, rather than the total number of land registry parcels. We get very similar results. The effect of one standard deviation in this alternative variable is comparable in magnitude to the one obtained when using our main ghost building intensity variable.[44]

1.6.3 Channels

This section elaborates on some of the potential channels through which the anti tax evasion program could increase voter support for the incumbent. Table 1.8 presents the results from the estimation of Equation 1.8. This step aims to show that the increase in tax enforcement induced by the program — the ghost buildings registration — drove the electoral response.

In Columns (1) and (2), we present the correlation between the registration rate and the incumbent reelection. We find that, controlling for ghost building intensity, a one standard deviation increase in ghost building registration rate raises reelection likelihood by 1.3 percentage points. The result holds when using both a) the April 2011 registration rate

[44] Appendix Figure A.2 also shows that the results are robust to changes in the election sample years.

Table 1.7: *Ghost Building Intensity and Incumbent Reelection: Robustness*

	(1) Baseline	(2) Region Interactions	(3) Trim top 1%	(4) Drops Small Towns	(5) Drop Touristic Towns	(6) Post 2007	(7) Alternative Normalization
Ghost Building Intensity*Post	1.042***	1.162***	1.255**	1.195***	0.911**	1.345**	0.119*
	(0.378)	(0.386)	(0.493)	(0.390)	(0.416)	(0.563)	(0.067)
Observations	25893	25893	25618	23337	22905	25893	25893

Notes: The dependent variable is a binary indicator equal to one if the incumbent mayor is reelected (mean 0.453). **Post** is a binary indicator equal to one if the election occurs after the Ghost Buildings program inception. In all the columns, *Post* is instrumented by *Province Post*, a binary indicator equal to one if the election occurs after the modal program inception year in the province. All the regressions include town fixed effects, election-year fixed effects and an interaction between macro-areas fixed effects and *Post*. In the specification **Region Interactions**, we add interactions between the post-program indicator and a regional fixed effect. In the specification **Trim top 1% Intensity**, we drop towns in the first percentile of ghost building intensity (0.0928). In the specification **Drop Small Towns**, we exclude from the sample towns within the bottom 10% of the distribution of land registry parcels (2633). In the specification **Drop Touristic Towns**, we drop towns that are classified as tourist destinations by the Italian Association of Cities, *Ancitel*. In Columns (1)-(3) and (5) *Post* is instrumented by *Province Post*, a binary indicator equal to one if the election occurs after the modal program inception year in the province. In Column (4) — **Post 2007**— we instrument it with a binary indicator equal to one if the election occurs after 2007, the start date for the first (and modal) round of the program in the country. **Ghost Building Intensity** is defined as the ratio between the number of land registry parcels with ghost buildings and the total number of land registry parcels. except in Column (5) — **Alternative Normalization** — where it is defined as the ratio between the number of land registry parcels with ghost buildings and the total number of buildings. All the regressions include town fixed effects, election-year fixed effects and an interaction between macro-areas fixed effects and *Post*. The regression sample includes all the elections between 1993 and 2011 in which the incumbent does not face a binding term-limit. Standard errors are clustered at provincial level. * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

with year fixed effects and b) the registration rate reached by the election year computed as described in Section 1.4. In Column (3), we show that adding the interaction between town- and mayor-level controls and the post-program indicator does not change the result. In Column (4), we show that the motivation for the instrumental variable strategy for the registration rate finds support in the data: years elapsed since the program start at election time are a good predictor of the registration rate at that time. In Column (5) we use the years elapsed since the program start as an instrument for registration rate.[45] In the IV specification, a one standard deviation increase in the registration rate (.079) raises the reelection likelihood by 4 percentage points in post-program elections. In Column (6) we show that the IV estimate is unchanged when adding town and mayor controls. Finally, Column (7) shows that an alternative computation of the registration rate imputable to the incumbent — assuming a constant growth rate of 50% in the registration levels — delivers similar results.

In Columns (5) and (7), we notice that the IV estimates are larger than the respective OLS estimates. This is relatively common (see, for example, the returns to schooling literature, reviewed in Card (2001)), and it can be explained either by OLS attenuation bias due to measurement error, or by the fact that in the set of cities affected by the IV — that is, cities where the registration activity depends on program duration — the political returns to registration might be bigger than in the rest of the cities (i.e., we are estimating a LATE). Even if our instrument is uncorrelated with any idiosyncratic city-specific characteristics, we are not able to rule out the possibility that having the program for longer time has an independent effect on its impact on the probability of reelection. While we acknowledge this possibility, we still believe that our instrument does a good job in addressing the main endogeneity concern for the registration efforts of the mayors (town-specific characteristics, such as the mayor's ability or effort).

[45] In our IV specification we do not control for year fixed effects. Three quarters of the post-program elections come from cities that started the program in 2007. Thus, we lose statistical significance when running this specification, though it is reassuring that the coefficient of interest remains of similar size. Results are available upon request.

Table 1.8: *Ghost Building Registration and Incumbent Reelection*

	(1) OLS	(2) OLS	(3) OLS	(4) 1st Stage	(5) 2SLS	(6) 2SLS	(7) 2SLS
Ghost Building Registration Rate*Post	0.173**	0.220***	0.199***		0.625***	0.654***	0.680***
	(0.072)	(0.082)	(0.076)		(0.196)	(0.197)	(0.214)
Years Elapsed since Program Inception				0.066***			
				(0.006)			
Ghost Building Intensity*Post	1.276***	1.271***	1.367***	-0.734**	1.561***	1.877***	1.616***
	(0.383)	(0.374)	(0.410)	(0.312)	(0.391)	(0.445)	(0.397)
Registration Rate	April2011	Constant	Constant	Constant	Constant	Constant	Growth
Extra Controls*Post			X			X	
Observations	25893	25893	25893	25893	25893	25893	25893

Notes: The dependent variable is a binary indicator equal to one if the incumbent mayor is reelected (mean 0.453), except in Column (3), First Stage, where it is the ghost building registration rate. **Registration Rate** refers to alternative definitions of the ghost-buildings registration rate we use in the analysis. The variable is equal to "April 2011" when we use the registration rate at April 2011 and introduce an election year fixed effects in the model. It is equal to "Constant" when we impute the registration rate at the time of the election using the April 2011 rate as a starting point and assuming a constant yearly registration rate. It is equal to "Growth" when we impute the registration rate at the time of the election using the April 2011 rate as a starting point and assuming yearly registration rate grows by 50% every year. The variable **Years Elapsed since Program Inception** takes value 0 if the elections occurs before the program. **Post** is a binary indicator equal to one if the election occurs after the Ghost Buildings program inception. In all the columns, *Post* is instrumented by *Province Post*, a binary indicator equal to one if the election occurs after the modal program inception year in the province. **Ghost Building Intensity** is defined as the ratio between the number of land registry parcels with ghost buildings and the total number of land registry parcels. **Extra Controls*Post** include *Geographic*, *Socio-Economic*, and *Mayor* Controls interacted with the *Post* dummy. Refer to Table 1.1 for a list of these variables. All the regressions include town fixed effects, election-year fixed effects and an interaction between macro-areas fixed effects and *Post*. The regression sample includes all the elections between 1993 and 2011 in which the incumbent does not face a binding term-limit. Standard errors are clustered at provincial level. *p<0.1, **p<0.05, ***p<0.01.

We then provide empirical support for two channels affecting the electoral response to the program. Our simple theoretical framework predicted that this should be higher in towns where the local government delivers public goods more effectively and where the non-monetary returns to tax enforcement are higher. We provide evidence about these hypotheses by estimating Equation 1.9. The coefficient of interest δ_3 captures the impact of a standard-deviation increase in the variables measuring either the municipal speed of public good provision or the tolerance for tax evasion on the electoral response to the program. Table 1.9 presents the results. In column (1) we find that a one standard deviation increase in the speed of public good provision increases the point estimate of the impact of ghost buildings on reelection by 0.63, and that this coefficient is statistically significant at the ten percent level. We then confirm that this interaction effect does not simply capture geographical variation in the responsiveness across different parts of Italy by adding triple interactions across the post-program indicator, the ghost building intensity, and the macro area dummies. The sign and economic significance of the coefficient is robust, though estimated less precisely (p=.137).

We then look at the role of tax culture. We exploit variations across regions in the extent to which respondents "justify tax cheating" in the *European Values Study*. These results provide clear evidence the tax culture matters. In Column (3), we show that a one standard deviation increase in the tolerance score reduces the point estimate of the impact of ghost buildings on reelection by .64, (significant at the ten percent level). Column (4) shows that the magnitude of the coefficient is stable, or if anything increases (in absolute value) when adding the triple interactions with macro-areas dummies. These results provide suggestive evidence that the positive effect of the Ghost Buildings program on incumbent reelection likelihood is larger in localities where voters have, on average, stronger preferences for tax enforcement and where the delivery of public goods is more effective.

Finally, Table 1.10 presents the results of the estimation of the baseline regression model in Equation 1.7, using the log of town-level government expenditures. Column (1) presents the reduced-form results, using the post-program indicator based on the provincial mode.

Table 1.9: *Ghost Building Intensity and Incumbent Reelection: Heterogeneity Analysis*

	(1)	(2)	(3)	(4)
Ghost Building Intensity*Post	1.183***	1.232*	1.063***	1.311**
	(0.394)	(0.684)	(0.380)	(0.668)
...*Speed of Public Good Provision	0.627*	0.591		
	(0.380)	(0.398)		
...*Justify Tax Cheating			-0.639*	-0.734*
			(0.364)	(0.404)
GBI*Macro Area*Post	No	Yes	No	Yes
Observations	25812	25812	25893	25893

Notes: The dependent variable is a binary indicator equal to one if the incumbent mayor is reelected (mean 0.453). **Post** is a binary indicator equal to one if the election occurs after the Ghost Buildings program inception. In all the columns, *Post* is instrumented by *Province Post*, a binary indicator equal to one if the election occurs after the modal program inception year in the province. **Ghost Building Intensity** is defined as the ratio between the number of land registry parcels with ghost buildings and the total number of land registry parcels. **GBI*Macro Area*Post** is the triple interaction among macro-areas fixed effect, ghost building intensity and *Post* All the regressions include town fixed effects, election-year fixed effects, interactions between macro-areas fixed effects and *Post*, and an interaction between the relevant heterogeneity variable for the column and *Post*. The regression sample includes all the elections between 1993 and 2011 in which the incumbent does not face a binding term-limit. Standard errors are clustered at provincial level. *p<0.1, **p<0.05, ***p<0.01.

The point estimate is .436 (significant at 10%). The coefficient is stable when instrumenting the post-program indicator with the provincial one and it is slightly larger when including interactions among controls and the post-program indicator.[46] While the effect of the program is statistically significant, we also note that it is fairly small. A one standard-deviation increase in ghost-building intensity increases expenditures by about 1%. We argue that it is unlikely that this effect explains the whole incumbent reelection effect we documented earlier in the paper. Consistently with the suggestive evidence provided by the heterogeneity by tax culture, we suggest that non-monetary factors (e.g., the direct utility non-evaders derive from catching of the shirkers) must play an important role.

[46] Appendix Figure A.1 shows that government expenditures satisfy the parallel trend assumption, and pre-program coefficients are not statistically different from zero.

Table 1.10: *Local Government Expenditures*

	OLS	2SLS	
	(1)	(2)	(3)
Ghost Buildings Intensity * Post		0.488*	0.607***
		(0.264)	(0.207)
Ghost Building Intensity*Post Province	0.436*		
	(0.252)		
Extra Controls*Post	No	No	Yes
Observations	74646	74646	74646

The dependent variables is the natural logarithm of town government expenditures.*Notes:* **Post** is a binary indicator equal to one if the election occurs after the Ghost Buildings program inception. In all the columns, *Post* is instrumented by *Province Post*, a binary indicator equal to one if the election occurs after the modal program inception year in the province. **Ghost Building Intensity** is defined as the ratio between the number of land registry parcels with ghost buildings and the total number of land registry parcels. **Extra Controls*Post** include *Geographic*, *Socio-Economic*, and *Mayor* Controls interacted with the *Post* dummy. Refer to Table 1.1 for a list of these variables. All the regressions include town fixed effects, election-year fixed effects and an interaction between macro-areas fixed effects and *Post*. Standard errors are clustered at provincial level. * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

1.6.4 Alternative Explanations

Finally, we use the entire set of our results to argue that the impact on incumbent reelection probability arising from the increase in tax enforcement more than offsets several alternative potential explanations about the impact of the Ghost Buildings program on voter support for the local incumbent.

According to the first of these alternative explanations, the publication of the number of ghost buildings generates information about the incumbent. We believe this to be both unlikely and inconsistent with our findings. First, the set of ghost buildings is a slow moving stock variable that is likely to have accumulated over decades, rather than a reflection of just the most recent years. Most of the buildings found by the *Agenzia del Territorio* were not newly constructed. The existence of a term limit, paired with the fact that the time to complete a building in Italy is generally longer than most of the other OECD countries, suggests that most of these buildings could not have been built while the incumbent was in office. Second, we notice that voters who could potentially receive information from the publication are most likely the ones who were not evading before the program, as evaders already knew about their own evasion.

Keeping this premise in mind, our results rule out this alternative explanation. In one version of this alternative story, voters, after learning about *low* levels of evasion detected by the program, reward the current mayor for having properly enforced tax payment in the past. This hypothesis predicts a *negative* impact of the detected ghost building intensity on incumbent reelection in post-program elections, and as such it is obviously inconsistent with our baseline results. In another version of this alternative explanation, voters reward an incumbent mayor for having allowed *high* levels of evasion in the past. First, this contradicts the intuition that non evaders, rather than those previously evading, are the ones who are potentially acquiring new information. Second, it is unlikely since the purpose of the program, and therefore the publication, was to shut down the evasion opportunity. Third, it is at odds with the fact that the positive impact of program intensity on incumbent reelection is lower in regions with higher tolerance for tax evasion. Fourth, it is inconsistent with our results showing that towns with higher registration levels are more likely, rather than less likely, to reelect an incumbent mayor.

In a second potential alternative explanation, the program gives an incumbent an electoral rent by allowing her *to not register* the targeted ghost buildings, for instance by reporting errors in the results generated by the mapping process.[47] If this were the case, we would expect the positive impact of the program to be stronger in regions with higher tolerance for tax evasion. We find the opposite to be the case. In addition, such an explanation is inconsistent with the result that a higher share of registered ghost buildings at the time of a local election increases reelection likelihood. The results of this section provide strong evidence that it is the additional tax enforcement induced by the program that drives the increase in the reelection prospects of the incumbent.

[47]For example, the press agency of one of the mayor of a city in our sample, Capaccio Paestum, explicitly criticized the "excessive media attention" to the program, indicating how the unregistered buildings in that city were due to "citizens' needs" (Comune di Capaccio Paestum, 2010).

1.7 Conclusion

A rapidly growing literature shows that interventions that improve the "technology" of tax enforcement — third party reporting, cross-checking, or better auditing algorithms — can substantially reduce tax evasion. Yet, political incentives to adopt these technologies are also of crucial importance. Policymakers will delay or prevent enforcement policies if they are bound to lose support from them. In spite of this, little is known about the electoral impact of fighting tax evasion. This paper provides evidence of a positive interaction between technological improvements in tax payer monitoring and political incentives. Specifically, local incumbents are shown to obtain positive political returns — an increase in their reelection likelihood — from the Ghost Building program, a nationwide anti tax evasion policy in Italy which was based on a new enforcement technology.

Underlying tax culture — broadly defined as the individual propensity and social norms determining evasion for a given level of technology — is another important determinant of tax compliance. It shapes the enforcement level a government can achieve for a given enforcement technology. We show that tax culture affects the political returns to undertaking anti tax evasion policies. The increase in incumbent reelection probabilities in response to the Ghost Buildings program is larger in areas with lower self-reported tolerance for tax evasion. Finally, we document that the political returns to enforcement policies are higher when the government in more efficient in public goods provision.

The findings of the paper have two important policy implications. First, they provide a framework for thinking about the political feasibility of policies that increase visibility of tax evasion, thus lowering the monitoring costs and increasing policymakers' incentives for raising enforcement. This has immediate relevance for special interest politics. Concentrated evader groups might effectively lobby to keep evasion hidden from the public. Yet, they are unlikely to be able to punish an incumbent who enforces tax compliance after the evasion becomes broadly visible.

Second, there is potential complementarity among anti tax evasion policies, government responsiveness, and social preferences for tax compliance. Governments that plan to implement novel enforcement policies should concurrently attempt to strengthen their capabilities, for instance by improving the speed at which they respond to citizen's needs, or by increasing social stigma associated with tax evasion. This complementarity will likely increase the returns politicians obtain from anti tax evasion policies and will thus make such policies more aligned with political agent incentives.

We are aware that using an identification strategy based on a specific natural experiment enhances internal validity of our study but might come at the price of lower external validity, for instance for extrapolating about similar programs in other countries or programs targeting other taxes. Yet, we speculate that evidence of positive political returns to anti tax evasion policies in Italy, a country often cited as an example of poor tax culture, will be a lower bound for other OECD countries. We believe an interesting goal for future work is to elucidate the potential non-linearity in the relation between tax evasion prevalence and political returns to enforcement policies. In addition, we believe that complementarity between enforcement policies and social norms on evasion could potentially be relevant for policy design in other regions of the world.

Another important dimension of external validity concerns enforcement policies targeting other types of evasion. One of the merits of the Ghost Buildings program is that it detected the entire stock of evasion. On the other hand, the effectiveness of policies targeting other tax-concealing activities might vary by the ability of the specific evader to hide. This might affect how the public would respond. We hope future work will shed light on the political returns to other enforcement policies around the world.

Chapter 2

Rural Roads and Intermediated Trade: Regression Discontinuity Evidence from Sierra Leone[1]

2.1 Introduction

In spite of the central relevance of infrastructure in the development policy debate and in foreign aid, there has been little rigorous research on the role of local infrastructure investments in developing countries, especially in the context of very poor economies. Over the last two decades, governments and donors in Sub-Saharan Africa have devoted considerable resources to rural road construction and rehabilitation, of which a large share was used to upgrade feeder roads that link up small localities with each other or to larger roads.[2] The rationale behind these investments is that good feeder roads, while expensive, are central to the integration of markets, primarily because, by reducing transport costs,

[1] Co-authored with Rachel Glennerster and Tavneet Suri

[2] Since 2000, major feeder roads rehabilitation programs have been implemented in Cameroon, the DRC, Ghana, Mozambique, Sierra Leone, and South Sudan, among other countries. Carruthers et al. (2010) document that it would cost about 2% of African GDP every year for ten years to reach some reasonable targets on improved transportation.

they improve the access of farmers to markets for their crops. The existing body of research on the impact of rural feeder roads is plagued by identification problems, and usually does not provide compelling causal evidence on the impact of infrastructure improvement (van de Walle (2009)).

In addition, the empirical literature on rural roads, for the most part, has not explicitly dealt with the role of the underlying market structure in shaping the response to these infrastructure investments. Agricultural markets throughout Sub-Saharan Africa are characterized by high levels of fragmentation and poor transport infrastructure, with intermediaries playing a central role (Fafchamps et al. (2004)). Different models of trader competition and intermediation generate different predictions about the price response to an improvement in rural road quality and how this response varies with market characteristics. The empirical analysis of these price responses can therefore provide a test for competing theoretical frameworks of the market structure in agricultural trade. Understanding which of these theoretical frameworks best explains the nature of competition amongst these agricultural intermediaries can shed light on the impact of other supply chain interventions, such as subsidies to different agents of the value chain or exporting promotion policies.

In this paper, we make two contributions. First, we provide empirical evidence on the impact of rural road rehabilitation on transport costs and rural market prices. We focus on a specific program in Sierra Leone where the road selection algorithm allows us to use a Regression Discontinuity Design (RDD), and we use novel data collection strategies to measure changes in prices and transport costs. Sierra Leone is an economy with low population density where feeder roads could potentially have large impacts. Rural markets in the economy are not well integrated, and farmers and traders both travel long distances to engage in trade. Transaction costs are extremely high but transport costs are only a part of these transaction costs. These characteristics are common to many rural African economies (Fafchamps and Hill (2008), Fafchamps et al (2004)) and other developing countries (World Bank, 2009). As feeder roads are commonly advocated as a key policy intervention in these contexts, our first objective is to understand their impact. Well identified measures of feeder road impacts are particularly important given the high costs of effectively connecting such dispersed populations.

As a second contribution, we use our empirical results to test between alternative models of competition in agricultural markets. We focus on Bertrand oligopsony, bilateral bargaining, Cournot oligopsony (with and without endogenous entry), and a basic Mortensen (2003) search model. For each of these frameworks, we derive the equilibrium price in rural markets and provide comparative statics on the price effects of a change in rural road transport costs. We then show how this price effect varies by market characteristics, such as the distance to the main urban centers and the agricultural productivity in the surrounding areas. To the extent that the four classes of models deliver different predictions on these comparative statics, we can use our empirical results to test between these models. It is an open question as to what the relevant model of market structure for traders in agriculture is in developing economies. As we describe below, market structure is important to understanding what policies are most effective at improving market access and, more broadly, welfare for farmers in Sub Saharan Africa.

Our empirical analysis focuses on a feeder road rehabilitation program in Sierra Leone that was financed by the European Union (EU) and implemented in four districts that cover 27% of the country's area and 30% of its population.[3] Sierra Leone's infrastructure is extremely poor: in 2002, at the end of a decade-long civil war, only 8% of the country's 11,300 km of roads were paved. Agriculture is the country's largest employer, with 64% of households farming, of which 87% produce rice, the main staple. The infrastructure network is generally described as inadequate to support well-integrated agricultural markets.[4]

The project targeted local dirt feeder roads, measuring 20 km on average, linking local markets to villages or to a more important road. They did not connect major cities in different parts of the country together. To structure the use of its funds, the EU created a priority ranking for each of 47 eligible roads based on an index of quantitative economic

[3] Sierra Leone National Census (2004).

[4] See the Agricultural Household Tracking Survey (AHTS) Report (2011) and Pushak and Foster (2011).

data. The highest-ranked roads were chosen in order until at least as close to a total of 150 km of roads had been assigned to each district. A total of 31 roads were ultimately selected. This method of fund allocation allows us to use a road-level RDD to study the impacts of the feeder road improvements.

Our empirical results can be summarized as follows. First, the rehabilitation program did improve road quality on the roads selected. Using data on transport fares and GIS video stream data collected on the roads, we show that transport costs fell significantly. Second, from trader surveys and monthly price surveys, we find that an improvement in rural road quality leads to a reduction in the prices of rice and cassava (the two main staples produced domestically) in rural markets along the rehabilitated roads.[5] Third, we find that this price reduction is stronger in markets that are farther away from main urban centers and is weaker in markets that are located in more productive areas.

Which of the standard classes of models best explains these findings? We look across a set of simple models which assume the city is a small open economy, traders are homogeneous and their transport costs are linear. In these models, rural roads can potentially affect transport costs for both buyers traveling from urban areas (primarily traders) and sellers (farmers) since both travel to rural marketplaces. In our analysis, we refer to the impact of the rural road rehabilitation on urban traders as the "demand effect", and the effect on producers as the "supply effect", with the former driving up prices in rural markets, and the latter driving them down. The relative importance of these two effects varies across the theoretical frameworks we present. In particular, both the Bertrand framework and the Cournot model with endogenous entry predict that reduced transport costs should be associated unambiguously with *increases* in prices in rural markets (i.e. the demand effect always dominates). However, this is inconsistent with our empirical findings - we find that improvements in road quality that reduce transport costs on average *reduce* prices. On the other hand, the remaining frameworks (bilateral bargaining, Cournot with an exogenous

[5]Rural markets in our sample are medium-sized marketplaces in which several crops are traded (rice, cassava, palm oil in particular). In Section 2.3 we describe in detail our sampling strategy for markets and traders within markets.

number of traders, and the search frictions framework) predict a role for both demand and supply effects.

Second, the models have different predictions regarding the heterogeneity in the price effect of reduced rural transports costs along two dimensions: (i) the distance between the rural markets and major urban areas, and (ii) farmer productivity. Both the bilateral bargaining model[6] and the Cournot oligopsony case with an exogenous number of traders predict that market characteristics have no impacts on these price effects, which is inconsistent with our results. On the other hand, a basic Mortensen search framework predicts that the interactions of the price effect with distance and productivity are non-zero because the magnitude of both the demand and the supply effect varies with these two market characteristics. Specifically, for the markets that are further away, or the markets in areas with lower productivity, the relevance of search costs in determining equilibrium prices is higher. In turn, this implies that the negative price effects of the improved roads should be stronger for these sets of markets. Overall, our empirical results are consistent only with a search framework and inconsistent with other models.

In support of this conclusion, we provide further evidence on search costs using data on cell phone penetration. We find that our price effects are muted in areas where there is better cell phone penetration, which is what one would expect in the presence of search costs. In particular, we expect that cell phone penetration lowers search frictions. Therefore, the prices responses in areas with higher cell phone penetration will be closer to the Bertrand case, i.e. less negative or even positive. We find evidence of this, consistent with the findings in Allen (2012), Aker (2010) and Jensen (2007).

These results have important policy-implications. First, the impact of improved road infrastructure varies by the characteristics of the road's location. Features such as productivity and linkages between urban consumers and traders can affect not just the magnitude of rural market price responses, but also their sign. This implies that the benefits of improvements

[6] The bilateral bargaining model represents the case where a given producer is locked into a relationship with one particular trader. It can also be interpreted as a model with search costs approaching infinity, so that producers have no outside option, outside of the existing relationship.

in infrastructure are heterogeneous, which is important for policy makers when deciding where to make such improvements. As in Suri (2011), where the returns to agricultural technologies are heterogeneous and some of this heterogeneity arises from differences in access to the rural road infrastructure network, returns to different road projects will vary sharply with underlying farmer heterogeneity. Second, the finding that agricultural intermediary markets in this setting are best characterized by a framework that includes search frictions has implications for other policies designed to influence agents in the agricultural value chain. These policies include price subsidies, agricultural export promotion interventions, credit-provision policies targeting traders, and more major road projects. Third, the empirical evidence in support of the presence of search frictions suggests a possible complementarity between hard infrastructure projects and other interventions aimed at reducing search costs, such as the introduction of marketing boards or the extension of mobile phone coverage in rural areas (for example, mobile phone penetration rates are only 36% in Sierra Leone (World Development Indicators, 2012)).

This paper's findings are consistent with recent empirical evidence on search frictions (see, for example, Jensen (2007), Aker (2010) and Goyal (2010)). While these studies focus on the price impact of a reduction in information frictions, we show that the response to improved transport infrastructure also depends on the presence of search frictions. From this standpoint, our approach is similar to that of Allen (2012) who shows that, in the presence of search costs, the elasticity of trade flows to destination prices depends on supply parameters, such as producer heterogeneity. In addition, our work draws from three other strands of literature. First, a small but influential set of papers has used rigorous identification strategies to shed light on the impact of large transport infrastructure improvements (examples include Michaels (2008), Donaldson (2012), Datta (2012), Faber (2012) and Banerjee, Duflo and Qian (2012)). We look at rural roads and not highways - van de Walle (2009) provides a good review of the literature on rural roads.[7] Second,

[7]Examples include Jacoby and Minten (2008), Dorosh et al. (2010), Gibson and Rozelle (2003), Ali (2011), Khandker, Bakht and Koolwal (2009), Khandker and Koolwal (2011), and Mu and van de Walle (2007).

a recent growing literature uses micro-data to estimate the degree of price pass-through internationally and domestically.[8] Third, our theory focus is motivated by the recent emphasis on trade intermediation (for example, Bardhan, Mookherjee, and Tsumagari (2012), Antràs and Costinot (2011)). In particular, we rely heavily on Chau, Goto and Kanbur (2009) in our theoretical setup and in the way we model search frictions.

The rest of the paper is structured as follows. In Section 2.2, we outline the theoretical frameworks and derive the comparative statics of interest. Section 2.3 describes some background on Sierra Leone, the EU road rehabilitation program and the data sources. Section 2.4 outlines the empirical strategy and tests the validity of the RDD for our setting. Section 2.5 presents the results of the empirical analysis, including a comparison of our results with the theoretical predictions of the different models of competition and results from a variety of robustness checks. Section 3.6 concludes.

2.2 Theoretical Frameworks

In this section, we present four different theories of trade intermediation between agricultural producers and traders. Traders play a central role in the rural economies of Sub-Saharan African by channeling product between crop-producing areas in the countryside and the consumers living in urban areas (Fafchamps et al. (2004)). We focus on rural markets where traders purchase and sell a number of different crops. Traders are mostly small scale, traveling from one market to another before transporting crops to the capital city. We describe these traders in more detail below.

We focus on four broad classes of theoretical models, representing different market structures: Bertrand competition, bilateral Nash bargaining, Cournot oligopsony (both with and without endogenous entry) and search frictions à la Mortensen (2003). The Bertrand case and the search frictions case, as well as the theoretical setup, are based on Chau, Goto and Kanbur (2009). We adapt their framework in three directions to fit our research questions.

[8] See, for example, Broda and Weinstein (2008), Burstein and Jaimovich (2009), Gopinath et al. (2011), Borraz et al (2012) and Atkin and Donaldson (2012).

First, we explicitly model the role of rural roads in linking producers and traders to rural market places. Second, we introduce a productivity parameter that varies across villages. Third, we model two additional cases: a specific case of bilateral bargaining between traders and producers and the case of Cournot competition (with and without endogenous entry).

The primary goal is to derive key comparative statics to be tested in our empirical framework. We focus specifically on the price effects of an improvement in rural road infrastructure, and on the heterogeneity in this effect across market characteristics.

2.2.1 Setup

We model the transactions that occur between traders and rural agricultural producers in rural markets to which both types of agents travel.[9] We make two important simplifying assumptions here. The first is that traders are homogeneous - this seems to be a reasonable interpretation of our data. For example, 97% of traders are male, 88% are from two main ethnic groups, 65% have no education, 86% are married, 96% report that they started trading on their own, 62% own a mobile phone, 96% own a radio and only 19 traders of local rice own their own mode of transport (motorbike, car or truck - unfortunately the survey did not distinguish between these). The second assumption we make is that utility is linear: both farmers and traders maximize their (expected) profits.

Around each market, there is an exogenous number of producers, N, who produce σ units of a certain crop. The opportunity cost of each unit (for instance the utility from consuming each unit) is c. In order to reach the local market, rural producers use rural roads to travel distance α. The unit transport cost on the rural road, τ, is therefore the inverse of a measure of road quality. When road quality improves, the unit transport cost on the road decreases.

City traders travel to the local market to purchase the crop which they can then sell in

[9]Throughout, we use the term producers to refer to either farmers or “aggregators", where aggregators refer to larger farmers who aggregate product from other farmers in a village to bring it to the market to sell. The theoretical models we present below all hold whether it is farmers themselves that come to the market or aggregators.

the urban areas at an exogenous price p^*. The exogeneity of the city price p^* is based on the intuition that the city receives the crop from many different markets across the country. Each market is thus a small open economy and changes in the transport costs between a given village and the surrounding urban areas are assumed not to affect the city price or prices in other rural markets. In order to reach the market, the traders bear two types of transport costs: first, a "major road" transport cost, x, and second, a rural road cost $\beta\tau$. The "net city price" is thus $p^* - x - \beta\tau$. Throughout, we assume this form of linear transport costs as well as a separation between the major road and the rural road transport costs for the trader. The linearity seems a natural assumption for the transport costs since the main form of transport for traders is renting motorbikes, which price a fare per distance travelled. The market for transport along these roads (major and rural) is reasonably competitive. Very few traders actually own a mode of transport - in our trader data only 19 traders owned any form of transport. There is therefore little scope for a (discrete) investment that simultaneously affects transport costs on both the main road and the rural road. In addition, the improvement in the rural roads due to the EU program causes a reduction in transport costs but not a big enough reduction to allow traders to invest in their own mode of transport. Finally, this assumption is consistent with Fafchamps et al (2004) who find that there are no returns to scale at this level of traders.

We denote the (endogenous) sale price in the rural market by p. In addition to city traders, we assume that farmers can sell to local consumers (non-farming rural households) at cost-recovery price $p^0 = c + \alpha\tau$. Throughout the paper, we assume there are gains from trade, that is: $p^* - x - \beta\tau > c + \alpha\tau$. The assumption that there are gains to trade implies that there is no extensive margin effect of changes in τ on traders or producers entering the market.

Below, we discuss equilibrium prices under four alternative market structures. For each of these, we are interested in three comparative statics:

1. The equilibrium price change in response to a change in rural road transport costs: $\frac{\partial p}{\partial \tau}$

2. Heterogeneity in this price response by the distance travelled on the major road, x:

$\frac{\partial^2 p}{\partial\tau\partial x}$

3. Heterogeneity in this price response by market-level productivity, σ: $\frac{\partial^2 p}{\partial\tau\partial\sigma}$

2.2.2 Bertrand Oligopsony

Bertrand competition is characterized by free entry of city traders into each village and perfect information. This implies that all the producers are matched to city traders and that competition drives up the equilibrium price, p^B to equalize the net city price:

$$p^B = p^* - x - \beta\tau \tag{2.1}$$

It is easy to see that producers appropriate all the gains from trade. Looking at the predictions of the Bertrand model with regards to the three comparative statics of interest, we obtain:

1. $\frac{\partial p^B}{\partial\tau} = -\beta < 0$: rural road transport costs enter the equilibrium price only via their effect on traders' transport costs.

2. $\frac{\partial^2 p^B}{\partial\tau\partial x} = 0$: as the trader cost function is separable between the rural road and the major road costs, the price response to rural road costs is not affected by the major road transport costs.

3. $\frac{\partial^2 p^B}{\partial\tau\partial\sigma} = 0$: productivity, or any other supply characteristic, does not affect the equilibrium price, and hence does not affect the price response to a reduction in rural transport costs.

2.2.3 Bilateral Bargaining with Lock In

We now present a simple model of prices under a specific model of bilateral bargaining in which each producer can only trade with one specific trader, with an outside option limited to the cost-recovery price $c + \alpha\tau$. In this scenario, the transactions equilibrium price, p^N, is assumed to be the generalized Nash-Bargaining solution:

$$p^N = \arg\max_p \{[\sigma(p^* - x - \beta\tau - p)]^\gamma * [\sigma(p - c - \alpha\tau)]^{1-\gamma}\}$$
$$= \gamma(c + \alpha\tau) + (1-\gamma)(p^* - x - \beta\tau), \tag{2.2}$$

where $\sigma(p^* - x - \beta\tau - p)$ is the trader's surplus, $\sigma(p - c - \alpha\tau)$ is the producer's surplus, and γ is the bargaining weight for the trader. The Bertrand outcome is a special case of the Nash Bargaining solution for $\gamma = 0$. In addition, the bilateral bargaining setup implies that, in the presence of a positive surplus from trade, all the producers are matched to traders.

The model delivers the following predictions with regards to our three key comparative statics:

1. $\frac{\partial p^N}{\partial \tau} = \gamma\alpha - (1-\gamma)\beta \gtrless 0$: with positive α and β, rural road transport costs now enter prices through their effect both on producers' and on traders' costs. The former drives down prices (what we term the supply effect), while the latter raises prices (what we term the demand effect).

2. $\frac{\partial^2 p^N}{\partial \tau \partial x} = 0$, as in the Bertrand case.

3. $\frac{\partial^2 p^N}{\partial \tau \partial \sigma} = 0$: the amount transacted, σ, does not affect how the price responds to a change in τ, even though σ does enter the equilibrium price.

2.2.4 Cournot Oligopsony

The third framework we present is one of Cournot competition. Here, we introduce an additional assumption to generate an upward sloping supply curve. We assume that the parameter α (the distance travelled by the producer on the rural feeder road) is a random variable uniformly distributed over $[\alpha^* - z, \alpha^* + z]$. Producers sell in the market only if the equilibrium price p is larger than their reservation price: $p > c + \alpha\tau$. This extensive margin on selling generates an upward sloping supply curve as only a share $F(\frac{p-c}{\tau}) = \frac{(p-c)/\tau - (\alpha^* - z)}{2z}$ of producers enters the market, where $F(\cdot)$ is the cdf of the uniform distribution.[10]

[10] We also studied the case where farmers are heterogeneous in the opportunity cost, c, instead of being heterogeneous in the parameter α. Since all the comparative statics for this case collapse to the Bertrand case, we do not report it here.

We study two versions of the Cournot model. First, we look at the case with an exogenous number of traders operating in the market, M. Second, we endogenize the number of entrants using a free entry condition, where we assume entrants have a fixed cost of entry K.

We focus on symmetric equilibria. With an exogenous number of traders, each trader chooses his optimal quantity given other traders' quantities which enter the profit functions through the (inverse) supply curve. After imposing symmetry, the equilibrium price, p^C is

$$p^C = \frac{c + \tau(\alpha^* - z) + M(p^* - x - \beta\tau)}{1 + M} \tag{2.3}$$

The Cournot model with exogenous entry delivers the following results with regards to our comparative statics of interest:

1. $\frac{\partial p^C}{\partial \tau} = \frac{\alpha^* - z - M\beta}{1+M} \gtrless 0$: both the supply and the demand effect of a change in τ enter the derivative of price with respect to transport costs.

2. $\frac{\partial^2 p^C}{\partial \tau \partial x} = 0$: while x shifts the net city price, it does not interact with τ in determining the equilibrium price.

3. $\frac{\partial^2 p^C}{\partial \tau \partial \sigma} = 0$: σ does not enter the solution for the equilibrium price, p^C.

We now extend the basic Cournot model above to allow for endogenous trader entry. By equating individual profits with the cost of entry, we find the equilibrium number of traders operating in the economy and then derive the corresponding equilibrium price, p^{CE}:

$$p^{CE} = p^* - x - \beta\tau - \sqrt{\frac{2\tau K}{Nz\sigma}}, \tag{2.4}$$

The comparative statics of interest are now:

1. $\frac{\partial p^{CE}}{\partial \tau} = -\beta - \sqrt{\frac{K}{2\tau Nz\sigma}} < 0$: τ affects p^{CE} through its impact on both demand and supply. The demand effect operates though β, as before. However, the reduction in aggregate supply generates a more than proportional reduction in the number of traders entering the economy. As a result, the price reduction in response to an increase in τ is stronger than in the benchmark Bertrand case. Intuitively, an increase

in τ has two effects. First, it reduces the "net city price", $p^* - x - \beta\tau$. All else equal, this leads traders to reduce p. Second, it reduces the elasticity of supply, $Q = \frac{\sigma N(p-c)}{\tau}$. All else equal, this induces traders to *increase* the mark down, thereby contributing further to a decrease in price.

2. $\frac{\partial^2 p^{CE}}{\partial\tau\partial x} = 0$: terms that only enter the net city price do not interact with transport costs in determining the equilibrium price.

3. $\frac{\partial^2 p^{CE}}{\partial\tau\partial\sigma} = \frac{\sqrt{K}}{2\sigma\sqrt{2rNz\sigma}} > 0$: by increasing aggregate supply, σ reduces the role of fixed costs in determining the equilibrium price. An economy with a higher productivity will be closer to the perfectly competitive case relative to one with a lower σ. Thus, an increase in σ implies a smaller in absolute value (i.e. closer to $-\beta$) price response to an increase in τ since the impact of τ on the elasticity of supply is lower.

Summarizing these results, the Cournot model with exogenous entry cannot be distinguished from the bilateral bargaining model based solely on the sign of the comparative statics. On the other hand, the version with endogenous entry delivers a unique set of predictions relative to the previous models.

2.2.5 Search Frictions

We now introduce search frictions in the model economy. Empirically, we think search costs may arise from two potential mechanisms. A first potential mechanism for search costs is that these markets are difficult to reach. Farmers and traders often travel on motorbikes to these markets and although transport is available, the timing is uncertain. This generates some search frictions as farmers and traders may not end up in a given rural market at the same time. There may therefore be waiting costs as well as uncertainty of what will be available in the market by the time the trader gets there. This is the type of search cost considered in Fafchamps and Hill (2008).

Second, farmers develop relationships with traders over time. There is a lot of evidence for such contracts both in Sierra Leone (see Casaburi and Reed (2012) for an example) as well as in other developing countries (see Deb and Suri (2012)). These relationships operate as search frictions in the sense that there are switching costs for farmers to shift across traders and this can generate market power for the traders. This is one explanation given by Fafchamps and Minten (2011) as to why a price information intervention did not have impacts on prices in India. These relational contracts may exist to enforce quality standards (see Bardhan, Mookherjee and Tsumagari (2012)) or due to the existence of trade credit as in the case of Casaburi and Reed (2012). The existence of such relationships between farmers and traders would generate search frictions of the sort we model here.

To construct our model, we closely follow Chau, Goto, and Kanbur (2009) and we refer readers to their paper for detailed derivations of the results. Their paper is in turn based on the static search framework derived by Mortensen (2003), which provides the key results. We choose a static framework to facilitate the introduction of search frictions in the previously described setup.

The interactions between traders and producers now occur in three stages. In the first stage, traders decide whether to enter a certain market. If they enter, they incur an entry cost, K, which includes bargaining time, travel fixed costs, waiting time and uncertainty. In the second stage, traders who entered the market choose a price offer. Due to the search frictions, only one random producer receives the offer. In the third stage, producers take up the best offer among those received. For a large enough market, the distribution of offers received by each producer follows a Poisson with mean $\lambda \equiv M/N$, where M is the endogenous number of traders entering the market:

$$Pr(z;\lambda) = e^{-\lambda}\frac{\lambda^z}{z!} \tag{2.5}$$

Thus, the cumulative distribution of the highest price offer received by each producer, or equivalently, from the trader's perspective, the probability of outbidding the second-best offer is:

$$H(p) = \sum_{z=0}^{\infty} Pr(z;\lambda)F(p)^z = \exp(1 - F(p)) \tag{2.6}$$

where $F(p)$ is the endogenous cdf of price offers made by the traders.

The trader's expected profit maximization implies that no symmetric pure strategy equilibrium exists. Rather, traders follow a mixed strategy where prices are in the support $[p_l, p_h]$, given by:

$$[c + \alpha\tau, (1 - e^{-\lambda})(p^* - x - \beta\tau) + e^{-\lambda}(c + \alpha\tau)] \tag{2.7}$$

and are drawn according to distribution:

$$F(p) = \frac{1}{\lambda} \ln \left(\frac{p^* - x - \beta\tau - c - \alpha\tau}{p^* - x - \beta\tau - p} \right) \tag{2.8}$$

Each price in the support leads to an equal expected profit. The free entry condition can therefore be written as $E[\pi(p_l)] = K$, which allows us to solve for the equilibrium ratio of traders to producers, λ:

$$\lambda^* = \ln \left(\frac{p^* - x - \beta\tau - c - \alpha\tau}{K/\sigma} \right) \tag{2.9}$$

We substitute λ^* into the optimal trader's price offer distribution, $F(p)$, in (8) and solve for both the expected price for farmers conditional on receiving at least one offer from traders and for the unconditional expected price, p^S. The latter can be written as:

$$p^S = p^* - x - \beta\tau - \frac{K}{\sigma} \left[1 + \ln(p^* - x - \beta\tau - c - \alpha\tau) - \ln \left(\frac{K}{\sigma} \right) \right] \tag{2.10}$$

We focus on the comparative statics from this expected unconditional price equilibrium. The model delivers the following predictions with respect to our three comparative statics:

1. $\frac{\partial p^S}{\partial \tau} = -\beta \left(1 - \frac{(K/\sigma)}{p^* - x - \beta r - c - \alpha\tau} \right) + \frac{(K/\sigma)\alpha}{p^* - x - \beta\tau - c - \alpha\tau} \gtrless 0$, where the first term represents the demand effect, and the second the supply effect, of a change in transport costs. The magnitude of the demand effect is lower in absolute value than the demand effect in the Bertrand case: traders' market power induced by search frictions implies an imperfect pass-through of traders' cost shocks.

2. $\frac{\partial^2 p^S}{\partial \tau \partial x} = \frac{K(\alpha + \beta)}{\sigma(p^* - x - \beta\tau - c - \alpha\tau)^2} > 0$: in locations that are farther away for the city, a lower net price induces lower entry, more monopsony power and stronger deviations from the Bertrand benchmark. Thus, if x is higher, the (negative) price effect of a higher transport cost for traders is weaker and the (positive) price effect of a higher transport

cost for producers is stronger.

3. $\frac{\partial^2 p^S}{\partial\tau\partial\sigma} = -\frac{K(\alpha+\beta)}{\sigma^2(p^*-x-\beta\tau-c-\alpha\tau)} < 0$: intuitively, an increase in σ lowers the "real" entry cost K/σ, bringing the economy closer to the benchmark competitive Bertrand case. A higher σ therefore moves the equilibrium both toward a stronger (negative) demand effect and toward a weaker (positive) supply effect of an increase in τ.

2.2.6 Summary of Theoretical Results

Table 2.1 summarizes the results from the models presented in this section. With regard to our comparative statics of interest, the models can be differentiated easily in two ways. First, some models predict that a decrease in τ will unambiguously lead to higher equilibrium prices, while others predict an ambiguous effect. This is due to the fact that a change in τ can induce either a demand effect alone, or both a demand and a supply effect. Second, the Cournot model with exogenous entry and the search friction frameworks are the only models predicting that the magnitude of the price response will depend on the market-specific features, x and σ. In section 2.5, we compare these predictions to our empirical results. In addition, we will test whether these price effects vary by cell phone penetration, something we may expect in a world with search frictions.

2.3 Background and Data Sources

2.3.1 The EU Rural Feeder Roads Rehabilitation program

Background and Implementation

The EU feeder roads rehabilitation program was designed to contribute to the reconstruction of Sierra Leone's infrastructure in the aftermath of a destructive civil war (1991-2002). The EU's program was implemented between 2009 to 2011 at a total cost of EUR 9.5 million (USD 13 million) - approximately EUR 16,000 per kilometer of rehabilitated road. The program targeted four districts in three different provinces: Kambia and Port Loko

Table 2.1: *Theory Comparative Statics and Empirical Results*

Comp. Statics	**Bertrand**	**Bilateral Bargaining**	**Cournot Exogenous M**	**Cournot Endogenous M**	**Search Frictions**	**Empirical Results**	**Regression Coefficient ($-\frac{\partial p}{\partial \tau}$)**
$\frac{\partial p}{\partial \tau}$	< 0	$\gtreqless 0$	$\gtreqless 0$	< 0	$\gtreqless 0$	> 0	Treat < 0
$\frac{\partial^2 p}{\partial \tau \partial x}$	$= 0$	=0	=0	$= 0$	> 0	> 0	Treat*Distance< 0
$\frac{\partial^2 p}{\partial \tau \partial \sigma}$	$= 0$	$= 0$	$= 0$	> 0	< 0	< 0	Treat*Harvest per hh> 0; Treat*% Selling> 0

(Northern Province), Kenema (Eastern Province) and Pujehun (Southern Province). The roads in the rehabilitation program primarily connected small towns and markets to one another or to a major road, rather than linking big cities and regions to one another. Because these feeder roads do not dramatically alter access to population centers, they are unlikely to modify migration patterns in the areas surrounding the rehabilitated roads. The roads are also unlikely to affect agricultural technology adoption because improved technologies such as fertilizer and improved seeds are rarely available even in urban centers, as reflected in extremely low baseline rates of technology adoption.[11] Rehabilitation work began in mid-2009 and ended by early 2011 for all the roads. In the final construction progress report, dated August 2011, 25 of the 31 roads selected for rehabilitation were described as fully rehabilitated and 6 roads as partially rehabilitated.[12]

Program Design

The rehabilitation program was designed in a way that allows for a road-level Regression Discontinuity Design (RDD) analysis. In 2003, field investigations with local stakeholders led to the identification of a base list of 47 rural roads eligible for rehabilitation, totaling about 800 km. The task of prioritizing 600 kilometers of roads for rehabilitation from amongst this larger sample was given to an external consultant, Edward Davies & Associates (EDA). EDA (2004) provides extensive details of the prioritization process that was used to decide which of these roads would be rehabilitated. The roads were ranked according to a score, which was a weighted average of five components:

[11] Roughly 5% of farmers use fertilizer nationally (AHTS (2011)) and much of this use is centered around the capital Freetown and the national agricultural research station. In separate work, Suri (2011) found that deficient infrastructure and differential access to good infrastructure were among the reasons explaining low technology adoption rates in Kenya. In our context, the small size of the roads considered makes it reasonable to assume that the supply of agricultural technology from the main urban centers to the countryside did not dramatically increase as a result of the program.

[12] The average completion rate for the six roads partially rehabilitated was 64%. This percentage does not reflect the fraction of the total length of the road that had been rehabilitated. Instead, it illustrates what fraction of the planned improvements were completed on the road i.e. the average completion across all specifications of the work - in many cases, improvements were implemented along the entire length of the road, but not all the planned improvements were fully completed.

i) **Economic Production per kilometer**, defined on the basis of survey measures of the volume of crops produced, income from economic activities and mode of transportation;

ii) **Population per kilometer** within the area of influence of each road;

iii) **Road Assessment**, a 1 to 5 score measuring the pre-existing condition of the selected roads, based on seven parameters (culverts, bridges, drainage, pavement surface, vertical alignment, horizontal alignment, and riding quality);

iv) **Social Value**, a 1 to 5 score, based on the number of schools, health centers, wells and toilets in the catchment area of the road;

v) **Length**.

The data on these components was compiled by EDA. Each component of the score was normalized by its district-level maximum value. The final weighted sum of each of the five components had weights of 0.4, 0.2, 0.2, 0.1 and 0.1 for components (i) to (v) above, respectively.[13] The decision rule was that, in each district, roads would be rehabilitated starting with the highest-priority one (based on the score), following in order of decreasing score until as close to 150 km of road as possible was allocated to be rehabilitated in each district. Since th roads could not be split to get exactly 150 km, this means that in some districts slightly over 150 km of road was rehabilitated, and in some districts just under 150 km was rehabilitated. Following this rule, 30 roads out of the eligible 47 should have been selected across the four districts. Figure **??** presents maps of the priority roads in the initial list across the four districts, as well as the connections between these roads and the major roads in the country.

After collecting data from the EU, EDA, and the Sierra Leone Roads Authority (SLRA) to verify that the RDD had been followed, we found evidence of two potential manipulations that occurred around the cutoff. In the district of Kenema, the final list of roads rehabilitated covered 184 km, while selecting one less road would have brought the total rehabilitated in this district to 153 km, closer to the approximate desired cutoff of 150 km. One additional road was therefore rehabilitated although it was not supposed to be. We treat this as a fuzzy

[13]Appendix B.1 provides the exact formula used to compute the score.

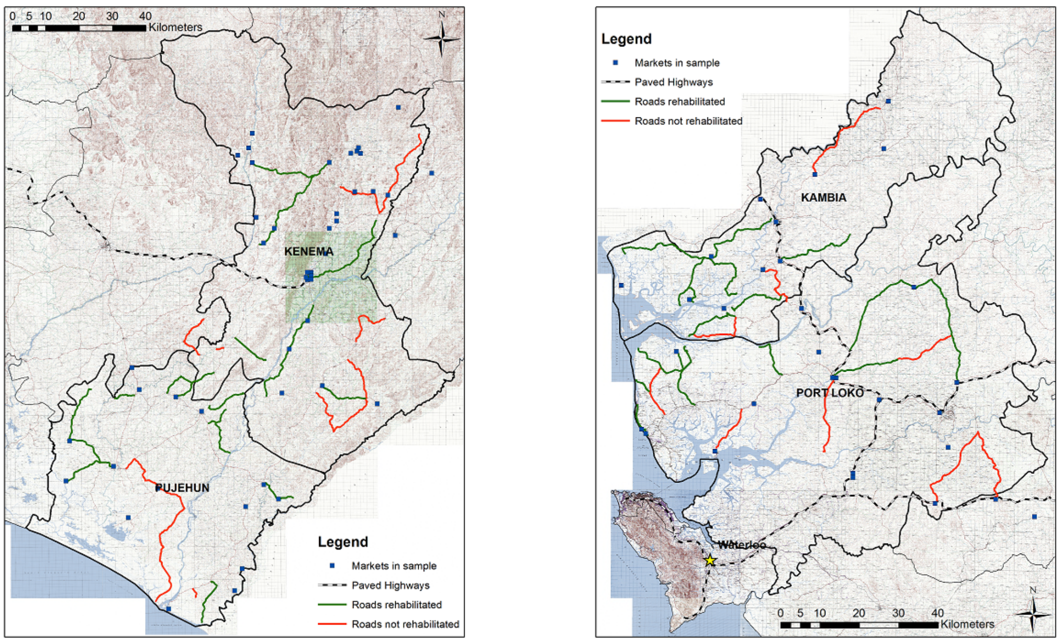

Figure 2.1: *District Maps of Sierra Leone*

RDD where the additional treated road is allocated to the control group. In Kambia District, two roads had the same score and were ranked exactly at the cutoff point. Only one of the two roads was rehabilitated to reach the appropriate total, but no rules were established for how roads should be selected in the case of a tie. In this case, we drop the two roads from our baseline specifications, though we check the robustness of our results to their inclusion. We discuss the implications of these issues for our identification strategy in more detail in Section 4 below.

2.3.2 Data

This section provides an overview of the data used in our empirical framework, ordered chronologically by collection date.

Baseline Roads Data (2003-2004)

To construct the scoring index, EDA collected data around all 47 eligible roads before any rehabilitation began. We use this data to check for baseline differences between the roads above and below the cutoff in our RDD framework. We include all five components of the score, as well as a number of road characteristics, such as the number of bridges, palmlog bridges[14] and culverts,[15] data for which was provided by the EU. Table 2.2 shows summary statistics on the baseline condition of the roads considered for rehabilitation. The average length of the roads was 21 km. In this table we also report the average rescaled score, which is the value of the index created by the EU, rescaled so that zero is the selection cutoff score between the first road below and the first road above in each district.

[14] A palmlog bridge is a makeshift bridge made from logs of palm trees. Regular cars are usually unable to cross such bridges, but four-wheel drive vehicles and motorbikes often can, although at times with considerable risk. During the rainy season, palmlog bridges get flooded, making the wood rot and sometimes causing the entire bridge to collapse.

[15] A culvert is a pipe or a drain that lets water flow below a road. Culverts are much smaller than bridges and are often used to create crossings over small waterways or as a way to improve drainage in locations where rainy seasons are extremely intense (as they are in Sierra Leone).

Table 2.2: *Baseline Summary Statistics*

	Mean	SD	Obs
EA-Level 2004 Census Data			
Proportion farming (all individuals over 10 years)	0.68	0.36	1373
Proportion farming (all households)	0.78	0.32	1373
Proportion literate (all individuals over 12 years)	0.31	0.19	1374
Years of education (HH heads)	2.86	2.26	1369
Roads Baseline Data			
Log population per km	6.53	0.93	45
Log economic production per km	18.60	1.01	45
Road assessment	17.24	3.81	45
Social value	9.56	2.15	45
Number of bridges (excluding palm logs)	3.21	4.07	45
Number of palm log crossings	2.80	5.01	45
Number of culverts	18.29	17.05	45
Road length (km)	20.61	14.72	45
Score rescaled at midpoint	0.08	0.17	45
Chiefdom-Level AHTS Data			
Proportion of HHs selling rice	0.20	0.13	34
Proportion of HHs selling cassava	0.37	0.20	34
Log HH rice harvest (kg)	6.15	0.51	34
Log HH cassava harvest (kg)	7.18	0.53	34
Roads Endline: First Stage			
Total travel fare (one way)	16100	12085	45
Travel fare per km (one way)	900	558	45
Average speed (kph)	36.26	10.42	44
Trader Data			
Local rice purchase price	774.84	90.08	149
Local rice sale price	931.14	105.46	118
Market-Level Price Survey Data			
Price of cassava/gari	339.71	99.92	787
Price of local rice	996.51	143.06	764
Cassava/gari available (proportion of mkts)	0.95	0.21	918
Local rice available (proportion of mkts)	0.84	0.37	918
Number of traders per market (cassava/gari)	9.37	11.31	918
Number of traders per market (local rice)	9.05	9.90	918
Log distance (km) to nearest large town	2.60	1.40	78

Note: All prices are in SLL (Sierra Leonean Leones). The relevant exchange rate was approximately $1=4300 SLL. Total travel fare is for a motorbike ride along the entire distance of the road. Outliers (top 1% of observations) are removed from price data. All prices are in cups unless otherwise indicated. Conversion rates are approximately 4 (5) cups of rice (cassava) per kg.

National Population Census (2004)

To test our RDD assumptions, we use data from the last population census conducted in Sierra Leone, which was completed in 2004. From the census data, our variables of interest are the fraction of households involved in crop farming, the fraction of individuals over ten years old involved in crop farming, the literacy rate and the years of education completed by household heads.

Agricultural Households Tracking Survey (2010)

To identify areas of high productivity or high surplus for different crops, we use measures of agricultural production at the road level, computed from the Agricultural Households Tracking Survey (AHTS). The AHTS was a nationally representative agricultural survey conducted in March-April 2010 with a sample of 8,803 farming households. Since the AHTS was potentially administered concurrently with the implementation of the EU program, we do not use it as baseline data. Instead we only use the household-level data to compute measures of local production around the markets for rice and cassava. We also use measures of the fractions of households selling rice and cassava as well as household cell phone ownership. We compute chiefdom-level averages, as discussed in more detail below.

Endline Roads Data (2011)

We designed novel GIS data collection strategies to verify whether the roads had indeed been rehabilitated, and to collect objective measures of transport costs. First, we drove a sport utility vehicle with a geo-referenced camera secured to the hood along each road. This exercise was completed for 46 out of the 47 eligible roads in November 2011.[16] The camera recorded the GPS position and speed of the vehicle every second, as well as collected a continuous stream of video along all roads. In the analysis presented below, we use average speed on the roads traveled as one of our outcomes of interest.

[16] One road was missed as part of this exercise because it could not be located. The road is 2.6 km long and was not selected for rehabilitation.

In another effort to gather evidence of the impact of the roads on transport, we collected data on transport fares along all 47 roads. We negotiated with motorbike taxi riders (locally known as *okadas*) for a route fare on every road,[17] and travelled with the *okada* on a random subset of these roads. *Okadas* are the most commonly available type of public transportation in rural Sierra Leone, for both people and goods, and unlike buses they can travel on most roads all year round. These data give us a measure of actual transport costs and freight rates on the sample of roads in the study.

We show summary statistics on these transport variables in Table 2.2. The average speed on the roads was about 21 km per hour and the motorbike fare was an average of 16,000 Leones per road ($3.72). Since the roads are of different lengths we report the fare per km travelled, which is about 900 Leones ($0.21) per km.

Trader Surveys (2011-2012)

To understand the effects of these road improvements on prices, we use data from two waves of a nationwide survey targeting rice traders in rural markets. The first wave, conducted in February/March 2011, targeted all the markets located within five kilometers of the 35 roads that were closest to the rehabilitation threshold in each of the four districts. In addition, random sampling of the remaining markets located in the rest of the country led to a sample of 54 markets located within 11 kilometers of the 47 roads. The second wave, conducted in January/February 2012, included the universe of markets (82) located within 11 kilometers of any of the 47 roads, including the 54 markets sampled for the first wave.[18] Within each market, we first listed all traders of local rice, then randomly sampled two traders to participate in the survey (or surveyed the unique trader when only one trader was present).

[17]This exercise was designed so that the surveyor bargaining with the taxi was the residual claimant.

[18]The threshold of 11km was defined before the second wave of the survey to maximize the number of markets surveyed under tight survey budget constraints. Within our sampling frame of markets, 54 were located within 5 kilometers of the nearest road, 78 were located within 10 kilometers, and 86 within 11 kilometers. Out of the 86 markets visited, 4 had no rice trader at the time of the survey, leading to the final number of 82 markets.

Table 2.2 shows summary statistics from these trader surveys. Sale prices from traders' most recent sales transactions are presented for local rice, but not for cassava, which was not covered by this particular survey. We report prices in cups, as this is the standard unit across all our price data.[19] The trader surveys were conducted during the harvest season (February/March), a period when local rice is plentiful and the price of local rice is lower than its annual average.

The GPS coordinates of the markets are used to compute as-the-crow-flies distances between rural markets and larger towns or urban centers in Sierra Leone. We define a large town to be either one of the six largest towns in the country based on the 2004 census (these towns were Freetown, Bo, Kenema, Makeni, Koidu, Waterloo and Lunsar, with a median population across these towns of 81,000 individuals) or any district headquarter town across all the 13 districts in the country. This gives us a total of 14 unique large towns, however, our markets only match to a subset of nine of those towns.

High Frequency Price Surveys (2011-2012)

Finally, between May 2011 and July 2012, we targeted the markets included in the trader survey sample to collect monthly price data via phone surveys. We think of the prices collected as the average price in the market for that month. The respondents were individuals identified as focal points for the markets during the data collection of the trader surveys. The number of markets for which we sought monthly pricing data increased after the tenth wave of calls to include the 28 additional markets targeted by the second wave of the trader survey. We focus on two main crops in our analysis of the high frequency data: local rice and a common type of processed cassava known locally as *gari*.[20] Table 2 shows summary statistics on the prices of these two crops, again reported in cups. The price of a

[19]There are roughly four cups of local rice per kilogram. Two varieties of local rice are sold by traders - parboiled and milled. These varieties are nearly identical to each other. We averaged these prices to obtain one local rice price.

[20]Gari is a form of processed cassava similar in aspect to kuskus or bulghur, common throughout most of West Africa. Gari is obtained from peeling and grating cassava tubers often using manual equipment.

cup of gari is about 340 Leones ($0.08) and that of a cup of rice about 1,000 Leones ($0.23).

2.4 Empirical Strategy

2.4.1 Regression Discontinuity Design

The design of the EU program allows us to use an RDD[21] to identify the causal effect of rural road rehabilitation on transport costs and rural market prices.[22] We have a small number of clusters for the RDD, though a large number of observations, especially when using the high frequency price survey data. This RDD therefore resembles the geographic or boundary RDD studies (for example, see Dell (2010)). Later, we also conduct a number of robustness checks to deal with the limitations on the number of clusters.

In what follows, we describe the identification problem using the standard Rubin (1974) potential outcomes framework. Define $y_i(1)$ as the potential outcome for road i when the matched road is rehabilitated by the EU program and $y_i(0)$ when it is not rehabilitated. Then the (sub) population parameter of interest is $\Delta = E[y_i(1) - y_i(0)]$, the average treatment effect (ATE). The basic identification problem arises because the econometrician only observes the realized outcome based on the actual rehabilitation status T_i of the road:

$$y_i = T_i \cdot y_i(1) + (1 - T_i) \cdot y_i(0) \tag{2.11}$$

An OLS framework would lead to an inconsistent estimate of Δ in the presence of unobservable covariates that are correlated both with the outcome and with the rehabilitation

[21] We refer the reader to Imbens and Lemieux (2007) and Lee and Lemieux (2010) for extensive reviews of the RDD.

[22] Our main outcomes of interest are measured either at the road level (transport costs) or at the market level (market prices). In addition, since we do not have baseline (pre-intervention) data for market prices, we use instead road level data collected from household surveys in Enumeration Area (EAs) around the project roads during the 2004 Census. We also use Chiefdom level data collected from an agricultural household survey for the heterogeneous treatment effects analysis. For the sake of clarity, in this section we only refer to our outcomes as being road-level or market-level. However, the reader should bear in mind that some of the variables used for baseline checks were actually collected at the household level and appropriately aggregated.

status. With multiple rounds of data, a difference-in-difference estimator would be biased if changes in the rehabilitation status are correlated with changes in other unobservable variables.[23]

Our RDD relies on the comparison between roads "just above" and roads "just below" the rehabilitation cutoff. In the case of perfect correspondence between the rehabilitation plan and the actual rehabilitation status, this empirical strategy would identify the average treatment effect (ATE) around the cutoff. This ATE is local in the sense that results around the treatment effect at the cutoff cannot be generalized to other points in the domain of the forcing variable. In the context of this study, we argue that our empirical design identifies the marginal effect of a program expansion, which is a policy relevant local average treatment effect.

We define S_i^N as the road score, normalized to zero at the mid-point between the first road below the cutoff and the first road above the cutoff.[24] Our empirical analysis starts with a graphical approach. For each of the variables of interest, we plot the bin level means of the outcome and include a linear fit of all the underlying data. We restrict the graphs to be within our chosen main specification bandwidth (see below) so that the graphs are easily comparable to the parametric results. For the parametric estimation, we restrict our attention to roads that are "close" to the cutoff, i.e. the sample of roads whose score is within h points from the cutoff ($S^N \in [-h, h]$). We use one main bandwidth as our preferred specification, $h = 0.15$, which delivers a subsample of 31 roads. This bandwidth of $h = 0.15$ is close to the Imbens and Kalyanaraman (2009) optimal bandwidth across a range of our outcomes at the road-level. As Lee and Lemieux (2010) mention, optimal bandwidth algorithms may suggest bandwidths that are larger than the rule of thumbs used by many researchers. Since we have a small sample size, we choose the optimal IK bandwidth as our main specification. We also report results in our main tables for two additional bandwidths, ($h = 0.075$ and

[23]Studies using a differences-in-differences approach include Ali (2011), Bakht, Khandker and Koolwal (2009), Khandker and Koolwal (2011), Bell and van Dillen (2012), and Mu and Van de Walle (2007)).

[24]Our results are similar if we normalize the cutoff to be at the first treated road in each district. These results are available upon request.

$h = 0.3$) which span half and double the preferred bandwidth, and deliver subsamples of 18 and 38 roads, respectively. For our transport cost variables, where our outcomes are at the road level, we do not report results for the $h = 0.075$ bandiwdth as these specifications only have 18 observations and are therefore not particularly meaningful.

For the preferred bandwidth and the two alternatives, we run a local-linear regression of the form:

$$y_{ik} = \gamma_0 + \gamma T_i + \gamma_R T_i * S_i^N + \gamma_L(1 - T_i) * S_i^N + \delta_k + \epsilon_{ik} \tag{2.12}$$

where y_{ik} is a road-level outcome, i denotes the road and k the district. We control for district fixed effects in the road level specifications and present heteroskedasticity robust standard errors.

For market-level outcomes (prices), the regression model is:

$$y_{ijkt} = \gamma_0 + \gamma T_i + \gamma_R T_i * S_i^N + \gamma_L(1 - T_i) * S_i^N + \delta_k + \mu_t + \epsilon_{ijkt} \tag{2.13}$$

where i denotes the road that the market is close to, j is the market, k is the district (we control for both the district in which the road is located as well as for the district the market is in, if different) and t is the survey round. For the trader survey, our regressions are at the trader level and we have multiple traders per market.

In addition to the local-linear regressions, we use the full sample of 47 roads and include a third-order polynomial approximation in S_i^N. We use third order polynomials though we also show robustness to this.

When running the market-level regressions, we adjust our standard errors for two way clustering, allowing for clustering at the road level as well as at the market level. This is important since some markets match to multiple roads. There are 82 unique markets in our surveys. On average, we have about 6 markets matching to each road. Approximately 48% of markets match to just one road. The others match to two or more roads, with the average being about 2.8 roads per market. We also weight the regressions by the inverse of the distance between the road and the market in question.

We look not only at the treatment effects of the improved roads but also the heterogeneity in these effects along four specific dimensions: distance between the market and the closest town, harvest of the relevant crop (rice or cassava), the fraction of households selling the relevant crop and cell phone penetration. For the heterogeneity analysis we use a dummy variable for whether the market is above the median value for the relevant variable. We therefore report results from these specifications as well as specifications where we control for these heterogeneity variables.

2.4.2 Testing the Validity of the Identification Strategy

The above RDD relies on the assumption that there is no manipulation of the theoretical rehabilitation status around the threshold. One potential challenge to identification therefore comes from the following two compliance issues briefly discussed in Section 2.3.1:

1. In Kenema district, there is a discrepancy between the rehabilitation plan and the actual status - an additional road was treated. This implies that the average treatment effect of the actual rehabilitation status cannot be estimated. Instead, we present the (local) Intent-To-Treat (ITT) estimator.[25]

2. In Kambia district, the same score was assigned to two roads, but only one road was rehabilitated without a clear rule for the event of a tie break. To complicate matters, across the whole sample, these are the two roads closest to the cutoff on either side, giving them the highest weight in a classic RDD. Since we lack a clear ranking protocol for this specific case, we present results after dropping these two roads from the sample.[26]

In addition, we adopt two standard strategies to test for the presence of manipulation. First, we test for discontinuity in the density of the forcing variable (McCrary (2008)). Second,

[25] LATE results are available on request.

[26] This is akin to the recent donut RDD (see, for example, Barreca et al. (2011)). Notice that we include these roads in one of the robustness checks later in Section 2.5.2.

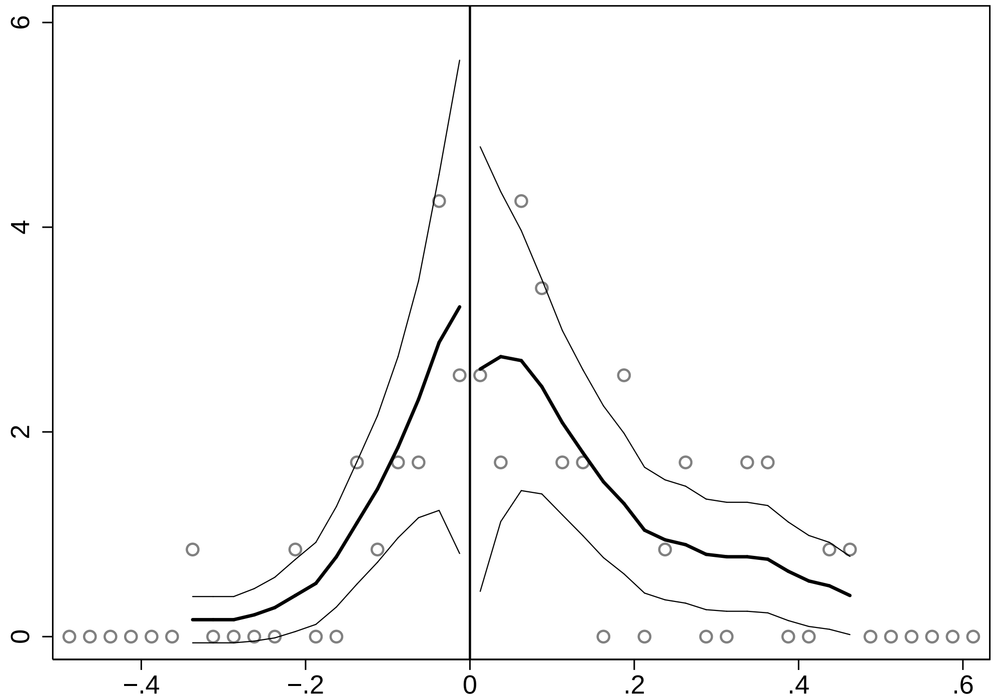

Figure 2.2: *McCrary (2008) Density Test*

we inspect whether baseline covariates show discontinuities around the score cuttoff. Figure **??** presents the results of the McCrary test.[27] The observed discontinuity is well within the confidence interval. In addition, if anything, the jump is downward while any manipulation would predict a jump upward. Finally, given the low number of observations at the extreme values or tails, the density is not well estimated at these tails. Our density test is, therefore, likely to have low power.

The analysis of covariate balance at the cutoff relies on three data sources. For a subset of variables (the most relevant ones), we also show graphical results. Table 2.3 focuses on variables collected by the EU as part of their baseline assessment of the roads and directly entering the scoring index. We look at just the population per km graphically - see Figure **??**. Both the figure and Table 2.3 show no significant discontinuity at the cutoff (only one

[27]For our test we use a binsize of 0.025 and a bandwidth of 0.15. Our estimate of discontinuity at the cutoff is -0.325 with a standard error of 0.684.

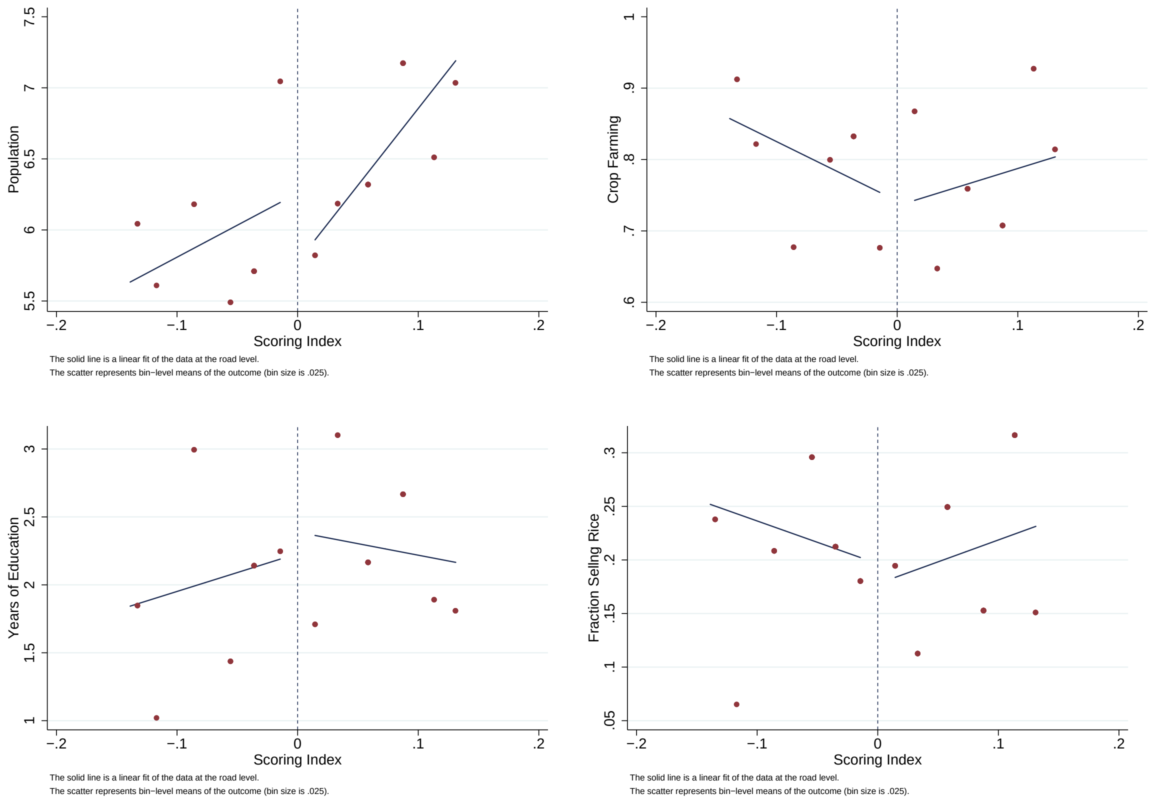

Figure 2.3: *Baseline Analysis*

coefficient of the twenty is significant).

Table 2.4 presents other baseline characteristics of the roads, including variables collected by EDA during the prioritization process that do not enter the score, and the agricultural and education variables collected during the 2004 Census. Unfortunately, the census has limited agricultural information but we present results for the available variables - the fraction of households involved in agriculture, the fraction of individuals involved in agriculture, literacy and education of the household head. We keep EAs within 2 km of the roads and we aggregate the data to the road level for these tests. We also graphically show results for the fraction of individuals in farming and the education of the household head in Figure **??**. Neither the figures nor the regressions show evidence of discontinuities at the cutoff - in Table 2.4, none of the 28 coefficients are significant.

Table 2.3: *Pre-Rehabilitation Analysis: Score Components*

	(1) Population per km	(2) Production per km	(3) Road Assessment	(4) Social Value	(5) Road Length
Preferred LLR (h=.15)					
Treatment	-0.537	0.342	-0.072	-1.008	15.694
	[0.409]	[0.328]	[2.264]	[1.599]	[17.433]
Mean for Control	5.977	17.863	17.143	8.857	23.199
Observations	31	31	31	31	31
LLR (h=.075)					
Treatment	-1.578***	0.006	2.897	-0.229	11.099
	[0.430]	[0.553]	[3.310]	[3.038]	[28.657]
Mean for Control	5.958	17.681	18.667	9.444	22.943
Observations	18	18	18	18	18
LLR (h=.3)					
Treatment	-0.213	0.232	1.641	-0.457	4.988
	[0.314]	[0.223]	[1.644]	[1.137]	[10.770]
Mean for Control	5.967	17.902	17.067	8.733	22.306
Observations	38	38	38	38	38
3rd Order Polynomial					
Treatment	-0.789	0.085	0.862	-1.065	24.238
	[0.530]	[0.488]	[2.929]	[2.356]	[23.118]
Mean for Control	5.935	17.889	16.750	8.750	21.074
Observations	45	45	45	45	45

Note: Road assessment is a 1-5 score assigned by EDA that measures the pre-existing condition of the selected road based on seven parameters (culverts, bridges, drainage, pavement surface, vertical/horizontal alignment, and riding quality).
Social value is a 1-5 score assigned by EDA based on the conditions of schools, health centers, wells and toilets in the catchment area of the road.
LLR is short for local linear regression.
All specifications include district fixed effects.
Heteroskedasticity-robust standard errors reported in brackets.
* $p<0.1$, ** $p<0.05$, *** $p<0.01$

Table 2.4: *Pre-Rehabilitation Analysis: Road Characteristics*

	(1) Number of Bridges	(2) Number of Palmlog Bridges	(3) Number of Culverts	(4) Prop of HHs Crop Farming	(5) Prop of Indiv Crop Farming	(6) Literate HH Head Dummy	(7) HH Head Years of Education
Preferred LLR (h=.15)							
Treatment	2.801	-2.242	4.647	-0.036	0.003	0.040	0.284
	[5.263]	[4.466]	[11.778]	[0.097]	[0.125]	[0.059]	[0.582]
Mean for Control	3.857	4.552	15.430	0.880	0.794	0.239	2.056
Observations	31	31	31	31	31	31	31
LLR (h=.075)							
Treatment	0.578	-5.051	-9.375	0.100	0.190	-0.059	-0.249
	[11.257]	[4.421]	[9.420]	[0.136]	[0.172]	[0.067]	[0.694]
Mean for Control	4.778	3.525	14.336	0.890	0.790	0.240	2.008
Observations	18	18	18	18	18	18	18
LLR (h=.3)							
Treatment	0.130	-3.645	2.613	-0.003	0.001	0.033	0.212
	[3.499]	[3.427]	[8.724]	[0.070]	[0.090]	[0.039]	[0.398]
Mean for Control	3.800	4.248	15.735	0.886	0.804	0.238	2.035
Observations	38	38	38	38	38	38	38
3rd Order Polynomial							
Treatment	2.287	3.863	-0.263	-0.065	-0.001	0.038	0.640
	[7.704]	[6.487]	[15.785]	[0.149]	[0.185]	[0.090]	[0.880]
Mean for Control	3.563	3.983	15.001	0.890	0.810	0.237	2.021
Observations	45	45	45	45	45	45	45

Note: Columns (4)-(7) based on 2004 Census data. A culvert is a pipe or drain allowing water to flow under a road.LLR is short for local linear regression. All specifications include district fixed effects. Heteroskedasticity-robust standard errors reported in brackets. * p<0.1, ** p<0.05, *** p<0.01

In Table 2.5, we focus on the variables used for the heterogeneity analysis. This includes four variables collected during the AHTS in early 2010 as proxies of σ, the distance from the market to the closest urban town (as a proxy for x in the heterogeneity analysis) and a measure of cell phone penetration, also from the AHTS. As proxies of σ, we use the average rice and cassava harvests per household and the fractions of households selling rice and cassava.[28] Although these variables were not collected at baseline in the strict sense, the harvest variables can be considered predetermined as the relevant farmer decisions (such as acreage planted and input use) were made the previous year, prior to any road rehabilitation. Because harvest data may be noisy, we also look at the extensive margin sale choice as a proxy for σ. Since the majority of the population relies on subsistence farming for the staple crops, the share of producers who enter the market economy is a reasonable measure of the average surplus of rice. As a measure of cell phone penetration, we use the fraction of households that own a cell phone. We use chiefdom-level averages of the agricultural and ell phone variables for the chiefdom in which a given market is located. Chiefdom-level data are less subject to endogeneity concerns since the roads typically considered affect a small fraction of the population of each chiefdom.[29]

Table 2.5 reports the results. In general, we do not find much evidence of discontinuities in these variables. For three of the six variables, one specification shows an effect that is both economically and statistically significant (for example, in the tight bandwidth linear specification, the fraction of households selling cassava is more than 50% lower in the treatment group than in the control group). The polynomial specification is responsible for two of the three problematic coefficients. Due to the small sample size in these cells, we think that the polynomial is not a good approximation to the data as it overfits leading to an overestimate of the magnitude of the coefficients. Graphically, Figure **??** shows results for the fraction of households selling rice with little evidence of a discontinuity.

[28] For this table, we use the continuous version of these variables. In the heterogeneity analysis, we use dummies for whether the value of the variable is above or below the median, as described later in the paper.

[29] Chiefdoms are the third-level administrative unit in Sierra Leone, coming after provinces and districts. There are 149 chiefdoms throughout the country.

Table 2.5: *Pre-Rehabilitation Analysis: Farming Activity and Harvests*

	(1) Log Market Distance to Big Town	(2) Chiefdom Frac of HHs Selling Rice	(3) Chiefdom Frac of HHs Selling Cassava	(4) Log Rice Harvest	(5) Log Cassava Harvest	(6) Chiefdom Phone Access
Preferred LLR (h=.15)						
Treatment	-0.518	-0.031	-0.063	-0.005	0.166	-0.062
	[0.341]	[0.027]	[0.061]	[0.168]	[0.112]	[0.054]
Mean for Control	9.707	0.218	0.389	6.197	7.307	0.353
Observations	104	65	65	65	65	65
LLR (h=.075)						
Treatment	-0.670	0.007	-0.256***	-0.197	0.060	-0.059
	[0.623]	[0.054]	[0.067]	[0.219]	[0.280]	[0.104]
Mean for Control	9.857	0.223	0.405	6.202	7.253	0.360
Observations	64	41	41	41	41	41
LLR (h=.3)						
Treatment	-0.702*	-0.027	-0.016	-0.040	0.081	-0.029
	[0.365]	[0.021]	[0.045]	[0.157]	[0.103]	[0.033]
Mean for Control	9.707	0.218	0.389	6.197	7.307	0.353
Observations	129	77	77	77	77	77
3rd Order Polynomial						
Treatment	-2.292**	-0.018	-0.165	-0.727**	-0.314	-0.066
	[1.058]	[0.036]	[0.168]	[0.285]	[0.322]	[0.091]
Mean for Control	9.707	0.218	0.389	6.197	7.307	0.353
Observations	156	91	91	91	91	91

Note: All data is from the 2010 AHTS.
Distance data (km) is at the market-road level. Harvest data (kg) is at the chiefdom-road level.
LLR is short for local linear regression.
All specifications include road-district and market-district fixed effects.
Standard errors two-way clustered by road and market reported in brackets in column (1).
Standard errors two-way clustered by road and chiefdom reported in brackets in columns (2)-(5).
* $p<0.1$, ** $p<0.05$, *** $p<0.01$

2.5 Results

We now discuss our empirical results. We first present estimates of the effects of the road improvements on transport costs and speed, and then for the comparative statics for rural market prices that we derived in Section 2.2. The comparative statics presented in the theory section focus on a change in crop prices with respect to a change in τ, the rural transport costs. The empirical results are for an *improvement* in road quality that *reduces* rural transport costs.

2.5.1 Speed and Transport Costs

Figure **??** presents a graphical analysis of the impact of the program on our two primary measures of transport costs: travel speed and (log) travel fares per km. As mentioned above, for the graphs, we show the data within the main bandwidth, $h = 0.15$ - we show bin level means for the data and a linear fit of all the data. Table 2.6 shows the results from the parametric regressions on these variables according to the specifications outlined in Section 2.4. However, the results from the $h = 0.075$ bandwidth for these road level outcomes are not meaningful as the sample size is too small and given the number of controls for the RDD specification, we are left with too few degrees of freedom. Looking at the impact of road quality improvements on speed, the bin means show a jump at the discontinuity (see Figure **??**), but the flexible polynomial specification is driven up at the left of the discontinuity. This is due to the fact that the last road on the left of the cutoff was rehabilitated even though it should not have been. The parametric analysis in Table 4 confirms this intuition: the wide ($h = 0.3$) bandwidth local linear specifications point at an increase in speed across the cutoff. The results are also robust to the inclusion of controls we later use for the heterogeneity analysis. The polynomial is less robust, but is likely to be overfitting the data. The discontinuity in transport fares is visible (though noisy) in Figure **??** and in Table 2.6: an improvement in road quality has a significantly negative impact on fares. In the local linear regression with the optimal bandwidth ($h = 0.15$), the road improvements lead to a 59% reduction in transport costs per km.

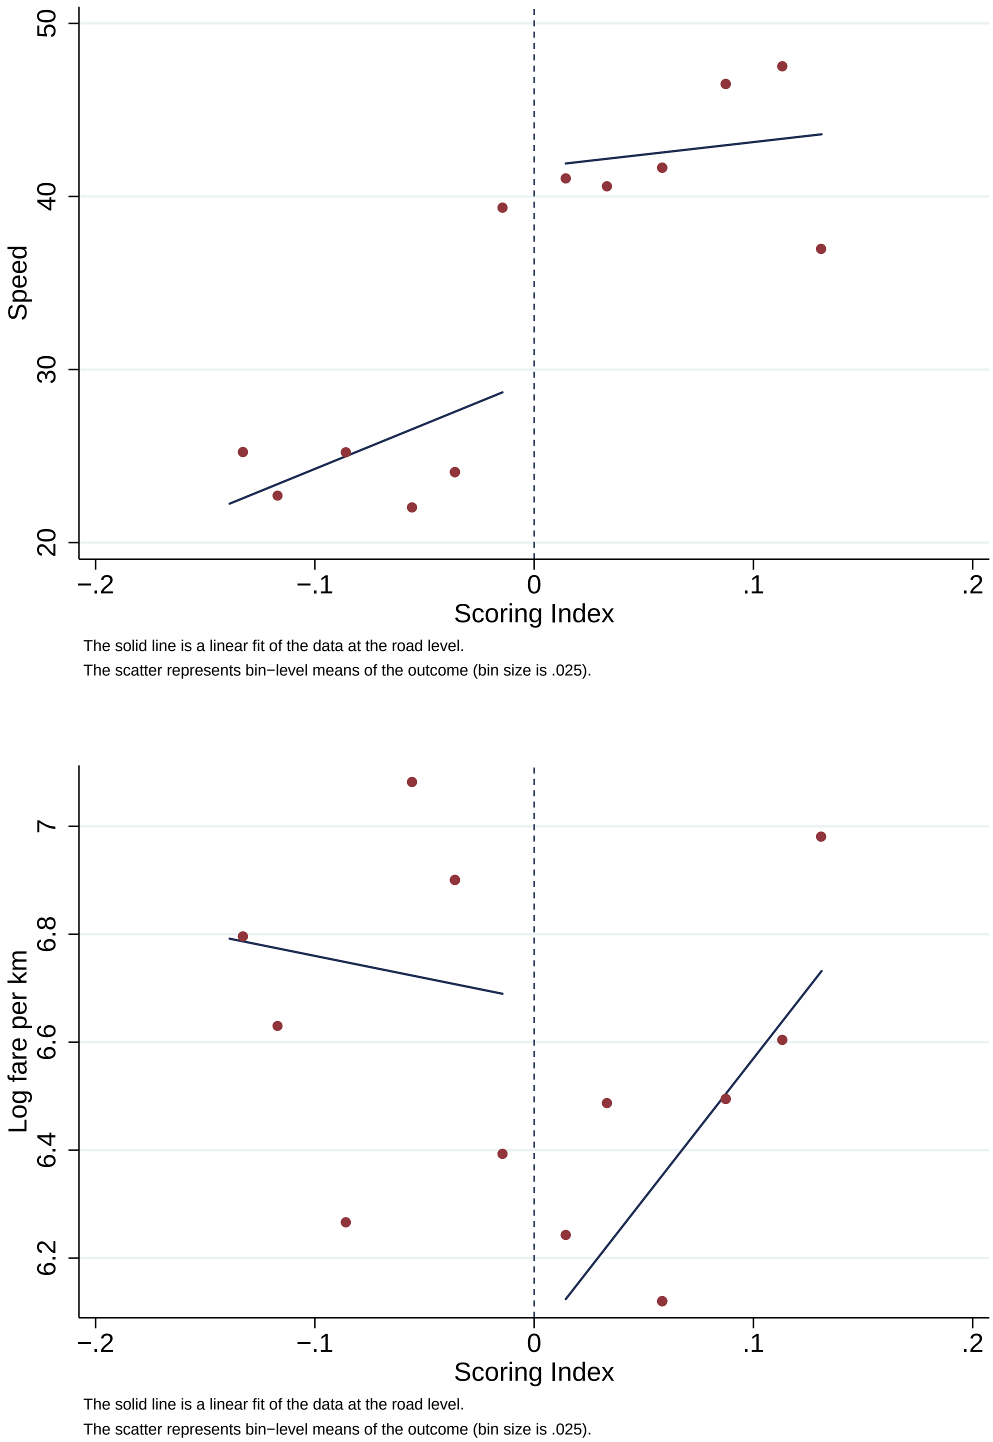

Figure 2.4: *Transport Costs Analysis: Motorbike Fares and Road Speed*

Table 2.6: *Transport Costs Analysis: Motorbike Fares and Road Speed*

	Average Speed (kph)		Log Fare/km	
	(1)	(2)	(3)	(4)
Preferred LLR (h=.15)				
Treatment	12.085**	12.769**	-0.594**	-0.610***
	[5.345]	[5.282]	[0.236]	[0.203]
Mean for Control	26.196	26.196	6.729	6.729
Heterogeneity Controls		X		X
Observations	31	31	31	31
LLR (h=.3)				
Treatment	16.122***	16.350***	-0.149	-0.178
	[4.881]	[4.704]	[0.214]	[0.215]
Mean for Control	26.732	26.732	6.789	6.789
Heterogeneity Controls		X		X
Observations	38	37	38	37
3rd Order Polynomial				
Treatment	-4.274	-3.674	-0.836**	-0.980**
	[8.509]	[8.498]	[0.325]	[0.419]
Mean for Control	26.732	26.732	6.805	6.805
Heterogeneity Controls		X		X
Observations	44	43	45	44

Note: Fares are one-way motorbike fares in SLL (Sierra Leonean Leones).
The relevant exchange rate was roughly $1=4300 SLL.
Average speed measured in kilometers per hour along the entire road.
Heterogeneity controls include cassava/rice sales indicators/harvests measured for EAs within 2km of the roads.
LLR is short for local linear regression.
All specifications include district fixed effects.
Heteroskedasticity-robust standard errors reported in brackets.
* p<0.1, ** p<0.05, *** p<0.01

2.5.2 Rural Market Prices

Basic Price Effects

We now look at the effect of the road improvements on crop prices. First, we relate our empirical approach to the theoretical framework above. The impact of rural road rehabilitation corresponds to the first comparative static of interest we presented in Section 2.2, $\frac{\partial p^S}{\partial \tau}$, in Table 1. Since the theoretical models derive predictions with respect to transport costs (not road quality), a negative $\frac{\partial p^S}{\partial \tau}$ implies that improvements in road quality have a positive effect on prices, and vice versa. If our empirical results are consistent with any given model, the sign of our regression coefficients should be the *opposite* of the theoretical predictions.

Our price results focus on the two main staple crops produced domestically, rice and cassava. In our data, rice and cassava are available 85% and 96% of the time, respectively, and both crops are sold in all markets at some point during the year. Before looking at prices, we first analyze the effects of the road improvement program on the extensive margin of crop availability in markets. The RDD treatment effect on availability of local rice is -0.10 with a standard error of 0.07. For cassava, the effect is -0.04 with a standard error of 0.06 (detailed results available upon request). We interpret this as evidence that market availability on the extensive margin was not different across treatment and control roads, although we do not have data on quantities traded that would allow us to further test this.

Figure **??** and Table 2.7 (columns (1) and (2)) show that road rehabilitation leads to a substantial reduction in the price of rice in the trader data. The graphical analysis shows a discontinuity in the price of rice at the cutoff. The price decrease is significant across all specifications (columns (1) and (2) of Table 2.7). In the rice price data from the high frequency surveys, evidence for a treatment effect in the form of a price reduction is much weaker (see Figure **??** and Table 2.7, columns (3)-(4)). While still negative in the tight bandwidth, the point estimates in these regressions are much smaller and mostly not significant. The results for rice from the high frequency data are different to those in the trader survey. One reason for this may be the difference in timing of the two surveys. The trader survey is conducted

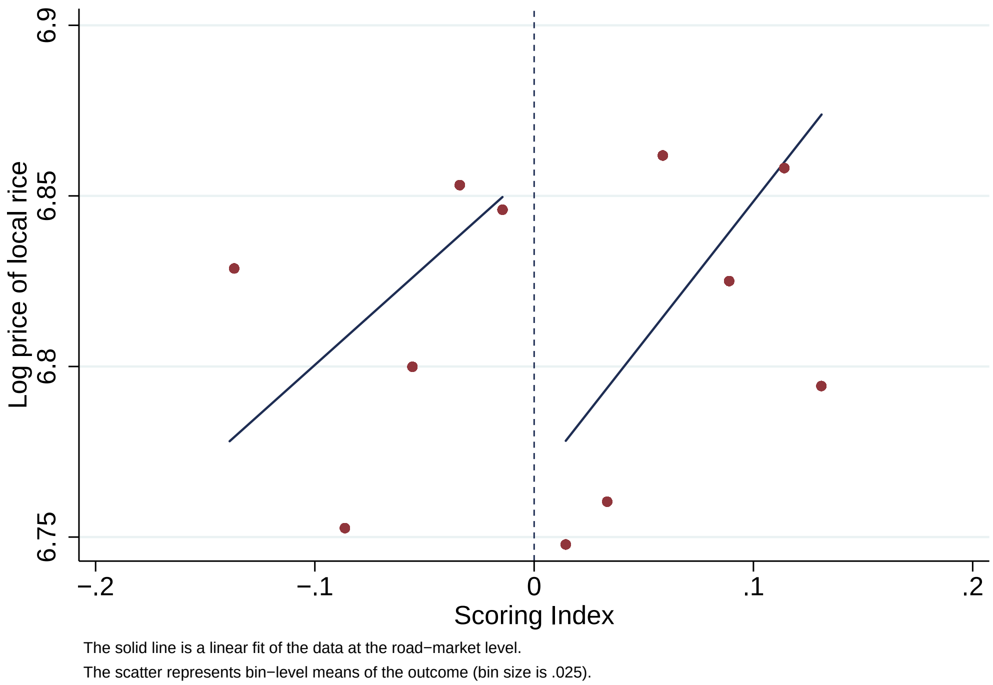

Figure 2.5: *Trader Survey: Log Price of Local Rice*

just after harvest, while the high frequency surveys are collected every month. We do not have enough sample size to look at just the harvest period in the high frequency data as we only have one year (hence one harvest season) of data. For cassava there is no specific harvest season. Cassava can remain in the ground for multiple seasons, can be harvested any time and is usually only harvested for immediate consumption or sale (it has to be processed immediately it is harvested).

Figure **??** and Table 2.7 (columns (5) and (6)) show price effects for cassava, based on the monthly price surveys only as the trader surveys did not cover cassava. Table 2.7 shows a large price drop in cassava. In the optimal bandwidth specification (top panel in column (5)), the road rehabilitation significantly reduces prices of cassava by 17.8%. This result is robust to different bandwidths, a polynomial specification, and to the inclusion of controls. The effect on cassava prices is larger than the effect on rice prices. We expect this to be the

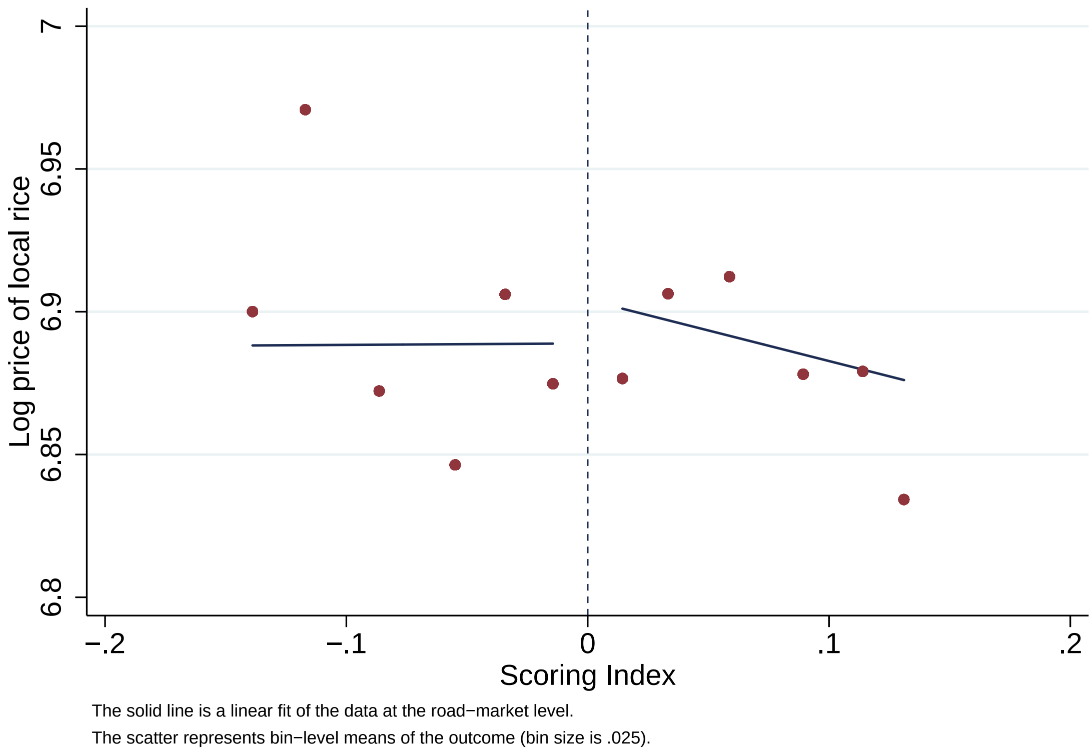

Figure 2.6: *High Frequency Price Survey: Log Price of Local Rice*

Table 2.7: *Effects on Prices from Trader and High Frequency Surveys*

	Log Local Rice Price		Log Local Rice Price		Log Cassava Price	
	(1)	(2)	(3)	(4)	(5)	(6)
Preferred LLR (h=.15)						
Treatment	-0.116**	-0.105***	0.006	0.005	-0.178**	-0.141**
	[0.058]	[0.029]	[0.029]	[0.016]	[0.086]	[0.058]
Mean for Control	6.831	6.831	6.889	6.889	5.703	5.703
Heterogeneity Controls		X		X		X
Observations	190	188	896	883	918	906
LLR (h=.075)						
Treatment	-0.178***	-0.174***	-0.030	-0.059**	-0.227*	-0.254***
	[0.066]	[0.047]	[0.048]	[0.027]	[0.125]	[0.074]
Mean for Control	6.843	6.843	6.887	6.887	5.728	5.728
Heterogeneity Controls		X		X		X
Observations	124	122	599	586	595	583
LLR (h=.3)						
Treatment	-0.074*	-0.069**	0.005	-0.004	-0.159*	-0.099*
	[0.045]	[0.031]	[0.021]	[0.021]	[0.089]	[0.052]
Mean for Control	6.831	6.831	6.889	6.889	5.703	5.703
Heterogeneity Controls		X		X		X
Observations	237	233	1099	1073	1112	1088
3rd Order Polynomial						
Treatment	-0.165***	-0.168***	-0.000	-0.058	-0.403**	-0.341***
	[0.058]	[0.044]	[0.055]	[0.052]	[0.170]	[0.063]
Mean for Control	6.831	6.831	6.889	6.889	5.703	5.703
Heterogeneity Controls		X		X		X
Observations	298	292	1314	1275	1338	1302

Note: Annual trader data (cols (1)-(2)); monthly high frequency price data (cols (3)-(6)). All prices in SLL (Sierra Leonean Leones) per cup. Heterogeneity controls include log distance to the nearest major town, and rice and cassava sales fractions and harvest amount medians. LLR is short for local linear regression. All specifications include round, road-district, and market-district fixed effects. Standard errors two-way clustered by road and market reported in brackets. * p<0.1, ** p<0.05, *** p<0.01

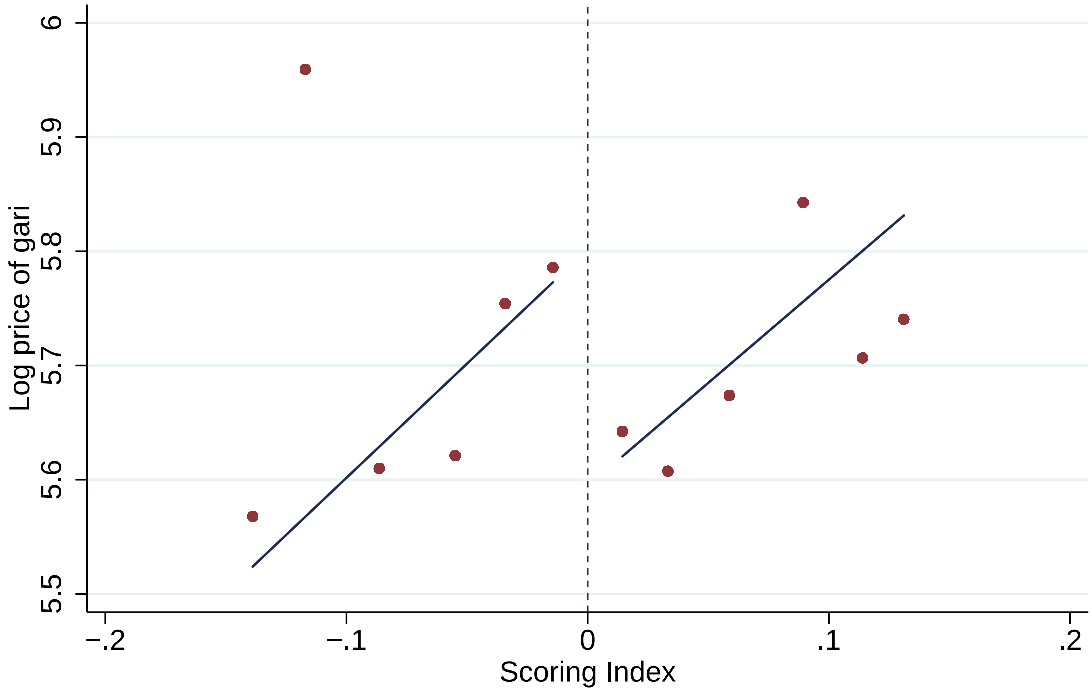

The solid line is a linear fit of the data at the road−market level.
The scatter represents bin−level means of the outcome (bin size is .025).

Figure 2.7: *High Frequency Price Survey: Log Price of Cassava (Gari)*

case as cassava is usually processed at the market site rather than on the farm,[30] and raw cassava is more expensive to transport than unprocessed (husk or threshed) rice.[31] Hence, cassava prices should be more responsive to changes in road infrastructure.

In summary, Table 2.7 provides evidence of price reductions for both rice and cassava in response to the road improvements. We interpret this as evidence that the improvements in infrastructure have primarily facilitated access to rural markets for producers, as opposed to

[30]In the AHTS (2011), only 2% of households and 8% of villages reported being equipped with at least one cassava grater; 0.05% of households and 8% of villages were equipped with a mechanized rice mill.

[31]We do not have accurate data on the transport cost of unprocessed cassava. However, raw cassava tubers are more voluminous, and hence more expensive, holding volume constant to transport than most other crops. In addition the processing rate to transform raw cassava into *gari* is low: 1 unit of unprocessed cassava yields 0.36 units of *gari* based on the AHTS (2011) data, while the typical milling rate for rice is 0.5. In other words, producers must transport larger volumes of raw cassava than rice to the marketplace to sell the same volume of the processed crop.

reducing transport costs on rural roads for traders. More formally, using the notation from our theoretical setup, these results imply that, in our sample, $\alpha > \beta$, which is consistent with the qualitative evidence that, on average, city traders' transport costs are primarily determined by the quality of the major roads, as opposed to rural roads. In our earlier terminology, the supply effect prevails over the demand effect. This importance of the supply effect (versus the demand effect) is inconsistent with a basic model of Bertrand competition. In the Bertrand framework, a change in transport costs affects the equilibrium price only through β, the costs for the trader. Similarly, the Cournot case with endogenous entry predicts that a reduction in transport costs would unambiguously increase prices. These predictions are not borne out in the empirics.

Heterogeneous Price Effects

We now focus on the heterogeneity in these price responses to the road improvements. The sign of the heterogeneity in these effects can shed light on the underlying market structure and allow us to empirically test between different classes of models. Using insights from the theory, our analysis of heterogeneous effects focuses on two market characteristics: the distance between the rural market and the closest urban center (the variable x in the theory) and crop-specific productivity around the markets (the variable σ in the theory). We proxy the latter with Chiefdom level measures of havests and sales. For this analysis, we work with dummy variables that indicate whether a certain market is above or below the median value of each of our heterogeneity variables. Using these binary indicators allows us to present the results in an intuitive way and to quantify the economic magnitude of the interactions easily.

In Table 2.8 (columns (1) and (2)), we present our heterogeneity analysis of price effects for based on the distance between markets and the nearest large town. The price reductions induced by road quality improvements are always stronger (i.e., more negative) for markets that are further away from large towns. The sign of the interaction term of treatment and the dummy indicating the market is far from urban centers is negative and statistically significant in all four specifications for rice, and all but the wide-bandwidth local linear regression for cassava. By contrast, the sign of the treatment coefficient is always positive, and significant in all cases for rice, but never for cassava. Based on this hypothesis, we also report one-sided p-values on the treatment interaction coefficient in each panel, all eight of which are significant at the 10% level (and seven are significant at the 1% level).

Next, we examine the heterogeneous impacts on crop prices by the chiefdom-level average household harvest as a proxy for productivity (columns (3) and (4)). For rice, we find that the treatment effect is negative in low-productivity areas, and in fact positive in the high-productivity areas, indicating a stronger demand effect in these areas. With respect to rice, the coefficient of treatment interacted with high-productivity areas is statistically significant in the optimal and tight bandwidths, but not significant in the other two specifications. We find similar results for cassava, with a strongly negative treatment effect in low-productivity markets (significant in the tight bandwidth and polynomial specifications) and a coefficient that is positive and statistically different from zero in two specifications in the high-harvest areas. The one-sided p values are significant in four of the eight specifications.

In columns (5) and (6), we present evidence of treatment effect heterogeneity by our second proxy of farmer productivity, the share of farmers selling any of their harvest. The price reduction induced by the improvements in road quality is significantly weaker in markets with a higher density of sellers. For rice, in markets with a high density of sellers, the price effects are positive. For both crops, the difference in the percentage effect is economically sizable, although it is larger for cassava. The significance of the interaction term, as well as the one-sided p-values, indicate that the effects across low and high seller density markets are significantly different from each other for all specifications for cassava, and all except the wide-bandwidth specification for rice.

Summarizing these results, we find evidence that market-specific characteristics, both on the demand side (distance from major market centers) and on the supply side (productivity), significantly affect the magnitude and the direction of the price effects of improvements in rural road infrastructure. Turning to the models presented in Section 2.2, these findings

Table 2.8: *Heterogeneity by Access to Towns, Agricultural Production and Density of Rice Sellers Near Roads*

	Distance		Harvest		Seller Density	
	(1) Log Rice Price	(2) Log Cassava Price	(3) Log Rice Price	(4) Log Cassava Price	(5) Log Rice Price	(6) Log Cassava Price
Preferred LLR (h=.15)						
Treatment	0.085***	0.002	-0.022	-0.115	-0.040	-0.275**
	[0.024]	[0.033]	[0.036]	[0.071]	[0.040]	[0.108]
Treat * Above Median	-0.162***	-0.345***	0.065*	0.064	0.095**	0.236**
	[0.047]	[0.104]	[0.039]	[0.080]	[0.039]	[0.113]
One-sided p-value	0.000	0.000	0.048	0.212	0.008	0.018
Mean for Control	6.889	5.703	6.889	5.703	6.889	5.703
Observations	896	918	883	906	883	906
LLR (h=.075)						
Treatment	0.107***	0.051	-0.138**	-0.404**	-0.121*	-0.476***
	[0.040]	[0.068]	[0.070]	[0.178]	[0.072]	[0.115]
Treat * Above Median	-0.210***	-0.535***	0.178**	0.410**	0.176**	0.490***
	[0.073]	[0.129]	[0.076]	[0.188]	[0.078]	[0.132]
One-sided p-value	0.002	0.000	0.009	0.015	0.012	0.000
Mean for Control	6.887	5.728	6.887	5.728	6.887	5.728
Observations	599	595	586	583	586	583
LLR (h=.3)						
Treatment	0.072***	0.002	-0.008	-0.254	-0.007	-0.278***
	[0.022]	[0.023]	[0.036]	[0.185]	[0.034]	[0.102]
Treat * Above Median	-0.128***	-0.186	0.019	0.219	0.040	0.308***
	[0.037]	[0.136]	[0.040]	[0.191]	[0.037]	[0.117]
One-sided p-value	0.000	0.086	0.318	0.126	0.136	0.004
Mean for Control	6.889	5.703	6.889	5.703	6.889	5.703
Observations	1099	1112	1073	1088	1073	1088
3rd Order Polynomial						
Treatment	0.135**	0.107	-0.053	-0.362***	-0.138***	-0.584***
	[0.055]	[0.075]	[0.065]	[0.095]	[0.033]	[0.101]
Treat * Above Median	-0.242***	-0.552***	0.065	0.420***	0.205***	0.556***
	[0.092]	[0.089]	[0.063]	[0.134]	[0.049]	[0.123]
One-sided p-value	0.004	0.000	0.151	0.001	0.000	0.000
Mean for Control	6.889	5.703	6.889	5.703	6.889	5.703
Observations	1314	1338	1275	1302	1275	1302

Note: All data is from high frequency price surveys. All prices in SLL (Sierra Leonean Leones) per cup. LLR is short for local linear regression. All specifications include round, road-district, and market-district fixed effects. One sided p-values are calculated on the interaction term (>0 for distance, <0 for harvest and seller density). Heteroskedasticity-robust standard errors for all narrow bandwidth (h=.075) regressions. Standard errors two-way clustered by road and market for cassava price regressions, except the h=.075 specification. For rice price (distance heterogeneity) regressions, we use heteroskedasticity-robust standard errors as there are too few clusters to allow for two-way clustering. Standard one-way clustering at either the road level or market level produces standard errors that are uniformly smaller than those reported. * p<0.1, ** p<0.05, *** p<0.01

are inconsistent with the bilateral bargaining framework and the Cournot model with an exogenous number of traders.[32] In the bilateral bargaining setting, the productivity parameter affects the size of the surplus but does not affect the optimal solution to the bargaining problem. In addition, the distance between the rural market and the urban center, x, affects the outside option of the city-traders but does not affect the impact of rural road quality on the bargaining solution. The Cournot model with an exogenous number of traders also predicts that neither x nor σ affect the price derivative with respect to rural transport costs.

Our empirical findings are therefore only consistent with a model with search frictions, as formalized by Chau, Goto and Kanbur (2009). In this framework, a reduction in rural road transport costs has both demand and supply effects on prices. It also predicts that more isolated markets have lower entry and hence more imperfect pass-through and that markets with higher producer productivity have lower real search frictions and hence stronger demand effects. We find evidence of both forms of heterogeneity and in the directions predicted by a search framework.

Additional Evidence in Support of Search Costs

To look more directly at search costs, we present additional results that make use of data on cell phone access. To measure access to cell phones, we use the average cell phone ownership in the Chiefdom in which the market is located. As with the earlier heterogeneity results, we use a dummy for whether the market has below or above median cell phone ownership. The intuition here is to test whether there is heterogeneity in our price effects by this measure of cell phone penetration. Cell phones would allow farmers and traders to gather information in prices, to coordinate on times to meet, and to generally explore outside options more easily. Therefore, we expect that higher cell phone prevalence would be associated with lower search costs and hence that our price effects would be muted or

[32] The Bertrand model and the Cournot model with an endogenous number of traders were ruled out on the basis of our main price effects in Table 2.7.

less negative in areas with better cell phone penetration.

These results are reported in Table 2.9. As a framework with search costs would predict, we find that the price effects are less negative (and in some cases even positive) in areas with better cell phone penetration, since the coefficients on the interaction are positive. We find this across specifications for cassava prices but less robustly across specifications for rice prices.

Robustness Checks

Given our small sample size and empirical design, we report three different sets of robustness checks on our main results: alternative specifications, alternative samples, and a small sample correction. We report all robustness check results for the main price effects in Table 2.10 (both with and without the heterogeneity controls) and for the heterogeneous effects in Table 2.11.

In Panel A of Tables 2.10 and 2.11, we present results from three additional specifications: (i) one that controls for second-order polynomials in the score, (ii) one that uses a triangle weight in the local linear regression (i.e., a weight that places less emphasis on observations further from the cutoff), and (iii) one that uses continuous heterogeneity variables (reported only in Table 9) instead of the 'greater than median' dummies we use in Table 2.8. As Panel A in each of Tables 2.10 and 2.11 shows, our main results are robust to these alternative specifications. In some cases we lose precision in our estimates, but the coefficients do not change significantly.

With respect to alternative samples, we present results from four additional specifications: (i) one that includes the two roads at the cutoff in Kambia District, (ii) one that restricts the sample to the closest market to each road and therefore has no weights, (iii) one that restricts the sample to markets within 5km of each road and does not weight by the market-distance to the road, (iv) one that excludes markets that are in more urban areas, and (v) one that excludes markets that match to more than four roads. Once again, looking at Panel B in Tables 8 and 9, we see that our results are robust across all these alternative samples though

Table 2.9: *Heterogeneity by Mobile Phone Penetration*

	(1) Log Rice Price	(2) Log Cassava Price
Preferred LLR (h=.15)		
Treatment	-0.038	-0.233**
	[0.033]	[0.091]
Treat * Above Median	0.084**	0.242**
	[0.040]	[0.106]
One-sided p-value	0.019	0.011
Mean for Control	6.889	5.703
Observations	883	906
LLR (h=.075)		
Treatment	-0.082	-0.386***
	[0.061]	[0.099]
Treat * Above Median	-0.069	0.675***
	[0.091]	[0.178]
One-sided p-value		
Mean for Control		
Observations	586	583
LLR (h=.3)		
Treatment	-0.027	-0.238***
	[0.024]	[0.087]
Treat * Above Median	0.036	0.346***
	[0.045]	[0.116]
One-sided p-value	0.210	0.001
Mean for Control	6.889	5.703
Observations	1073	1088
3rd Order Polynomial		
Treatment	-0.071	-0.451***
	[0.056]	[0.145]
Treat * Above Median	-0.131	1.297***
	[0.193]	[0.442]
One-sided p-value	0.751	0.002
Mean for Control	6.889	5.703
Observations	1275	1302

Note: All data is from high frequency price surveys. All prices in SLL (Sierra Leonean Leones) per cup. LLR is short for local linear regression. All specifications include round, road-district, and market-district fixed effects. One sided p-values are calculated on the interaction term (>0 for distance, <0 for harvest and seller density). Heteroskedasticity-robust standard errors for all narrow bandwidth (h=.075) regressions. Standard errors two-way clustered by road and market for cassava price regressions, except the h=.075 specification. * p<0.1, ** p<0.05, *** p<0.01

Table 2.10: *Robustness Checks: Main Effects on Prices*

	Log Local Rice Price		Log Local Rice Price		Log Cassava Price	
	(1) Basic	(2) Het Controls	(3) Basic	(4) Het Controls	(5) Basic	(6) Het Controls
A: Alternative Specifications						
Paper Spec: 2nd Poly						
Treatment	-0.140**	-0.120***	0.005	-0.020	-0.218	-0.113
	[0.066]	[0.044]	[0.036]	[0.039]	[0.135]	[0.074]
Observations	298	292	1314	1275	1338	1302
Triangle Weight						
Treatment	-0.082	-0.074***	0.027	0.033**	-0.079	-0.077**
	[0.050]	[0.028]	[0.032]	[0.016]	[0.067]	[0.036]
Observations	190	188	896	883	918	906
B: Alternative Samples						
All Roads						
Treatment	-0.100**	-0.073**	0.015	0.024	-0.169***	-0.095*
	[0.048]	[0.031]	[0.023]	[0.016]	[0.063]	[0.050]
Observations	203	201	970	957	995	983
Closest Market (unw)						
Treatment	-0.123*	-0.179**	0.030	0.027	-0.179**	-0.161***
	[0.074]	[0.073]	[0.026]	[0.027]	[0.077]	[0.061]
Observations	56	56	282	282	284	284
Markets w/in 5km (unw)						
Treatment	-0.085	-0.063*	0.032	0.018	-0.122*	-0.080
	[0.061]	[0.033]	[0.028]	[0.027]	[0.073]	[0.061]
Observations	82	82	438	438	459	459
Drop Town Markets						
Treatment	-0.131**	-0.127***	-0.002	-0.000	-0.198**	-0.172***
	[0.056]	[0.022]	[0.030]	[0.018]	[0.091]	[0.052]
Observations	163	161	733	720	753	741
Drop Multimatch (>4) Markets						
Treatment	-0.131**	-0.127***	0.001	0.010	-0.168*	-0.132**
	[0.056]	[0.022]	[0.029]	[0.016]	[0.092]	[0.058]
Observations	163	161	778	765	780	768
C: Small Sample Corrections						
One-way Cluster (Road)						
Treatment	-0.116**	-0.105***	0.006	0.015	-0.178**	-0.126**
	[0.051]	[0.027]	[0.028]	[0.018]	[0.080]	[0.053]
Wild Bootstrap Coeff 95% CI	-0.21, -0.04	-0.15, -0.06	-0.05, 0.06	-0.02, 0.05	-0.33, -0.03	-0.23, -0.02
Observations	190	188	896	883	918	906

Note: All specifications are local linear regressions at the preferred (h=.15) bandwidth unless otherwise indicated.Wild bootstrap coefficient confidence intervals based on 400 iterations.

Table 2.11: *Robustness Checks: Heterogeneity*

	Distance		Harvest		Seller Density	
	(1) Rice Price	(2) Cassava Price	(3) Rice Price	(4) Cassava Price	(5) Rice Price	(6) Cassava Price
A: Alternative Specifications						
Paper Spec: 2nd Poly						
Treat * Above Median	-0.156***	-0.369**	0.088	0.466*	0.101*	0.564***
	[0.057]	[0.149]	[0.056]	[0.277]	[0.060]	[0.126]
Triangle Weight						
Treat * Above Median	-0.160***	-0.266***	0.073**	0.019	0.101**	0.173*
	[0.042]	[0.086]	[0.031]	[0.045]	[0.039]	[0.089]
Continuous Het Vars						
Treatment Interaction	-0.019	-0.121**	-0.047	0.123***	0.085 s	0.505*
	[0.014]	[0.056]	[0.035]	[0.035]	[0.165]	[0.287]
B: Alternative Samples						
All Roads						
Treat * Above Median	-0.148***	-0.304***	0.023	0.007	0.080*	0.189*
	[0.044]	[0.099]	[0.045]	[0.070]	[0.042]	[0.108]
Closest Market (unw)						
Treat * Above Median	-0.096	-0.252**	0.035	0.139	0.023	0.181*
	[0.059]	[0.121]	[0.063]	[0.322]	[0.065]	[0.109]
Markets w/in 5km (unw)						
Treat * Above Median	-0.183***	-0.369***	0.154***	0.126*	0.135***	0.220**
	[0.050]	[0.118]	[0.018]	[0.070]	[0.031]	[0.105]
Drop Town Markets						
Treat * Above Median	-0.165***	-0.352***	0.054	0.136*	0.084**	0.273***
	[0.051]	[0.108]	[0.036]	[0.076]	[0.039]	[0.105]
Drop Multimatch (>4) Markets						
Treat * Above Median	-0.177***	-0.431***	0.053	0.088	0.082**	0.261**
	[0.054]	[0.066]	[0.046]	[0.078]	[0.041]	[0.114]
C: Small Sample Corrections						
One-way Cluster (Road)						
Treat * Above Median	-0.162***	-0.345***	0.065*	0.064	0.095**	0.236**
	[0.047]	[0.088]	[0.038]	[0.086]	[0.042]	[0.102]
Wild Bootstrap Coeff 95% CI (Lower)	-0.21, -0.12	-0.51, -0.19	-0.01, 0.14	-0.11, 0.23	0.02, 0.17	0.05, 0.41

Note: All specifications are local linear regressions at the preferred (h=.15) bandwidth unless otherwise indicated. Two-way clustered standard errors except rice price/distance heterogeneity regressions and one-way cluster spec. All one-sided p-values are for the interaction term (>0 for distance, <0 for harvest and seller density). Wild bootstrap interaction coefficient confidence intervals based on 400 iterations.

we lose power in the specifications with low sample sizes.

Finally, Panel C in Tables 2.10 and 2.11 includes a robustness check to address the sample size and low number of clusters. Only 47 roads were considered for rehabilitation, but since we have market-level data, our results are two-way clustered by market and road. Our concern is the road-level clustering. From the applied literature (see Cameron, Gelbach and Miller (2008) and Angrist and Pischke (2008)), it seems that 47 clusters is not small enough to be a concern, and neither is 31 (the number of roads in our preferred bandwidth specification). However, as a check, we implement the wild bootstrap - Cameron, Gelbach and Miller (2008) propose this for small numbers of lusters. Since our concern is clustering at the road-level, as a preliminary step, we show that one-way clustering by road does not change our results for the preferred bandwidth specification (see Panel C of Tables 8 and 9). We then report the 95% confidence interval from a wild bootstrap routine based on clustering by road only. The wild bootstrap confidence bands support our main results.

2.6 Conclusion

Road infrastructure projects represent a large share of current foreign aid and government expenditure in Sub-Saharan Africa. Promoters of these investments argue that enhancing market access is a prerequisite for agricultural development and economic growth. As such, the assessment of the impact of local improvements in rural road quality has immediate policy relevance. In addition, any empirical study focusing on these policies should explicitly assess the role of market structure and hence the heterogeneity in impacts across markets with different characteristics, such as the proximity of the market to urban centers, agricultural productivity and cell phone access.

In this paper, we make two contributions relative to the existing literature. First, we use the specific design of a rural road rehabilitation program in Sierra Leone to estimate the causal impact of road quality on rural market prices of rice and cassava, the two most important staple crops in the country. We also estimate how this response varies with specific market characteristics. Second, we use these empirical findings to test between alternative stylized frameworks of intermediated trade which describe the interaction between rural producers, intermediaries and final consumers. Specifically, we model responses to a change in rural road transport costs under the following four settings: (i) Bertrand traders' oligopsony, which leads to perfect competition outcomes; (ii) bilateral bargaining, which models the case of rural monopsonies; (iii) Cournot oligopsony, both with and without an endogenous number of traders; and (iv) a basic search framework à la Mortensen (2003). We show that these frameworks deliver different predictions on how prices respond to a reduction in rural transport costs. The predictions for how these price effects vary also differ. Comparing these theoretical predictions to our empirical findings, we find support for the model with search frictions but not for the other frameworks of trader competition. In addition, we present additional evidence in support of search costs using data on cell phone penetration.

The evidence on search frictions is consistent with a recent literature on the effects of information on prices (see Jensen (2007), Aker (2010), and Goyal (2010)). However, relative to this literature, we show that search frictions which determine the nature of traders' competition also have an effect on how the local economy responds to improvements in infrastructure. More generally, we point out that the structure of the market for agricultural commodities will affect the direction and the magnitude of price response to such investments, thus affecting the size of the gains for different types of agents. We believe our results could inform the decisions of policymakers considering similar improvements in infrastructure as well as a wide range of other policies that would impact agricultural trade, such as price subsidies, agricultural export promotion interventions, credit-provision policies targeting traders, and more major road projects. In addition, finding evidence consistent with the presence of search frictions implies that there may be strong complementarity between roads (or other infrastructure projects) and policy interventions that may reduce search frictions, such as the expansion of mobile phone coverage.

Possible extensions of this research would look at the longer-term effects of the rural roads improvements. Our study focused on the short-term impacts of these improvements since our surveys were conducted between one and two years after the completion of the rehabilitation program. Given the characteristics of these roads, i.e. that they are small dirt roads, we potentially expect the infrastructure to degrade over time and hence the impact of the rehabilitation to be more limited as this degradation occurs.

Chapter 3

Interlinked Transactions and Pass-Through: Experimental Evidence from Sierra Leone[1]

3.1 Introduction

Rural areas of developing economies often lack formal financial institutions. In their absence, agents in the rural supply chain have emerged as a substitute source of credit for producers and households. For instance, an intermediary buying agricultural produce for a wholesaler may provide payment in advance to the farmer for output, allowing the farmer to smooth consumption.[2] A long tradition in development economics has emphasized that relationships such as these lead to transactions that are *interlinked*: the price at which output is purchased is determined jointly with the terms of the credit contract, and vice versa (e.g. Bardhan, 1980, Braverman and Stiglitz, 1982, Bell, 1988, Grosh, 1994, Deb and Suri, 2012). A corollary to this observation is that product market conditions may affect the supply of

[1] Co-authored with Tristan Reed

[2] It has been observed more generally in the finance literature that a particular form of credit given through the supply chain—trade credit—becomes more prominent when financial institutions are weak (e.g. Petersen and Rajan, 1997; Fisman and Love, 2003).

credit. If the wholesale value of produce for the intermediary rises, so might the credit supplied to the farmer.

The presence of such interlinkages complicates our understanding of how prices pass through the supply chain. The rate of pass-through is an important parameter, as it is output prices, along with costs, that ultimately determine a farmer's investment and the decision to adopt technologies; the benefit to growth offered by trade liberalization, for instance, may be attenuated to the extent that international price signals fail to reach domestic producers. If intermediaries pass through some of the value of world prices in credit, however, there may be in fact more transmission of incentives than would be observed if one only looked at farm gate prices. A large literature in international economics (for a review see Burstein and Gopinath, 2012) explains low price pass-through alternatively with price rigidity, imperfect competition, and distribution costs. The presence of interlinkages may provide another important explanation, particularly in remote areas of developing economies, where a low rate of price pass-through has been observed (Fafchamps and Hill, 2008; Atkin and Donaldson, 2013; Mitra et al.,2013).[3]

Building on these observations, our paper makes three contributions. First, we discuss the results of a randomized experiment in a set of agricultural markets in sub-Saharan Africa designed to elucidate the separate margins of pass through. The experiment is set in the cocoa industry of Sierra Leone where interlinked transactions such as the one described above are common. We pay a treatment group of intermediaries a per-unit bonus for delivering cocoa (above a certain quality standard) to wholesalers. Using detailed data on the prices and credit supplied to farmers, we show that although average pass-through of the bonus is small in terms of prices, it is substantial in terms of credit outlay. The experiment confirms the two conjectures above: product market conditions faced by the intermediary affect substantially the supply of credit to farmers, and the total pass through of the cocoa's value is masked when one only observes the price at which the cocoa is

[3] It has been alleged, correctly or not, that in the wake of cocoa and coffee market liberalization, intermediaries have exerted market power over farmers, dampening the transmission of price signals, and muting their incentives to invest (Oxfam, 2002).

transacted.

Second, we develop a simple theoretical model that illustrates this point. In our model, changes in the price paid to intermediaries for output shift the share of producers engaged in interlinked transactions as opposed to simply selling on a spot market. In the interlinked transaction, intermediaries pay the producer in advance for the good, a form of forward credit that the producers use to smooth consumption.[4] This credit is paid back in the form of a lower output price. The average rate of price pass-through is determined by the measure of producers who endogenously switch into (or out of) interlinked transactions. In response to an increase in the price they receive from wholesalers, intermediaries may choose to give credit to more producers. As these producers move from the spot market to the interlinked transaction, the observed price they receive falls; indeed they have already been paid in advance. While farmers benefit from credit provision, this switching between contracts drives down the average rate of price pass-through further than the rate that would obtain if the intermediary were simply an oligopsonist on the spot market. We test empirically the core prediction of the model, that across markets, price and credit pass-through are substitutes. Using an analysis of heterogeneous treatments effects in our experiment, we show that those markets that experience a stronger credit response show a lower rate of price pass-through. The markets in which the credit response is stronger and the price response weaker are precisely those markets in which interlinked transactions were more common in the baseline.

Third, we calibrate our model and show that a model that focuses only on price pass-through will substantially underestimate the change in producer welfare derived from an increase in the buyer price relative to one that accounts for interlinkages. The result is robust to a wide range of parameter values. This insight also speaks to recent literature that uses the price pass-through rate as a tool to infer the shares of surplus captured

[4]Intermediaries have also been observed to write contracts that transfer risk from farmers to traders, providing insurance against adverse price and productivity shocks that may affect the farmer. While we acknowledge this is another important margin on which intermediaries may pass through value to farmers, it is not common in our setting, and we leave it to be studied in others.

by producers, consumers, and intermediaries in the economy (Fabinger and Weyl, 2013; Atkin and Donaldson, 2012). In the presence of interlinked credit and output markets, welfare analysis needs to be based on "credit pass-through", as well as the standard price pass-through.

The presence of many layers of intermediation is a defining characteristic of agricultural markets in Sub-Saharan Africa (Bauer, 1954; Fafchamps, 2004). Besides simply transporting goods, intermediaries also provide services such as information, insurance, and, as in our context, credit. While the literature has acknowledged this role, there is little quantitative evidence about how traders' own incentives affect the level of service provision. This paper shows they do, and that the magnitude of this effect can be substantial. Our work then supports a view of intermediaries as productive members of the supply chain who undertake value-enhancing investments, as opposed to a view in which they are simply arbitrageurs. In this sense our work is related to that of Rubenstein and Wolinsky (1987) and Antras and Costinot (2012), who develop models in which traders provide a service to the market by alleviating search frictions. It is also related to the work on micro-finance by Maitra, Mitra, Mookherjee, Motta, and Visaria (2012) who identify another way in which traders may add value. The authors argue that given the strength of traders' relationships with clients, traders may have more information about default risk and be able higher quality clients to financial institutions.

Given the context, our work also contributes to the extensive literature on agricultural traders in Africa in particular, initiated by Bauer (1954) and Hill (1961) and continued by Fafchamps (2004), Fafchamps, Gabre-Madhin and Minten (2005), Osborne (2005), Fafchamps and Hill (2008) and Casaburi, Glennerster and Suri (2012), among others. More broadly, we also add to the literature studying the importance of inter-firm relationships in developing countries (McMillan and Woodruff, 1999; Banerjee and Duflo, 2000, and Macchiavello and Morjaria 2012).

The paper proceeds as follows. In section 3.2 we describe our experiment and provide summary statistics on traders and the markets used in the study. Section 3.3 discusses our experimental results. In section 3.4 we present a model of pass-through in interlinked transactions. Section 3.5 tests further implications of the model and presents a calibration and a discussion of its welfare implications. Section 3.6 concludes.

3.2 An Experiment in the Sierra Leone Cocoa Industry

In order to elucidate the multiple margins through which intermediaries may pass value to producers in response to a change in their price incentives, we run an experiment in a set of African agricultural markets in which interlinked transactions including credit are common, within the cocoa industry of Sierra Leone.[5] As the summary statistics presented below will show, the provision of loans by traders to farmers is a defining characteristic of this industry, making the context similar to those in other developing economies discussed in the papers cited in the introduction. In the course of an ongoing business relationship, traders will offer farmers credit, and then allow farmers to repay the loan in cocoa, by accepting from the trader a below market price for subsequent sales until the loan has been repaid. This credit could be productive, and allow the farmer to invest in post-harvesting quality-enhancing processing, or could be simply a payment advance that the farmer uses for consumption.

We paid a bonus of 150 Leones—5.6% of the average wholesale price—for high quality cocoa to randomly selected traders, who themselves buy directly from farmers. We then measure how this bonus affects prices and credit delivered to farmers across the different villages in which the traders operate. By estimating heterogeneous treatment effects across villages, we will be able to test a core prediction of our model, that price and credit pass-through are substitutes from the perspective of the trader.

The bonus in our experiment was designed to model fluctuations in the market price

[5]West Africa produces two-thirds of the world cocoa supply. Though given its small size Sierra Leone accounts for only a small share of this total, cocoa is important nationally. The crop comprised 8.6% of exports in 2009, and is by far the country's largest export crop by value, according to the UN COMTRADE database. The industry has also grown tremendously in the last decade, with the value of exports growing ten-fold between 2009 and 2001, when the country's decade long civil war came to an end.

received by traders, who themselves sell to wholesalers. The within-country cocoa trade in Sierra Leone is highly fragmented across many traders, and the supply chain has many links, similar to other agricultural markets in developing countries (for examples in Africa see Fafchamps, et. al. 2005 and Osborne 2005).[6] Farmers sell to traders, who sell to wholesalers in small towns, who in turn sell to exporters in larger towns, who in turn sell to buyers at the port. While the study of pass-through is surely relevant at each of these links in the supply chain, we focus on the final link closest to production, and leave the examination of other levels for future research. Working at this level is not only the most feasible from a cost-effectiveness perspective, but it also allows us to examine heterogeneity in pass-through across many different locations. As one moves further down the supply chain, the number of origin locations for cocoa necessarily falls quite quickly.

The cocoa season in Sierra Leone lasts from the beginning of the rainy season in June until February. Our experiment covers the months of September-December of 2011. The experimental approach allows us to exploit price variation across traders throughout the study time span. Prices also vary because of international market fluctuations. However, these changes are infrequent, making it difficult to estimate an effect of pass-through unconfounded by other seasonal variables that may affect supply. This concern motivates an experiment in preference to an observational study. In the following two subsections we discuss respectively the traders and the villages they operate in. Throughout both we discuss the experimental design, the data collected, and summary statistics from these data.

3.2.1 Traders and the Treatment

We developed our experiment in partnership with five privately owned wholesalers in Sierra Leone's Eastern cocoa producing district, three in the town of Segbwema, and

[6]Though Sierra Leone does have an official marketing board, the organization has been defunct since the war, and the government is responsible for a negligible share of purchases. The discussion in Gilbert (2009) suggests that Sierra Leone's market is similar to those in Nigeria, Cameroon and Côte d'Ivoire, all of which liberalized during the 1990s, and became similarly fragmented. A potential explanation for the lack of vertical integration in the market in the absence of a strong marketing board are the stringent legal restrictions on the transaction of land discussed in Acemoglu, Reed and Robinson (2013). These, along with weak legal institutions more broadly, would make vertical integration of the supply chain difficult, if not impossible.

one each in the towns Pendembu and Kailahun.[7] These wholesalers collect cocoa in their warehouses, and then sell it on to exporters in the provincial capital of Kenema. Our sample of 80 traders, henceforth study traders, comprise almost the complete set of traders who do business regularly with these wholesalers.[8]

Treatment and Random Assignment

Treatment traders were paid a bonus of 150 Leones per pound of Grade A cocoa. Though this bonus was only a 5.6% share of the price traders received before the treatment, it was a large, approximately 60% increase over average baseline margins. Traders were not told when the treatment price increase would end, though some certainty was given by our research team that it would last for at least a few months.

As emphasized by Atkin and Donaldson (2012), it is important to measure pass-through only for narrowly defined homogenous goods, as one must not conflate pass-through and changes in the composition of quality. The quality of cocoa is indeed heterogeneous, and market prices depend on a variety of characteristics including moisture content, mold, germination, lack of fermentation and a discoloration known as slate. Though there is no official measure of quality in the market, wholesalers and traders agree on broad determinants of quality that are consistent with international standards (see CABISCO, 2002). In order to measure pass-through for given classes of quality, we worked with wholesalers to develop a quality grade that correlates well with baseline prices. When traders arrive at the warehouse, inspectors hired by the research team sampled 50 beans

[7]These towns are now quite remote, accessible only by unpaved roads that can become impassible in the rainy season. During the colonial period, however, Pendembu was a prosperous trading town and the final stop on the Sierra Leone Railroad, which was dismantled and sold by the government of Siaka Stevens in the 1970s. The decline in the country's cocoa industry since then can be observed at the massive abandoned produce warehouse where the end of the tracks once lay. Exporters we visited in 2011 joked with some cynicism that the cocoa stocks of the largest wholesalers in Pendembu could not come close to filling it.

[8]In a census of regular business partners of the wholesalers, we counted originally 84 traders. Two were outliers with respect to baseline quantity. They were not matched and thus de facto dropped from the randomization. One other trader was lost due to attrition–he did not return after the census and no follow up data on either credit or prices could be collected. Since all of the analysis is done within matched pairs, his pair is also dropped from the analysis. Given the pairwise randomization, this attrition is not a threat to the internal validity of our study.

from each bag, and used them to create an index of quality—grades A, B or C—which was then applied to each bag. In Appendix C.1, we discuss in greater detail the construction of the grades, and their relationship to wholesale prices and international standards of cocoa quality.

To improve the statistical power of our experiment, we follow a pairwise randomization strategy (for a review, see Bruhn and McKenzie, 2009). We first match traders within wholesalers according to a self-reported estimate of the number of grade A bags that they had sold since the beginning of the cocoa season, a plausible proxy for the scale of their business. We felt this a useful proxy for similarity in capacity for price and credit pass-through, since the ability to give credit will be a function of the total wealth of the trader, which, given constant or increasing margins, should rise with the scale of business. Having matched the traders, we assigned treatment and control within pairs using a random number generator.

Trader Data and Summary Statistics

Over the course of the experiment we collect a variety of data from traders. Summary statistics are presented in Table 3.1. At baseline, we interviewed each trader about their experience in the industry, and basic demographic indicators. These results are presented in Panel A of Table 3.1. Traders operate at a small scale in terms of value. At average cocoa prices and 2011 exchange rates, the self-estimate of bags sold reports sales per trader at since the beginning of the season (at that point about half completed) at approximately $4,360.[9] Traders are experienced, with an average of 8.5 years in the industry and 6.5 years selling to the wholesaler. Their average age is 38. 58% have a cement or tile floor (as opposed to dirt or thatch) in their and 92% own a mobile phone versus 8% in the same survey. 83% have access to a storage facility. The third column of Table 3.1 shows that treatment and control are balanced on all trader-level covariates.

[9]This is calculated as the control group's average number of bags, 30.3, times the approximate pounds per bag, 180 times the average dollar price of cocoa over this period, Le. 3,200, divided by 4,000, the nominal exchange rate.

Table 3.1: *Trader summary statistics*

Covariate	Treatment	Control	Treatment - Control
Panel A: Baseline Interview			
Self-estimate bags sold in 2011	32.8	30.3	2.5
			(6.7)
Self-estimate grade A bags sold in 2011	20.0	18.6	1.4
			(5.4)
Age, years	38.2	36.9	1.3
			(2.1)
Years trading cocoa	8.1	8.9	-0.8
			(1.2)
Years selling to study wholesaler	5.7	7.3	-1.4
			(1.1)
Cement or tile floor in house $\in \{0,1\}$	0.53	0.62	-0.09
			(0.1)
Mobile phone owner $\in \{0,1\}$	0.90	0.93	-0.03
			(0.06)
Access to storage facility $\in \{0,1\}$	0.88	0.78	0.10
			(0.09)
Panel B: Pre-treatment shipment data			
Cocoa (pounds) sold during pre-treatment	2940	3180	-240
			(750)
Per pound farmer price for Grade A (Leones)[a]	3,121.7	3,121.1	-0.6
			(47.6)
Per pound farmer price for Grades B or C (Leones)[b]	3,045.6	3,040.3	-5.3
			(32.2)
Panel C: Baseline farmer listing			
Villages operating in	4.3	4.9	-0.6
			(0.4)
Number of farmers buying from	23.3	28.4	-5.1
			(3.5)
Mean number of farmers per village	5.8	5.6	0.2
			(0.8)
Share of farmers given credit since March	0.72	0.68	0.04
			(0.07)
Number of observations	40	40	

Notes: Standard errors allowing for unequal variance between groups in parenthesis. Treatment and control assigned randomly within pair of matched on self-estimates of grade A bags sold in 2011. [a] There are only 22 treatment observations of the grade A price in pre-treatment shipments, and 24 control. [b] There are only 30 treatment observations of a grade B or C price in pre-treatment shipments, and 34 control.

During the experiment, when traders arrived at the warehouse, our inspectors measured the quality of their shipment and, if the trader had been assigned to the treatment group, he or she was paid a bonus for grade A. Bags were then weighted and quantities recorded. We collected these data for three weeks before treatment assignments were announced. Panel B of Table 3.1 shows deliveries from this period. These results confirm that treatment and control traders are balanced on the volume of their business: treatment traders sold on average 2,940 pounds and control traders sold 3,180.

In these shipment data, we also collect the price per pound paid to farmers, and the name of the village in which the cocoa was purchased. Traders typically mix cocoa from different farmers of the same village in the same bag. Questions about which farmers in these villages the bags were purchased from would have been too time consuming to collect, and so farmer prices reported are the average purchase per unit purchase price in a village. Given concerns about measurement error in the farmer per pound price, we also collected information on the total amount paid to farmers for a given bag. Dividing this by its weight, we create an alternative measure of the per unit price for use in our results on price pass-through. Farmer prices reported in baseline in Panel B show that traders in treatment and control were balanced on the prices they paid to farmers, and confirm that average prices of grade A cocoa are larger than for grades B and C: in the control group the average price paid for grade A is Le. 3,121 and the average price paid for B or C is Le. 3,040.

Finally, in the baseline we asked traders to list each farmer they buy from and all of the villages in which they buy. For each farmer, we asked whether they had given the trader had given the farmer a loan. These results are shown in Panel C. The average trader operates in 4.6 villages, and buys from 25.9 farmers, on average 5.7 per village. Importantly, traders have given at least one loan to on average 70% of their clients since the beginning of the season. We repeated this listing exercise one and two months after the treatment began, in the final round asking the amount of the last loan. We will use these data to estimate our treatment effects on credit supply.

3.2.2 Villages

Study traders reported ever purchasing cocoa in 165 villages. Of these villages, 80 are used in the analysis, because these are the ones for which we have at least one observation of the grade A during the study period. Figure 3.1 presents a map of these villages along with the major towns, and the road network, which is unpaved. Panel A of Table 3.2 presents summary statistics from this sample of villages. On average, each village has 3.2 study traders, and 1.5 treatment traders. 34 of our 80 sample villages have at least one treatment and one control trader. This merits concern about spillover effects between treatment and control. We address this in Section 3.3.

As shown in Figure 3.1 the average road distance from a village to the nearest town is 9.6 miles using Dijsktra's minimum distance algorithm along the road network. Importantly, over all study traders, on average 65% of farmers selling to study traders have been given credit by at least one trader over the last year, highlighting the importance of interlinked transactions in this industry. Traders have multiple clients (famers) within each village, on average 6.2, and on average 18.7 clients per village. The population of these villages, calculated using the 2004 census, is substantially greater, at 494 people on average, with a large standard deviation of 753. Though we lack a direct measure of market share sizes, markets appear not to be greatly concentrated. In the baseline, study traders report having on average 5.7 regular competitors in a village. Randomization across traders randomly allocates treatment traders to sets of villages conditional on the number of study traders in the village. Since we will estimate heterogeneous treatment effects across villages, it is important to check whether villages are balanced in the composition of treatment and control traders. Panel B of Table 3.2 presents the coefficients of a regression of a village level covariate on the number of treatment traders and number of study traders as a test of balance. In all cases, the coefficient on the number of treatment traders is statistically insignificant at standard levels.

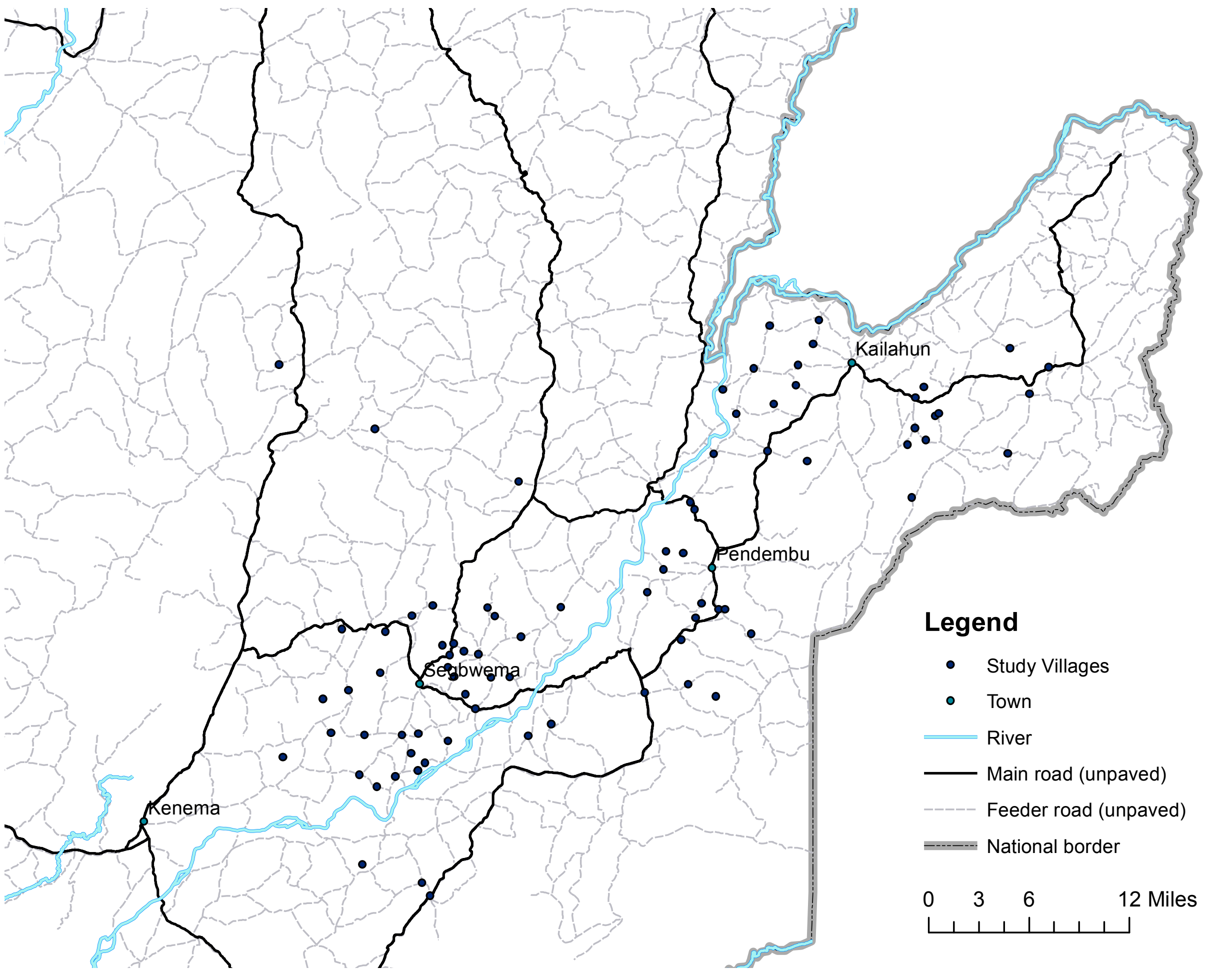

Figure 3.1: *Map of study villages*

Table 3.2: *Village summary statistics*

Village covariate	# of study traders	# of treatment traders	Miles to nearest town	Baseline credit share	Ave. # of farmers per trader	Ave. # of competitors per trader	Pop. per square mile	Total farmers reported by traders
Panel A: Sample averages								
Mean	3.2	1.5	9.6	0.65	6.2	5.7	494	18.7
	(2.4)	(1.5)	(5.7)	(0.29)	(4.3)	(3.9)	(753)	(16.4)
Number of observations	80	80	80	80	80	80	80	80
Panel B: Balance in count of treatment traders across sample villages								
# of treatment traders			0.79	0.01	0.39	0.41	74.39	0.02
			(0.66)	(0.03)	(0.39)	(0.42)	(62.86)	(1.50)
# of study traders			-0.28	0.02	-0.30	-0.03	-32.51	5.24
			(0.41)	(0.02)	(0.30)	(0.25)	(45.87)	(1.03)
Number of observations	80	80	80	80	80	80	80	80

Notes: Panel A shows means of study sample of villages with standard deviations in parenthesis. Panel B shows the coefficients in a regression of the covariate on the number of treatment traders, the number of study traders and a constant. Robust standard errors are presented in parentheses in panel B. Miles to nearest town calculated using Dijkstra's minimum distance algorithm along the network of rural feeder roads. Baseline credit share and total number of clients reported in a baseline listing of all clients in each village. Average number of clients and competitors reported by traders in a baseline survey. Population per square mile calculated over all 2004 census enumeration areas within a 1/5th mile radius of the GPS coordinates of the village. Cost constraints prohibited collecting GPS coordinates for the population of villages.

3.3 Experimental Results

In this section, we present the average treatment effect results from our experiment. We first document the negligible effect of the bonus on prices paid to farmers and show that the lack of price pass-through cannot be explained by increasing marginal costs of transport. We then show that the traders respond to the bonus by increasing credit provision to farmers. In section 5, we complement these results with an analysis of heterogeneous treatments across villages motivated by the theoretical framework we develop in section 4. Throughout the paper, the standard errors we report are robust to heteroskedasticity and are estimated with two non-nested clusters that allow for arbitrary correlation across observations from a given village, and across observations from a given trader.[10]

3.3.1 Price Pass-Through

To estimate pass-through in prices, we run the following regression, where an observation is a shipment k delivered by trader i of pair p, from village v in week t:

$$y_{kipvt} = \alpha_{ip} + \tau_t + \theta^p(\text{Bonus}_i) + \mathbf{X}_i'\boldsymbol{\beta}_{\boldsymbol{x}} + \mathbf{W}_v'\boldsymbol{\beta}_{\boldsymbol{w}} + \epsilon_{kiptv} \tag{3.1}$$

The term α_{ip} is a fixed effect for each matched pair in the randomization. Since pairs were matched within wholesalers, this effectively controls for the town in which the trader sells his cocoa. The term τ_t is a week fixed effect, to capture time varying factors in supply, such as weather, as well as any variation in the expectation of the wholesaler price generated by the international market. The vector $\mathbf{X}_i'$, used in some specifications, includes the trader-level covariates of age, years working with wholesaler and dummies for ownership of a mobile phone and a concrete floor. The latter term is a useful proxy for wealth in our context. The vector $\mathbf{W}_v'$ includes the village-level covariates of minimum road distance to nearest town, number of study traders, number of other treatment traders, and baseline share of farmers having ever received credit from a study trader. Bonus_i is a dummy equal to one if trader i

[10]This clustering approach follows Cameron, Gelbach and Miller (2006).

is assigned to treatment, and so θ^P is the average treatment effect. The term ϵ_{iptv} is an error. Pairwise randomization motivates the assumption that $E[\epsilon_{kiptv}|\text{Bonus}_i, \alpha_{ip}, \tau_t, \mathbf{X}_i', \mathbf{W}_v'] = 0$.

The term θ^p is the coefficient of interest. We will pick y_{kipvt} to be the level of prices so if $\theta^p = 150$, we have perfect pass-through, as the treatment traders will have increased the price paid to farmers by the full amount of the bonus. Table 3.3 presents estimates of θ^p. In the basic specification in column 1, with no village or trader covariates, pass-through is statistically indistinguishable from zero with a point estimate of $\theta^p = -5.4$ (s.e. = 14.9). Even at the upper bound of a 95% confidence interval, pass-through would be just 24 Leones, less than one fifth of the amount of the bonus, 150 Leones. This extremely low level of pass-through to farmgate prices is consistent with evidence from West Bengal in Mitra et. al (2013).

Given that some villages contain both treatment and control traders, we are concerned about spillovers between groups. It may be that Bertrand competition between treatment and control traders drives up the price offered by control traders, so that there is no difference between the prices offered by both groups. We test for this by adding the number of other study traders and other treatment traders in column 2. This is a specification similar to the one developed to test for externalities by Kremer and Miguel (2004). The number of other treatment traders in the market does not significantly affect the price paid to farmers. This is confirmed by column 3, where we interact the two market-level variables with the treatment status of the trader. We find some evidence that treatment traders pay a higher price when there are other treated competitors in the market, but the coefficient on the interaction is not significant at conventional levels. Column 4 confirms the results using the entire vector of village-level and individual-level controls. Column 5 presents the same regression using as our outcome an alternative measure of price taken by dividing a trader's total expenditure by its weight. We can again reject perfect pass-through and cannot reject that pass-through is equal to zero. This provides reassurance that our price results are not driven by measurement error in prices. In columns 6 and 7, we also present results for grade B and C cocoa. For grade B, in column 6, pass-through is positive and statistically

Table 3.3: *Farmer price response*

Variable	(1) Price of Grade A	(2) Price of Grade A	(3) Price of Grade A	(4) Price of Grade A (alternative)	(5) Price of Grade A	(6) Price of Grade B	(7) Price of Grade C
Bonus $\in$ {0,1}	-5.4	-5.5	-11.0	-4.1	-3.6	37.95***	4.64
	(14.9)	(13.8)	(19.5)	(11.1)	(15.1)	(12.71)	(19.54)
# Other Bonus Traders		8.8	3.3	6.7			
		(7.4)	(10.0)	(6.3)			
Bonus*# Other Bonus Study Traders			7.7				
			(10.7)				
# Other Study Traders		-10.7*	-9.8	-7.9			
		(5.7)	(7.4)	(5.3)			
Bonus*# Other Study Traders			-1.6				
			(6.6)				
R-squared	0.9	0.9	0.9	0.9	0.9	0.88	0.86
Observations	1,090	1,090	1,090	1,090	1,090	527	226
Village Controls	NO	NO	NO	YES	NO	NO	NO
Trader Controls	NO	NO	NO	YES	NO	NO	NO

Notes: Robust standard errors allowing for two-way clustering at the village and trader level are shown in parenthesis. All specifications include calendar week and randomization pair fixed effects. Village controls include minimum road distance to nearest town, number of study traders, number of other treatment traders, and baseline share of farmers having ever received credit from a study trader. Trader controls are age, years working with wholesaler, and dummies for ownership of a mobile phone and a concrete floor. An observation is a shipment delivered to a wholesaler, and prices are per pound in Leones. The bonus is an increase in the trader resale price of 150 Leones per pound for grade A only, and so perfect pass-through would imply a coefficient of 150 on the bonus indicator in columns 1, 2 and 3. The alternative measure of price in column 5 is the total price paid to the farmer divided by weight of shipment. There were approximately 4,000 Leones to the U.S. dollar at the time of the study. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

significant at 1 percent, with $\theta^p = 49.41$ (s.e. = 11.76). At first pass this is surprising. How could pass-through for grade A be zero, and pass-through for grade B be so much higher? We argue that this is consistent with Type I error on the part of traders, who observe quality imperfectly. The bonus has increased the expected price for grade A quality cocoa relative to grade B. Even if pass-through is zero for a given quality, if quality is imperfectly observable traders will now be more willing to pay the grade A price premium for cocoa that has some probability of being grade A. That in column 7 we see pass-through is again reduced for grade C, which is much less likely to be mistaken for grade A, confirms this intuition.

In a perfectly competitive model of spatial arbitrage, the difference in price between two locations that trade will equal the marginal cost of transport. Thus, our lack of pass-through could be explained by the fact that marginal costs of transport are increasing rapidly among treatment traders, who now bring more cocoa per shipment. Table 3.4 presents estimates of equation (3.1), with outcomes related to cost in the place of y_{ipvt}, using all grades of cocoa shipped. Columns 1 and 2 show that the weight per shipment—scale—is increasing significantly by the treatment. In the preferred specification with village and trader controls we have that treatment traders increase their shipment volume by 8.20 (s.e. = 2.37). In columns 3 and 4, we see that unit costs reported by the trader are also falling. When shipments arrived, we recorded the total cost of transporting of that shipment to the wholesaler. Unit costs are obtained by dividing this by weight. In the preferred specification with village and trader controls we have that the treatment effect is -11.38 Leones (s.e. = 2.38). This implies that in addition to the bonus of 150 Leones per unit, traders also received a gain in the form of 11 Leones lower transport costs per pound shipped. Finally column 5, which amounts to a linear probability model in which the outcome variable is a dummy indicating that a truck and not a motorcycle was used to transport the cocoa, shows that this cost result is being driven by a change in transport technology. As shown in Figure 3.2, which plots unit cost by transport technology, trucks have consistently lower per unit costs. It is our understanding from conversations with traders that trucks are easier to hire if one has a larger total shipment size, as the truck driver himself must amortize the fixed

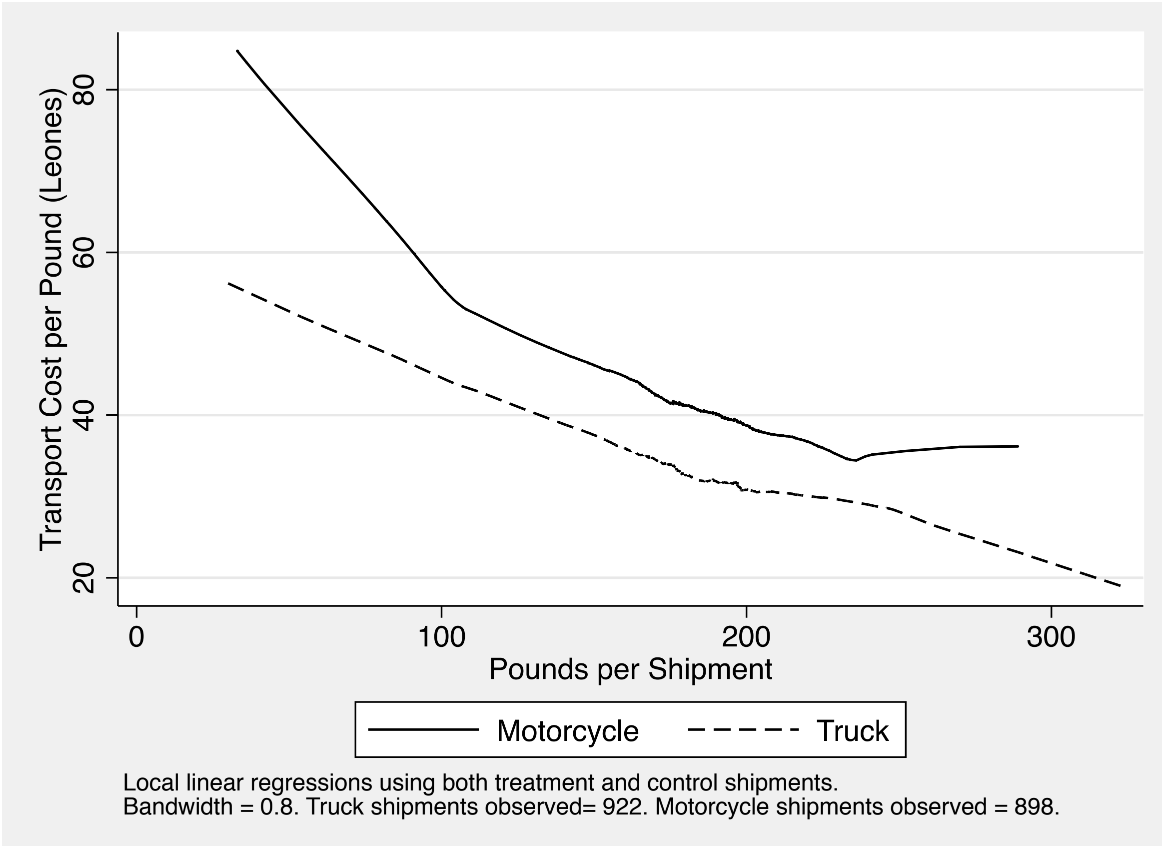

Figure 3.2: *Declining unit costs by transport technology*

cost of driving to the pick up location. These results show that the lack of pass-through cannot be explained by increasing marginal costs. They also confirm those of Fafchamps et. al. (2005) for Africa, who find evidence of increasing returns to scale in transport in Malawi and Madagascar, though not in Benin.

3.3.2 Credit

To investigate the effects of the bonus on credit, we estimate the following regression, which is a modified version of (3.1):

$$y_{fipv} = \alpha_{ip} + \theta^c(\text{Bonus}_i) + \mathbf{X}_\mathbf{i}'\boldsymbol{\beta}_x + \mathbf{W}_\mathbf{v}'\boldsymbol{\beta}_w + \varepsilon_{\mathbf{fipv}} \qquad (3.2)$$

Table 3.4: *Transport cost and technology choice response*

VARIABLES	(1) Pounds per shipment	(2) Pounds per shipment	(3) Unit cost	(4) Unit cost	(5) Truck use $\in \{0,1\}$
Bonus $\in \{0,1\}$	7.75***	8.20***	-10.26***	-11.38***	0.21***
	(2.18)	(2.37)	(2.58)	(2.38)	(0.05)
R^2	0.21	0.22	0.44	0.47	0.46
Number of observations	1,837	1,837	1,837	1,837	1,837
Trader Controls	NO	YES	NO	YES	YES
Village Controls	NO	YES	NO	YES	YES

Notes: Robust standard errors allowing for two-way clustering at the village and trader level are shown in parenthesis. An observation is a shipment delivered to a wholesaler. All specifications include calendar week and randomization pair fixed effects. Village controls include minimum road distance to nearest town, number of study traders, number of other treatment traders, and baseline share of farmers having ever received credit from a study trader. Trader controls are age, years working with wholesaler, and dummies for ownership of a mobile phone and a concrete floor. The control mean of pounds per shipment is 178 (s.d. = 36), Leones unit cost, 45 (s.d. = 27), and likelihood of truck use, 0.34. *** p$<$0.01, ** p$<$0.05, * p$<$0.1

This is a farmer-level regression. Specifically, y_{fipv} is an indicator of whether farmer f in village v was given credit by trader i of pair p. The term θ^c is the treatment effect estimator, and ν_{fipv} is an error term. All other terms are as in (3.1), except that we have removed the week fixed effect, since we only observe the outcome variable in one cross section. Pairwise randomization again motivates the assumption that $E[\nu_{fipv}|\text{Bonus}_i, \alpha_{ip}, \mathbf{X}'_i, \mathbf{W}'_v] = 0$.

Table 3.5 presents estimates of the θ^c in equation (3.2). The treatment effect on credit is substantial: in the month before the interview, farmers reported as regular suppliers by treatment traders in the baseline listing are 14 percentage points more likely to receive credit from these traders relative to a control mean of 11 percentage points. In column 1 we run a linear probability model where the outcome is a dummy equal to one if credit was provided to a farmer in the last month. Columns 2 and 3 control again for potential spillovers. In particular, in column 3, the coefficient θ^c captures the difference between the credit provision likelihood between a treatment and a control trader, in a market with no other competitor. Column 4 shows that the results are robust when adding trader and village controls. In column 5 we see this in terms of Leones. Here, traders were asked after two months of treatment the amount of the loan last given to the farmer, if any was given in the past month. Those that did not give any were set equal to zero. We see that traders are more than doubling their credit outlay, with $\theta^c =$ Le. 9.771 (s.e. = 5,209), off a control group mean of Le. 18,908.

3.4 A Model of Pass-through in Interlinked Transactions

In this section, we develop a simple model of interlinked transactions between intermediaries and producers, who we will call traders and farmers to match the context of our experiment. We derive the equilibrium output prices paid to farmers and the level of credit provision. We show how these variables respond to an increase in the wholesale price—the price at which traders resell their purchases—and conclude by studying the average welfare effects of such a change on farmer welfare in this model. We contrast them with those derived from a benchmark model that does not account for interlinkages.

Table 3.5: *Credit response*

Variable	(1) Lent $\in \{0,1\}$	(2) Lent $\in \{0,1\}$	(3) Lent $\in \{0,1\}$	(4) Lent $\in \{0,1\}$	(5) Loan Amount
Bonus ∈ {0,1}	0.14***	0.14***	0.11*	0.12***	9,771*
	(0.03)	(0.03)	(0.07)	(0.03)	(5,209)
# Other Bonus Traders		-0.01	-0.00	-0.01	
		(0.01)	(0.02)	(0.01)	
Bonus × # Other Bonus Study Traders			-0.01		
			(0.02)		
# Other Study Traders		0.01	0.00	0.01	
		(0.01)	(0.01)	(0.01)	
Bonus × # Other Study Traders			0.01		
			(0.02)		
R-squared	0.31	0.31	0.31	0.32	1
Observations	1,541	1,541	1,541	1,541	1,541
Trader Controls	NO	NO	NO	YES	NO
Village Controls	NO	NO	NO	YES	NO
Control Mean	0.119	0.119	0.119	0.119	18908

Notes: Robust standard errors allowing for two-way clustering at the village and trader level are shown in parenthesis. An observation is a farmer listed by the trader in the baseline. The dependent variable in columns 1 - 4 is an indicator for whether the trader had lent to the farmer since the treatment began, asked a month after the treatment began. The control mean of this dummy was 0.12. The dependent variable in column 5 is the amount of money lent in the past month. The control mean of this amount was 18,908 (s.d. = 52,597). All specifications include randomization pair fixed effects. Village controls include minimum road distance to nearest town, number of study traders, number of other treatment traders, and baseline share of farmers having ever received credit from a study trader. Trader controls are age, years working with wholesaler, and dummies for ownership of a mobile phone, and a concrete floor. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

3.4.1 Environment

The economy is composed of M isolated markets. Each market m is populated by I_m farmers and J_m traders. Interactions between the two agents can occur under two alternative contracts. In *spot markets*, the trader and farmer transact only at harvest time. This contract can be viewed as the benchmark model, with no interlinkages. In *interlinked transactions* (ILT), a trader provides credit before harvest and then purchases output from the farmer at harvest, unless the farmer chooses to default.

Traders are homogeneous. Farmer i in market m is endowed with an amount of land l_i. This varies across farmers. If the farmer does not receive credit, the land produces quantity q_{im}, which we will write as q_i to reduce notation. If the farmer receives credit, the land is assumed to produce $q_i(1+r)$. We can also interpret q_i as quality-adjusted quantity. In this case, which is particularly relevant for our setting, credit allow farmers to improve quality of their output, through post-harvesting processing. We assume farmers and traders are risk-neutral and maximize their (expected) profits. We now discuss the payoff structure under the two tyes of contracts

Spot Markets

Spot market prices in market m, p_m^S, are equal to the wholesale price over a constant markdown,

$$p_m^S = \frac{w_m}{\mu}, \quad \mu \geq 1 \tag{3.3}$$

The mark down is a reduced form capturing both trading costs proportional to value and any market power the trader may have. The utility of farmer i farmer transacting on the spot market is:

$$u_{im}^S = p^S q_i = \frac{w_m}{\mu} q_i \tag{3.4}$$

On the spot market, we assume that each trader has an equal probability to secure cocoa from a given farmer in market m equal to $1/J_m$. Therefore, the expected utility for trader j

transacting with farmer i in market m is:

$$v^S_{jim} = \frac{1}{J_m}(w_m - p^S_m)q_i = \frac{1}{J_m}\frac{\mu - 1}{\mu}w_m q_i \tag{3.5}$$

Interlinked Transactions

In interlinked transactions, a trader provides credit before harvest, and subsequently purchases output at a pre-determined price. We assume that traders and farmers are randomly matched and that only one trader can offer credit to a given farmer.[11] In order to provide credit, the trader incurs a fixed cost f. This can be interpreted as the minimum amount of screening and monitoring that trader needs to undertake, independent of the amount of credit outlay (for a review of these issues, see Banerjee, 2002). We assume that if the two parties enter an interlinked transaction, the trader provides a fixed amount of credit per each bag denoted by c. As we mentioned above, this raises the farmer output by a factor $(1 + r)$.

In addition to the impact of credit on production levels, we assume that the farmer marginal utility from consuming (part of the) loan c before harvest is $c_F = \lambda \cdot c$, with $\lambda \geq 1$. This is a reduced form for the increased utility of the farmer from extra pre-harvest consumption, which is assumed to be weakly larger than the trader's utility cost of disbursing the loan. One way to think about this is that the farmer experiences a higher marginal utility per unit of income relative to the trader in the pre-harvest season.

After receiving credit, the farmer decides whether to stick to the terms of the contract or to undertake a strategic default. If the farmer respects the contract, he receives a farmer-specific contract price p^C_{im}. We describe how this price is determined in equilibrium in section 3.4.2. When he does not default, the utility of farmer i under ILT is

$$u^{CN}_{im} = (p^C_{im}(1 + r) + c_F)q_i \tag{3.6}$$

[11]This is consistent with anecdotal evidence in our setting that farmers mostly receive advances from one trader. However, we note that the static nature of the game obviously neglects the relational component of credit provision, which is potentially very relevant for our setting.

and the utility for trader j in an ILT contract with farmer i is

$$v_{jim}^{CN} = ((w_m - p_{im}^{C})(1+r) - c)q_i - f \tag{3.7}$$

Note here that the trader's utility no longer includes the term $1/J_m$, since in the interlinked transaction contract, he is now certain to get the farmer's output (if the contract is enforced).

The benefit of strategic default for the farmer depends on the underlying contracting institutions in market m. Specifically, we assume that, if the farmer defaults, he loses a share γ_{im} of his output. He then sells $(1-\gamma_{im})q_i(1+r)$ on the spot market. The parameter γ_{im} is a reduced form measure of contracting institutions capturing market characteristics that could shape the cost of default, including trader monitoring costs, proximity to law enforcement, and social norms specific to farmer i. Thus, another benefit of the credit provision for the trader, in addition to the higher quantity to purchase, is that the credit contract "locks in" the farmer's supply by imposing a penalty γ to the farmer if he side-sells to other traders.[12]

In the strategic default scenario, the utility of farmer i in the ILT contract is

$$u_{im}^{CD} = \left(p_m^S(1-\gamma_{im})(1+r) + c_F\right) q_i = \left(\frac{w_m}{\mu}(1-\gamma_{im})(1+r) + c_F\right) q_i, \tag{3.8}$$

while the utility of trader j is

$$v_{jim}^{CD} = -c \cdot q_i - f \tag{3.9}$$

3.4.2 The Equilibrium Contract

In this subsection, we describe the conditions under which farmers and traders will opt to transact on the spot market or in an ILT, and the conditions under which ILT will persist as a subgame-perfect Nash-equilibrium. The timing of the game is as follows: in the first stage, the trader is randomly matched to the farmer, and decides whether to offer credit. He also decides the terms of the contract—the contract price at which output will be sold after harvest, p_{im}^{C}. If the trader does not offer credit they transact on the spot market. If the

[12]The main insights of the framework hold if γ is modeled as a constant loss, as opposed to a proportion of the output.

trader does offer credit, they proceed to the second stage and the farmer decides whether to accept or not. In the third stage, the farmer decides whether to default or not, conditional on having accepted the ILT. We solve the model by backwards induction. In the third stage, the farmer decides not to default if

$$u_{im}^{CN} \geq u_{im}^{CD} \Rightarrow p_{im}^{C} \geq \frac{w_m}{\mu}(1-\gamma_{im}). \tag{3.10}$$

This is the farmer's incentive compatibility constraint. In order not to induce default, the trader must offer a large enough contract price to satisfy it.

In the second stage, if the trader offers credit and a contract price in the first stage, the farmer must decide whether to accept it. The farmer accepts credit if

$$\max\left(u_{im}^{CN}(p_{im}^{C}), u_{im}^{CD}\right) \geq u_{im}^{S}, \tag{3.11}$$

which highlights the fact that the decision to participate in credit depends on the proposed contract price, p_{im}^{C}. In order to simplify the presentation, we assume that the farmer always prefers to accept a contract and default than to reject a credit offer. This amounts to assuming that $\gamma < \frac{r}{1+r} + \frac{c_F \mu}{w} \equiv \gamma_A$. We argue that the case where contracting institution quality are low enough — relative to the consumption smoothing and productive role of credit — is the relevant one for our setting.[13]

In the first stage, the trader decides whether to offer credit and, if so, the contract price to offer. For simplicity, we focus on the equilibrium where the trader extracts all the ex-post surplus from the relation by setting a price such that the farmer's incentive compatibility constraint (3.10) binds with equality.[14] Thus, the price of the ILT contract is:

$$p_{im}^{C^*} = \frac{w_m}{\mu}(1-\gamma_{im}), \tag{3.12}$$

which is decreasing in the quality of contracting institutions. Given this optimal price, he

[13]The insights from the comparative static analysis hold when we look at the case of $\gamma > \gamma_A$

[14]The main comparative statics results described in Section 3.4.3 do not depend on this particular equilibrium selection criterion

decides to offer credit to farmer i if

$$v_{jim}^{CN}(p_{im}^{C^*}) \geq v_{jim}^{S} \tag{3.13}$$

Therefore, interlinked transaction contracts arise as the equilibrium contractual form if the inequality in equation 3.13 holds. The equilibrium price in the interlinked transaction contract is described by Equation 3.12.

To build intuition for our empirical results, we consider the case where farmers vary by their production-level under the no-credit case, q_i, while contract institutions vary only at the market level, denoted γ_m, and the spot-market markdown, μ, is constant across markets. In this case, farmer i and trader j in market m enter an ILT arrangement if: i) $q_i \geq q_m^* \equiv \frac{fJ_m\mu}{w_m(J(1+r)-1)(\mu-1)+J(1+r)\gamma w_m - cJ_m\mu}$; ii) $\gamma > \frac{cJ\mu-(J(1+r)-1)w(\mu-1)}{J(1+r)w}$ Intuitively, traders provide credit only to those farmers whose quantities are large enough that the increase in revenues for the traders to provide credit more than offset the fixed cost of credit provision f. In addition, observe that, under the parameter restriction on γ_m imposed by (ii) above, the minimum production volume a farmer needs to produce to access credit is decreasing in J_m and γ_m, and increasing in f. Intuitively, credit provision increases when the relative benefit for the trader from interlinked transactions increases. This occurs when: a) the number of competitors increases (and thus the expected profit from spot market interactions decreases); b) the quality of contracting institutions increases (and thus the contract price the trader has to offer to induce no-default falls); c) the fixed cost from credit provision decreases. The equilibrium contractual form determines the price each farmer faces. Farmers on the spot market sell at $p_i = p_m^S$. Farmers in an ILT arrangement sell at $p_i = p_{im}^{C*}$

3.4.3 Comparative Statics

In this section, we study the impact of an increase in the wholesale price at which the trader can resell the output bought from farmer, w. The difference between the trader profit under interlinked transaction and spot markets, $v_{jim}^{CN} - v_{jim}^{S}$, is easily shown to be increasing in w_m. Those farmer-trader pairs for which the difference changes from negative to positive,

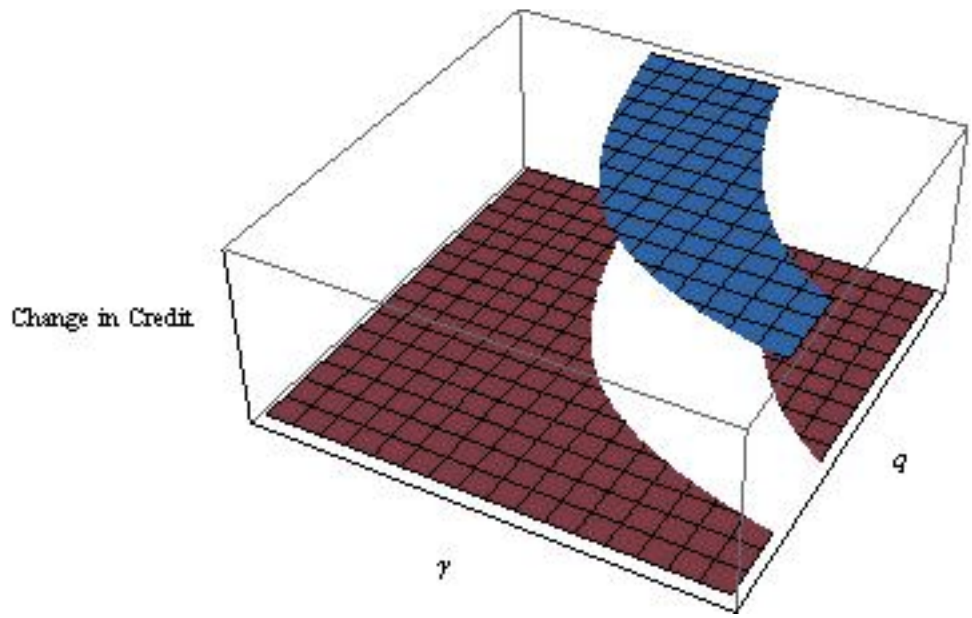

Figure 3.3: *The impact of a change the wholesale price on credit provision. Red indicates negative values, blue positive.*

will switch from spot market to interlinked transactions

We consider again the case where q_i is the only parameter varying across farmers in a given market. The threshold level of production required for the farmer to receive credit in a given village, q_m^*, decreases in response to an increase in w_m. As w_m increases, a greater share of farmers switch to ILT. Figure 3.3 shows the increase in credit status in response to a discrete increase in the wholesale price, as a function of γ and q.

We now consider how p_i changes in response to a change in w_m from w_m^0 to $w_m = w_m^0 + \Delta$, where Δ is some constant. Observe that the direction and the magnitude of the change depends on whether the farmers are on the spot market, in an ILT, or whether they switch into ILT in response to the change in w_m. First, farmers who remain on the spot market experience an increase in their price, p_m^S, of $\frac{\Delta}{\mu}$. Second, farmers who were in ILT contracts both before and after the change in w_m experience an increase in their contract price, p_{im}^C, of $\frac{\Delta(1-\gamma_m)}{\mu}$. Third, farmers who enter an ILT contract in response to the increase in w_m face the following change in price: $\frac{\Delta(1-\gamma_m)-\gamma_m w_m^0}{\mu}$. The last result shows that farmers switching into

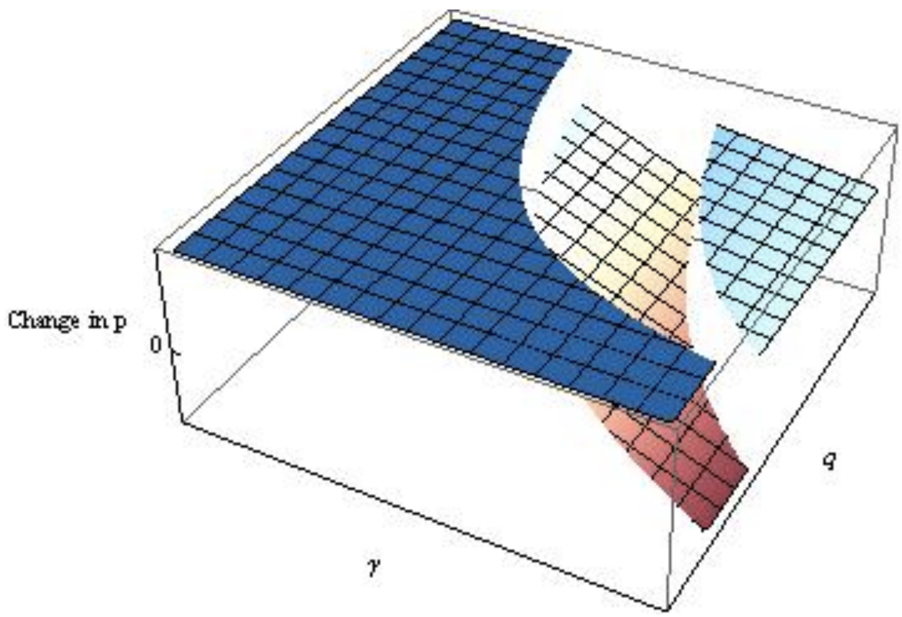

Figure 3.4: *The impact of a change in the wholesale price on farmer prices. Red indicates negative values, blue positive.*

ILT in response to the change experience a decrease in price, since they switch from the spot price to the contract. They are however still better off, because they are now receiving credit.

Figure 3.4 summarizes the price response for different combinations of γ and q. The graph points at the main result from the theory. The credit and pricing pass-through in response to a change in the wholesale prices are substitutes. Those pairs (γ, q) that switch into ILT are exactly the ones experiencing a decrease in output prices. For the remaining pairs, who stay in the same contract form, the price goes up.

3.4.4 Welfare Analysis under Interlinked Transactions

In this section we use our model to measure the effect of a change in wholesale prices on average farmer welfare and compare this estimate to that of a benchmark model without interlinked transactions. We simplify the exposition by assuming contract institutions are fixed across farmers at the market level, γ_m, and that q_{im}, the farmer-level output produced *if the farmer does not receive credit*, is distributed with cumulative distribution function $G(q_{im})$

defined over the support $[q_{Lm}, q_{Hm}]$. Thus there exists a $q^*_m(w_m)$ which represents the production volume threshold above which farmers enter ILT contracts in market m. The insights of the analysis are similar when we relax the assumption of a constant γ_m in the market. Below, we consider a single market and simplify notation by omitting the subscript m.

The average farmer welfare in the market is equal to:

$$W^{ILT} = \int_{q_L}^{q^*} q\frac{w}{\mu}g(q)dq + \int_{q^*}^{q_H} q\left(\frac{(1+r)(1-\gamma)w}{\mu} + c_F\right)g(q)dq, \tag{3.14}$$

which is simply a weighted average of the welfare of farmers under the spot market and ILT contracts. The first term represents the welfare of farmers transacting on the spot market and the second is the welfare for farmers who are in ILT contracts. To find the average welfare effects of an increase in the wholesale price w, we apply Leibniz's rule to obtain

$$\begin{aligned} \frac{dW^{ILT}}{dw} &= G(q^*)E[q|q < q^*]\frac{1}{\mu} + (1-G(q^*))E[q|q \geq q^*]\frac{(1+r)(1-\gamma)}{\mu} \\ &\quad - \left(c_F + (r(1-\gamma)-\gamma)\frac{w}{\mu}\right)q^*g(q^*)\frac{dq^*}{dw} \end{aligned} \tag{3.15}$$

Intuitively, the change in welfare is the sum of three terms: i) the average change in welfare for farmers on the spot market, weighted by the share of farmers on the spot market; ii) the average change in welfare for farmers on ILT contracts, weighted by the share of farmers in ILT; iii) the change in welfare for farmers that switch into ILT in response to the change in welfare. Notice that $\frac{dq^*}{dw} < 0$, and so the last term is positive.

We compare these welfare results with those of a benchmark model, where farmers only transact on the spot market. In this benchmark case average welfare is

$$W^S = \int_{q_L}^{q_H} q\frac{w}{\mu}g(q)dq, \tag{3.16}$$

and the change in average welfare in response to a change in wholesale price is

$$\frac{dW^S}{dw} = E[q]\frac{1}{\mu} \tag{3.17}$$

Equation (3.17) shows that, in the absence of credit provision, price pass through is the key variable that shapes predicted welfare changes. This simple result obviously relies on the assumption of perfectly inelastic supply at the farmer level, and our constant markdown, μ. A more general model would have additional predictions about how the welfare response to an increase in the wholesale price varies with the supply elasticity.

3.5 Model Testing and Calibration

3.5.1 The Substitutability of Price and Credit Pass-through

The core testable prediction of our model is that price pass-through and credit pass-through (on the extensive margin: the share of farmers receiving credit) are substitutes across markets. In this section we provide evidence on this substitutability using data from the markets in our experiment. A key result will be that while price pass-through is positive in some markets, it is low and even negative in others. This heterogeneity is driven by the baseline intensity with which interlinked transactions are used in these markets. This finding emphasizes the core result of our paper, that price-pass-through is diminished in the presence of interlinked transactions. An increase in the wholesale price leads to a reduction in farmer prices for those farmers who switch into interlinked transactions as a result of the increase. We conclude this section with a welfare calibration of our model, showing that the effects of ignoring the margin of credit pass-through can be substantial for welfare analysis.

First, we document that several predictions of the model are supported by baseline correlations in the data. Table 3.6 presents these results. Each column shows a regression of the village outcome on the proxies listed in the rows and a constant. These indicators come from the pre-treatment period, when our inspectors were collecting data on quantity, prices, and quality of cocoa delivered to wholesalers, but treatments had not been assigned. Our model predicts a negative cross-sectional correlation between the level of prices and the supply of credit at the village level. This is confirmed in column 1 of Table 3.6. Moving from zero credit share to full credit share decreases the price paid in the cross section by 4.2 percent (s.e. = 2.3). In column 2, we present a regression where the outcome variable

Table 3.6: *Substitutability, baseline correlations*

VARIABLES	(1) Baseline Credit Share	(2) ln(Price of grade A)	(3) Share of grade A
Quantity Cocoa/Farmers	0.077***		
	(0.027)		
Number of Study Traders in Village	0.024*	-0.001	-0.025*
	(0.013)	(0.002)	(0.014)
Number of Study Traders	-0.005	-0.001	0.003
	(0.005)	(0.001)	(0.007)
Baseline Credit Share		-0.042*	0.339**
		(0.023)	(0.139)
R-squared	0.08	0.09	0.11
Observations	80	44	75

Notes: Robust standard errors in parentheses. Observations in columns are limited by the number of villages in which we observe one price in the pre-treatment period, and include those villages for which we did not observe prices post-treatment *** p<0.01, ** p<0.05, * p<0.1

is the village-level share of farmers receiving credit from study traders. This outcome is positively correlated with the total quantity supplied to traders, which is calculated as total quantity supplied in baseline, divided by the number of study traders. This confirms the model's insight that traders are more likely to extend credit to villages from which they can get higher volumes of cocoa. In addition, village-level credit share is also positively (and significantly) correlated with the number of traders in the economy, confirming another prediction of our model that markets with more traders, and thus a greater threat of loosing output to a competitor, are related to more intensive use of interlinked transactions. Finally, column 3 shows that the ratio between the volume of a high quality cocoa and the total volume of cocoa produced in the village is increasing in the level of credit provision. This is consistent with the model insight that loans from traders can be used to increase the value of output through quality-enhancing post-processing activities.

We then test the prediction that price and credit pass-through are substitutes across markets by studying across-village heterogeneous treatment effects of the experiment. We modify equations (3.1) and (3.2) respectively to allow for heterogeneous treatment effects across villages by specifying the two regression equations

$$y_{kipvt} = \alpha_{ip} + \tau_t + \theta^P(\text{Bonus}_i) + (\text{Bonus}_i \times \mathbf{W}'_v)\boldsymbol{\theta}^{\boldsymbol{P}}_{\boldsymbol{w}} + \mathbf{W}'_v\boldsymbol{\beta}_{\boldsymbol{w}} + \mathbf{X}'_i\boldsymbol{\beta}_{\boldsymbol{x}} + \epsilon_{kiptv} \tag{3.18}$$

$$y_{fipv} = \alpha_{ip} + \theta^C(\text{Bonus}_i) + (\text{Bonus}_i \times \mathbf{W}'_v)\boldsymbol{\theta}^{\boldsymbol{C}}_{\boldsymbol{w}} + \mathbf{W}'_v\boldsymbol{\beta}_{\boldsymbol{w}} + \mathbf{X}'_{\mathbf{i}}\boldsymbol{\beta}_{\boldsymbol{x}} + \,^{\circ}\,_{\mathbf{fipv}} \tag{3.19}$$

where $\mathbf{W}'_v$ is some vector of village covariates as before. For any village v then we have an estimator for the pass-through in that village given by

$$\hat{\rho}^s = \mathbf{W}'_v\boldsymbol{\theta}^{\boldsymbol{s}}_{\boldsymbol{w}} + \theta^s \tag{3.20}$$

where $s \in \{P, C\}$, price or credit. Our model predicts that $\hat{P} = Corr(\hat{\rho}^P, \hat{\rho}^C) < 0$ across villages with different market size, the level of competition, and the ease of enforcing contracts.

The most natural starting point for an element of $\mathbf{W}'_v$ is a baseline measure of credit supply at the village level, which proxies directly for the prevalence of interlinked transactions in a given market. As discussed above, we are able to generate such a measure from our baseline survey of traders. For all the farmers listed by study traders in the baseline, we calculate the within-village share of farmers that has been given credit by a study trader since March, shortly before the cocoa season begins. As shown in Table 3.2, the mean of this variable across study villages is 65%, with a standard deviation of 29 percentage points. To gain more variation we will add to $\mathbf{W}'_v$ the minimum road distance to the nearest town, a proxy for ease of enforcement of contracts, as well as factors that may affect per unit transport costs, and the number of study traders in the village, a proxy for market size.

We proceed to estimate to estimate (3.18) and (3.19) using seven different specifications of $\mathbf{W}'_v$, and to report $\hat{P}_r = Corr(\hat{\rho}^P_r, \hat{\rho}^C_r)$ for each specification r. The regression coefficients $\boldsymbol{\theta}^{\boldsymbol{s}}_{\boldsymbol{w}}$ and θ^s for each model are presented in Table 3.7, and the estimates of $\hat{P}_r$ are presented in Table 3.8. Standard errors for each $\hat{P}_r$ are constructed by resampling with replacement both the price and credit sample 500 times, each time reestimating (3.18) and (3.19). For each iteration, we estimate $\hat{\rho}^s_r$ for each s, and hold the correlation in memory. The standard error is the standard deviation of these correlations, and the test statistic is the $\hat{P}_r$ divided by its standard error.

The models, or different specifications of $\mathbf{W}'_v$, are as follows. Model A includes only baseline credit share. In columns 1 and 2 of Table 3.7 we see that the interaction terms have the opposite sign, as expected. In markets with higher baseline credit share, price pass-through is lower, though not significantly so, and credit pass-through is higher, and significant at 1 percent. Since there is only one variable in this model, the correlation coefficient presented in Table 3.8 is necessarily -1.0. Model B is our baseline linear model, including interactions terms with baseline credit share, the number of study traders, the number of treatment traders and minimum distance to the nearest town. Again, the signs on baseline credit share are the opposite sign, as expected. The correlation coefficient across villages generated by this model is -0.88. With a standard error of 0.24, this is significant at 1%. A scatter plot and best fit line indicating the relation between the two predicted treatment effects is presented in Figure **??**. The negative relationship is clear. Model C includes the same covariates, but includes quadratic and cubic terms for each to allow for non-linearity in the treatment effects as they vary with the covariates. In this model, the correlation coefficient is smaller, at -0.44, but still significant at 6 percent with a standard error of 0.23. Adding the quadratic and cubic terms increases the variance of the overall predicted effects, but does not change the overall pattern.

We note that the slope of the graph is quite consistent with the baseline correlation between prices and village-level share of farmers who receive credit. Specifically, a one-tenth increase in the treatment effect on the share of credit provision reduces the treatment effect on cocoa prices by about twenty Leones. The baseline correlation, based on pre-experiment data, would predict that such an increase in credit share should reduce prices by about fifteen Leones.

The rest of the models drop covariates other than baseline credit share to show the robustness of the relationship. Model D includes only baseline credit share and minimum road distance, and model E includes these variables with quadratic and cubic terms. Again

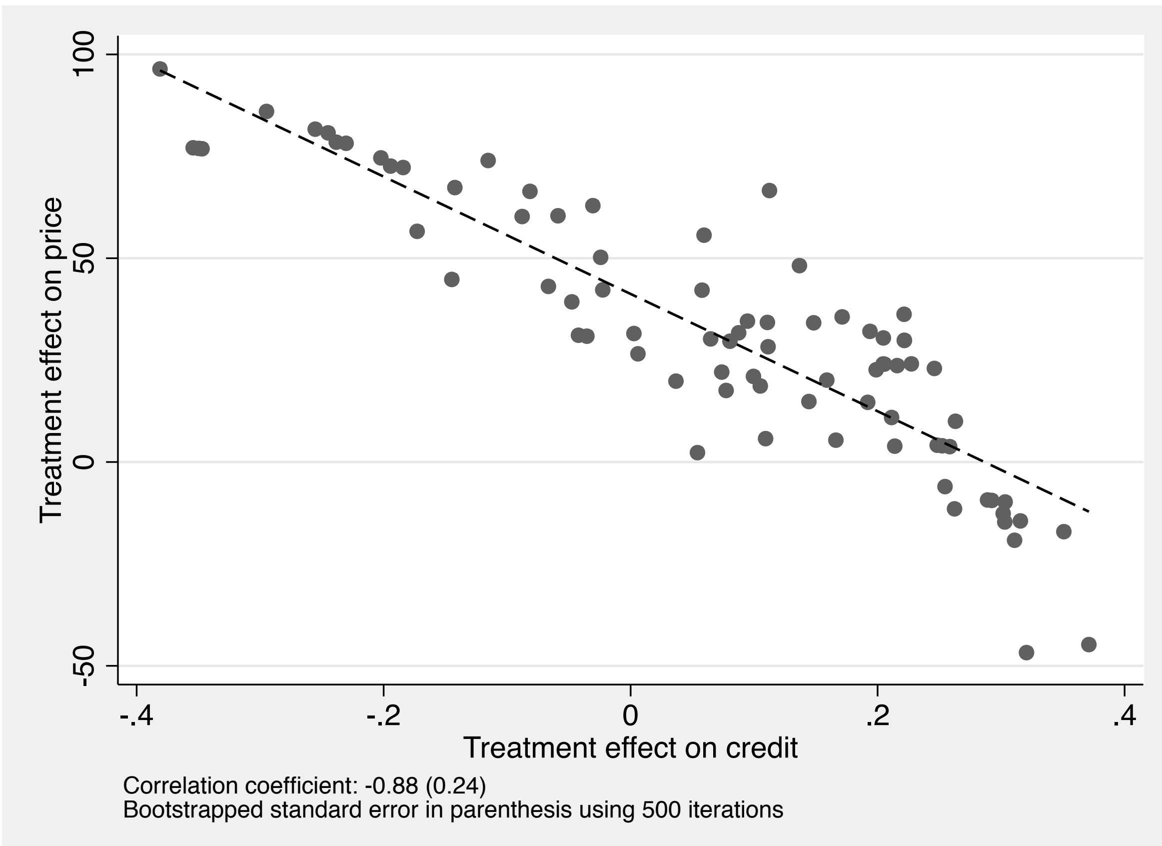

Figure 3.5: *Substitutability, correlation of predicted treatment effects in baseline linear model (Model B).*

Table 3.7: *Substitutability, regression coefficients*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Model	A		B		C		D		E		F		G	
Outcome	Price	Credit	Price	Credit	Price	Credit	Price	Credit	Price	Credit	Price	Credit	Price	Credit
Bonus $\in \{0,1\}$	54.61	-0.32**	89.63**	-0.37**	-69.98	-0.15	70.97	-0.31*	88.34	-0.11	75.66**	-0.38***	24.24	-0.76***
	(39.30)	(0.14)	(37.89)	(0.16)	(169.63)	(0.35)	(43.81)	(0.16)	(137.03)	(0.34)	(38.33)	(0.14)	(85.40)	(0.25)
Interaction Terms														
Baseline credit share	-63.89	0.60***	-72.61	0.59***	522.41	0.46	-78.95	0.60***	80.01	0.94	-54.07	0.58***	576.12	-0.49
	(50.95)	(0.21)	(49.01)	(0.22)	(476.57)	(1.49)	(52.21)	(0.21)	(528.11)	(1.28)	(48.73)	(0.22)	(460.93)	(1.32)
Baseline credit share^2					-1,327.95	3.24			-558.78	2.23			-1,446.56	4.48
					(1,208.54)	(3.57)			(1,265.06)	(3.19)			(1,084.05)	(3.29)
Baseline credit share^3					799.21	-3.11			402.44	-2.49			896.40	-3.53
					(782.44)	(2.34)			(787.57)	(2.10)			(683.43)	(2.16)
# of study traders			-13.55	0.04*	-83.09	0.27					-16.21*	0.04*	-105.26	0.34
			(9.44)	(0.02)	(65.81)	(0.20)					(9.21)	(0.02)	(75.02)	(0.23)
# of study traders ^{2}					8.49	-0.06							13.15	-0.07
					(12.91)	(0.04)							(14.58)	(0.05)
# of study traders ^{3}					-0.41	0.00							-0.68	0.00*
					(0.75)	(0.00)							(0.84)	(0.00)
# of treatment traders			19.88	-0.04	273.66**	-0.23					23.59	-0.04	257.86**	0.02
			(14.49)	(0.03)	(118.45)	(0.25)					(14.91)	(0.03)	(111.83)	(0.24)
# of treatment traders^2					-71.61*	0.07							-67.61*	-0.01
					(36.89)	(0.07)							(34.70)	(0.07)
# of treatment traders^3					5.90*	-0.01							5.65*	-0.00
					(3.27)	(0.01)							(3.08)	(0.01)
Miles to nearest town			0.06	-0.00	16.01	-0.13	-0.08	-0.00	-18.61	-0.14				
			(2.88)	(0.01)	(40.35)	(0.09)	(2.47)	(0.01)	(38.17)	(0.09)				
Miles to nearest town^2					-0.92	0.01			2.57	0.01				
					(3.45)	(0.01)			(3.18)	(0.01)				
Miles to nearest town^3					0.02	-0.00			-0.07	-0.00				
					(0.08)	(0.00)			(0.07)	(0.00)				
R^2	0.89	0.33	0.89	0.34	0.90	0.38	0.89	0.33	0.89	0.37	0.89	0.34	0.90	0.36
Number of observations	1,090	1,541	1,090	1,541	1,090	1,541	1,090	1,541	1,090	1,541	1,090	1,541	1,090	1,541

Table 3.8: *Substitutability, predicted treatment effect correlations*

Model	A	B	C	D	E	F	G
$\hat{P}$	-1.0	-0.88	-0.44	-0.99	-0.68	-0.78	-0.25
	–	(0.24)	(0.23)	(0.33)	(0.23)	(0.29)	(0.25)
Test Statistic	-2.27	-3.56	-1.93	-3.03	-2.95	-2.68	-0.93

Notes: $\hat{P}$ gives the correlation coefficient between the predicted treatment effect on price and on credit for each model estimated using weights equal to the inverse number of observations per village. Bootstrapped standard errors estimating using 500 iterations, resampling price and credit samples with replacement, are in parenthesis. The test statistic is $\hat{P}$ divided by its standard error. Model A includes only baseline credit share; Model B includes baseline credit share, number of study traders, number of treatment traders, and minimum road distance to nearest town. Model D includes only baseline credit share and minimum road distance, and Model F includes only baseline credit share, number of study traders and number of study traders. Model C includes a cubic for all terms in B, E includes a cubic for all terms in D, and G includes a cubic for all terms in F.

the coefficients remain negative, at -0.99 (s.e. = 0.33) and -0.68 (0.23) respectively. Model F includes only baseline credit share, number of study traders and number of study traders, and model G includes these in addition to quadratic and cubic terms as before. Here the relationship is negative, at -0.78 (s.e. = 0.29) and -0.25 (s.e. = 0.25) respectively in each model, though it is not significant at standard levels in model G. The best fit lines for these models are summarized in Figure **??**. In sum, the results show a robust and strong negative relationship between credit and price pass-through across markets.

3.5.2 Calibration and Welfare Analysis

Our results emphasize that a focus only on prices may obscure the total value passed through to farmers. A calibration of our model provides an illustration of the extent to which not accounting for interlinked credit transactions may underestimate the effect of a change in the wholesale price of output on the average farmer's welfare, and under which parameter values the bias will be most severe. These calculations rely heavily on the model, which makes several simplifying functional form assumptions, and the precise magnitudes

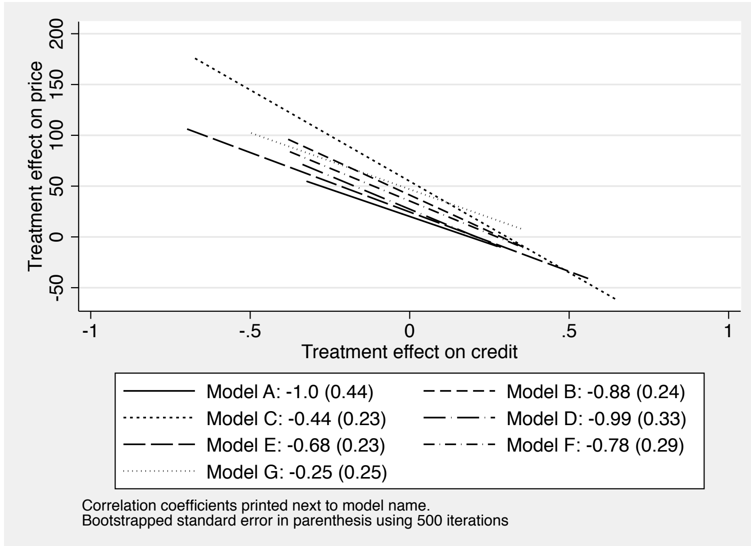

Figure 3.6: *Substitutability, best linear fit between price and credit treatment effects estimated from each model A-G.*

of our results should be interpreted with caution. Nevertheless, we believe that this exercise is useful, and helps illustrate the point that the inferences one draws about welfare from low pass-through in a market with interlinked transactions can be substantially different from those one would draw using a standard pass-through model.

We analyze the change in welfare with respect to a marginal increase in the wholesale price in a benchmark model without interlinkages and in our model. In the benchmark model, the change is given by (3.17). Since farmers receive $1/\mu$ of the wholesale price, the markdown is equal their welfare change per unit sold. We multiply by the average farm size to obtain an estimate of the change in average welfare in the economy.

In our model of interlinked transactions, the effect on welfare of a price is a weighted average of three terms, the effect on those who still receive no credit, those who still receive credit after the change and those who switch into the interlinked transaction. This result is summarized by (3.15), restated here:

$$\frac{dW^{ILT}}{dw} = G(q^*)E[q|q < q^*]\frac{1}{\mu} + (1 - G(q^*))E[q|q \geq q^*]\frac{(1-\gamma)}{\mu} - (c_F - \gamma\frac{w}{\mu})q^*g(q^*)\frac{dq^*}{dw}$$

In our calibration we will estimate all terms except μ and γ from the data, and then compare

(3.17) and (3.15) for different values of the two parameters. It remains then to obtain values for $G(q*)$, $E[q|q < q*]$, $E[q|q > q*]$, c_F, w, $q*$, and $g(q*)\frac{dq^*}{dw}$.

We begin with c_F, the utility gain to the farmer from the credit relation per unit sold. A natural starting point is the cash value of the average loan given to farmers, divided by the average quantity sold by a farmer. Our data on loan size come from the endline survey, asked in late December, 2011. Traders were asked if they had given a loan in the last thirty days, and if so, how much. The average loan size was Le. 177,297 (s.d. = 143,594), approximately 45 US dollars. To arrive at a per unit quantity, we must assume a period over which the loan was amortized. To be conservative, we will choose two months. Choosing a shorter period would make our estimates of the marginal utility from a loan higher. To calculate per farmer quantity sold during that period, we calculate total quantity of grade A purchased by all traders in each village in the month before the question was asked. We then divide this by the number of suppliers in the village, and take the average, arriving at 96 pounds. Assuming that the trader, when giving the loan expected to receive this flow for four months, the implicit per unit credit outlay is 923 Leones. We assume that λ, the marginal utility for the farmer of a one-dollar credit provision, valued at the time of the harvest sale, is 1.3,[15] so we have $c_F = 1,201$. Current prices at this time are Le. 2,666 per pound so we take that to equal W_0. Finally, in the calibration exercise, we conservatively assume that r, the return of credit on production quantity/quality, is zero. Thus, our estimates will be a lower bound on the impact of a change in trades resale price on farmer welfare.

We then estimate the share of farmers in the credit relation, and those who are not. Since the data on credit provision were collected referencing the last month, we will use this same time frame to estimate the share of farmers currently in an interlinked transaction. This is just then the control group's share of people who have received credit in the last month, 12%. We use this value for $G(q*)$. To estimate the expected farmer loan size for traders in

[15]This is a conservative estimate. Several studies on similar populations find *weekly* discount factors around 0.75 (e.g. Schaner, 2013)

and out of the credit relation, we return to baseline data on quantities sold by each village to all of our study traders, so that our estimates of quantities are not influenced by the treatment. We use the baseline share of farmers who have received credit in that village since the beginning of the season as our proxy for the level of the credit in that village. The median value of this variable across villages is 63%. We code villages above this median as being in the credit relation, and those below this median as being outside of it, and then calculate average baseline quantity sold during the 4 pre-treatment weeks of the experiment per farmer in the village for the two groups. From this we arrive at $E[q|q < q*] = 36$ and $E[q|q > q*] = 95$. For the value of $q*$, we simply take expected per farmer quantity for the median village with respect to the credit share, getting $q* = 67$.

Finally, we require a term for $g(q*)\frac{\partial q*}{\partial W}$, the change in the threshold value of $q*$ times its density. This we can estimate from our treatment effects. If each unit of q represents a farmer, a decrease in $q*$ is just the negative of the change in the share of people getting credit. If we assume this term is constant, it is just equal to the negative of the slope of our treatment effect on credit with respect to the increase the bonus, or $-0.12/150 = -0.0008$. Finally, to estimate the change in the standard model, we just require expected quantity per farmer, which in the baseline equals 66 pounds.

Figure **??** summarizes the results of our calibration of the welfare impact of the experiment with and without accounting for the credit margin. Specifically, the figure shows the ratio between the average welfare changes estimated using the interlinked transaction model (3.15) and using the benchmark framework (3.17), for a range of values of γ and μ, and the calibrated values discussed above. The ratio is larger than one in most of the examined range. Intuitively, the inclusion of the credit channel raises the estimated welfare impact of a change in trader prices relatively more when i) μ is high, since a higher markdown reduces the price pass-through, which is the sole driver of welfare change in the benchmark model; ii) γ is low, since better contract enforcement institutions increase the amount of rent traders can extract from farmers when they provide them with credit. In conclusion, the exercise shows that neglecting the role of interlinkages can substantially affect the estimation of the

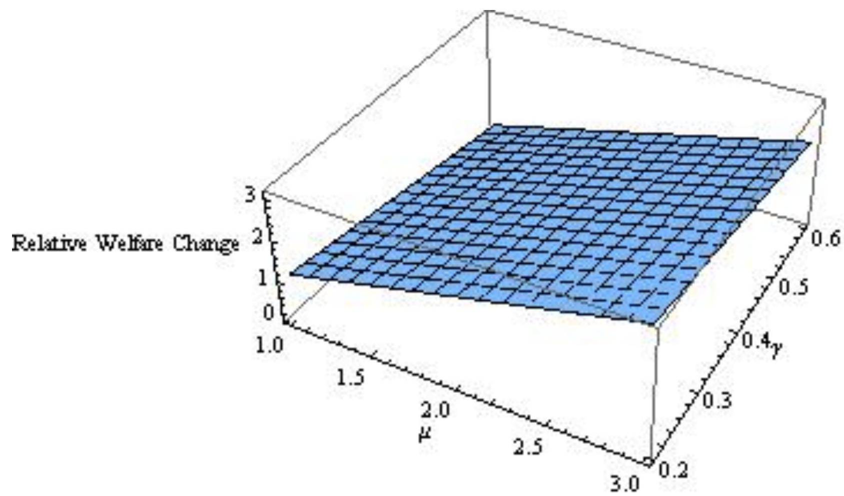

Figure 3.7: *A welfare change in response to a marginal change in the intermediaries price in the ILT model relative to in the benchmark model.*

welfare consequences of a change in wholesale prices.

3.6 Concluding Remarks

The theory and evidence presented in the paper show that, in the presence of interlinked transactions, low price pass-through may obscure other channels through which value is passed from end-buyers to producers. Price pass-through and credit pass-through are shown to be substitutes across markets in an agricultural setting in sub-Saharan Africa that has similarities to markets in other developing economies. This finding is shown to have substantial implications for the analysis of the impact of changes in downstream prices on farmer welfare. Indeed, in our context, low price pass-through can result from high responsiveness of interlinked rural capital markets, as well as from intermediaries' market power.

Our study complements existing work on the welfare implications. Interlinked transactions along the value chain are common in developing economies and we expect our framework to be particularly valuable in these settings. More broadly, interlinkages play a role in a wide range of transactions. For instance, trade credit is a major source of finance for firms in industrialized countries (Petersen and Rajan, 1997; Fisman and Love, 2003) and cash-in-advance plays an important role in international trade contracts (Antràs and Foley, 2011). Our argument that, in the presence of interlinked transactions, one needs to assess the multiple margins through which value passes-through the value chain is thus a general one.

While our empirical work studies pass-through at the level of the intermediary buying from a producers, the issues described here naturally extend to other links in the supply chain, for instance from international buyers to exporters. We thus believe the question of how interlinkages affect the transmission of price signals through entire supply chains is an exciting one for future research.

References

Chapter 1

Alesina, Alberto (1988). "Credibility and Policy Convergence in a Two-Party System with Rational Voters." *American Economic Review*, 78, pp.796-806.

—, Reza Baqir and Caroline Hoxby (2004). "Political jurisdictions in heterogeneous communities." *The Journal of Political Economy*, 112(2):348-396.

—, Dorian Carloni, and Giampaolo Lecce (2011). "The Electoral Consequence of Large Fiscal Adjustments." Working paper.

—, William Easterly and Janina Matuszeski (2011). "Artificial States." *Journal of the European Economic Association*, 9(2): 246-277.

— and Enrico Spolaore (1997). "On the number and size of nations." *The Quarterly Journal of Economics*, 112(4):1027-1056.

Alm, James, Betty Jackson, and Michael McKee (1992). "Estimating the Determinants of Taxpayer Compliance with Experimental Data." *Economic Development and Cultural Change*, March, 39(4): 107-114.

Almunia, Miguel and David Lopez-Rodriguez (2012). "The Efficiency Costs of Tax Enforcement: Evidence from a Panel of Spanish Firms." University of California Berkeley, working paper.

Andreoni, James, Brian Erard and Jonathan Feinstein (1998). "Tax Compliance." *Journal of Economic Literature*, 36(2):818-860.

Aparicio, Gabriela (2012). "Monitoring and its Interaction with Punishment in Tax Enforcement: Evidence from a Regression Discontinuity Design." Georgetown University, working paper.

Artavanis, Nikolaos T., Adait Morse, and Margarita Tsoutsoura (2012). "Tax Evasion Across Industries: Soft Credit Evidence from Greece." Chicago University Booth School of Business, working paper.

Banks Jeffrey and Rangarajan Sundaram (1998). "Optimal Retention in Agency Problems." *Journal of Economic Theory*, 82(2). 293-323.

Barca, Marcello (2008). “I primi cittadini della montagna commentano i dati apparsi sul sito dell’Agenzia del territorio sugli immobili non denunciati.” *L’Informazione Provincia Montagna*, March 21st.

Barro, Robert (1973). “The Control of Politicians: An Economic Model.” *Public Choice*, 14, 1942.

Bernardini, Daniele (2011). “Il Comune va a caccia delle case fantasma.” *La Nazione (2011)*, March 18th.

Besley, Tim and Andrea Pratt (2006). “Handcuffing the Grabbing Hand?: Media Capture and Government Accountability.” *American Economic Review*, 96(3). 720-736.

—, and Stephen Coate (1997). “An Economic Model of Representative Democracy.”*The Quarterly Journal of Economics*, 112, pp.85-114.

—, Jose G. Montalvo, and Marta Reynal-Quenol (2012). “Do Educated Leaders Matter?” *Economic Journal*, forthcoming.

Bird, Richard M. and Eric M. Zolt (2008). “Technology and Taxation in Developing Countries: From Hand to Mouse.” *National Tax Journal*, Vol. LXI, No. 4, Part 2.

Brender, Anton (2003). “The Effect of Fiscal Performance on Local Government Election Results in Israel: 1989-1998.” *Journal of Public Economics*, 87 (9): 2187-2205.

— and Allan Drazen (2008). “How Do Budget Deficits and Economic Growth Affect Reelection Prospects? Evidence from a Large Panel of Countries.” *American Economic Review*, 98(5): 2203-20.

Bubb, Ryan (2008). “Blame It On The Rain? Voter Rationality and Exogenous Economic Shocks.” New York University, working paper.

Card, David (2001). ”Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems.” *Econometrica*, 69, pp. 1127-1160.

Carpenter, Jeffrey, Samuel Bowles, Herbert Gintis, and Sung-Ha Hwang (2009). “Strong Reciprocity and Team Production: Theory and Evidence.” *Journal of Economic Behavior and Organization*.

Carrillo, Paul E., Dina Pomeranz, and Monica Singhal (2012). “Tax Me If You Can: Information Cross-Checks and Evasion Substitution.” Working paper.

Casari, Marco and Luigi Luini (2009). “Group cooperation under alternative punishment institutions: an experiment.” *Journal of Economic Behavior and Organization*, 71(2). 273-282.

Caselli, Francesco, Tom Cunningham, Massimo Morelli, and Inés Moreno-de-Barreda (2013). “The Incumbency Effects of Signalling." *Economica*, forthcoming.

Cavallaro, Francesco (2011). “Sindaci Pronti alla Lotta all’Evasione." *Il Sole 24 Ore*. February 16th.

Chetty, Raj, John Friedman and Emmanuel Saez (2011). "Using Neighborhood Effects to Uncover the Impacts of Tax Policy: The Effect of the EITC on Earnings." Working paper.

Cole, Shawn, Andrew Healy, and Eric Werker (2012). "Do Voters Demand Responsive Governments? Evidence from Indian Disaster Relief."*Journal of Development Economics*, no. 97 (2012): 167-181.

Comune di Capaccio Paestum (2010). "Comunicato Stampa, n. 134/10", July 29th.

Corriere della Citta' (2012). "Crisi? Parla il sindaco De Fusco." March 28th.

Dell'Oste, Cristiano and Gianni Trovati (2011). "Il Tesoro delle Case Fantasma." *Il Sole 24 Ore*, January 31st.

Ferejohn, John (1986). "Incumbent performance and electoral control." *Public Choice*, 50: 5-26.

Ferraz, Claudio, and Frederico Finan (2008) "Exposing Corrupt Politicians: The Effects of Brazil's Publicly Released Audits on Electoral Outcomes." *Quarterly Journal of Economics*, 123 (2): 703-745.

Fisman, Raymons and Shang-Jim Wei (2004). "Tax Rates and Tax Evasion: Evidence from Missing Imports in China." *Journal of Political Economy*, 112 (471-500).

Gagliarducci, Stefano and Daniele Paserman (2012). "Gender Interactions within Hierarchies: Evidence form the Political Arena." *Review of Economic Studies*, 79(3). 1021-1052.

—. and Tommaso Nannicini (2013). "Do Better Paid Politicians Perform Better? Disentangling Incentives from Selection" , *Journal of the European Economic Association*, forthcoming

Gazzetta del Mezzogiorno (2012). " Caccia aperta ai furbetti proprietari di case-fantasma. Il Comune vuole censire quelle abitazioni che sfuggono alle tasse." January 11th.

Gordon, Roger H. and Wei Li. (2009). "Tax Structure in Developing Countries: Many Puzzles and a Possible Explanation." *Journal of Public Economics*, 93:7-8 (855-866).

Grembi, Veronica, Tommaso Nannicini and Ugo Troiano (2013). "Policy Responses to Fiscal Restraints. A Difference-in-Discontinuities Design", Working paper

Holmstrom, Berndt (1982). "Managerial incentive problems: A dynamic perspective." In *Essays in Economics and Management in Honor of Lars Wahlbeck.* Helsinki: Swedish School of Economics; reprinted in *Review of Economic Studies 66 (1999)*: 169-182.

Kleven Henrik J., Martin b. Knudsen, Claus T. Kreiner, Soren Pedersen , and Emmanuel Saez (2011). "Unwilling or Unable to Cheat? Evidence from a Randomized Tax Audit Experiment in Denmark." *Econometrica*, 79(3). 2011, 651-692.

Kumler, Todd, Eric Verhoogen and Judith A. Frías (2011). "Enlisting Workers in Monitoring Firms:Payroll Tax Compliance in Mexico." Working paper.

Lee, David. (2001). "The Electoral Advantage to Incumbency and Voters' Valutation of Polticians' Experience: a Regression Discontinuity Analysis of Elections to the U.S. House." NBER Working Paper #8441.

Litschig, Stephan and Kevin Morrison (2012). "Government Spending and Re-election: Quasi-Experimental Evidence from Brazilian Municipalities." Barcelona GSE, working paper.

Marino, M. Rosaria and Roberta Zizza (2008). "The personal income tax evasion in Italy." *Bank of Italy*, working paper.

Mian Atif and Amir Sufi (2012). "The Effects of Fiscal Stimulus: Evidence from the 2009 Cash for Clunkers Program." *Quarterly Journal of Economics*, forthcoming.

Miguel, Edward A., Marco Manacorda and Andrea Vigorito (2011). "Government Transfers and Political Support." *American Economic Journal: Applied Economics*, 3(3). 128.

Nannicini, Tommaso, Andrea Stella, Guido Tabellini, and Ugo Troiano (2012). "Social Capital and Political Accountability." *American Economic Journal: Economic Policy*, forthcoming.

Nordhaus, William (1975). "The political business cycle." *Review of Economic Studies*, 42:169-190.

Olson, Mancur (1965). "The Logic of Collective Action: Public Goods and the Theory of Groups." Harvard University Press.

Ouss Aurelie and Alex Peysakhovich "When Punishment Doesn't Pay: The 'Cold Glow' Heuristic and Decisions to Punish." Harvard University, working paper.

Pande, Rohini (2011). "Can Informed Voters Enforce Better Governance? Experiments in Low-Income Democracies."*Annual Review of Economics*, Vol. 3: 215-237.

Pomeranz Dina (2012). "No Taxation without Information: Deterrence and Self-Enforcement in the Value Added Tax." Harvard Business School, working paper.

Putnam, Robert (2001). "Social Capital: Measurement and Consequences." *ISUMA - Canadian Journal of Policy Research*, 2, 1, 41-51.

Quattrone, George A. and Amos Tversky. (1988). "Contrasting Rational and Psychological Analyses of Political Choice." *The American Political Science Review*, Vol. 82, No. 3., pp. 719-736.

Rogoff, Kenneth, and Anne Sibert (1988). "Elections and macroeconomic policy cycles." *Review of Economic Studies*, 55:1-16.

Rothstein, Bo (2000). “Trust, Social Dilemmas and Collective Memories.” *Journal of Theoretical Politics*, vol. 12, 477-501.

Saez, Emmanuel (2010). “Do Taxpayers Bunch at Kink Points?” *American Economic Journal: Economic Policy*, 2, pp. 180-212.

Saiz, Albert (2010). “The Geographic Determinants of Housing Supply.” *The Quarterly Journal of Economics*, vol. 125(3). pages 1253-1296.

Slemrod, Joel (2003). “Trust in Public Finance.” In S. Cnossen and H. Werner (eds.) *Public Finance and Public Policy in the New Century*, 49–88. Cambridge, MA: MIT Press.

— (2007). “Cheating Ourselves: The Economics of Tax Evasion.” *Journal of Economic Perspectives*, 21(1): 25-48.

—, Marsha Blumenthal, and Charles Christian (2001). “Taxpayer Response to an Increased Probability of Audit: Evidence from a Controlled Experiment in Minnesota." *Journal of Public Economics* 79(3): 455- 483.

— and Shlomo Yitzhaki (2002). “Tax avoidance, evasion, and administration.” Handbook of Public Economics, in: A. J. Auerbach and M. Feldstein (ed.). *Handbook of Public Economics*, edition 1, volume 3, chapter 22, pages 1423-1470.

Spicer, Michael W., and Lee A. Becker (1980). “Fiscal Inequity and Tax Evasion: An Experimental Approach.” *National Tax Journal*, June, 33(2):171-175.

Torgler, Benno (2003). “Tax Morale, Rule-Governed Behaviour, and Trust.” *Constitutional Political Economy*, June, 14(2): 119-140.

— (2007). *Tax Morale and Tax Compliance: A Theoretical and Empirical Analysis*, Cheltenham U.K.: Edward Elgar.

Tullock, Gordon (1959). “Some problems of majority voting.” *Journal of Political Economy*, 67: 571-579.

Wolfers, Justin (2009). “Are Voters Rational? Evidence from Gubernatorial Elections.” University of Pennsylvania, working paper.

Chapter 2

Aker, Jenny (2010). “Information From Markets Near and Far: The Impact of Mobile Phones on Grain Markets in Niger", *American Economic Journal: Applied Economics*, 2(1): 46-59.

Ali, Rubaba (2011). “Impact of Rural Road Improvement on High Yield Variety Technology Adoption: Evidence From Bangladesh", Working Paper, University of Maryland.

Allen, Treb (2012). “Information Frictions in Trade", Working Paper, Princeton University and Northwestern University.

Antràs, Pol and Arnaud Costinot. "Intermediated Trade", *Quarterly Journal of Economics*, 126(3): 1319-1374.

Atkin, David and Dave Donaldson (2012). "Who's Getting Globalized? Intranational Trade Costs and World Price Pass-Through", Working Paper, Stanford University.

Banerjee, Abhijit, Esther Duflo, and Nancy Qian (2012). "On the Road: Access to Transportation Infrastructure and Economic Growth in China", NBER Working Paper 17897.

Bardhan, Pranab, Dilip Mookherjee, and Masatoshi Tsumagari (2012). "Middlemen Margins And Globalization," Working Paper, Boston University.

Barreca, Alan, Melanie Guldi, Jason Lindo, and Glen Waddell (2011). "Saving Babies? Revisiting the Effect of Very Low Birth Weight Classification", *Quarterly Journal of Economics*, 126(4): 2117-2123.

Bell, Clive, and Susanne van Dillen (2012). "How Does India's Rural Roads Program Affect the Grassroots? Findings from a Survey in Upland Orissa", Working Paper, University of Heidelberg.

Borraz, Fernando, Alberto Cavallo, Roberto Rigobon, and Leandro Zipitria (2012). "Distance and Political Boundaries: Estimating Border Effects under Inequality Constraints", NBER Working Paper No. 18122.

Broda, Christian and David E. Weinstein (2008). "Understanding International Price Differences Using Barcode Data," NBER Working Paper No. 14017.

Burstein, Ariel and Nir Jaimovich (2009). "Understanding Movements in Aggregate and Product-Level Real Exchange Rates", Working Paper, UCLA.

Cameron, A. Colin, Jonah Gelbach, and Douglas Miller (2008). "Bootstrap-Based Improvements for Inference with Clustered Errors", *Review of Economics and Statistics*, 90(3).

Carruthers, Robin, Ranga Rajan Krishnamani, and Siobhan Murray (2009). "Improving Connectivity: Investing in Transport Infrastructure in Sub-Saharan Africa", Africa Infrastructure Country Diagnostic Background Paper 7 (Phase II), World Bank.

Casaburi, Lorenzo, and Tristan Reed (2012). "Incomplete Pass-Through and Interlinked Transactions: Experimental Evidence from Sierra Leone", Working Paper, Harvard.

Chau, Nancy H., Hideaki Goto, and Ravi Kanbur (2009). "Middlemen, Non-Profits, and Poverty", Working Paper, Cornell University.

Datta, Saugato (2012). "The Impact of Improved Highways on Indian Firms", *Journal of Development Economics*, 99(1): 46-57.

Deb, Rahul, and Tavneet Suri (2012). "Endogenous Emergence of Credit Markets: Contracting in Response to a New Technology in Ghana", Forthcoming, *Journal of Development Economics*.

Donaldson, Dave (2012), “Railroads of the Raj: Estimating the Impact of Transportation Infrastructure", Working Paper, MIT.

Dorosh, Paul, Hyoung-Gun Wang, Liang You, and Emily Schmidt (2010). “Crop Production and Road Connectivity in Sub-Saharan Africa", Policy Research Working Paper 5385, World Bank.

EDA (2004). "Identification and Prioritization of Rural Roads in Kambia, Port Loko, Kenema and Pujehun Districts. Final Report, Volume 1".

Faber, Benjamin (2012). “Trade Integration, Market Size, and Industrialization: Evidence From China’s National Truck Highway System", Working Paper, LSE.

Fafchamps, Marcel, and Ruth Vargas Hill (2008). “Price Transmission and Trader Entry in Domestic Commodity Markets”, Working Paper, University of Oxford.

—, Eleni Gabre-Madhin, and Bart Minten (2004). “Increasing Returns and Market Efficiency in Agricultural Trade", *Journal of Development Economics*, 78(2).

Gibson, John and Scott Rozelle (2003). “Poverty and Access to Roads in Papua New Guinea", *Economic Development and Cultural Change*, 52(1): 159-185.

Gopinath, Gita, Pierre-Olivier Gourinchas, Chang Tai-Hsieh, and Nicholas Li (2011). “International Prices, Costs, and Markup Differences", *American Economic Review*, 101(6).

Goyal, Aparajita (2010). “Information, Direct Access to Farmers, and Rural Market Performance in Central India," *American Economic Journal: Applied Economics*, 2(3): 22–45.

Imbens, Guido and Karthik Kalyanaraman (2009). “Optimal Bandwidth Choice for the Regression Discontinuity Estimator", NBER Working Paper No. 14726.

— and Thomas Lemieux (2007. “Regression Discontinuity Designs: A Guide to Practice", *Journal of Econometrics*, 142(2): 615-635.

Jacoby, Hanan and Bart Minten (2008). “On Measuring the Benefits of Lower Transport Costs", *Journal of Development Economics*, 89(1): 28-38.

Jensen, Robert (2007). “The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector", *Quarterly Journal of Economics*, 122(3): 879-924.

Khandker, Shahidur, Zaid Bakht, and Gayatri Koolwal (2009). “The Poverty Impact of Rural Roads: Evidence From Bangladesh", *Economic Development and Cultural Change*, 57(4): 685-722.

Khandker, Shahidur and Gayatri Koolwal (2011). “Estimating the Long-Term Impacts of Rural Roads", Policy Research Working Paper 5867, World Bank.

Lee, David and Thomas Lemieux (2010). “Regression Discontinuity Designs in Economics", *Journal of Economic Literature*, 48(2): 281-355.

McCrary, Justin (2008). "Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test", *Journal of Econometrics*, 142(2): 698-714.

Michaels, Guy (2008). "The Effect of Trade on the Demand for Skill: Evidence from the Interstate Highway System", *Review of Economics and Statistics*, 90(4): 683-701.

Mortensen, Dale (2003), *Wage Dispersion: Why Are Similar Workers Paid Differently?*, MIT Press, Cambridge.

Mu, Ren and Dominique van de Walle (2007). "Rural Roads and Local Market Developments in Vietnam", Policy Research Working Paper 4340, World Bank.

Pushak, Nataliya and Vivien Foster (2011). "Sierra Leone's Infrastructure: A Continental Perspective", *Africa Infrastructure Country Diagnostic Country Report*, Washington, D.C.: The World Bank.

Rubin, Donald (1974). "Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies", *Journal of Educational Psychology*, 66(5), 688-701.

Statistics Sierra Leone (2004). "2004 Population and Housing Census".

— (2011). "Agricultural Households Tracking Survey Report".

Suri, Tavneet (2011). "Selection and Comparative Advantage in Technology Adoption", *Econometrica*, 79(1): 159-209.

Van de Walle, Dominique (2009). "Impact Evaluation of Rural Road Projects", *Journal of Development Effectiveness*, 1(1): 15-36.

World Bank (2009). *World Development Report 2009: Reshaping Economic Geography*, Washington

Chapter 3

Antràs, Pol and Arnaud Costinot (2011). "Intermediated Trade," *Quarterly Journal of Economics*, 126(3): 1319-1374.
—, and C. Fritz Foley (2011). "Poultry in motion: a study of international trade finance practices," NBER working paper No. 17091.

Acemoglu, Daron, Tristan Reed, and James A. Robinson (2013). "Chiefs: Elite Control of Civil Society and Economic Development in Sierra Leone," NBER working paper No. 18691.

Atkin, David and Dave Donaldson (2013). "Who's Getting Globalized? The Size and Nature of Intranational Trade Costs," MIT mimeo.

Banerjee, Abhijit (2002). "Contracting Constraints, Credit Markets and Economic Development." *Advances in Economics and Econometrics* 3.

— and Esther Duflo (2000). "Reputation Effects and the Limits of Contracting: A Study of the Indian Software Industry", *Quarterly Journal of Economics* 115: 989-1017.

Bardhan, Pranab K. (1980). “Interlocking factor markets and agrarian development: A review of issues.” *Oxford Economic Papers* 32(1): 82-98.

Bauer, Peter T. (1954). *West African Trade: A Study of Competition, Oligopoly and Monopoly in a Changing Economy*, Cambridge: University Press.

Bell, Clive (1988). “Credit Markets and Interlinked Transactions,” in Hollis Chenery and T.N. Srinivasan eds. *Handbook of Development Economics*, Volume 1, Chapter 16, Amsterdam: Elsevier.

Braverman, Avishay and Joseph E. Stiglitz (1982). “Sharecropping and the Interlinking of Agrarian Markets,” *The American Economic Review*, 72(4): 695-715.

Bruhn, Miriam and David McKenzie (2009). "In Pursuit of Balance: Randomization in Practice in Development Field Experiments,” *The American Economic Journal: Applied Economics*, 1(4): 200-232.

Cameron, A. Colin, Jonah B. Gelbach and Douglas L. Miller (2006). “Robust Inference with Multi-way Clustering,” NBER technical working paper No. 327.

CAOBISCO: Chocolate, biscuits & confectionery of Europe (2002). “Cocoa quality,” Memorandum: Brussels, R’ef: 166-58sj-Rev.5, March 13.

Casaburi, Lorenzo, Rachel Glennerster and Tavneet Suri (2012). “Rural Roads and Intermediated Trade: Regression Discontinuity Evidence from Sierra Leone,” MIT Sloan mimeo.

Deb, Rahul and Tavneet Suri (2012). “Endogenous Emergence of Credit Markets: Contracting in Response to a New Technology in Ghana,” *Journal of Development Economics*, forthcoming.

Fabinger, Michal and E. Glen Weyl (2013). “Pass-through as an Economic Tool,” *Journal of Political Economy*, forthcoming.

Fafchamps, Marcel (2004). *Market Institutions and Sub-Saharan Africa: Theory and Evidence*, Cambridge: MIT Press.

Fafchamps, Marcel, Eleni Gabre-Madhin, and Bart Minten (2005). “Increasing returns and market efficiency in agricultural trade.” *Journal of Development Economics* 78(2): 406-442.

Fafchamps, Marcel and Ruth Hill (2008). “Price Transmission and Trader Entry in Domestic Commodity Markets,” *Economic Development and Cultural Change*, 56(4): 729-766.

Feenstra, Robert C. (1989). “Symmetric pass-through of tariffs and exchange rates under imperfect competition: an empirical test,” *Journal of International Economics*, 27(1-2), 25-45.

Fisman, Raymond and Inessa Love (2003). “Trade Credit, Financial Intermediary Development, and Industry Growth," *Journal of Finance,* vol. 58(1), pages 353-374, 02.

Gilbert, Christopher L. (2009). “Cocoa Market Liberalization in Retrospect,” *Review of Business and Economics*, 54(3): 294-312.

Goldberg, Pinelopi Koujianou, and Rebecca Hellerstein (2008). "A structural approach to explaining incomplete exchange-rate pass-through and pricing-to-market." *The American Economic Review* 98(2): 423-429.

Grosh, Barbara (1994). “Contract Farming in Africa: An Application of the New Institutional Economics,” *Journal of African Economies*, 3(2): 231-61.

Hill, Polly (1961). “The Migrant Cocoa Farmers of Southern Ghana,” *Africa: Journal of the International African Institute*, 31(3): 209-230.

Kremer, Michael and Edward Miguel (2004). “Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities” *Econometrica,* 72(1): 159-217.

Macchiavello, Rocco and Ameet Morjaria (2012). “The Value of Relationships: Evidence from a Supply Shock to Kenyan Flower Exports,” Warwick University mimeo.

Maitra, Pushkar, Sandip Mitra, Dilip Mookherjee, Alberto Motta, and Sujata Visaria (2012). “Intermediated Loans: A New Approach to Microfinance,” Boston University mimeo.

Mitra, Sandip, Dilip Mookherjee, Maximo Torero, and Sujata Visaria (2013). “Middleman margins and Information: An Experiment with Potato Farmers in West Bengal,” Boston University mimeo.

McMillan, John, and Christopher Woodruff (1999). “Interfirm relationships and informal credit in Vietnam.” *The Quarterly Journal of Economics* 114(4): 1285-1320.

Osborne, Theresa (2005). “Imperfect competition in agricultural markets: evidence from Ethiopia,” *Journal of Development Economics* 76(2): 405-428.

Oxfam (2002). “Mugged: Poverty in Your Coffee Cup”, Oxfam (2002).

Petersen, Mitchell A. and Raghuram G. Rajan (1997). “Trade Credit: Theories and Evidence.” *Review of Financial Studies* 10 (3): 661-691

Poulton, Colin, Andrew Dorward, and Jonathan Kydd (1998). “The revival of smallholder cash crops in Africa: public and private roles in the provision of finance.” *Journal of International Development* 10(1): 85-103.

Rubinstein, Ariel and Asher Wolinksy (1987). “Middlemen,” *The Quarterly Journal of Economics*, 102(3): 581-594.

Sonii, David (2005). “Learning about sustainable cocoa production: a guide for participatory farmer training. 1. Integrated crop and pest management,” Sustainable Tree Crops Program, International Institute of Tropical Agriculture, Yaoundé, Cameroon.

Appendix A

Appendix to Chapter 1

A.1 Appendix Figures

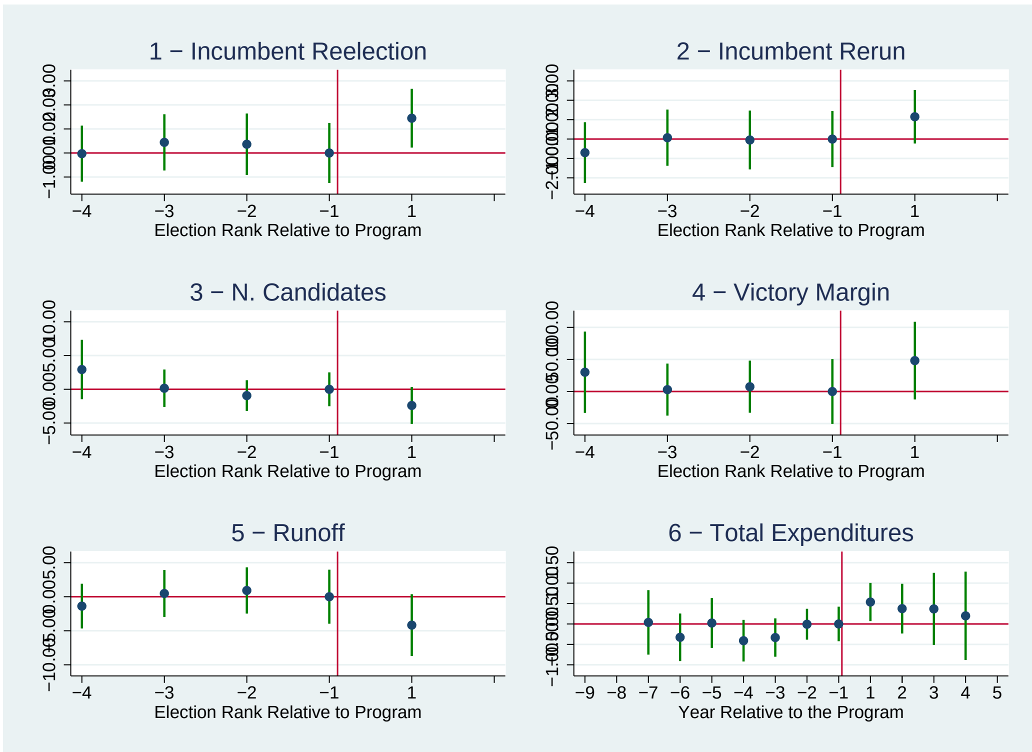

Notes: **Graphs 1 to 5** report the coefficients on the ghost building intensity *for each election* before and after the beginning of the Ghost Buildings program. The modal number of years between elections is five years between 1993 and 2001, and four afterwards. On the x-axis, elections are ranked based on their occurrence relative to the program. **Graph 6** reports the coefficients on the ghost building intensity *for each calendar year* before and after the beginning of the Ghost Buildings program. The dependent variable is in natural logarithm. On the x-axis, years are ranked based on their occurrence relative to the program. We drop the year of program inception due to its ambiguous treatment status. In all the graphs, the regression includes town and year fixed effects. We report the point estimate and the 95% confidence interval. The last election/year before the program ("-1") is the omitted category. The coefficient on ghost building intensity for this group is normalized to zero. Confidence interval width for this election is obtained as the mean of the confidence interval width in election/year -2 and election/year +1.

Figure A.1: *Ghost Building Intensity Coefficient by Election Pre/Post Program*

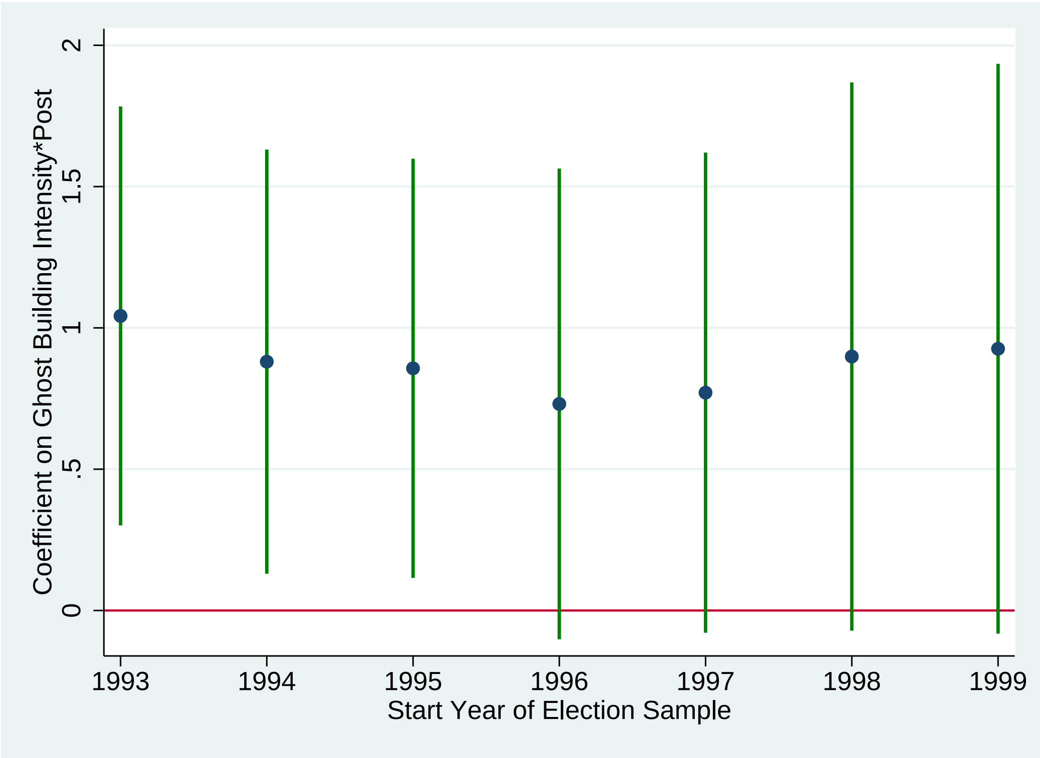

Notes: The figure presents robustness of the results on incumbent reelection to changes in sample years. The y-axis shows the coefficient (and 95% confidence intervals) on the interaction between ghost building intensity and the post-program indicator as estimated in our baseline specification (Table 1.4). The x-axis is the start year of the alternative election samples we use (the final year is 2011 for all the samples). The first sample, 1993-2011, is the baseline sample. In 1993, an electoral law reform introduced mayor individual ballot election.

Figure A.2: *Robustness to changes in election sample time span*

A.2 Appendix Tables

Table A.1: *Variables' description and sources*

Variable	Definition and measure	Sample	Source
Ghost Building Intensitity	Ghost Building Intensity Ration between the number of land registry parcels with ghost buildings and the total number of parcels.	Program inception	ATD
Registration Rate	Registration Rate Percentage of ghost building parcels that get registered by the April 2011	2011	ATD
Total expenditures	Total local expenditures Per-resident	2000-2011	IMI Financial reports, *Quadro 3*
Town Area Size	Area Size of the town, in square km	2001	SAIM
Altitude	Altitude Altitude of the city, in meters	2001	SAIM
Population	Population Population, in thousand of inhabitants	2001	Census
Disposable Income per capita	Disposable income per capita at the municipal level, in thousand of euros	2005	SAIM
Urbanization Index	Index is equal to one if density is less than 100 people per sq. km; it is equal to two if density is between 100 and 500 people per sq. km; it is equal to three if density is above 500 people per sq.km.	2001	SAIM
Justify Tax Cheating	Answers to the question "Do you Justify Tax Cheating?", originally coded on a scale 1 (never justifiable) to 10 (always justifiable). Normalized variable.	1981-2008	EVS
Speed of Public Good Provision	Speed of current expenditures Ratio between paid over committed current expenditures	2005-2006	IMI

Notes: ATD stands for Agenzia del Territorio Database. IMI stands for Italian Ministry of the Interior. SAIM stands for Statistical Atlas of Italian Municipalities. EVS stands for European Values Survey.

Table A.2: *Political Variables' description and sources*

Variable	Definition and measure	Sample	Source
Mayor Age	Age of the mayor Age of the mayor, in number of years	Program inception	IMI
Mayor Education	Education of the mayor Categories: Primary Education, High school education, University degree, Postgraduate professional schooling, GED equivalent schooling, Vocational schooling	Program inception	IMI
Mayor Born Same City	Place of birth of the mayor Dummy variable equal to 1 if mayor is born in the same city	Program inception	IMI
Mayor Term Number	Tenure of the mayor Number of the mayoral's term, in number of years	Program inception	IMI
Mayor Woman	Gender of the mayor Dummy variable equal to 1 if the mayor is a woman	Program inception	IMI
Incumbent Reelection	Incumbent mayor is reelected Equal to 1 if mayor is reelected	1993-2011	IMI
Incumbent Rerun	Incumbent mayor decides to run for office again Equal to 1 if mayor re-runs	1993-2011	IMI
N. Candidates	Number of candidates Number of candidates	1993-2011	IMI
Victory Margin	Margin of victory Margin of victory of the winning candidate	1993-2011	IMI
Runoff	Election has a runoff Runoff	1993-2011	IMI

Notes: IMI stands for Italian Ministry of the Interior.

Appendix B

Appendix to Chapter 2

B.1 Formula for Road Score

The score of road i in district k is given by:

$$score_i = 0.4 * \frac{ep_i}{ep\_max_k} + 0.2 * \frac{pd_i}{pd\_max_k} + 0.2 * \frac{ra_i}{ra\_max_k} + 0.1 * \frac{sv_i}{sv\_max_k} + 0.1 * \frac{l_i}{l\_max_k} \tag{B.1}$$

where:

ep_i is economic productivity per kilometer of road i,

ep_max_k is the maximum economic productivity in district k,

pd_i is population density around road i,

pd_max_k is the maximum population density in dsitrict k,

ra_i is the baseline road condition assessment of road i,

ra_max_k is the maximum baseline road condition in district k,

sv_i is the social value indicator of road i,

sv_max_k is the maximum social value in district k,

l_i is the length of road i,

l_max_k is the maximum road length in district k.

Appendix C

Appendix to Chapter 3

C.1 Cocoa Quality

Both international and local cocoa prices vary with quality. Factors contributing to poor quality cocoa are high moisture content, mold, germination, a lack of fermentation and slate, a discoloration signaling poor flavor. There is wide agreement on these standards internationally. For a discussion see CAOBISCO (2002) and for a manual specific to West Africa on how to improve cocoa at the farm level see Sonii (2005). Other dimensions of quality affecting price on the international market are various fair-trade and environmental certifications. Such certification generally requires that beans can be verifiably traced to individual producers. In our market, there is not yet the infrastructure to do such tracing, and so this quality dimension does not apply.

Table C.1 shows the average quality and wholesale prices of cocoa bags from the experiment, before the November fall in the international price. As can be seen, moisture content has the highest price elasticity—price falls by 0.32% with a one percentage point increase in moisture. Moisture is an important variable in our market, because wet cocoa rots in storage, destroying value. At an average 11% moisture content, cocoa in our market is substantially wetter than export grade, which requires a maximum moisture content of 7%. For this reason, many exporters maintain large drying facilities. There is an efficiency cost to this organizational structure, as some cocoa that is not dried at the farm gate will be lost to rot in transport.

In our grading system, members of our research team stayed in the warehouses of wholesalers and tested a sample of 50 beans from each bag of cocoa as it arrived. Moisture was measured using Dickey John MiniGAC moisture meters, two of which were generously donated by the manufacturer. Other defects were spotted by eye, after cracking beans open with a knife. Grade A beans have no more than average 11.5% moisture, no more than 2% mold (1 bean of 50), and no less than 72% beans with no defect (36 beans of 50). Grade B beans have no more than 22% moisture, 4% mold (2 beans of 50) and no less than 52% good beans (27 beans of 50). Grade

Table C.1: *Appendix, Cocoa Quality*

Defect	Average per shipment	Price elasticity	Average per pound price by tercile of defect (Le.) 1	2	3
Moisture Content	11%	-0.32%	3,384	3,297	3,263
Mold	2%	-0.02%	3,308	3,353	3,241
Germinated	3%	-0.01%	3,309	3,313	3,298
Under-fermented	15%	-0.02%	3,345	3,333	3,228
Slate	7%	-0.01%	3,323	3,304	3,279

Notes: Data from 916 treatment and control transactions observed before the November decrease in the international price. Elasiticty gives the percentage reduction in price for a 1 percentage point increase in the defect.

C applies to any bean failing to be grade A or B. At baseline, quantities supplied by traders were approximately one third of each.