The London School of Economics and Political Science

Essays on Public Service Delivery and Agricultural Development

Florian Michael Blum

A thesis submitted to the Department of Economics of the London School of Economics for the degree of Doctor of Philosophy. London, June 2017.

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that this thesis consists of approximately 35,000 words.

Acknowledgments

I am indebted to a number of people for their insights, support and encouragement while writing this thesis. Most importantly, I wish to thank my supervisor Oriana Bandiera for her invaluable advice and trust. My advisor, Henrik Kleven, and faculty at the LSE and elsewhere, especially Robin Burgess, Gharad Bryan, Alessandro Gavazza, Greg Fischer, Maitreesh Ghatak, Gerard Padró i Miquel, Rohini Pande and Silvana Tenreyro, also deserve gratitude for their insights.

My work in Tanzania would not have been possible without the support of two exceptional people. Ugo Pica-Ciamarra's ability to build relationships, solve problems and manage stakeholders has taught me more about the world of policy making than I could have hoped for. Longin Nsiima's management and communication skills bridged the gap between the Tanzanian government and academic research. I am deeply indebted to both of them.

Clare Balboni, Kilian Huber, Stephan Maurer, Gustav Meibauer, Tom O'Keeffe, Julia Philipp, Francesco Sannino and Danielle Stein are remarkable colleagues and friends who, through collaboration, advice and friendship, have made my years at the LSE as special as they were. I also owe gratitude to members of the research community at the LSE and the Suntory and Toyota International Centre for Economics and Related Disciplines for their support.

I am deeply indebted to my wonderful family. Birgit, Jürgen, Inge, Hermann, Paul, Renate and Heimfried have, throughout my life, provided the trust, support and encouragement that made this possible. Thank you.

Abstract

This thesis consists of three chapters that study public service delivery, nutrition and agricultural productivity in developing countries.

The first chapter investigates whether imposing price-caps on frontline service delivery agents enhances welfare. I implement a field experiment in which I randomize whether public extension agents are subject to a price-cap or not. I find that while price-caps are effective in enhancing the affordability of extension services and increasing recipients' surplus, they also reduce the geographic coverage of services. This suggests that price-cap regulation creates a tension between making services affordable and providing incentives for agents to serve remote recipients. I then show that the marginal welfare effect of reducing discretion over prices can be expressed as a function of two sufficient statistics: the elasticity of geographic service coverage with respect to the price-cap and the price elasticity of demand. Calculating the welfare effects, I find that any reduction of agents' discretion reduces social welfare.

The second chapter is concerned with contract design in public service delivery when delivery agents are boundedly rational. A theoretically efficient contract that minimizes moral hazard costs and avoids behavioural distortions charges agents a fixed fee for the usage of public assets and makes them residual claimants on its returns. I investigate whether such contracts are indeed efficient in practice by investigating whether imposing lump-sum fees on livestock extension agents distorts their choices. Using a field experiment, I first show that, contrary to classic economic theory, levying a fixed fee on agents leads them to increase user fees for a livestock vaccine and induces demand effects that reduce quantities. To understand the mechanisms underlying this result, I implement a series of lab-in-the-field experiments with a subset of the field-experimental participants. The results suggest that instead of setting prices for user fees as mark-ups over marginal costs agents use simplified rules-of-thumb that anchor pricing decisions on aggregate profits. The results highlight that boundedly rational behavior can reduce the effectiveness of adopting fixed fee contracts.

The third chapter investigates whether improvements to agricultural production technology, a common response to undernutrition, can enhance food security and improve nutrition. In India, groundwater irrigation using tube wells has long been promoted as a means to reduce rainfall-dependence and enhance food security. The merits of adopting tube wells have, however, been debated widely, with opponents fearing a deprivation of smaller farmers and impoverishment of rural laborers. To evaluate the causal effects of tube well adoption on nutrition, I employ an instrumental variable framework that exploits variation in land suitability for deep groundwater irrigation caused by differences in hydrogeological structures. I find that groundwater irrigation significantly improves nutrition across the income spectrum: a one standard deviation increase in the proportion of cropped area irrigated with tube wells increases calorie intake by 770 to 915 calories per day. In addition, groundwater irrigation generates positive spillovers on the calorie intake of urban populations and households not employed in agriculture. I present additional evidence which suggests that these effects are driven by increases in agricultural productivity that reduce staple prices and raise wage rates. The findings thus highlight the value of groundwater irrigation in fighting undernutrition and promoting agricultural development.

Contents

List of Tables

List of Figures

Chapter 1

The Value of Discretion - Price-Caps and Public Service Delivery

It is often argued that price-caps – a ceiling on prices charged by monopolistic suppliers - are necessary to redistribute surplus and make essential goods and services affordable. I explore whether price-caps lead to welfare improvements through a field experiment in which I randomize whether public livestock extension agents are subject to a price-cap or not. This intervention has three effects. First, conditional on being served, the treatment increases the consumer surplus available to recipients: the price-cap reduces average prices by 17% and the within-agent standard deviation of prices by 42%. Second, the intervention increases the affordability of extension services: the price-cap increases the share of previously unserved and needy customers in the beneficiary pool by 15% and 9%, respectively. Third, the price-cap reduces the geographic coverage of services by decreasing the likelihood that agents will serve remote villages by 25%. This suggests that price-cap regulation creates a tension between making services affordable and providing incentives for agents to serve remote recipients. In light of this trade-off, I show that the marginal welfare effect of reducing discretion over prices can be expressed as a function of two sufficient statistics: the elasticity of geographic service coverage with respect to the price-cap and the price elasticity of demand. Calculating the welfare effects, I find that any reduction of agents' discretion reduces social welfare.

1.1 Introduction

A key function of governments is to assure that citizens are provided with essential services for welfare and economic development. When agents responsible for supplying these services have market power, decentralized market outcomes can generate an inefficient allocation of services and allow agents to extract surplus from recipients.[1] These undesirable outcomes threaten the ability of governments to provide for their citizens. Government imposed price-caps, which aim to reduce agents' discretion over prices, are one commonly proposed solution to this challenge.[2] Theoretical work suggests that such interventions may be welfare-enhancing, as they can redistribute surplus and improve aggregate efficiency (Laffont and Tirole, 1993; Laffont, 2005). In addition, there is evidence that charging for essential services in a developing country setting can adversely affect peoples' willingness to use these services (Cohen and Dupas, 2010; Ashraf, Berry, and Shapiro, 2010). Those findings highlight a potential value of low price-caps and suggest that using price-caps as an instrument to reduce agents' discretion over prices may be welfare-enhancing.

In many situations, unregulated monopolistic agents responsible for supplying public services do not, however, only hold decision power over prices. Rather, these agents also have discretion over which and how many recipients to serve.[3] In such situations, a theoretical literature argues that allocating discretion over pricing to agents can be welfare-enhancing (Grossman and Hart, 1986; Aghion and Tirole, 1997). In particular, when governments are unable to impose and enforce universal service obligations, discretion over prices allows agents to extract surplus, creating an incentive for them to extend service coverage.[4] Price-caps reduce agents' discretion and, by limiting their ability to extract surplus, can reduce incentives to extend

[1] For example, Sappington and Weisman (2010) in their review of evidence on regulatory policy state: *"When competition is unable to impose meaningful discipline on incumbent suppliers of essential services, regulation can be employed as an imperfect substitute for the missing market discipline."*

[2] Recent newspaper headlines include: *"Hillary Clinton Proposes Cap on Patients' Drug Costs"* (New York Times, 2015); *"FCA proposes price-cap for payday lenders"* (Financial Conduct Authority, 2014); *"Kenya to cap interest rates on bank loans"* (Financial Times, 2016).

[3] Chaudhury et al. (2006), for example, document an inability of governments to limit absenteeism among teachers and health workers in developing countries.

[4] In Tanzania, the setting of this study, the government cites the binding fiscal constraints and high monitoring costs to explain the absence of universal service obligations.

service coverage.[5] Capping prices for essential services can therefore also reduce welfare. The question of whether to cap prices or not hence has to solve a trade-off between making services affordable and providing incentives for agents to extend service coverage.

In light of this trade-off, evaluating whether reducing agents' discretion over prices improves welfare requires knowledge of two parameters: the price elasticity of demand (*"intensive margin elasticity"*) and the impact of price-cap regulation on the number of markets served (*"extensive margin elasticity"*). While empirical estimates of demand elasticities are available for many goods and services, empirical evidence on the latter parameter is limited. As a result, little is known about whether price-cap regulation enhances or reduces welfare.

This paper provides experimental evidence to address this gap. In particular, I estimate the intensive and extensive margin elasticities and examine the trade-off between discretion and price-cap regulation for an essential public service provided in developing countries: agricultural extension.[6] As the basis of this investigation, I focus on the provision of I-2 poultry vaccines that protect against Newcastle Disease (ND). ND is highly prevalent in East Africa and is lethal for infected birds, therefore posing a substantial economic risk for populations dependent on agriculture.[7] I-2 is the primary ND vaccine used by rural farmers in Tanzania. It is exclusively produced by the government and distributed directly to farmers by specialized government service delivery agents. These agents receive the vaccine at a subsidized price from the government and distribute it by travelling to farmers in a geographically defined service area, where they perform veterinary examinations and apply the vaccine.[8] Each service area has exactly one agent tasked with providing

[5] The incentive value of discretion provides a possible explanation for an observed correlation between worker autonomy and project completion in the public sector (Rasul and Rogger, 2016). Duflo, Greenstone, et al. (2014) also study how agents' discretion affect the effectiveness of public operations but, instead of focusing on incentives, highlight that discretion allows public agents to use private information.

[6] On average, African countries spend 5% of their budget on agricultural extension, compared to 8% spent on primary health services.

[7] Poor farmers are particularly affected by ND, as their livelihoods and asset holdings are especially dependent on livestock.

[8] Private provision is widely viewed as sub-optimal to public provision, as effective ND vaccination requires adequate handling and application of the vaccine as well as the ability to detect preexisting poultry infections. A central concern is that private market provision would compromise service quality which risks reducing the acceptance and adoption of I-2 vaccinations. For example, the formal instruction sheet for I-2 reads: *Never vaccinate chickens which are already unhealthy. If they die, the owner will blame you.*

vaccines.[9] Agents collect user charges from farmers in order to finance the vaccine delivery and application.[10] While there exist privately supplied imperfect substitutes for I-2, access to them is limited to urban areas and larger scale farmers, and trust in the effectiveness of privately supplied vaccines is low.[11] There is no formal service charge schedule in place that regulates prices for vaccinations. This lack of competition and price regulation thus gives substantial price discretion to agents.

To investigate whether capping prices for ND vaccinations enhances welfare, I collaborate with the Tanzanian government. Working with the Ministry of Agriculture, Livestock and Fisheries, I allocate 550 wards, the main administrative units at which agents are organized, to one of two experimental groups.[12] 274 of these wards are assigned to the control group that mirrors the Tanzanian status quo. Agents assigned to these wards have full discretion over pricing and typically charge recipients either 50 or 100 Shillings (\$0.023 or \$0.046) per vaccination. In contrast, 276 wards are assigned to the treatment group, in which agents face a price-cap of 80 Tanzanian Shillings (\$0.035) per vaccination. Agents in both experimental groups have full discretion over which and how many recipients are served and which geographic areas are targeted.

The evaluation of this intervention yields three main results. First, the price-cap is effective in reducing agents' discretion over prices. I find that the intervention reduces average prices charged by agents by 17% and the within-agent standard deviation of these prices by 42%. I also show that farmers whose main source of livelihood is agriculture and smallholders were paying higher prices in the status quo and benefit particularly from the price-cap. Second, the price-cap affects the composition of the beneficiary pool: I find that the cap increases the share of previously unserved recipients in the beneficiary pool by 14%. In addition, the price-cap increases the share of farmers whose livelihood depends on agriculture and smallholders in the

[9] Agents are, on average, responsible for 4 villages. The government's ability to introduce competition *within* the current system is constrained by a lack of qualified extension staff.

[10] Financing service delivery agents through user charges is common in developed and developing countries. See section 1.2.2 for an overview.

[11] 11% of recipients in my sample report having access to an imperfect substitute. Such recipients are predominantly located in urban areas. The absence of a quality control system for privately provided vaccines makes I-2 delivery through public agents the only available option to obtain access to a certified ND vaccine.

[12] While the Newcastle Disease program covers most of Tanzania, this study focuses on four regions: Dodoma, Iringa, Morogoro and Tanga. The sample includes all wards in those regions that were assigned a public service delivery agent at the time of the study.

recipient pool by 8% and 9%, respectively. Third, the price-cap reduces the geographic coverage of services by reducing the proportion of villages that are served in each service area by 12%. This reduction is driven by a reduced likelihood of agents visiting remote villages: Agents are 25% less likely to visit the furthest half of villages and 26% less likely to visit the most remote village in their area of responsibility. Taken together, the evaluation suggests that the price-cap induces agents to shift towards comparatively closer but previously unserved recipients. These results therefore present direct evidence in support of the trade-off discussed previously: On the one hand, price-caps can enhance welfare by making services affordable for previously unserved recipients in markets that are comparatively cheaper to serve. In addition, price-caps can redistribute surplus to recipients. On the other hand, price-caps can adversely affect welfare by reducing agents' incentive to extend services to remote markets that are costlier to serve.

Given the countervailing effects induced by the price-cap, it is not clear whether it is welfare-enhancing for the government to restrict the agents' discretion over prices and, if so, to what extent. To assess this, I develop a model of monopoly regulation that is consistent with the empirical results. Assuming that the government maximizes a weighted sum of consumer and producer surplus, this model allows me to express the aggregate social welfare effect of marginally reducing prices below the full discretion level as a function of two sufficient statistics: First, the elasticity of the proportion of villages served with respect to the price-cap, which acts as a sufficient statistic for the welfare loss incurred through the reduced number of markets served. Second, the price elasticity of demand, which captures the welfare benefits associated with reducing prices.

The relationship between price-caps and service allocation evaluated as part of the field experiment is sufficient to estimate the elasticity of the proportion of villages served with respect to the price-cap. However, the price-cap treatment is insufficient to estimate the elasticity of demand for I-2 vaccinations.[13] To address this, I induce additional experimental variation in costs. In particular, I independently assign agents in 273 wards to a second treatment that requires them to contribute a fixed fee of 25,000 Tanzanian Shillings (approximately \$11.40) to cover the cost of the vaccine if they agree

[13] This is because disentangling supply and demand responses using one source of variation is difficult.

to participate in the program.[14] Agents in the remaining 277 wards receive the vaccine for free. This intervention generates variation in prices because, as I show in a lab-in-the-field experiment, agents choose prices based on average instead of marginal costs.[15] Using this variation, I obtain an estimate of the aggregate elasticity of demand of -1.22.

I then combine these estimates to evaluate the effect of marginally reducing agents' discretion over prices using the sufficient statistics formula. I find that marginally reducing prices below the monopoly level induces a social welfare *loss* of between 3% and 12% of total sales revenue per agent. In contrast, a counterfactual scenario that ignores the supply effects induced by the intervention suggests that price-cap regulation would lead to moderate welfare increases of between 0.5% to 2.5% of total sales revenue per agent on the margin. Taken together, this result highlights the importance of studying the effect of price-interventions on service coverage: While price-caps can increase welfare by making services more affordable, reducing agents' discretion over prices can be counterproductive when agents also have discretion over which markets to serve.

This paper provides evidence on the welfare effects induced by imposing price-caps on monopolistically supplied services. By doing this, it complements a literature on pricing for public services. Theoretical work in this area has investigated how to optimally regulate providers of public services. Laffont and Tirole (1986), show that optimal level of regulation creates a tension between rent-extraction and efficiency. In deriving the optimal regulation mechanism, they show that price-caps can be optimal for highly efficient firms. While they also show that price-caps might allocate excessive rents to less efficient firms compared to alternative contracts, both Laffont (2005) and Alonso and Matouschek (2008) argue that in the absence of transfers and verifiable information on costs, price-caps remain the optimal regulatory policy.[16] This paper provides direct and causally identified empirical evidence on the effectiveness of price-cap regulation and the rent-extraction versus efficiency trade-off.

[14] To avoid challenges associated with liquidity constraints, this fee was collected after the vaccination campaign had ended.

[15] The intervention and the lab-in-the-field experiment are explained and evaluated in chapter 2.

[16] Empirical evidence that tests those theories to date has primarily focused on the comparison of various regulatory mechanisms with each other. For example, the review by Abel (2000) highlights that incentive regulation through price-caps in the telecommunication market reduces prices and provides incentives to invest in infrastructure compared to rate of return regulation.

Empirical evidence on pricing for essential services in developing countries has focused primarily on demand effects induced by prices. For example, Cohen and Dupas (2010) investigate how charging for insecticide treated bed nets affects demand and use. Similarly, Ashraf, Berry, and Shapiro (2010) show that prices for water-purifies can act as a screening device for high-use customers. This paper complements this work by highlighting the importance of considering the interaction between supply and demand decisions as the basis of setting prices for public service provision.

More broadly, this paper provides field-experimental evidence on the welfare effects induced by allocating discretion to agents. Thereby it complements a new and rapidly growing literature that investigates the organization of the public sector in developing countries.[17] While previous empirical work has focused on understanding how contracts for public agents can be designed to encourage effort (e.g. Khan, Khwaja, and Olken, 2016b; Muralidharan and Sundararaman, 2011 and Olken, Onishi, and Wong, 2014), a more recent stream of theoretical literature has begun to take into account that contracts in practice are rarely complete.[18] This poses a central question as to who should have decision rights in cases not covered by the employment contract. While theoretical work highlights that allocating discretion to agents in situations not covered by contracts can act as an incentive (Aghion and Tirole, 1997), empirical evidence on this mechanism is rare. Previous empirical work has primarily focused on documenting a correlation between employee autonomy and public project completion (Rasul and Rogger, 2016) as well as improved information use as a result of regulatory discretion (Duflo, Greenstone, et al., 2014). This paper is, to the extent of my knowledge, the first to show the existence of incentive effects as a result of discretion.

The remainder of this article is structured as follows. Section 1.2 describes the relevant features of the setting in which the study takes place. Section 1.3 outlines the experimental design and section 1.4 presents the results. Section 1.5 theoretically conceptualizes the effects of the price-cap intervention and derives a sufficient statistics formula to evaluate the welfare impact of capping prices. Section 1.6 presents the welfare analysis. Section 1.8 concludes.

[17] The review by Finan, Olken, and Pande (2015) provides a comprehensive overview of this literature.

[18] See Grossman and Hart (1986) for the theoretical foundations of this argument.

1.2 Setting

This project explores the effect of price-caps in the context of agricultural and livestock extension services in Tanzania, which is a public service delivery program administered by local governments and coordinated nationally by the Ministry of Agriculture, Livestock and Fisheries.[19] The service aims to subsidize animal health and production services to make them available to small-scale rural farmers who are excluded from private input markets. It hence provides crucial economic infrastructure in an economy in which over 60% of households depend on livestock for their livelihoods, and where livestock is the primary asset held by rural households.

In light of their importance, livestock and agricultural extension services are one of the key services provided in developing countries (Swanson, Farner, and Bahal, 1989; Feder, Willett, and Zijp, 1999). According to the United Nations Food and Agricultural Organization, governments in Africa spend on average 5% of their total annual expenditure on agricultural services, which is only slightly lower than their expenditure on health (8%), with a total global spending on extension services estimated at $31 billion in 2008. In total, 10,891 extension agents were employed in Tanzania in 2012, comprising approximately 5% of local government staff. This leaves extension workers as the third highest proportion of government employees, after education and health services.[20]

1.2.1 Agents

The agents responsible for delivering livestock extension services in Tanzania are para-veterinarians, who are employed by local governments. Agents have advanced professional qualifications and typically hold a diploma from specialized training institutes in animal health, animal production or general agriculture. They are responsible for 1 to 12 villages, averaging 4 villages per agent, and typically operate in areas where the private coverage of livestock services is low. Agents work by themselves and have their own geographically defined area of responsibility, in which they face no competition from other public providers. The main organizational unit of agents at the local level are wards, which are accumulations of roughly 8 villages. There are, on average, two agents per ward. Agents in the same ward interact on a daily basis

[19] This division of responsibilities is the common organizational form for extension in Africa (Crowder et al., 2002) and Latin America (Wilson, 1991).

[20] Total numbers of local government staff in 2007 were 224'114, with 148'607 being teachers and 39'217 being health workers.

and, while maintaining their own geographic areas, typically coordinate their work.[21] There was no entry or exit of agents during my study period.

Agents' primary task is to travel to farmers in order to provide services. Around 25% of agents have access to a government motorcycle, whereas the remainder travels by foot and uses public motorcycle taxis and buses to reach farmers.[22] The government does not provide any reimbursements of travel costs. Instead, service delivery is completely funded through user fees.

1.2.2 User Fees

A key component of agents' contracts is their compensation structure. Delivery agents are compensated through a mix between government wages and user fees. Specifically, agents receive a flat compensation of around $200 per month.[23] In addition, agents can collect user fees from farmers, which cover delivery costs and act as performance pay. In the status quo local governments have allocated full discretion over pricing to the agents.

Existing evidence shows that the rationale for employing user fees in extension is threefold. First, user fees are equivalent to commissions for private sellers, which provide high-powered incentives for agents to exert effort. This addresses one of the key challenges in public provision of extension services, which have traditionally suffered from a lack of mechanisms to induce providers to exert effort (Howell, 1986; Farrington et al., 2002). Second, high monitoring costs make performance pay schemes, in which payments from the government are linked to output, infeasible. User fees reduce monitoring costs as they delegate monitoring responsibilities to the recipients of services, who can directly observe output (J. R. Anderson and Feder, 2007; Kidd et al., 2000). In the presence of high monitoring costs, the alternative to user fees are therefore fixed wages, which provide no incentives to agents. Third, user fees are a cost-sharing device that reduce pressure on local governments' budgets to fund service delivery (e.g. Cary, 1998).

While cost-sharing schemes are common for agricultural and livestock extension programs (Rivera and Gustafson, 1991; Dancey, 1993), compensation schemes that partially rely on user payments are also present for other public

[21] To avoid spillovers and interaction, treatments are assigned at the ward level. See also section 1.3.1.

[22] This process of service delivery is the modus operandi for a number of key service delivery programs in developing countries. For example, Ashraf, Bandiera, and Lee (2016) study community health workers in Zambia. Community health workers are expected to devote 80% of their time to household visits and are hence required to incur similar costs as the agents in my setting.

[23] Wages vary across local government administrations.

service delivery schemes, such as health services and food distribution.[24] On the one hand, there is a substantial amount of evidence documenting bribes paid to public service delivery agents, which, while illicit, play a similar role to a user fee.[25] More formally, a number of countries and organizations have, either temporarily or permanently, switched to a system that relies on user fees to cover expenditure of health facilities.[26] For example, Deserranno (2016) studies community health workers recruited by BRAC, an international NGO, in Uganda. Those workers are tasked with providing basic health services to local residents and are compensated through medication sales to service recipients. Similarly, food distribution systems such as Solidaridad in the Dominican Republic and Raskin in Indonesia rely on co-payments to finance distributors and local government agents, respectively (Busso and Galiani, 2014; Banerjee, Hanna, Kyle, et al., 2015).

1.2.3 Services

Agents provide a range of services to recipients. This includes preventive animal health treatments such as vaccinations, deworming procedures and reactive treatment to address common livestock diseases. Unregulated user fees are charged for all animal health services.[27] As part of this project I focus on the provision of I-2 vaccines as one dimension of service provision.[28] I-2 is a thermotolarent vaccine for poultry that protects against Newcastle Disease (ND), a viral disease that is transmitted between birds and leads to almost 100% mortality in affected and unvaccinated chicken. Estimates from Tanzania suggest that more than 30% of chicken die from ND every year, leading to an annual cost of up to $78 Million (Msami, 2007). As part of an I-2 vaccination program, agents receive subsidized vaccines from the

[24]Countries that have, among others, implemented such schemes include Cameroon, Chad, Mali, the Central African Republic, India, Kenya, Nicaragua, China, New Zealand, the Netherlands, Chile, Australia as well as most OECD countries. See the comprehensive reviews by Haan et al. (2001) and J. R. Anderson and Feder (2007) for details.

[25]See, for example, Deininger and Mpuga (2005) for an overview from Uganda.

[26]Examples include Burkina Faso, Kenya, Papua New Guinea , Uganda, South Africa, Colombia, Sudan and Lesotho. For an overview of those experiences, see the review paper by Lagarde and Palmer (2008).

[27]The program also aims to provide recipients with traditional extension services, such as advice regarding animal husbandry practices and information on optimal feed composition. As the extent to which agents engage in providing those services and the prevalence of user fees for such services is more heterogeneous, the main focus of this investigation is on animal health services.

[28]By focusing only on I-2, I ignore possibly compounded adverse effects of the intervention induced by an additional reduction in other services provided to remote villages. Estimates are therefore a lower bound.

government and then travel to recipients in order to apply vaccinations to farmers' livestock.

Four characteristics of I-2 service provision make it particularly suitable for my study. First, the Tanzanian government is the only producer and provider of I-2 vaccinations. This gives delivery agents market power and allows them to extract surplus from recipients.[29] In addition, this characteristic, together with the fact that there is no competition between public providers, simplifies the interpretation of my results as it alleviates concerns that the treatment shifts market shares between different providers. Second, the public provision of I-2 is based on a vaccination calendar which requires a coordinated vaccination effort on a four-monthly basis.[30] During such campaigns, agents' primary task is the provision of vaccinations. As this study focuses on vaccination periods, this reduces the concern that the intervention might affect the effort agents exert on alternative tasks. Third, in order to eradicate ND, vaccination levels in the poultry population need to be maintained at at least 85% (Boven et al., 2008). Given high turnover rates of flocks, the fact that an important transmission channel of ND is through non-domesticated birds, and the low coverage of vaccination programs, Tanzania's system is unlikely to eradicate ND in the near future. In light of this argument, I simplify the analysis by abstracting from externalities and focus on consumer surplus in my welfare analysis.[31] Finally, I-2 is the only public animal health service provided to poultry keepers. Although I do not observe prices for other services provided by agents to I-2 recipients, this property allows me to investigate potential cross-price effects by investigating whether the treatment induces agents to target more non-poultry farmers.

1.3 Experimental Design and Data

The experiment discussed in this paper examines how price-cap regulation affects public service delivery in the context of the public provision of I-2

[29]Private provision is widely viewed as sub-optimal to public provision, as effective ND vaccination requires adequate handling and application of the vaccine as well as the ability to detect preexisting poultry infections. A central concern is that private market provision would compromise service quality which risks reducing the acceptance and adoption of I-2 vaccinations.

[30]Vaccination campaigns follow regional rainfall patterns and typically take place in January, May and September. Campaigns last three weeks before the lack of cooling renders the vaccine unusable.

[31]A complementary argument in favor of focussing on I-2 relates to the common criticism of public extension services in Africa that the services provided are largely ineffective and add little to farmer productivity (e.g. Dejene, 1989; Gautam, 2000). Focusing on a vaccine for which effectiveness has been medically proven alleviates this concern.

Newcastle Disease vaccination. This section explains the experimental design before describing the data used as the basis of this evaluation.

1.3.1 Experimental Design

This paper examines the welfare effect of an intervention that imposes and enforces a maximum price ("price-cap treatment") for I-2 vaccinations. To avoid spillovers resulting from coincidental interactions between agents, treatment assignment was performed at the ward instead of at the individual level. I randomly and independently assign each of the 550 wards in the study to either the control or the treatment group. Table 1.1 displays the basic experimental design. Group allocation was stratified by 108 strata, where each stratum was defined by a district identifier and two binary variables, indicating whether all agents in the ward had specialized in general agriculture and whether only one agent was assigned to the ward.

This intervention was carried out during the first I-2 vaccination campaign of 2016.[32] The study covers the time period between January and February 2016 and enumerated the universe of agents in four of Tanzania's 30 regions (Dodoma, Iringa, Morogoro and Tanga). Figure 1.1 shows a map of the study regions and the wards included in this study. All 27 districts in the four regions were included in the study. The study area was chosen to include a wide variety of agricultural environments while assuring geographic proximity to the ministry headquarter in Dar Es Salaam. From each study district, I obtained administrative records of all employed agents, detailing their name, specialization, ward of responsibility and telephone number. In total, I collected this information for all 990 agents registered in the four regions, which forms the provisional sample of this study. 832 of those agents attended the training and participated in the vaccination campaign.

1.3.2 Implementation Procedures

All participants were invited to attend a 90-minute meeting at the district headquarter at the beginning of the campaign to collect the vaccine and receive instructions. Agents who attended this meeting received a show-up fee to cover their transport expenditure. Payments varied between 10,000 and 50,000 Tanzanian Shillings ($4.50 to $22), depending on the distance and available transport methods. The specific instructions were announced to participants only after they had arrived for vaccine collection at the district headquarters but before they departed to the field again. Thus the decision

[32] The timing is described in detail in the appendix to this paper.

whether to attend the vaccine collection should be viewed as exogenous with respect to the experiments. Trainings and surveys were conducted on different days for the different experimental groups and districts to avoid spillovers.

During this meeting, agents in the control group were reassured that they were allowed to collect fees from farmers which they could keep for themselves. Agents were specifically encouraged to profit financially from the transaction, stating that the government viewed user fees as a way to motivate employees and compensate them for good performance. In addition, the instructions reiterated that agents were allowed to charge farmers any price they chose and that it was acceptable to charge different prices to different farmers.

Agents were then informed that the ministry wanted to keep better records of how many chickens were vaccinated and that therefore reporting procedures during this campaign would differ slightly from the status quo. In particular, a condition of participation in the vaccination campaign was that agents would issue formal receipts to every farmer served and submit the receipt information directly to a central headquarter using a phone based reporting system. Agents were specifically told that the ministry would contact farmers to verify that the information provided on receipts was correct. In order to assure compliance with this reporting system and encourage effort, the ministry offered a bonus payment of 60 Tanzanian Shillings (approximately $0.025) for every verified vaccination.

After the instructions, training staff collected data on demographics, work history and workplace characteristics of participants. Ministry staff then distributed the vaccines to agents, supplying agents with as many doses as they requested for their area of responsibility and informing them that more doses would be stored at the district headquarter where they could be picked up in case of additional demand.

1.3.3 Price-Cap Treatment

Compared to the control group, the instructions given to agents in the price-cap treatment differed only with regards to the rules on pricing. In particular, ministry officials informed participants that they were free to set any price up to 80 Tanzanian Shillings (approximately $0.035) per vaccination. This cap was calibrated to balance two considerations. On the one hand, it had to be sufficiently low to be binding in order to affect agents' pricing and allocation behavior. On the other hand, it had to be sufficiently high to allow agents to cover their marginal costs. To achieve this balance, the maximum price of 80 Tanzanian Shillings was chosen after careful consultations with experts from the Tanzanian Veterinary Laboratory Agency (the vaccine's main producer), MALF, local governments and international academics. In addition, this decision was also based on a mixed methods pilot study conducted by the author that analyzed pricing behavior during previous I-2 campaigns. To avoid setting a price-cap that would not allow agents to recover their marginal costs of applying the vaccine, the cap was conservatively set to bind only for comparatively high prices.

As price-caps are only effective if they can be enforced, I took the following measures to ensure compliance with the price-cap: First, the receipts that are normally employed during campaigns were amended to contain the national emblem of the United Republic of Tanzania, transforming them into official government documents. As receipts require the delivering agent's signature, forging them is equivalent to tempering with official government documents which is punishable by law and can lead to dismissal. Anecdotal evidence suggests that this incentive mechanism was taken seriously: local government level supervisors requested detailed information on verified compliance behavior by their employees in the aftermath of the intervention to discipline non-compliant employees.

Second, MALF conditioned the bonus payment of 60 Tanzanian Shillings per vaccination on compliance with the price-cap. This scheme makes it incentive compatible to comply with the price-cap as long as deviation yields a price lower or equal to 140 Tanzanian Shillings per vaccination and the detection probability is sufficiently high. Given that 99% of transactions in the control group were conducted at user charges below this threshold, compliance was incentive compatible for the vast majority of transactions.

1.3.4 Data

Data used as the basis of this paper was collected from two different sources: administrative government receipts and a survey of service recipients. I designed and conducted the recipient survey specifically as part of this project. In addition, I implemented a new procedure of reporting service provision receipts via text message to increase accuracy and usability of the administrative data. I complement this data using information from a baseline survey of agents, described in detail in the appendix.

The information provided on official government receipts and the number of receipts issued constitutes my provisional outcome data. The receipts detail each recipient's name, contact number, village, the date of the visit, the total user fee collected and the number of vaccinations applied. After issuing the receipt, agents electronically transmitted the receipt information to a government database using a text-message template.[33]

Using the receipt data, I construct two unverified, and therefore provisional, outcome measures. First, the total user fee collected divided by the number of vaccinations applied gives a direct measure of the per unit price charged to farmers. Second, the total number of farmers served can be measured through the total number of receipts submitted.

After the end of the vaccination campaign, I administered a survey to service recipients. The survey was conducted over a period of six weeks between March and April 2016 and sampled a randomly selected fraction of 15% of all receipts submitted, selected randomly and stratified by agent. This led to a total sample of 4'516 receipts selected for surveying and verification,[34] 80% of which were successfully contacted.[35] The farmer survey collected detailed information on the service provision and on recipient characteristics, thereby verifying that the service was actually provided and collecting verified information on user fees.

I use the information obtained from this survey to construct my main outcome measures. In order to arrive at a measure of the total number of farmers served, I multiply the number of verifiable receipts per agent by the agent-specific sampling weight of each receipt. I repeat the same procedure for the average price, total revenue collected and the total number of chickens vaccinated. In order to analyze outcome measures related to service allocation, I use farmer survey data on farmer demographics, asset holdings distance between farmers' home and the agent's headquarter and farmers' sources of livelihoods.[36]

[33]In total, agents issued 31,657 valid receipts, accounting for 702,762 animals vaccinated.

[34]Rounding errors induced by the stratification led to a sample that is slightly smaller than 15% of 31657.

[35]The procedures to contact farmers are described in detail in the appendix. Among the farmers not reached, enumerators were unable to reach 42% because of incorrect or invalid contact details. In total, phone survey procedures therefore were able to assess the validity of almost 90% of receipts sampled. I treat the remaining receipts as unverifiable and hence incorrect.

[36]While it would have been optimal to conduct a detailed consumption survey as part of this exercise in order to obtain a more precise measure of farmers' livelihoods, budget limitations rendered this option infeasible.

1.4 Results

This section presents the empirical methodology and results from the evaluation of the price-cap treatment. Table 1.2 presents summary statistics and a balance check using baseline characteristics of agents participating in the campaign. All characteristics in the table were chosen prior to estimating the balance checks. These results suggest that experiment participants are similar across the treatment and the control group. Panel A considers agent level characteristics while panel B investigates differences in workstation characteristics. None of the 28 differences are statistically significant at the 5% level, which confirms balance at baseline.

The presentation of the results proceeds in two steps. Subsections 1.4.1 and 1.4.2 present results that investigate how the price-cap affects transaction prices. Subsection 1.4.3 then focuses on the allocation and extension of services to show the central trade-off between affordability and coverage necessary for welfare analysis.

1.4.1 Impact of the Price-Cap on Prices

I begin the evaluation by estimating the impact of the price-cap treatment on user fees charged over the course of the vaccination campaign. As treatment assignment was randomized, the empirical methodology is straightforward. I estimate Ordinary Least Squares (OLS) equations of the following form:

$$y_{iwd} = \beta_0 + \beta_1 PriceCap_{wd} + \beta_2 X_{wd} + \gamma_d + \epsilon_{iwd}$$

where y_{iwd} is the outcome of interest for participant i in ward w and district d, $PriceCap$ is a binary variable that indicates whether agents' wards were assigned to the price-cap treatment, and X_{wd} denotes a vector of ward-level stratification variables.[37] The coefficient of interest is β_1. I also include district level fixed-effects (γ_d), as the assignment lottery was stratified by these strata. As the treatment is perfectly correlated within wards, every specification reports robust standard errors clustered at the ward level.

[37] X_{wd} also contains an indicator for whether a ward was assigned to a cross-cutting treatment ('Fixed Fee Treatment') explained in more detail in section 1.6 and in chapter 2. Note that, as I find no evidence of interaction effects between the price-cap and the cross-cutting treatment, I treat both as separate experiments. Given that the two treatments were assigned as part of a cross-cutting design, treatment effects of the price-cap intervention should therefore be interpreted conditional on 50% of the sample being assigned to the cross-cutting treatment.

I first investigate the effect of the price-cap treatment on the distribution of prices. Panel A in figure 1.2 plots a histogram that visualizes the distribution of prices in the status quo, using farmer survey data from the control group.[38] Prices follow a bimodal distribution with peaks at 50 and 100.[39] Panel B overlays the distribution of prices in the treatment group over the histogram from the control group. Significant bunching at 80 suggests that the price-cap was binding and effectively reduced the level of prices.[40] Columns 1 and 2 in table 1.3 confirms this finding by showing that the price-cap reduced average prices by approximately 17%, which is statistically significant at the 1% level.[41] Column 5 confirms the visual impression of bunching at 80 by showing that this intervention increased the fraction of transactions per agent where a price of 80 Shillings was charged by a factor four, from around 5% to 20%.

I then investigate the impact of the price-cap on the within-agent variation of prices. To do this, I calculate the residuals of a regression of prices on agent fixed effects. Figure 1.3 presents a box-plot of the residuals, separated between treatment and control group, to visualize the effect of the treatment on price variation. The height of the box corresponds to the difference between the 25^{th} and 75^{th} percentile of residuals, and the whiskers correspond to the 10^{th} and 90^{th} percentile, respectively. I find that the intervention reduced price disparities between recipients. In particular, the figure shows that within-agent price variation is substantially lower in the treatment group than in the control group.

Columns 3 and 4 in table 1.3 show that this reduction in variation is also statistically significant. Column 3 reports the estimate of the treatment effect on the within-agent standard deviation of prices using farmer survey data, whereas column 4 repeats the same analysis using receipt data. The results suggest that the treatment reduced the within agent variation of prices by 42

[38]To improve the visualization, the histograms are truncated at 200 Tanzanian Shillings, which excludes less than 1% of all observed transactions.

[39]As expected when designing the intervention, the price-cap hence only binds for a subset of transactions. The histogram also shows that less than 1% of transactions in the control group occur at prices above 140 Tanzanian Shillings, which assures that complying with the price-cap is incentive compatible.

[40]Figure 1.2 also suggests that the price-cap intervention increased the mass of the price distribution for prices significantly *below* the cap. Anecdotal evidence suggests that this is driven by difficulties with calculating multiples of 80 and a tendency to round down to the nearest 1000 for the total price.

[41]The estimate using the farmer survey data is slightly lower than the estimate obtained from the receipt data, which is partially driven by (detected) under-reporting of prices on the receipts.

to 44%, which is statistically significant at the 1% level. The results do not differ substantially between the receipt and the farmer survey data.

1.4.2 Who benefits from the Price-Cap?

The bimodal distribution of prices in the control group and the evidence of within-agent price variation suggest that different farmers, even if they are served by the same agent, pay different prices. To gain an understanding of the impact of the price-cap on welfare, it is important to understand which farmers pay higher prices and thus are more likely to be affected by the cap. To do this, I present two pieces of evidence.

First, table 1.4 presents the correlates of unit prices in the control group. Column 2 shows that agents offer lower per-unit prices for larger flocks. In particular, every additional vaccination applied by the agent is associated with a quantity discount of 0.12 Shillings. In addition, indicators of education and asset holdings suggest that more disadvantaged recipients also, ceteris paribus, pay higher vaccination prices. Column 4 shows that recipients who have more than primary education pay 9 Shillings less per vaccination, and columns 5 to 7 suggest that increased holdings of non-poultry livestock assets are associated with lower per vaccination prices. The estimates of the partial correlations remain significant even after controlling for self-reported travel time to the recipient. Finally, while there is some evidence that prices are related to travel and application costs, it is unlikely that this variation can conclusively explain the observed variation in prices. Column 1 in table 1.4 first shows that, while not statistically significant, an additional minute of walking to the recipient's village, as measured through self reported walking distance, is associated with marginally higher vaccination prices. I then investigate the extent to which measures of cost, in particular walking and motorcycle travel times and the number of vaccinations applied, can explain the within-agent price variation. I find that while individually the indicators show associations with prices, controlling for cost measures only explains 7% of the within-agent price variation. Taken together, this is suggestive that agents use discretion over prices in the status quo to extract surplus.

Second, in light of the previous results, it is instructive to investigate who benefits most from the price-cap. To investigate this, I run regressions at the transaction level that take the following form:

$$Price_{fwd} = \beta_0 + \beta_1 PriceCap_{wd} + \beta_2 K_{fwd} + \beta_4 PriceCap_{wd} \times K_{fwd} + \beta_5 Z_{fwd} + \beta_6 X_{wd} + \gamma_d + \epsilon_{fwd}$$

where $Price_{fwd}$ denotes outcome variables for recipient f in ward w and district d, K_{fwd} denotes a characteristic of the recipient and Z_{fwd} denotes control variables at the recipient level and X_{wd} controls at the agent level. Control variables at the farmer level include measures of travel distance. Control variables at the agent level contain stratification variables and indicators for the fixed fee treatment. Farmer level regressions are weighted to obtain equal weights for each service delivery agent. Standard errors are again clustered at the ward level.

Table 1.5 presents the results of this exercise. In column 1 I estimate the effect of the intervention on transaction prices and allow the treatment effect to vary depending on whether the recipients' livelihoods depend on agriculture. Approximately 80% of households in the control group match this definition. The results show that while the point estimate for the treatment effect on prices is negative for all farmers, it is small in absolute terms and statistically insignificant for non-agricultural households but approximately 50% larger and statistically significant at the 10% level for households whose main livelihood is derived from agriculture.

Column 2 shows that agents price differentially based on not only farmers' livelihood characteristics, but also on the number of chickens vaccinated per farmer. In particular, agents offer lower per-unit prices for larger flocks. Table 1.5, column 3 shows that farmers who own fewer than 11 chickens on average pay 14 Tanzanian Shillings (or 18%) more per vaccination than farmers with larger flocks.[42] Implementing a price-cap not only reduces average prices for all recipients by 12% but also eliminates this quantity discount. Taken together, the price-cap intervention appears to particularly benefit agriculturally dependent households and smallholders, who are likely to be poorer and hence more susceptible to shocks to livestock holdings.

1.4.3 Impact on Service Allocation

The previous section has shown that price-caps affect prices. While this directly affects the distribution of surplus, price-caps' primary welfare implications operate through their effect on the allocation of services. This section highlights two channels through which reducing discretion over prices affects service allocation. First, price-caps increase the affordability of services. This increases the likelihood of agents extending services to

[42] When asked about the motivation for this pricing strategy, agents mentioned that quantity discounts were needed to convince larger flock holders to bear the higher total cost of the service.

previously unserved recipients and, in light of the evidence on differential pricing presented in the previous section, to agricultural households and smallholders. Second, price-caps reduce agents' expected profits from serving a given village, which in turn reduces their incentives to incur the travel costs associated with travelling to remote villages.

Impact of Price-Cap on Composition of Beneficiary Pool

I first investigate how price-caps affect the composition of the beneficiary pool. As the previous section has shown, price-caps reduce prices on average and do so in particular for agricultural households and smallholders. Price-caps therefore not only redistribute surplus but also increases the affordability of services. The intervention therefore should increase the share of previously unserved beneficiaries in the recipient pool. Consistent with this, columns 1 and 2 in table 1.6 show that the price-cap indeed increases the proportion of previously unserved recipients in the beneficiary pool by 12% to 15%.

I then investigate how this price-cap affects the share of farmers in the recipient pool who benefited particularly from the price-cap, namely those whose main source of income is derived from agricultural production and those with comparatively small chicken flocks. As discussed in section 1.4.2 and shown in table 1.5, the price-cap reduces transaction prices more for such recipients. Columns 3 and 4 in table 1.6 suggest that this price-adjustment indeed leads to a positive demand effect, as households whose main source of income stems from agriculture are 6% more likely to be served in response to the price-cap treatment, conditional on the size of the recipient pool. While not statistically significant, the point estimate in columns 5 and 6 in table 1.6 suggest that smallholders are 9% more likely to be served in the price-cap group.

Taken together, the results presented in this section suggest that in the absence of price-caps agents use their discretion to extract rents from service recipients. Capping prices redistributes surplus to recipients, crowds in previously unserved farmers and makes services more affordable for recipients in need.

Price-Caps reduce Geographic Coverage

While the previous sections have shown that price-caps reduce prices and increase the proportion of new recipients and recipients with a high need for the service in the beneficiary pool, it is unclear how capping prices affects the aggregate coverage of services. In particular, it is possible that implementing a price-cap reduces agents' incentives to extend services to markets that are more costly to serve. For the setting studied in this paper, capping prices might reduce their incentives to extend services to remote villages. Price-caps can therefore reduce welfare by discretely eliminating aggregate surplus obtained from serving a given market.

This section provides evidence in support of this mechanism. To do this, I merge information on villages and travel distances with the farmer survey and the receipt data. In particular, a list of all villages in their area of responsibility and the approximate travel time by foot from their headquarter to each village was collected from agents during the baseline survey. I use the data on travel times to rank the villages by their distance to the agent's headquarter. I then match the village information provided during the farmer survey and on the receipts to the village list collected during the baseline survey, to obtain information on whether agents visited a given village.[43]

Table 1.7 shows how capping prices affects which villages agents visit. Column 1 shows that while agents in the status quo visit approximately 37% of villages that they are assigned to, the price-cap reduces this proportion by 4.5 percentage points. Columns 2 and 3 confirm that this reduction is driven by a reduced likelihood of agents visiting remote villages: they are 25% less likely to visit villages whose distance from their headquarter is above median, and 26% less likely to visit the furthest village in their area of responsibility. Taken together, this suggests that price-cap regulation reduces agents' incentives to serve more remote markets.[44]

Impact on Total Number Served

Given the countervailing forces discussed previously, it is unclear whether the price-cap will increase or decrease the total number of farmers served. Figure 1.4 investigates this question and shows little evidence of the price-cap affecting the total number of farmers served. The figure separately plots the daily number of farmers served for the treatment and the control group using receipt data. This shows that the difference between the daily number of farmers served is statistically indistinguishable from zero for 18 out of the 21 days of the campaign. Column 1 in table 1.8 confirms this impression:

[43] Approximately 11% of receipts were unmatchable to villages. This can either be because the information provided in the surveys or on the receipts was incorrect or because recipients live outside of the formal villages. Reassuringly, the likelihood of an agent visiting an "unmatched" recipient is uncorrelated with the treatment.

[44] In the appendix to this paper I present additional evidence on the trade-offs generated by price-caps by focusing on availability of I-2 substitutes.

agents in the price-cap treatment serve an average of 3.6 fewer farmers than agents in the control group. This difference is statistically insignificant.

I conduct two robustness checks to verify this result. First, a possible concern is that the result is a composite effect between a participation response on the extensive margin and an effort response on the intensive margin. To address this, I restrict the sample to agents who verifiably served at least one farmer, therefore ruling out responses on the extensive margin. Column 2 in table 1.8 confirms that the result is robust to this restriction: ruling out extensive margin responses, agents in the treatment group serve on average 5 fewer farmers than agents in the control group, which remains statistically insignificant. Second, I consider the impact of the treatment on the total number of vaccines applied. Column 3 in table 1.8 shows that while the point estimate for the treatment effect is negative for the number of farmers served, it flips sign for the total number of chickens vaccinated while remaining insignificant. Taken together, those results suggest that while the price-cap affected the *types* of farmers served, I do not detect an effect on the total number of recipients served.

1.5 Sufficient Statistics Model for Welfare Analysis

The previous section has shown that price-caps have three key effects on service provision. First, they *reduce* average prices and the within-agent variation of prices. Second, they *increase* the proportion of new recipients and recipients who were paying higher prices in the status quo in the recipient pool. Third, they *reduce* the likelihood of agents visiting remote villages. To conceptualize these effects, and to understand their effect on welfare, this section develops a model that is consistent with the empirical results and allows for the estimation of welfare effects through a sufficient statistics formula.

I model I-2 provision as a slot assignment problem in which slots are assigned through two allocation mechanisms. First, suppliers choose which villages to visit. While agents are responsible for a given service area, the model takes into account that visiting a given village requires agents to pay a travel cost. Agents' willingness to pay this travel cost then determines service allocation *between* villages in a given service area. Second, agents choose prices that determine which recipients are served within a village, conditional on the agent visiting their village.

1.5.1 Model Setup

This model considers a situation in which a monopolistic agent is supplying services to a population of potential customers, the size of which is normalized to 1. Customers are defined by their valuation of the service, which is denoted by v_i. I assume that v_i is a continuous random variable drawn from a distribution $F(v)$. I further assume that recipients' elasticity of demand is given by ϵ_D^i which can either be high or low: $\epsilon_D^i \in \{\epsilon_D^L; \epsilon_D^H\}$. Suppose that a fraction μ of recipients has ϵ_D^L, whereas everyone else has ϵ_D^H. Customers do not only differ with regard to their valuation, but also in their location. Specifically, I assume that recipients live in a continuum of villages with differing travel distances to their agent's headquarter. I further assume that the distribution of valuations is the same for every village.

Agents in this model face two sequential choices. First, they decide which markets to serve by determining the allocation of services *between* villages. I assume that travelling to village j requires paying a cost of c_j which agents can choose to either pay or not. I assume that c_j is drawn from a distribution with c.d.f. $M(c)$ defined over $[0; c_{max}]$ with $c_{max} < \infty$. Second, agents decide on prices, which determine service allocation *within* a village, conditional on the village being served.[45] For simplicity, I model this by assuming that agents offer a take-it-or-leave-it price based on observable recipient characteristics. While I allow for price-discrimination, I assume that agents do not observe v_i and instead only learn about ϵ_D^i.[46] If customers accept the agent's offer, the agent receives the agreed sum and delivers the service at a constant cost τ. If the recipient rejects the price offer, no transaction takes place, but the agent still has to pay the travel cost to the village.

1.5.2 Theoretical Effect of Price-Cap on Observables

To aid the interpretation of the empirical results, this section shows that the implications of the model are consistent with the three main empirical results. In particular, I show how agents' pricing decisions as well as the allocation of services in the model are affected by the price-cap intervention. Section 1.5.3

[45] This implicitly assumes that travel costs within a village are 0. It is straightforward to relax this making travel costs farmer specific.

[46] Anecdotal evidence is consistent with this assumption. Agents report that negotiations with farmers regularly break down and that they are unable to charge similar farmers different prices. They do, however, mention that it is possible for them to give discounts based on observable characteristics, such as household wealth and on the number of chickens held by the household. While this description is in line with my model, it is inconsistent with alternative bargaining models, such as uniform pricing, first-degree price-discrimination and Nash bargaining.

then uses the empirical estimates to evaluate the marginal welfare effect of the price-cap.

Effect of Price-Cap on Prices

To begin the analysis of the model, note first that travel costs to villages are sunk at the time of price setting. This allows me to investigate the agent's two decisions separately. To understand the effect on prices, notice that agents face a monopoly trade-off: raising prices increases profit from a transaction but reduces the likelihood that farmers will accept the price. Formally, visiting a recipient of type i yields the following expected profit:

$$\pi_i = \left[1 - F\left(p_i | \epsilon_D^i\right)\right] [p_i - \tau]$$

Pointwise maximization of the objective function yields the standard monopoly pricing solution:

$$p_{Dec}\left(\epsilon_D^i\right) = \tau \left(\frac{\epsilon_D^i}{1 + \epsilon_D^i}\right) \tag{1.1}$$

Agents set prices based on a mark-up over marginal costs, with a low-elasticity of demand leading to high mark-ups. This model therefore gives rise to a bimodal price distribution in which the price-cap is more likely to bind for recipients with a low elasticity of demand. To explore the effect of the cap, assume that, consistent with the empirical design, the price-cap only binds for recipients with a low elasticity of demand. Prices under a price-cap $\bar{p}$ are hence given by:

$$p^{Reg}\left(\epsilon_D^L\right) = \bar{p} \text{ and } p^{Reg}\left(\epsilon_D^H\right) = \tau \left(\frac{\epsilon_D^H}{1 + \epsilon_D^H}\right)$$

This directly implies that the price-cap mechanically reduces the agent's ability to price-discriminate between recipients with a high and low elasticity of demand. Price-caps hence reduce the standard deviation and average of prices in equilibrium and particularly benefit recipients with a low demand elasticity.

Effect of Price-Cap on Beneficiary Pool

In order to understand how price-caps affect service allocation, I now investigate how price-caps affect the distribution of the different elasticity types in the recipient pool. This requires me to investigate recipients' acceptance decisions, conditional on their village being served. After receiving a price offer, recipients decide whether to accept or reject it. In equilibrium, recipients accept every offer that does not exceed their willingness to pay. The distribution of types in the recipient pool is then:

$$\mu D\left(p\left(\epsilon_D^L\right)|\epsilon_D^L\right)+(1-\mu)\,D\left(p\left(\epsilon_D^H\right)|\epsilon_D^H\right) \tag{1.2}$$

where $D\left(p\left(\epsilon_D^i\right)|\epsilon_D^i\right)$ denotes the demand for customers of type i.

Note also that here the demand curve is downward sloping. If the price-cap binds only for recipients with a low elasticity of demand, it is hence straightforward to see that price-cap regulation will increase their demand. This increases the likelihood of serving previously unserved recipients and shifts the distribution of types in the direction of recipients with a low elasticity of demand.

Effect of Price-Cap on Village Choices

To understand the agent's coverage decision, I investigate the effect of the price-cap on the choice of villages conditional on a price vector $\mathbf{p}$. In the status quo, agents decide to visit village j if the expected profit exceeds the associated costs:

$$\mu\pi_{\epsilon_D^L}\left(\mathbf{p}\right)+(1-\mu)\,\pi_{\epsilon_D^H}\left(\mathbf{p}\right)\geq c_j$$

where $\pi_{\epsilon_D^i}\left(\mathbf{p}\right)$ denotes the expected profit obtained from recipients with elasticity equal to ϵ_D^i. This defines a cut-off value for c_j, denoted by $c^*\left(\mathbf{p}\right)$, which is the highest cost village visited by the agent. The proportion of villages visited in the status quo is hence given by:

$$\sigma = M\left(c^*\left(\mathbf{p}\right)\right) \tag{1.3}$$

Regarding allocation between villages, agents are hence more likely to serve a larger proportion of villages when (i) the expected profit is higher and (ii) the proportion of recipients with a high elasticity of demand in each village is lower.

To understand how price-caps affect the allocation of services between villages, denote by π_j^{Dec} and π_j^{Reg} the agent's expected profit from visiting village j in the status quo and under the price-cap, respectively. As agents maximize their profits with respect to prices in the status quo, it follows that:

$$\pi_j^{Dec}\geq\pi_j^{Reg}$$

Profits play a dual role in this model. First, profits allow agents to extract surplus. Second, however, profits also compensate agents for the travel costs to remote villages incurred. This is necessary because travel costs are sunk when price offers are made. Pricing decisions therefore do not ensure that the agent breaks even in remote areas. Differentiating equation 1.3 with respect to prices shows how a price-cap affects the targeting of remote areas:

$$\frac{\partial \sigma}{\partial p_{\epsilon_D^L}} = M'\left(c^*\right)\mu \frac{\partial \pi_{\epsilon_D^L}}{\partial p_{\epsilon_D^L}} < 0$$

This implies that the price-cap reduces the proportion of villages visited. This is because reducing discretion reduces the amount of surplus agents can extract from remote villages, which lowers the highest travel cost they can pay and still break even. Taken together, the model shows that price-cap regulation generates a tension between preventing surplus extraction to make services affordable for recipients with a low elasticity of demand and providing incentives to serve remote villages.

1.5.3 Effect of Price-Cap on Welfare

The previous section has shown that price-caps crowd in previously unserved recipients and redistribute surplus at the expense of remote farmers. Given those countervailing forces, it is not clear whether it is welfare improving for the government to cap prices, and, if so, to what extent. To address this, this section first investigates the government's policy decision to motivate the choice of price-caps as a regulatory instrument. It then solves the government's objective function to derive an expression for the marginal effect of capping prices as a function of empirically estimatable sufficient statistics.

Regulatory Policy

The government's objective is to maximize social welfare. Denoting by g_{c_H}, g_{c_L} and g_a the government's welfare weight on high elasticity customers, low elasticity customers and the agent, social welfare for a generic price vector is given by:

$$
\begin{aligned}
&SWF(\mathbf{p}) = \\
&\qquad M(c^*(\mathbf{p}))\left[g_{c_L}\mu\int_{p(\epsilon_D^L)}^{\infty} v_i - p\left(\epsilon_D^L\right)dF\left(v|\epsilon_D^L\right)\right] \\
&\qquad + M(c^*(\mathbf{p}))\left[g_{c_H}(1-\mu)\int_{p(\epsilon_D^H)}^{\infty} v_i - p\left(\epsilon_D^H\right)dF\left(v|\epsilon_D^H\right)\right] \\
&+ g_a M(c^*(\mathbf{p}))\left[\mu\left(p\left(\epsilon_D^L\right)-\tau\right)D\left(\mathbf{p}|\epsilon_D^L\right) + (1-\mu)\left(p\left(\epsilon_D^H\right)-\tau\right)D\left(\mathbf{p}|\epsilon_D^H\right)\right] \\
&\qquad - g_a\int_0^{c^*(\mathbf{p})} c_j dM(c)
\end{aligned}
\tag{1.4}
$$

Social welfare therefore consists of a weighted sum between consumer surplus and the agent's profit. To build intuition, it is instructive to define the first best regulatory policy for when the government cares equally about producers and consumers ($g_{c_H} = g_{c_L} = g_a = 1$). Suppose that in the first best the government can make costless transfers to the agent and enforce which villages the agent serves. Denoting transfers by t, it is straightforward to see that the optimal regulatory contract implements the following policies:

$$
\begin{aligned}
\bar{p} &= \tau \\
c^* &= c^{max} \\
t &= \int_0^{c^*(\mathbf{p})} c_j dM(c)
\end{aligned}
$$

As denoted above, the optimal regulatory policy in the first best scenario sets prices equal to marginal costs, mandates the agent to serve all villages and uses transfers to reimburse the agent's travel costs.

In reality, fiscal constraints prevent governments from paying transfers to the agents.[47] In addition, governments are constrained by moral hazard, which limits their ability to mandate which villages the agent visits. Under those circumstances, Laffont (2005) notes that the optimal regulatory policy includes price-caps that rule out agents' most opportunistic choices.

[47] In the absence of costless transfers the government will choose a transfer level that equates the marginal social welfare gain of raising transfers to the marginal cost of public funds.

Sufficient Statistics Formula for Uniform Pricing

The central objective of this paper is to understand the effect of the price-cap on welfare. To achieve this, I take a sufficient statistics approach to determine the welfare effect of marginally lowering prices below the full discretion level. This approach has three key advantages. First, by expressing welfare effects as a function of reduced form parameters, it allows me to use the empirical results to evaluate the optimal price-cap for the service I study. Second, my approach requires me to make no structural assumptions on agents' and recipients' behavior. Third, it uses estimates from the non-marginal experimental intervention to investigate the effect of *marginally* lowering prices below the full discretion level. This makes the welfare results less dependent on the chosen value of the price-cap.

For tractability, this section will derive the sufficient statistics formula for the case of uniform pricing and then postulate the appropriate extension to third-degree price-discrimination discussed previously. I present the derivation for this extension in the appendix. Governments choose the price-cap to maximize social welfare. Analogous to equation 1.4, social welfare for the uniform price case is given by a weighted sum between consumer and producer surplus:

$$SWF\left(p\right) = M\left(c^{*}\left(p\right)\right)\left[\int_{p}^{\infty} v_i - pdF\left(v_i\right)\right] + gM\left(c^{*}\left(\bar{p}\right)\right)\left(p-\tau\right)D\left(p\right) - g_a \int_{0}^{c^{*}(p)} c_i dM\left(c_i\right) \quad (1.5)$$

where g denotes the welfare weight on agents relative to recipients. This formulation allows me to consider the welfare effect for scenarios in which governments value only consumer surplus ($g = 0$) and in which governments take into account aggregate surplus ($g = 1$).

Starting from unregulated prices, the marginal welfare effect of lowering prices has three first order effects on welfare. First, on the extensive margin, marginally lowering prices reduces the fraction of villages served, which leads to a discrete loss in consumer surplus. Second, on the intensive margin, lowering prices reduces the monopoly distortions within a village, as it closes the gap between prices and marginal costs τ. Third, reducing prices redistributes surplus from agents to consumers, which has a direct effect on social welfare if the government values surplus accruing to recipients more than surplus accruing to agents.[48] Taking derivatives of equation 1.5 and using the definition of $c^*(p)$, the marginal effect on welfare is given by:

$$\begin{aligned} \frac{\partial SWF(p)}{\partial p} &= \frac{\partial M(c^*(p))}{\partial p}\int_p^\infty v_i - pdF(v_i) \\ &- M(c^*(p))\,g(p-\tau)\frac{\partial D(p)}{\partial p} \\ &- M(c^*(p))\,D(p)(1-g) \end{aligned}$$

The first and second term capture the extensive and intensive margin effects, respectively. The third term captures the redistributive effect. To derive a formula based on sufficient statistics, it is useful to define two parameters. First, I denote by θ the extensive margin elasticity of village visits with respect to the price-cap. Formally:

$$\theta = \frac{\partial M(c^*(p))}{\partial p}\frac{p}{M(c^*(p))} \tag{1.6}$$

Second, I denote by ε_D the price elasticity of demand:

$$\varepsilon_D = \frac{\partial D(p)}{\partial p}\frac{p}{D(p)} \tag{1.7}$$

Finally, notice that consumer surplus at price p is given by:

$$CS(p) = \int_p^\infty v_i - pdF(v_i) \tag{1.8}$$

The above definitions, together with the fact that total number of farmers served is given by $N(p) = M(c^*(p))\,D(p)$, yields the following proposition:

Proposition 1. *The welfare effect of marginally reducing prices below the uniform monopoly pricing level can be estimated using θ and ϵ_D as sufficient statistics:*

$$\frac{\partial SWF(p)}{\partial p} = \theta N(p)\frac{CS(p)}{pD(p)} + \varepsilon_D g N(p)\frac{p-\tau}{p} - (1-g)\,N(p) \tag{1.9}$$

To understand the intuition behind this formula, consider two scenarios. First, suppose there are no distortions associated with the exploitation of

[48] In addition to those effects, there are also two second-order effects. First, reducing prices increases demand, which has a second order effect on welfare because buyers on the margin were indifferent between purchasing and not-purchasing in the first place. Second, reducing prices reduces providers' profit from the villages that are no longer visited. This effect is second order because the expected profit from the marginal village was 0 in expectation during the status quo.

market power. In this case $p = \tau$ and the intensive margin benefit of capping prices disappears. Second, suppose the government puts equal weight on surplus accruing to agents and customers. In this case $g = 1$ and the last term, which captures the redistributive effect of the price-cap, disappears.

It is straightforward to extend this analysis to price-discrimination when there are two types of buyers in the market: One for whom the elasticity of demand is high and one for whom it is low. I denote by ε_D^L and ε_D^H the demand elasticities of the low and high elasticity customers, respectively. In addition, I denote by μ the share of low elasticity customers in the market. The following proposition then describes the sufficient statistics formula that allows for the estimation of welfare effects:

Proposition 2. *The welfare effect of marginally capping prices for consumers with a low-elasticity of demand under third-degree price-discrimination is given by:*

$$
\begin{aligned}
\frac{\partial SWF(\mathbf{p})}{\partial p\left(\varepsilon_D^L\right)} = \quad & \theta N(\mathbf{p}) \frac{g_{c_L} \mu CS\left(p|\varepsilon_D^L\right) + g_{c_H}(1-\mu) CS\left(p|\varepsilon_D^H\right)}{p\left(\varepsilon_D^L\right)\left(\mu D^L(\mathbf{p}) + (1-\mu) D^H(\mathbf{p})\right)} \\
+ \quad & \varepsilon_D^L g_a \mu N^L(\mathbf{p}) \frac{p\left(\varepsilon_D^L\right) - \tau}{p\left(\varepsilon_D^L\right)} \\
- \quad & \left(g_{c_1} - g_A\right) N^L(\mathbf{p}) \mu
\end{aligned}
\tag{1.10}
$$

Here $N(\mathbf{p})$ denotes the total number of services provided and $N^L(\mathbf{p})$ the number of services provided for recipients with a low elasticity of demand. Further, g_{c_L}, g_{c_H} and g_a denote the government's welfare weights on low elasticity customers, high elasticity customers and the agent, respectively.

1.6 Welfare Analysis

The previous discussion has shown that estimating the marginal welfare effect of price-cap regulation requires knowledge of two parameters: The extensive margin elasticity of village visits with respect to the price-cap and the intensive margin elasticity of demand. This section first discusses the estimation of the elasticity of demand before using those estimates to evaluate the welfare effect of capping prices.

1.6.1 Estimation

While the price-cap treatment allows me to estimate the extensive margin elasticity of village visits with respect to the price-cap,[49] I require additional

[49] Alternatively, one could allow travel costs to vary at the farmer level. In this case, the responses would not be calculatable directly from the price-cap treatment, but could be

variation in prices to estimate the price elasticity of demand. This is because the variation induced by the price-cap generates responses both on the extensive and the intensive margin. Estimating demand elasticities traditionally would require experimental variation in marginal costs. In my case, this would imply generating variation in vaccination costs. Introducing such variation in the context I study is challenging as it generates incentives for agents to report fewer vaccinations than were actually conducted.[50] To overcome this challenge, I leverage the findings from a another paper. In chapter 2 I present evidence from a lab-in-the-field experiment which shows that agents choose prices based on a simplified heuristic that is affected by average costs. Furthermore, chapter 2 shows that this behavior makes it possible to induce price variation through a treatment that varies fixed fees, which does not generate any incentives to incorrectly report vaccination. This allows me to consistently estimate the elasticity of demand.

Fixed Fee Treatment and Elasticity Estimates

While the main results and design are discussed in chapter 2, I briefly describe the induction of additional experimental variation in costs here for convenience.

As part of the fixed fee variation, I assign agents in 273 wards to a cross-cutting treatment that requires them to contribute a fixed fee of 25,000 Tanzanian Shillings (approximately \$11.40) to cover a portion of the vaccine cost if they agree to participate in the program. To avoid concerns about liquidity constraints, the fee was collected after the completion of the vaccination campaign. Agents in the remaining 277 wards received the vaccine for free. The randomization was designed so that the probability that each ward received a given treatment was always held constant, regardless of what stratum the village was in and whether the price-cap treatment had been cross-randomized. The probability of being in the price-cap group is therefore orthogonal to having to pay a fixed fee.

The fixed fee treatment successfully induced variation in prices. Table 1.9 presents the treatment effect estimates from the intervention. Column 1 shows that imposing a fixed fee raises average prices charged by around

backed out using the elasticity of total recipients served with respect to the price-cap and the price elasticity of demand. My results are qualitatively robust to this alternative.

[50] As vaccines expire after 3 weeks without cooling, there is no formal system in place that requires agents to return unused vaccines to the headquarter.

11%.[51] Column 2 in table 1.9 shows that this increase in prices reduced the average number of farmers served by around 12%. As column 3 shows, the fixed fee treatment therefore only reduced collected revenue by 3%, which is statistically insignificant. I then use this variation to estimate the price elasticity of demand. To do this, I run agent level instrumental variable regressions in which I regress the log of the total number of farmers served on log prices. I instrument for prices using an indicator for an agent's assignment to the fixed fee treatment. Column 4 in table 1.9 presents the results from this exercise. The estimate of the price elasticity of demand is -1.223, making ND vaccinations a fairly elastic good. This estimate is consistent with the literature, which has estimated demand elasticities around -1.5 for Newcastle Disease Control products (Fisher, 2014).

Welfare Wedges

In addition to knowledge of the two sufficient statistics, θ and ε_D, which were estimated through the field experiment, the welfare analysis requires estimates of the welfare wedges $\frac{CS(p)}{pD(p)}$, which depends on consumer surplus, and $\frac{p-\tau}{p}$, which depends on marginal application costs.

To estimate consumer surplus, I assume that demand follows a constant elasticity demand function. When demand is $D(p) = D_0 p^{\varepsilon_D}$, consumer surplus is then:

$$CS(p) = \int_p^{\infty} D_0 x^{\varepsilon_D} dx$$

which can be calculated directly.

Regarding estimates of $\frac{p-\tau}{p}$, I take two approaches. First, I obtain an estimate of τ from the monopolists' pricing problem. In particular, when profit maximizing monopolists set uniform prices, they maximize $\pi = (p-\tau) D(p)$. The solution to this problem yields the first way to estimate mark-ups:

$$\frac{p-\tau}{p} = \frac{-1}{\varepsilon_D}$$

The evidence obtained in chapter 2 suggests, however, that agents do not set prices optimally. In light of this, I also bound the estimates by setting $\tau = 0$.

One additional complication arises, as the price-cap reduces the prices *per vaccination* and recipients typically purchase more than one vaccination.

[51] In chapter 2 I show that this is not a selection effect, as the fixed fee treatment did not affect participation in the vaccination campaign.

Assuming, for simplicity, a constant number of vaccinations per recipient, $N(p)$ then refers to the total number of vaccinations applied.

For the case of price-discrimination, calibration of the sufficient statistics formula requires three further parameters. First, I require separate estimates of demand elasticities for recipients with a high and a low elasticity of demand. To estimate those, I define smallholders, farmers in areas without private providers and households whose livelihood depends on agriculture as low-elasticity households. I then obtain demand elasticities by separately estimating demand functions for the two populations, using the fixed fee treatment as an instrumental variable. The results of this estimation are presented in column 5 and 6 in table 1.9. The estimated elasticities of demand are -0.41 for the low elasticity types and -3.83 for the high elasticity types. Second, I require knowledge of μ, the share of low elasticity households. As I cannot obtain this directly from the data, I bound my estimates by setting μ to either 0, 0.5, or 1.

1.6.2 Results

Table 1.10 presents the results from the calibration of the sufficient statistics formula. Panel A shows the calibrated marginal welfare effects, whereas panel B shows the welfare effects for a counter factual scenario in which extensive margin effects are absent. Three results are worth noting. First, the calibrated marginal welfare effects are negative across the board. Panel A shows that marginally reducing prices below the full discretion level leads to a welfare loss to the magnitude of between 3% and 11% of total sales revenue per agent. This suggests that the adverse effects of price-cap regulation on the extensive margin are so strong that any deviation from full discretion leads to a welfare loss. This directly implies that, for the setting I study, any form of price-cap regulation will cause welfare to decrease. Second, the implied welfare losses are substantially larger for the price-discrimination case as compared to the no price-discrimination case. This is natural, because the benefits of price-cap regulation under price-discrimination only accrue to a subset of a given village market compared to uniform pricing, whereas the adverse extensive margin effects affect the whole village market. Finally, panel B shows that, in the absence of extensive margin responses, price-cap regulation can lead to moderate welfare increases ranging from 0.4% to 2.5% of total sales revenue per agent on the margin. Taken together, the results therefore conclusively show that price-cap regulation generates a tension between intensive margin demand effects that increase welfare, and extensive margin effects that reduce welfare. For ND vaccinations in Tanzania, extensive margin effects are sufficiently strong to lead to a net-welfare loss, making price-cap regulation counter-productive.

1.7 Caveats

While the experiment and data collection procedures were designed to estimate the channel of interest, some caveats to the analysis exist that make alternative explanations possible. First, while all available agents in the enumeration region were assigned the vaccination task, some of them failed to attend the necessary training. There were several reasons for this: some were on annual leave, sick, on professional training, or were assigned other long-term duties. This attendance gap poses a challenge to experimental validity if attendance rates differ between treatment and control groups. In the appendix I alleviate this concern by showing that on average 83% of agents attended training. This result does not differ significantly between treatment and control groups.

Similarly, while all agents who attended training also collected vaccines, some of them failed to serve any farmers. A concern is that this failure to participate is a response to the treatment. In the appendix I again alleviate this concern by showing that the treatment did not affect the participation decisions: among the 832 agents who attended training, 82% submitted receipts in the control group, compared to 84% in the price-cap group and 82% in the fixed fee group. This difference is not statistically significant. I therefore conclude that the treatments did not induce any response on the participation margin.

Second, a concern is that the price-cap generated incentives for selective reporting. In particular, agents might be tempted to report only transactions that comply with the price-cap, while not reporting transactions whose value exceeds the price-cap. The experimental design addresses this concern through the bonus payment, which assures that it is always incentive compatible to report transactions, as only formal reporting generates eligibility for the bonus of 60 Tanzanian Shillings per vaccination. Consistent with this assertion, the farmer survey detected limited non-compliance with the price-cap: for only 4.5% of participants in the treatment group did farmers report paying prices that exceed 80 Tanzanian Shillings per vaccination.[52] To

[52] While this figure is small, it is still key to notice that, even under lower compliance levels, rules can still improve outcomes by assuring that those with a high-cost of non-compliance comply (see also Banerjee, Hanna, and Mullainathan (2013) for a discussion of this).

further validate this point, I investigate whether vaccine loss differs between the experimental groups. While the logistics of the vaccine distribution and storage render it infeasible to track every dose, I can proxy for leakage using the ratio between confirmed number of vaccinations and the initially distributed amount of vaccine doses.[53] Evaluating this proxy suggests that leakage rates were generally low, as the average proxy value is 96% in the control group. More importantly, this figure does not differ systematically between treatment and control. It is therefore unlikely that systematic leakage and misreporting is influential enough to drive my results.

Third, it is possible that the treatment induced agents to report receipts for which no service was provided in order to receive the bonus payment. To investigate this possibility, the appendix shows that on average 69% of transactions reported by agents could be verified. This figure does not differ significantly between treatment arms.[54]

Fourth, although the data verification procedures are reassuring in interpreting the observed price effects as a real transfer of surplus, one potential concern is that these impacts might be due to undetected misreporting. A particular concern is collusion between the agent and the farmer in generating inaccurate receipts. While it is not possible to conclusively rule out this possibility, the experimental design requires a high level of trust to make collusion profitable. To see this, notice that if agents decide to misreport, they face a lottery which pays the unconstrained revenue plus the bonus payment if they remain undetected and only the unconstrained revenue if the fraud is detected. If agents choose to report correctly, they receive the constrained revenue plus the bonus payment in every state of the world. Assuming risk-neutrality to obtain an upper bound, the largest possible detection probability agents are willing to accept is given by the expected increase in revenue from misreporting divided by the bonus payment. The experimental data suggests that non-compliance on average yields an additional revenue of around 11,000 Tanzanian Shillings (Table 1.8, column 5) while detection would lead to the loss of approximately 70,000 Tanzanian Shillings in bonus payments. Agents therefore decide to misreport if their detection probability is lower than 15%

[53]Notice that this measure allows for fractions that exceed 1, as agents might have collected additional vaccines from the storage locations at later stages of the vaccination campaign.

[54]A similar concern is that misreporting is distributed unevenly across agents, implying a heterogeneity between honest reporters and employees who misreport their performance. The appendix addresses this possibility by investigating how the fraction of verifiable receipts varies across individuals and showing that the inability to verify receipts is evenly distributed between respondents.

for all farmers. With 50 farmers served on average, this implies that collusion is profitable if the probability that every farmer honors the agreement is above 99%. Taken together, the experimental design therefore generates very small incentives for non-compliance that are unlikely to justify large-scale misreporting.

Fifth, while agents complied with the price-cap for the vaccination service, they might have increased prices on other services in response to the treatment. While data on prices for such transactions is not available, two factors make it unlikely that this mechanism is driving my results. First, fewer than 1% of respondents in the farmer survey reported paying a transport and consultancy fee in addition to the vaccination charge, which suggests that transactions on top of the user fees are rare. Most importantly, this figure does not differ systematically between treatment and control group. Second, I-2 vaccinations are the only large-scale profitable service that agents provide for poultry farmers. Instead, their main profit raising activities accrue from services for large ruminants, especially cattle. Any cross-price effects would therefore have to raise prices for cattle-related services. On the one hand, this implies that in the presence of cross-price effects agents in the price-cap treatment should be more likely to serve poultry farmers that also hold cattle, as this allows them to mitigate the effect of the price-cap. Data from this study rejects this hypothesis. While 29% of service recipients report owning at least one cow, this does not differ between treatment and control group. Having ruled out selection effects, I also investigate whether excluding cattle owners, and therefore potential cross-price effects, qualitatively changes my main results. This is not the case. The treatment still reduced average prices, and the aforementioned composition effects in the recipient pool remain even when excluding cattle owners, although the reduced sample size has made the estimates less precise. Taken together, this evidence makes it unlikely that cross-price effects substantially challenge the presented interpretation of the results.

1.8 Conclusion

In this paper I evaluate whether capping prices for public services increases or reduces welfare. I combine administrative government data with survey data to evaluate a field experiment which investigates how capping prices for public livestock vaccinations in Tanzania affects service delivery. The evaluation yields three main results. First, the price-cap reduces average prices charged by agents. Second, the price-cap affects the composition of the beneficiary pool by increasing the proportion of previously unserved recipients. Third, price-caps reduce the proportion of remote villages served.

I then employ a model of monopoly regulation to derive a sufficient statistics formula which allows me to evaluate and decompose the welfare effect of the intervention. This analysis yields two findings. The first finding is that the decision whether to introduce price-caps has to address a trade-off between demand and supply considerations. On the one hand, price-caps can enhance the provision of services by making services more affordable. This redistributes surplus to recipients and crowds in new recipients in markets that are comparatively cheaper to serve. On the other hand, price-caps can also harm social welfare, as such interventions reduce agents' incentive to extend service to markets that are costlier to serve. The second finding is that for public livestock vaccine service provision in Tanzania, the introduction of any form of price-cap regulation reduces welfare compared to the status quo.

A central contribution of the paper is to highlight the importance of incentive effects induced by capping prices. One implication of this finding is that public regulation which mandates that essential services should be provided for free or at very low prices can be suboptimal when governments cannot control or incentivize the agent to maintain a sufficient coverage of services. More broadly, this paper shows that it can be optimal to allocate discretion over prices to agents when contracts are incomplete, even though this allows agents to extract surplus. In addition to being informative about pricing policies for public services, this findings therefore also has broader implications for organizational design.

One limitation of this paper is that the experimental setting does not allow me to study long-term effects of the price-cap, such as possible dynamic effects induced by changes to the competitive structure and demand responses. Another limitation is that the experimental design does not allow for the comparison between alternative regulatory contracts. I hope to address those shortcomings in future work.

Tables

Table 1.1: Treatment Groups

Price-Cap	
No	**Yes**
274 wards (410 agents)	276 wards (422 agents)

Table 1.2: Summary Statistics and Balance Table

	Price-Cap Experiment		
Panel A: Agent Level	Control	Treatment	P-Value of Difference
Tenure	12.466	12.198	0.743
	(0.602)	(0.556)	
Ward Level agent	0.663	0.642	0.531
	(0.127)	(0.144)	
Number of Villages	4.022	3.974	0.804
	0.040	0.019	
Animal Health Specialist	0.434	0.476	0.241
	(0.025)	(0.025)	
Main Income Earner	0.866	0.864	0.938
	(0.017)	(0.018)	
Uses Motorcycle	0.446	0.400	0.204
	(0.024)	(0.027)	
Secondary Income Source	0.659	0.604	0.129
	(0.026)	(0.025)	
Acting Village Leader	0.144	0.152	0.753
	(0.017)	(0.018)	
Raises Livestock	0.798	0.796	0.934
	(0.021)	(0.020)	
Panel B: Work Station			
Rural	0.844	0.820	0.443
	(0.021)	(0.023)	
Average Travel Time	80.015	90.842	0.366
	(4.757)	(10.990)	
Private Veterinarian	0.156	0.197	0.176
	(0.020)	(0.022)	
Private Drug Seller	0.029	0.033	0.763
	(0.008)	(0.010)	
Poultry Area	0.076	0.062	0.433
	(0.013)	(0.012)	
Observations	410	422	

Notes: The sample includes all agents who agreed to participate in the experiment. Standard errors (clustered at the ward level) are reported in brackets. Travel time is reported in walking minutes.

Table 1.3: Effect of Price-Cap on Price Variation and Levels

Outcome:	**Mean Price**		**Within Agent Price Variation**		**% at Price-Cap**
	(1)	(2)	(3)	(4)	(5)
Price-Cap Treatment	-12.82***	-14.95***	-8.848***	-8.706***	0.200***
	(3.148)	(2.177)	(2.027)	(1.896)	(0.0212)
Observations	679	769	679	768	679
Data Source	Farmer Survey	Receipts	Farmer Survey	Receipts	Farmer Survey
District FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
Control Mean	76.28	73.01	20.53	20.80	0.051
Control St. Dev.	50.29	39.16	31.96	35.96	0.16

Notes: Standard Errors are clustered at the ward level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. All specifications control for stratification variables and district fixed-effects. Columns 1 and 2 present coefficient estimates of a regression of the within-agent standard deviation of prices on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Columns 3 and 4 present coefficient estimates of a regression of average price per chicken charged per agent on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. The outcome variable for columns 3 and 4 is the average price charged by participants. Column 5 presents coefficient estimates of a regression of the fraction of all transaction at 80 Tanzanian Shillings on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Columns 1, 3, and 5 employ farmer survey data whereas columns 2 and 4 use the receipt data.

Table 1.4: Recipient level correlates of prices in the control group

	Outcome Variable: Price per vaccination						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Walking time to village	0.0712						
	(0.0463)						
# Vaccinations		-0.115***					
		(0.0298)					
Main livelihood is agriculture			5.386				
			(4.689)				
More than primary education				-9.033*			
				(4.607)			
# Cattle owned					-0.176*		
					(0.0907)		
# Sheep owned						-0.283***	
						(0.0976)	
# Goats owned							-0.196**
							(0.0970)
Observations	1,556	1,562	1,554	1,552	1,549	1,551	1,550
Data Source	Survey	Survey	Survey	Survey	Survey	Survey	Survey
District FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Standard Errors are clustered at the ward level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. All specifications control for district fixed-effects and employ farmer survey data from the control group. Column 1 presents coefficient estimates of a regression of the price per vaccination on a continuous variable measuring the walking distance from the agent's headquarter to the farmer's home. Column 2 presents coefficient estimates of a regression of the price per vaccination on the number of vaccinations applied. Column 3 presents coefficient estimates of a regression of the price per vaccination on a binary variable indicating whether a household's main livelihood is derived from agriculture. Column 4 presents coefficient estimates of a regression of the price per vaccination on a binary variable indicating whether the household head has received more than primary education. Column 5 presents coefficient estimates of a regression of the price per vaccination on a continuous variable for the number of cattle owned by the farmer. Column 6 presents coefficient estimates of a regression of the price per vaccination on a continuous variable for the number of sheep owned by the farmer. Column 7 presents coefficient estimates of a regression of the price per vaccination on a continuous variable for the number of goats owned by the farmer.

Table 1.5: Heterogeneous Effects on Prices by Elasticity

Interaction Variable:	**Main Livelihood is Agriculture**	**Farmer is a Smallholder**
Outcome: Price	(1)	(2)
Price-Cap Treatment	-6.347	-10.27***
	(4.223)	(3.589)
Interaction Var.	5.978	13.82***
	(4.414)	(5.008)
Price-Cap $\times$ Interaction Var.	-9.382*	-13.61**
	(5.546)	(5.690)
Observations	3,043	3,045
Data Source	Farmer Survey	Farmer Survey
District Fixed Effects	Yes	Yes
Agent Controls	Yes	Yes
Farmer Controls	Yes	Yes

Notes: Standard Errors are clustered at the ward level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. All columns present regressions of transaction level prices on an indicator variable for the treatment, district fixed effects, ward-level stratification and control variables, a proxy for farmers' elasticity of demand and the interaction of this proxy with the treatment indicator. The proxy variables are binary variables indicating whether the recipient's main source of income is agriculture and whether households are smallholders who own fewer than 11 chickens.

Table 1.6: Effect of Treatment on Composition of Recipients, by Elasticity

Recipient Characteristic:	**Not served before**		**Main Livelihood is Agriculture**		**Farmer is a Smallholder**	
	(1)	(2)	(3)	(4)	(5)	(6)
Price-Cap Treatment	0.0565**	0.0612**	0.0543**	0.0470**	0.0150	0.0248
	(0.0250)	(0.0260)	(0.0272)	(0.0215)	(0.0137)	(0.0205)
Observations	832	3,095	832	3,096	832	3,098
Observation Level	Officer	Transaction	Officer	Transaction	Officer	Transaction
Data Source	Survey	Survey	Survey	Survey	Survey	Survey
District Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Agent Controls	Yes	Yes	Yes	Yes	Yes	Yes
Farmer Controls	N.A.	Yes	N.A.	Yes	N.A.	Yes
Control Mean	0.38	0.51	0.59	0.79	0.14	0.28
Control St. Dev.	0.38	0.50	0.40	0.40	0.19	0.45

Notes: Standard Errors are clustered at the ward level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Columns 1 presents coefficient estimates of a regression of the proportion of farmers who have not received services before on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Columns 2 presents coefficient estimates of a regression of a binary variable indicating whether a farmer has received services before on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Columns 3 presents coefficient estimates of a regression of the fraction of farmers served per agent whose main livelihood comes from agriculture on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Columns 4 presents coefficient estimates of a regression of a binary variable indicating whether a recipient's main source of income is from agriculture on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Columns 5 presents coefficient estimates of a regression of the fraction of farmers served that own fewer than 11 chickens on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Columns 2 presents coefficient estimates of a regression of a binary variable indicating whether a farmer owns fewer than 11 chickens on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables.

Table 1.7: Effect of Treatment on Village Choices

Village Level Outcome:	**Proportion Visited**	**Above Median Distance**	**Furthest Village**
	(1)	(2)	(3)
Price-Cap	-0.0445**	-0.0966***	-0.0598**
	(0.0219)	(0.0292)	(0.0303)
Observations	832	832	832
District Fixed Effects	Yes	Yes	Yes
Controls	Yes	Yes	Yes
Control Mean	0.37	0.37	0.23
Control St. Dev.	0.25	0.48	0.42

Notes: Standard Errors are clustered at the ward level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Column 1 presents coefficient estimates of the proportion of villages visited in the agent's area of responsibility on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Column 2 presents coefficient estimates of a regression of a binary variable indicating whether agents visited a village that was further than the median distance of all villages to their headquarter on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Column 3 presents coefficient estimates of a regression of a binary variable indicating whether agents visited the furthest away village in their area of responsibility on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables.

Table 1.8: Effect of Price-Cap on Quantities and Revenue

Outcome:	**# Farmers Served**		**# Vaccinations**	**Revenue**
	(1)	(2)	(3)	(4)
Price-Cap Treatment	-3.450	-4.759	13.54	-11,675*
	(3.708)	(4.209)	(121.6)	(6,372)
Observations	832	679	832	832
Data Source	Farmer Survey	Farmer Survey	Farmer Survey	Farmer Survey
District Fixed Effects	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Cond. on Participation	No	Yes	No	No
Control Mean	50.66	62.62	1’154.47	76’118.21
Control St. Dev.	56.27	56.21	1’614.46	96’206.8

Notes: Standard Errors are clustered at the ward level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. All specifications control for stratification variables and district fixed-effects. Columns 1 and 2 present coefficient estimates of a regression of the number of farmers served on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Columns 3 presents coefficient estimates of a regression of the number of vaccinations applied on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Columns 4 presents coefficient estimates of a regression of total revenue collected on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. All specifications employ farmer survey data.

Table 1.9: Fixed Fee Treatment Effects on Prices, Quantities and Revenue

Outcome:	**Price**	**# Farmers**	**Participation**	**Log(Q)**	**Log(Q)**	**Log(Q)**
	(1)	(2)	(3)	(4)	(5)	(6)
Fixed Fee	7.155**	-6.693*	-0.00405			
	(3.106)	(4.009)	(0.0267)			
Log(Price)				-1.223	-0.413	-3.834
				(0.907)	(0.607)	(38.663)
Observations	679	679	832	679	594	395
Recipients	All	All	All	All	Low ϵ_D	High ϵ_D
Estimation Method	OLS	OLS	OLS	IV	IV	IV
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Standard Errors are clustered at the ward level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Column 1 presents coefficient estimates of a regression of the average price per vaccination on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Column 2 presents coefficient estimates of a regression of the total number of farmers served on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Column 3 presents coefficient estimates of a regression of the revenue collected on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Columns 4 to 6 presents coefficient estimates of an instrumental variable regression of the log number of farmers served on the log of average prices charged, using the fixed fee treatment as an instrument. Column 4 presents the coefficient estimates for the whole sample, whereas columns 5 and 6 present the results separately for high and low elasticity recipients.

Table 1.10: Estimates of Welfare Effect of Marginally Reducing Prices

	Uniform Pricing		**Price Discrimination**					
	$g = 0$	$g = 1$	$g_a = 0, g_{c_1} = g_{c_2} = 1$			$g_a = g_{c_1} = g_{c_2} = 1$		
			$\mu = 0$	$\mu = 0.5$	$\mu = 1$	$\mu = 0$	$\mu = 0.5$	$\mu = 1$
Panel A: Estimates	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\tau = -\frac{1}{\epsilon_D}$	-4.62%	-4.62%	-3.33%	-10.55%	-17.77%	-3.33%	-11.57%	-19.82%
$\tau = 0$	-4.62%	-4.17%	-3.33%	-10.47%	-17.62%	-3.33%	-11.50%	-19.67%
Panel B: Counterfactual								
$\tau = -\frac{1}{\epsilon_D}$	2.05%	2.05%	0%	1.37%	2.74%	0%	0.34%	0.69%
$\tau = 0$	2.05%	2.50%	0%	1.45%	2.89%	0%	0.42%	0.84%

Notes: This table presents the results from the calibration of the sufficient statistics formulas. All estimates are expressed as percent of total sales revenue per agent. Columns 1 and 2 show the results for uniform pricing. Columns 3 to 8 show the welfare estimates measured using the sufficient statistics formula extended to third-degree price-discrimination. Panel A considers the aggregate welfare effects, whereas Panel B considers a counterfactual in which I ignore extensive margin responses.

Graphs

Figure 1.1: Map of Study Area

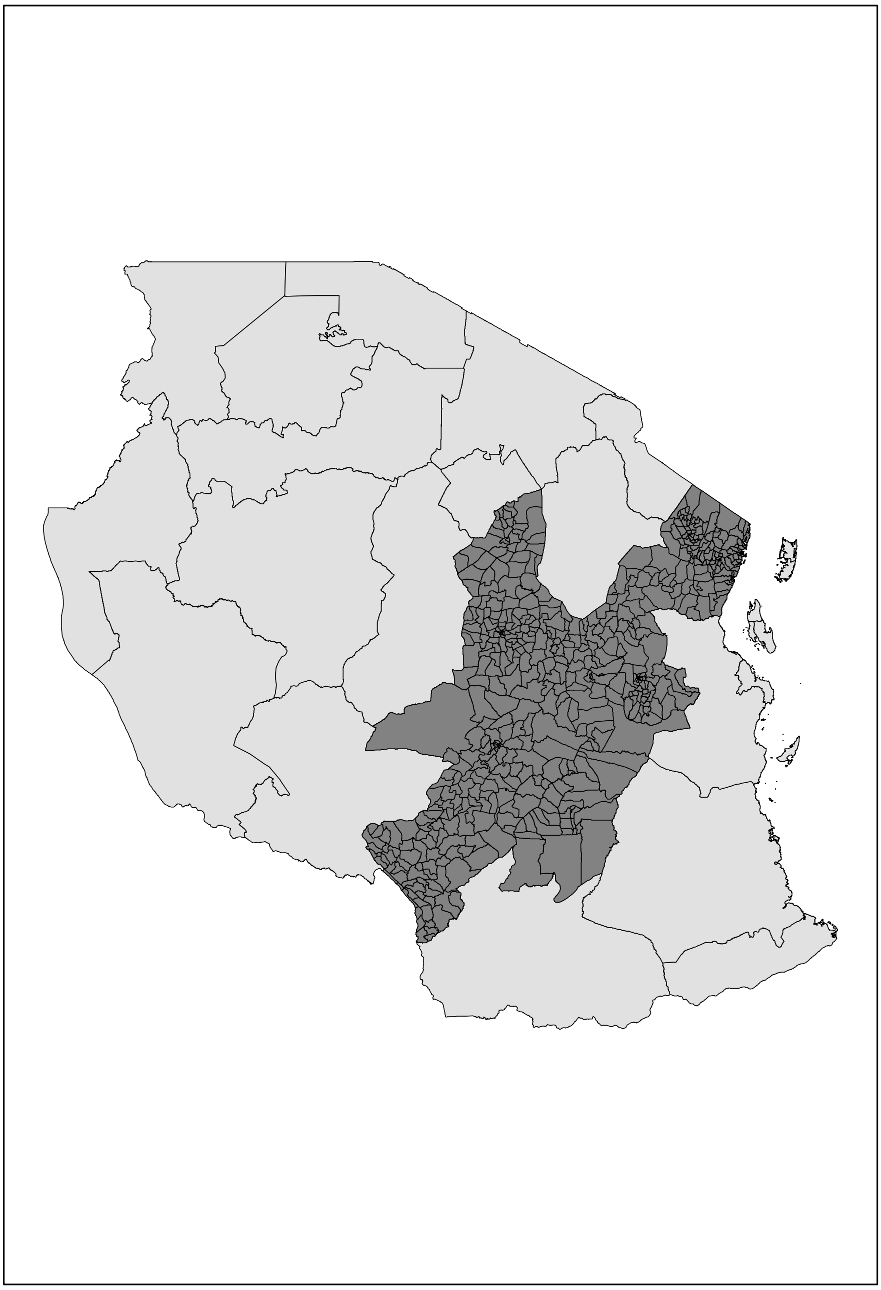

Notes: The figure shows a map of Tanzania. The study area, including the 550 sample wards, are shaded in darker grey.

Figure 1.2: Effect of Price-Cap on Price Distribution

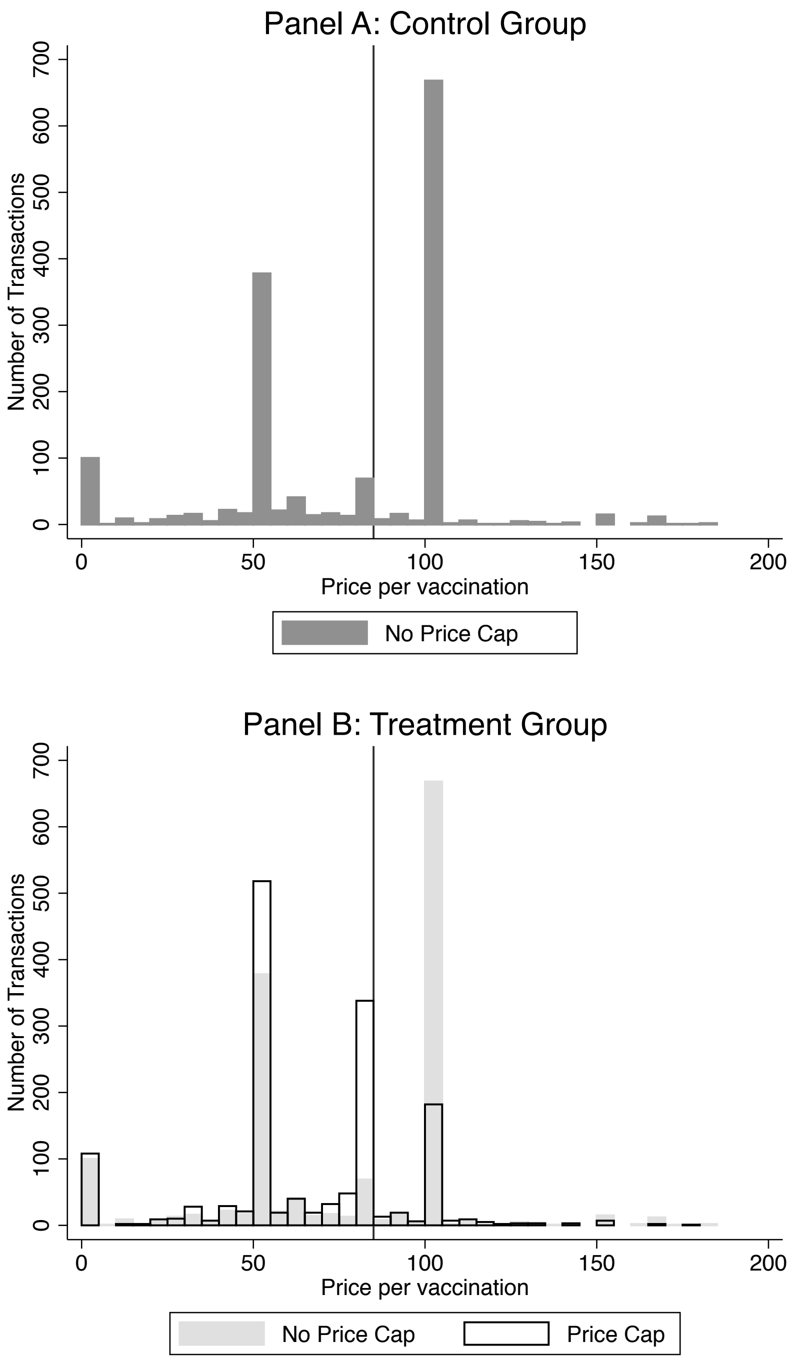

Figure 1.3: Box Plot of Price Variation

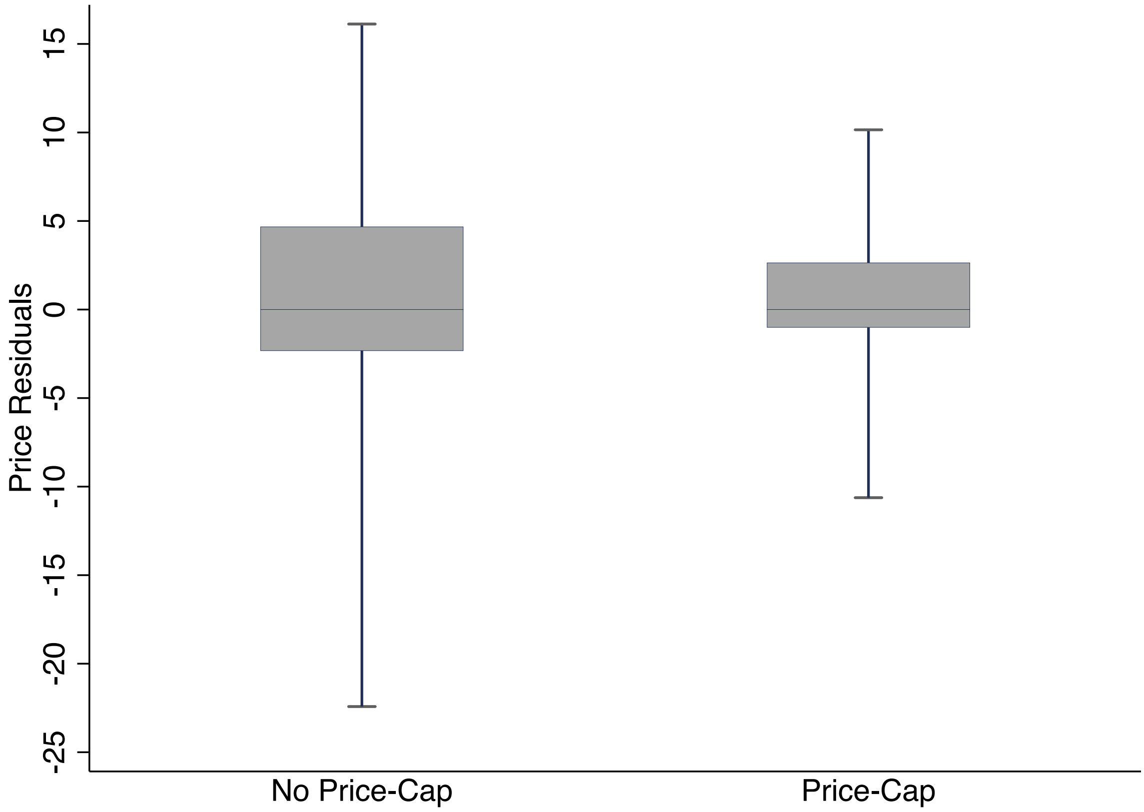

Notes: The figure shows box plots of residuals of a regression of prices on agent fixed effects. The regressions are estimated using receipt data. The box denotes the distribution of observations between the 25th and 75th percentile. The whiskers denote the length between the 10th and the 90th percentile of the price distribution. The vertical bar denotes the mean which, by construction of residuals, is at 0.

Figure 1.4: Effect of Price-Cap on Number of Farmers Served

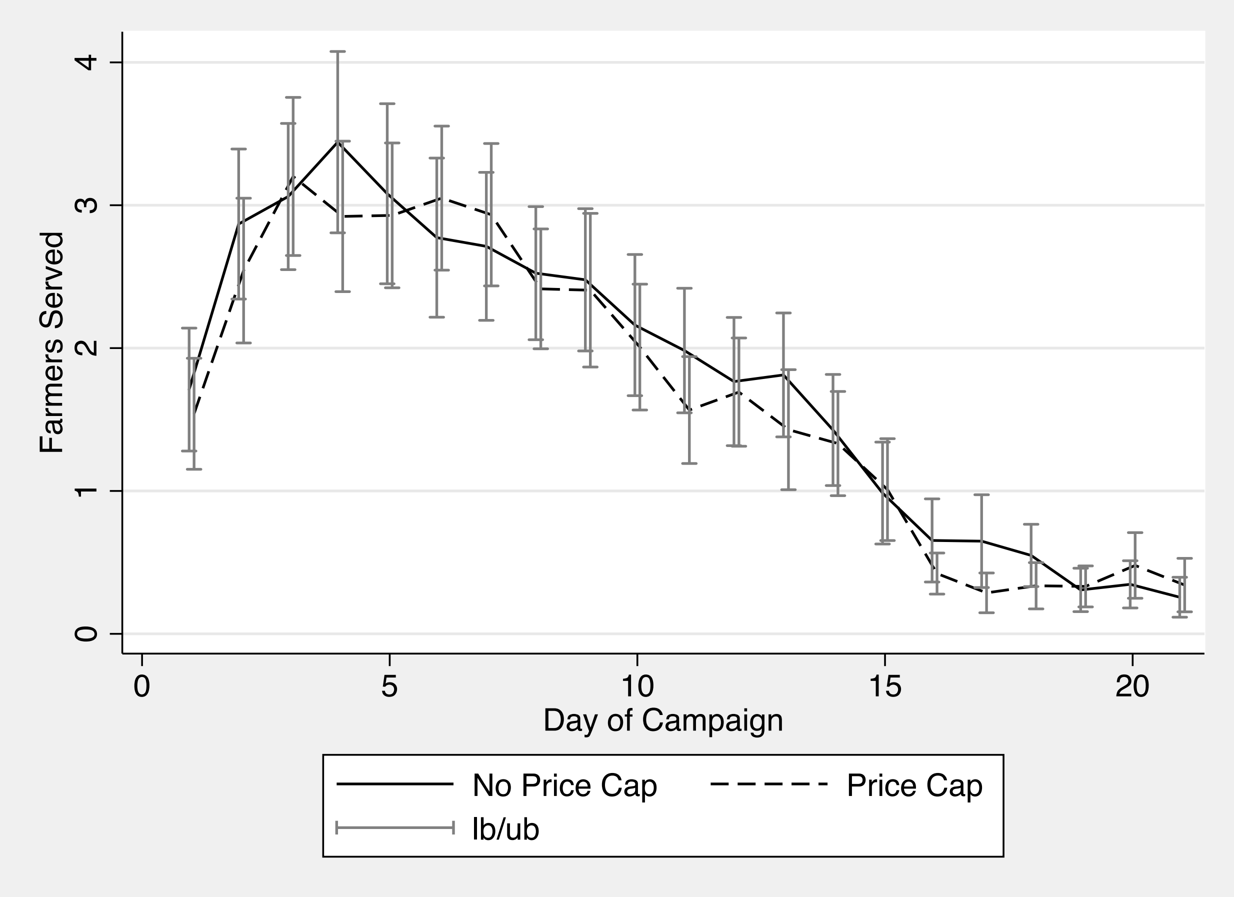

Notes: The figure shows the daily number of farmers served for every day of vaccination campaign, separated by treatment and control group. The error bars denote 95% confidence intervals. The figure uses receipt data.

Chapter 2

Are Fixed Fees Distortionary? Experimental Evidence

A theoretically efficient contract for public service delivery agents that minimizes moral hazard costs and avoids behavioral distortions charges agents a fixed fee for the usage of public assets and makes them residual claimants on its returns. I investigate whether such contracts are indeed efficient in practice by investigating whether imposing lump-sum fees on livestock extension agents distorts their choices. Using a field experiment, I first show that, contrary to classic economic theory, levying a fixed fee on agents leads them to increase user fees for a livestock vaccine and induces demand effects that reduce quantities. To understand the mechanisms underlying this result, I implement a series of lab-in-the-field experiments with a subset of the field-experimental participants. The results suggest that instead of setting prices for user fees as mark-ups over marginal costs agents use simplified rules-of-thumb that anchor pricing decisions on aggregate profits. The results highlight that boundedly rational behavior can reduce the effectiveness of adopting fixed fee contracts.

2.1 Introduction

In many contexts, principals delegate the use of their assets to an agent. Such relationships are particularly common in the public sector, where governments hire agents (such as bureaucrats and frontline workers), task them with the delivery of public services and supply them with the assets necessary to carry out this task.[1] A central challenge when designing contracts to govern such relationships is moral hazard: when the principal cannot observe effort, agents might shirk on the job and fail to use assets efficiently.

To overcome this problem, fixed fee contracts are often employed in practice. As part of such contracts, agents pay a fixed fee to the principals who in turn makes agents residual claimants on their assets' returns. For example, in many instances public service delivery agents in developing countries receive supplies (such as vaccines or bed nets) from governments in exchange for a fixed fee, and are then allowed to sell the goods to service recipients.[2] Such arrangements are attractive because they provide strong incentives to agents and thus reduce moral hazard costs.

A central assumption underlying the effectiveness of such contracts is that agents treat the fixed fee that is used to extract rents from them as sunk when deciding how to use the asset. Under this assumption, variation in the fixed fee does not distort agents' choices, so that asset returns are maximized. This feature is particularly attractive for public service delivery as it minimizes moral hazard costs and thus implies that tax revenue is used effectively and aggregate returns to public investment are maximized.

It is, however, not clear whether agents really treat the fixed fee component of such contracts as sunk. For example, Liebman and Zeckhauser (2004) suggest that individuals might either treat average as marginal prices

[1]The separation between ownership and usage of assets is also common in other contexts. For example, in sharecropping relationships land-owning principals delegate cropping and harvesting to tenants.

[2]Arrangements that rely on user fees that accrue private benefits for service delivery agents are common in agricultural and livestock extension (e.g. Rivera and Gustafson, 1991, Dancey, 1993). Fixed fee contracts are also increasingly used for health service provision. For example, Deserranno (2016) studies community health workers recruited by BRAC, an international NGO, in Uganda. Those workers are compensated by allowing them to sell medication, supplied by BRAC, to service recipients. Similarly, distributors and local government agents in public food distribution systems, such as Solidaridad in the Dominican Republic and Raskin in Indonesia, receive public supplies and are residual claimants on their sales (Busso and Galiani, 2014, Banerjee, Hanna, Kyle, et al., 2015.) Fixed fee contracts are also common in other contexts, such as land tenancy arrangements and in the taxi industry.

(*"ironing"*) or focus on local prices and ignore full price schedules (*"spotlight"*). In addition, the work by Tversky and Kahneman (1974), and the behavioral economics literature building on it, points out that people might resort to simplified heuristics, or rules-of-thumb, instead of identifying optimal solutions.[3] If agents fail to perfectly optimize their choices and instead rely on heuristics that are affected by sunk fees, employing theoretically optimal fixed fee contracts can induce distortions that reduce their effectiveness in practice.

This paper investigates empirically whether agents treat fixed fees as sunk. I study the behavior of public livestock extension workers in Tanzania who are tasked with delivering poultry vaccines and focus on understanding whether introducing a fixed fee component to their contracts affects user fees charged by agents and the equilibrium number of recipients served. As the basis of this study, I investigate the publicly subsidized provision of I-2 poultry vaccines that protect against Newcastle Disease (ND). ND is highly prevalent in East Africa and is lethal for infected birds, therefore posing a substantial economic risk for populations dependent on agriculture. I-2 is the primary ND vaccine used by rural farmers in Tanzania. It is exclusively produced by the government and distributed to farmers by public livestock extension agents. In the status quo, agents receive the vaccine for free from the government and distribute it by travelling to farmers in a geographically defined service area, where they perform veterinary examinations and apply the vaccine. Agents have a local monopoly over the provision of the I-2 vaccine, which they are allowed to exploit by collecting user charges from farmers. They are thus residual claimants on the returns to the vaccine sales.

To identify whether agents treat fixed fees as sunk, I collaborate with the Tanzanian government. I first provide evidence from a field experiment which shows that adding a fixed fee component to agents' contracts in order to extract rents from them affects equilibrium prices and quantities. Working with the Ministry of Agriculture, Livestock and Fisheries (MALF), I allocate 550 wards, the main administrative units at which agents are organized, to one of two experimental groups. 277 of these wards are assigned to the control group in which agents receive the vaccine for free from the government. This group is designed to mirror the status quo. In contrast, 273 wards

[3]This notion is similar to the idea by Baumol and Quandt (1964), who point out that *"the more refined a decision-making process, the more expensive it is likely to be, and therefore [...] no more than an approximate solution may be justified."* Conlisk (1996) summarizes a literature which argues that such behavior can be rationalized in the presence of deliberation costs. Pingle and Day (1996) provide lab-experimental evidence that empirically links decision cost to such approximate behavior. See also the summary of these arguments in Kahneman (2002).

are assigned to the treatment group, in which agents are required to remit a fixed fee of 25,000 Tanzanian Shillings (approximately \$11.40). Agents in both experimental groups have full discretion over which and how many customers are served, which geographic areas are targeted, and whether they choose to participate in the vaccination campaign or not.

The evaluation of this intervention yields three results. First, contrary to classic economic theory, implementing a fixed fee induces agents to raise prices. I find that the intervention increases average prices charged by 11% and induces demand effects that reduce the aggregate quantity of customers served by 12%. Second, variation in the fixed fee does not affect agents' participation in the vaccination campaign: I find that implementing a fixed fee decreases participation by 0.005%, which is statistically insignificant. The observed increase in prices in response to implementing a fixed fee is thus not driven by selective participation decisions. Third, the fixed fee only marginally reduces revenue earned by agents. I find that the intervention reduces revenue by 3%. Back-of-the-envelope calculations suggest that for reasonable values of service delivery costs, agents make a larger profit before remitting the fixed fee in the treatment than in the control group. Taken together, these results therefore present evidence in support of the hypothesis that agents are not choosing prices based on marginal costs. Instead, they appear to base prices on heuristics that are affected by variation in sunk costs, which means that the fixed fee induces a behavioral response that distorts prices and quantities compared to the status quo.

The evidence from the field experiment shows that fixed fee contracts induce distortions. The results are, however, insufficient to provide conclusive evidence on the underlying mechanisms. To understand the reasons for the observed response to the introduction of the fixed fee, I implement a series of lab-in-the-field experiments with a randomly selected subset of field-experimental participants.[4]

The design of the experiments is guided by a simple theoretical model which predicts comparative statics that allow me to distinguish between three possible pricing strategies. First, I consider a baseline scenario in which respondents are unconstrained in their profit maximization and chose prices optimally based on the available information. Second, Liebman and

[4]The combination of field- with lab-experiments is an increasingly adopted technique to identify mechanisms while showing the economic significance of effects (e.g. Kessler (2017)) and to provide evidence on the external validity of lab-experimental results (e.g. Rondeau and List (2008)).

Zeckhauser (2004) introduce the term *"ironing"* to refer to a situation in which individuals perceive a multipart price schedule only as the average price at the point of consumption. Consistent with this notion, an empirical literature documents individuals' confusion between average and marginal prices.[5] To address this possibility, I consider a scenario in which respondents attempt to optimize prices, but confuse average and marginal costs. Third, a literature in behavioral economics argues that instead of solving optimization problems, individuals might resort to intuitive rules-of-thumb to avoid deliberation costs.[6] The design of my experiments allows me to consider and evaluate whether respondents use a rule-of-thumb that anchors prices on the aggregate level of profits.

The main lab-in-the-field experiment I consider is a financially incentivized experiment which simulates the task that agents encounter in the field. Respondents are tasked with providing services to four potential customers, each of whom has a different marginal service delivery cost. For each customer, I elicit agents' perceived reservation price, i.e. the smallest price at which respondents are willing to serve the customer, using a multiple price list mechanism.[7] The experimental design provides exogenous variation in fixed fees and revenue for every customer, which allows me to separately identify their effect on reservation prices.[8]

If agents chose their reservation prices optimally, they would equal marginal service delivery cost for every customer, and would thus not be affected by variation in fixed fees and revenue. In contrast, I find that variation in both variables affects respondents' reservation prices: a 1 unit increase in average costs caused by the introduction of a fixed fee, conditional on marginal costs, leads to a 0.4 unit increase in reservation prices, while a 1 unit increase in revenue causes reservation prices to decrease by 0.07 units. Both effects are statistically significant at the 1% level. By comparing those

[5]See, for example, Feldman, Katuscak, and Kawano (2016) and Ito (2014).

[6]For example, Simon (1978) introduces the distinction between substantive and procedural rationality to argue that in complicated environments individuals might use rules to reduce costly optimization. This notion is supported by empirical evidence (e.g. Pingle and Day (1996) and Kahneman and Frederick (2002)) and its implications have been evaluated theoretically both at the micro (e.g. Ellison and Fudenberg (1993), Ellison and Fudenberg (1995) and Spiegler (2006)) and the macro level (e.g. Sims (2003) and Reis (2006)).

[7]The mechanism I employ is a simplified and context-adapted version of the one proposed by Becker, DeGroot, and Marschak (1964).

[8]The lab-experiment is similar in spirit to the one employed in a literature on dynamic stochastic decision making for saving and investment decisions. See, for example, Hey and Dardanoni (1988), Ballinger, Palumbo, and Wilcox (2003) and Oprea, Friedman, and S. T. Anderson (2009).

comparative statics to the implications of my model, I find that the observed responses to changes in revenue are inconsistent with marginal and average cost pricing, but are consistent with a rule-of-thumb which leads respondents to target an aggregate level of profits and adjust their pricing strategies to achieve this target. As aggregate profits are a function of sunk costs, the reliance on rules-of-thumb explains the observed effect of the fixed fee on prices.

To further support the notion of such rule-of-thumb pricing, I develop and implement a second lab-in-the-field experiment with the same sample of respondents. In contrast to the previous experiments, this simulation asks respondents to determine aggregate quantities of production. To this end, I implement an experiment in which respondents face increasing marginal costs and stochastic returns to production and must decide dynamically whether to continue production. The experimental design again generates exogenous variation in fixed fees and revenue.

If agents were optimally choosing their production quantities on the margin, they would continue producing until the expected return of producing an additional unit equals its marginal cost. Contrary to this assertion, I find that charging a fixed fee significantly reduces production, while increases in revenue significantly increase the aggregate quantity produced: introducing a fixed fee leads to a 16% reduction in the aggregate quantity produced, while a 1% increase in revenue causes a 7% to 13% increase. Both effects are statistically significant at the 1% level. This result is again inconsistent with profit maximization and a confusion of average with marginal costs, but is consistent with a model in which agents make economic decisions based on rules-of-thumb that are anchored on aggregate profits.

Taken together, this paper thus provides evidence that agents do not treat fixed fees as sunk in practice, which reduces the efficiency of fixed fee contracts. By doing this, this article complements two strands of literature. First, I add to a literature in behavioral economics that documents sub-optimal consumer and producer behavior.[9] Altomonte, Barattieri, and Basu (2015) present survey evidence which suggests that firms set prices as mark-ups over total instead of marginal costs.[10] Oprea (2014) provides evidence from a

[9]Primarily theoretical contributions of the effects of bounded rationality include Sah and Stiglitz (1986), Haltiwanger and Waldman (1985) and De Long et al. (1990) and are summarized in Conlisk (1996).

[10]The literature that investigates the optimality of price choices by firms is primarily descriptive. The finding Altomonte, Barattieri, and Basu (2015) builds on a classical debate on

lab-experiment to show that when calculating the profit maximizing strategy is complicated, respondents rely on statistics that are easier to observe but might lead to sub-optimal outcomes. Ito (2014) studies non-linear electricity pricing and finds that consumers respond to variation in average instead of marginal electricity prices. Feldman, Katuscak, and Kawano (2016) show that taxpayers' earnings respond to predictable lump-sum changes in tax liability, which suggests that households misperceive complicated tax schedules. This paper adds to this literature by providing causal evidence on the mechanisms underlying sub-optimal price setting behavior.[11]

Second, this study complements a literature on contract design for public sector workers. Existing work in this area has focused on the role of monitoring (e.g. Rasul and Rogger (2016)), team work (e.g. Chan (2015)), as well as financial (e.g. Ashraf, Bandiera, and Jack (2014), Khan, Khwaja, and Olken (2016b)) and non-financial (e.g. Khan, Khwaja, and Olken (2016a)) incentives. I add to this stream of literature by evaluating the practical value of fixed fee contracts in providing incentives and maximizing the returns to public investment.

The remainder of this article is structured as follows. Section 2.2 describes the relevant features of the setting in which the study takes place. Section 2.3 provides a theoretical model to guide the empirical analysis and derives comparative statics that allow to distinguish between different pricing strategies. Section 2.4 outlines the design of the field experiment and presents the results. Sections 2.5 and 2.6 outline the design of the two lab-in-the-field experiments and present their results. Section 2.7 concludes.

2.2 Setting

The setting of this study is the same as that described in chapter 1. I repeat the key characteristics here for convenience.

This project explores the response to fixed fees in the context of agricultural and livestock extension in Tanzania. Extension services are a publicly subsidized service delivery program administered by local governments and coordinated nationally by the Ministry of Agriculture, Livestock and Fish-

full cost pricing, which started with Hall and Hitch (1939) and was summarized in Faulhaber and Baumol (1988) and Mongin (1992).

[11] In a broader sense, the results in this paper appear inconsistent with the *"poor but rational"* view which suggests that the poor maximize profits subject to financial constraints (e.g. Duflo (2006)).

eries.[12] As part of this project I focus on the provision of I-2 vaccines as one dimension of services. I-2 is a thermotolarent vaccine for poultry that protects against Newcastle Disease (ND), a viral disease that is transmitted between birds and leads to almost 100% mortality in affected and unvaccinated chicken. Estimates from Tanzania suggest that more than 30% of chicken die from ND every year, leading to an annual cost of up to $78 Million (Msami, 2007).

The vaccine is delivered to farmers by public livestock extension agents. They have advanced professional qualifications and typically hold a diploma from specialized training institutes in animal health, animal production or general agriculture. As part of an I-2 vaccination program, they receive vaccines from the government and then travel to customers in order to apply vaccinations to farmers' livestock. During the study period, vaccine delivery was the primary task assigned to the agents. They are responsible for 1 to 12 villages, averaging 4 villages per agent, work by themselves and have their own geographically defined area of responsibility, in which they face no competition from other public providers. The main organizational unit of agents at the local level are wards, which are accumulations of roughly 8 villages. There are, on average, two agents per ward. Agents in the same ward interact on a daily basis and, while maintaining their own geographic areas, typically coordinate their work. There was no entry or exit of agents during my study period.

Delivery agents are compensated through a mix between government wages and user fees. Specifically, agents receive a flat compensation of around $200 per month.[13] In addition, agents can collect user fees from farmers, which cover delivery costs and provide incentives.[14] User fees are necessary as only around 25% of agents have access to a government motorcycle, whereas the remainder travels by foot and uses public motorcycle taxis and buses to reach farmers.[15] Local governments have allocated full discretion

[12]This division of responsibilities is the common organizational form for extension in Africa (Crowder et al., 2002) and Latin America (Wilson, 1991).

[13]Wages vary across local government administrations.

[14]This dual compensation structure is common for agricultural extension, and publicly subsidized service delivery programs in developing countries in general. See, for example Rivera and Gustafson, 1991, Dancey, 1993, Haan et al. (2001), J. R. Anderson and Feder (2007) and Lagarde and Palmer (2008).

[15]This process of service delivery is the modus operandi for a number of key service delivery programs in developing countries. For example, Ashraf, Bandiera, and Lee (2016) study community health workers in Zambia. Community health workers are expected to devote 80% of their time to household visits and are hence required to incur similar costs as the agents in my setting.

over pricing for user fees to the agents.[16] As the Tanzanian government is the only producer and provider of I-2 vaccinations, agents have substantial market power, which limits the extent to which competition can discipline agents' discretion over prices.[17] In the status quo, agents are not required to submit a fixed payment in exchange for the right to sell vaccines.

Given this market structure, prices for user fees are a central choice variable that determines the allocation of services. The focus of this study is on understanding how pricing strategies for such user fees are affected by the addition of a fixed fee component to agents' contracts.

2.3 Conceptual Framework

This section provides a framework that conceptualizes the implications of different pricing strategies for the response of prices to the introduction of a fixed fee component to agents' contracts. To this end, I extend the model developed in chapter 1 to derive comparative statics for three possible pricing strategies. The comparison of the theoretical comparative statics with empirically estimated elasticities provides indications regarding the pricing strategies used by agents. Throughout this analysis I focus on deriving three comparative statics: the response of prices to changes in marginal costs, average costs and revenue.

I focus on three possible pricing strategies. First, I derive comparative statics assuming that agents are choosing prices optimally. In this case, prices respond only to marginal costs. Second, I evaluate the implications of a situation where agents set prices based on average instead of marginal costs. In this scenario, I assume that agents solve the profit maximizing optimization problem but confuse average and marginal costs. Finally, I investigate a situation of bounded rationality. In this scenario, agents do not solve an optimization problem but instead rely on simplified heuristics that approximate the profit maximizing pricing choice. In particular, I consider a situation in which prices are chosen to target a pre-defined level of aggregate profits. In this scenario, reductions in the stock of profits can lead agents to raise prices, as they attempt to realign their aggregate profits with their

[16] In a separate intervention, discussed in chapter 1, I evaluate the effect of introducing regulation to reduce the discretion over prices.

[17] Private provision is widely viewed as sub-optimal to public provision, as effective ND vaccination requires adequate handling and application of the vaccine as well as the ability to detect preexisting poultry infections. A central concern is that private market provision would compromise service quality which risks reducing the acceptance and adoption of I-2 vaccinations.

targets. This section first presents the general set-up of the model before considering the three possible scenarios and deriving the corresponding comparative statics.

2.3.1 Set-Up

I consider a situation where one agent is responsible for supplying a service to a continuum of customers, the size of which is normalized to 1. Each customer is indexed with the subscript i. Serving customer i requires the agent to pay a travel cost c_i and a service application cost τ. Travel costs are drawn from a continuous distribution $M(c)$ defined on $[0, c^{max}]$. Denote by $c^*(p_i)$ the highest cost customer that agents are willing to visit at price p_i. In addition, suppose that market participation requires agents to pay a fixed fee K before any services are delivered. Customers have an idiosyncratic valuation v_i for the service, which is drawn from a homogeneous distribution $F(v)$.[18]

The timing of the interaction between agents and service recipients is as follows. Agents first decide whether to participate in the market and pay the fixed fee K. They then choose whether to travel to service recipient i. After paying the travel cost c_i, agents meet with the service recipient and propose a price p_i. Customers then decide whether to accept or reject the price offer. If they accept, agents provide the service at a cost τ. If customers reject, no transaction takes place and agents move on to the next customer.

For ease of exposition, assume that agents visit customers one after the other, starting with the customer with the lowest travel cost. Denote by R_{i-1} the aggregate revenue earned before deciding whether to visit customer i. Similarly, denote by C_{i-1} the total cost paid by agents before deciding whether to visit customer i. In the following analysis I will focus on the considerations the agent takes into account when proposing a price to the customer.

2.3.2 Choosing Prices to Maximize Profits

I begin by considering a situation in which agents choose prices optimally. Their aggregate profits from the vaccination task are given by:

$$\Pi = \int_{c_i \leq c^*(p_i)} \left[(p_i - \tau)(1 - F(p_i)) - c_i \right] di - K$$

Aggregate profits are additively separable for all customers that are visited. Optimizing agents thus set prices to maximize profits separately for every

[18]Note that, for simplicity and in contrast to the model presented in chapter 1, I assume that the elasticity of demand is constant.

service recipient. Profits at the time when prices are set for customer i are given by:

$$\pi(p_i) = R_{i-1} - C_{i-1} + (p_i - \tau)(1 - F(p_i)) - c_i$$

where $F(p_i)$ denotes the probability that individual i will accept the price offer p_i. Optimizing profits with respect to prices yields the well known formula for a uniform pricing monopolist:

$$p_i = \tau \left(\frac{\eta}{1+\eta} \right) \tag{2.1}$$

where η denotes the price elasticity of demand.

The objective of this section is to derive three comparative statics that can then be compared to the data. Using formula 2.1, it is straightforward to see that:

$$\frac{\partial p_i}{\partial \tau} > 0, \frac{\partial p_i}{\partial C_{i-1}} = 0, \frac{\partial p_i}{\partial R_{i-1}} = 0 \tag{2.2}$$

Intuitively, the separability of the objective function assures that prices that maximize profits are chosen exclusively based on marginal costs. While variation in travel costs affect whether a customer is served at all, all costs apart from the service application cost τ are sunk at the time of pricing and are thus not considered by the agent when proposing prices.

2.3.3 Confusion of Marginal and Average Costs

While choosing prices based on marginal costs maximizes profits, it is not clear whether agents actually employ this pricing strategy. Liebman and Zeckhauser (2004) introduce the term *"ironing"* to refer to a situation in which individuals perceive a multipart price schedule only as the average price at the point of consumption. Consistent with this notion, an empirical literature documents individuals' confusion between average and marginal prices.[19] Consistent with these ideas, it is hence possible that service delivery agents base their pricing decisions on average instead of marginal costs.[20]

To formalize this scenario, consider a situation in which, instead of observing τ, agents perceive marginal costs to be $\tilde{\tau}$. Suppose further that $\tilde{\tau}$ is a weighted average between true marginal costs τ and the average of total costs incurred at the time of visit to customer i: $\tilde{\tau} = \omega\tau + (1-\omega)\, ac_i$,

[19] E.g. Feldman, Katuscak, and Kawano (2016) and Ito (2014).

[20] This distinction has important implications for equilibrium supply. For example, when marginal costs are constant an agent choosing prices based on average costs will consistently set prices above the monopoly level, thus increasing the distortions created by the uniformly pricing monopolist further.

where $ac_i = \frac{C_{i-1}}{\gamma} + \tau + c_i$ and γ denotes the historical acceptance rate. Under this scenario, formula 2.1 defines the prices chosen by the confused agent p_i^C:

$$p_i^C = \widetilde{\tau}\left(\frac{\eta}{1+\eta}\right) = \left\{\tau + (1-\omega)\left(\frac{C_{i-1}}{\gamma} + c_i\right)\right\}\left(\frac{\eta}{1+\eta}\right) \tag{2.3}$$

Confusion between average and marginal costs has implications for price responses to exogenous shocks. In particular, it is straightforward to see that pricing based on average costs implies the following comparative statics:

$$\frac{\partial p_i^C}{\partial \tau} > 0, \frac{\partial p_i^C}{\partial C_{i-1}} > 0, \frac{\partial p_i^C}{\partial R_{i-1}} = 0 \tag{2.4}$$

Intuitively, confusion between average and marginal costs implies that prices will not only respond positively to changes in marginal costs, but also to increases in fixed costs K and travel costs incurred before the time of pricing, both of which are captured in C_{i-1}. When marginal service application costs are constant, increasing fixed costs under average cost pricing will reduce aggregate profits, as it will lead agents to set prices even further above the profit maximizing level.

2.3.4 Rules-of-Thumb

The previous scenario assumes that agents maximize an objective function, even though their confusion between average and marginal costs prevents them from optimizing profits. A large literature in behavioral economics has, however, pointed out that individuals might rely on simplified heuristics instead of maximizing utility. To explore the implications of this possibility, I focus on a specific rule-of-thumb: profit targeting. In this scenario, agents are unaware of the expected maximum amount of profits they are able to achieve, but instead have a target aggregate amount in mind that they aspire to. Prices are then not chosen to maximize an objective function but merely aim to minimize the distance between the target and the actual amount. I consider two scenarios of such behavior.

I first consider a simplified situation in which agents are unaware of demand effects. In this scenario, agents assume a locally linear and increasing relationship between prices and profits, and prices are chosen to achieve a target level of profits. Their perceived profit obtained from customer i is thus:

$$\widetilde{\pi}_i = (p_i - \tau)\, P\,(acc) - c_i$$

where $P(acc)$ denotes the acceptance probability which, in this scenario, is independent of prices. While agents are thus aware of the stochastic nature of customers' acceptance decisions, they fail to internalize that lowering prices shifts the odds in favor of acceptance. By considering a situation of profit targeting, this scenario further assumes that agents choose a sequence of prices $\mathbf{P} = \{p_1, .., p_i, ...\}$ to achieve a pre-defined aggregate level of profits π^T. This implies a sequence of target profits $\left\{\pi_1^T(p_1), ..., \pi_i^T(p_i), ...\right\}$ which agents define ex-ante and is used to adapt pricing decisions to the stochastic realization of customer acceptances. In particular, for every service customer i agents set prices to solve the following equality:

$$\pi_i^T = R_{i-1} - C_{i-1} + (p_i - \tau) P(acc) - c_i$$

This directly defines the prices p_i^{ROT} chosen by the agent:

$$p_i^{ROT} = \frac{\pi_i^T - (R_{i-1} - C_{i-1})}{P(acc)} + \frac{c_i}{P(acc)} + \tau \tag{2.5}$$

An agent who uses rules-of-thumb based on profit targeting thus sets mark-ups over marginal costs as a function of the target profits, average costs and average revenue. In contrast to an agent who is merely confused between average and marginal costs, agents relying on the proposed heuristic thus also consider their aggregate earnings when making pricing decisions. This implies that agents lower prices in response to positive shocks to revenues. Taken together, this implies the following comparative statics:

$$\frac{\partial p_i^{ROT}}{\partial \tau} > 0, \frac{\partial p_i^{ROT}}{\partial C_{i-1}} > 0, \frac{\partial p_i^{ROT}}{\partial R_{i-1}} < 0 \tag{2.6}$$

Profit targeting behavior also means that it is unclear whether increases in fixed costs bring prices closer to the true profit maximizing level, as the implied consideration of average costs raises prices above marginal costs, but profits that are higher than expected can counteract this increase. The relationship between profits and marginal increases in fixed-costs is thus ambiguous.

The comparative statics also apply to situations in which agents are aware of (and internalize) demand effects. To see this, assume for simplicity that customers' valuation is distributed such that their acceptance curve becomes linear: $P(acc|p_i) = 1 - \beta p_i$. Agents then choose prices to satisfy the following equation:

$$\begin{aligned}\pi_i^T &= R_{i-1} - C_{i-1} + (p_i - \tau)(1 - \beta p_i) - c_i \\ &= -\beta p_i^2 + p_i(1 - \tau\beta) + R_{i-1} - C_{i-1} - \tau - c_i\end{aligned}$$

This is a quadratic in prices and thus has at most two solutions:

$$p_{1/2}^{ROT} = \frac{\tau\beta - 1 \pm \sqrt{4\beta\left((R_{i-1} - C_{i-1}) - \pi_t^T\right) + \beta\left(\tau^2\beta - 6\tau - 4c_i\right) + 1}}{-2\beta} \quad (2.7)$$

In contrast to the simpler exhibition in which agents do not internalize the demand curve, it is now ambiguous whether increases in the stock of earnings lead to an increase in prices. This is because agents can choose to either set low prices to achieve a high acceptance rate or to achieve their profit target through high prices with a comparatively lower acceptance rate. If agents choose the former approach, an unexpected decrease in the stock of already collected revenue leads to an increase in prices. In contrast, if agents choose the latter approach an unexpected decrease in the stock of already collected revenue leads to a decrease in prices. To distinguish between the two approaches, it is useful to note that they also have differing implications in terms of agents responses to an increase in average costs. In the low-price case, prices respond positively to increases in average costs whereas in the high-price case agents lower prices if average costs increase. When agents take into account demand responses, two sets of comparative statics are thus possible:

$$\frac{\partial p_i^{ROT}}{\partial \tau} \lesseqgtr 0, \frac{\partial p_i^{ROT}}{\partial C_{i-1}} > 0, \frac{\partial p_i^{ROT}}{\partial R_{i-1}} < 0 \quad (2.8)$$

or

$$\frac{\partial p_i^{ROT}}{\partial \tau} \lesseqgtr 0, \frac{\partial p_i^{ROT}}{\partial C_{i-1}} < 0, \frac{\partial p_i^{ROT}}{\partial R_{i-1}} > 0 \quad (2.9)$$

2.3.5 Summary of Model Implications

Taken together, this section developed a theoretical framework that investigates pricing behavior under three different scenarios: Profit maximization, confusion between average and marginal costs, and rule-of-thumb pricing. The conceptual framework allows me to compare the implied comparative statics with empirically observable parameters to understand which pricing strategy is consistent with real world behavior. The conceptual framework generates the central empirical implication that under profit maximization agents' prices should only respond to variation in marginal costs. In contrast, if agents confuse marginal and average costs, prices should respond positively to increases in both marginal costs and fixed fees. Finally, if agents use rules-of-thumb and thus anchor their pricing strategies on a target level of profits, prices should also respond to variation in revenue. Table 2.1 summarizes the implied conclusions. The next sections will present evidence that is consistent with the latter interpretation but inconsistent with marginal and average cost pricing.

2.4 Evidence from a Field Experiment

The experiments discussed in this paper are set in the context of the publicly subsidized provision of I2 Newcastle Disease vaccination in Tanzania. This section explains the experimental design of a field experiment that assesses whether introducing a fixed fee into delivery agents' contracts distorts prices for user fees in this context.

2.4.1 Experimental Design

This project was carried out during the first I2 vaccination campaign of 2016. The study covers the time period between January and February 2016 and enumerated the population of agents in four of Tanzania's 30 regions (Dodoma, Iringa, Morogoro and Tanga).[21] In total, 832 agents participated in the study.

All participants in the field experiment were invited to attend a 90-minute meeting at the district headquarter at the beginning of the vaccination campaign to collect the vaccine and receive instructions on procedures.[22] During this meeting, agents were informed that, similar to normal procedures, they were allowed to collect a user fee from farmers which they could keep for themselves. Agents were specifically encouraged to profit financially from the transaction, stating that the government viewed user fees as a way to motivate employees and compensate them for good performance. In addition, it was reiterated during the meeting that agents were allowed to charge farmers any price they chose and that it was acceptable to charge different prices to different farmers.

[21] All 27 districts in the enumeration regions were included in the study. The study area was chosen to include a wide variety of agricultural environments while assuring geographic proximity to the ministry headquarter in Dar Es Salaam.

[22] Agents who attended this meeting received a show-up fee to cover their transport expenditure. Payments varied between 10'000 and 50'000 Tanzanian Shillings ($4.50 to $22), depending on the distance and available transport methods.

Agents were then informed that the ministry wanted to keep better records of how many animals were vaccinated and that hence reporting procedures during this campaign would differ slightly from the status quo. In particular, a condition of participation in the vaccination campaign was that agents would issue formal receipts to every farmer served and submit the receipt information directly to the ministry using a phone based reporting system. Agents were specifically told that the ministry would contact farmers to verify that the information provided on receipts was correct. In order to assure compliance with this reporting system, the ministry offered a bonus payment of 60 Tanzanian Shillings (approximately $0.025) for every verified vaccination.

The different treatments were announced to participants only after they had arrived for vaccine collection at the district headquarters but before they departed to the field again. I examine an intervention that charges agents a fixed fee ("fixed fee treatment"). The instructions given to agents in the fixed fee treatment differed from the ones given to the control group only with regards to vaccine distribution. In particular, ministry officials informed participants that, in contrast to normal campaigns, vaccines would not be provided for free and instead participation in the vaccination campaign required a fixed payment of 25,000 Tanzanian Shillings (approximately $11.40) to cover parts of the vaccine production cost. To avoid concerns about liquidity constraints, the ministry allowed agents to cover this fee through charges from farmers and collected the funds after the completion of the vaccination campaign. Agents were explicitly given the choice whether to accept the cost, perform vaccinations and collect user fees and bonus payments, or to reject participation without any obligation to pay. The ministry repeatedly emphasized that there would be no repercussions from refusing participation.

The fixed fee is comparatively small relative to the expected revenue from participation, with average revenue in the control group exceeding 70,000 Tanzanian Shillings. However, 25% of total earnings (revenue plus bonus payments) fall below the fixed fee. For agents with low revenue potential it may hence make sense to reject this proposal, but for those with sufficient business potential the expected return from accepting the fixed fee appear substantial.

I randomly and independently assigned each of the 550 wards in the study to either the control or the treatment group. Table 2.2 displays the basic experimental design. Group allocation was stratified by 108 strata, where each stratum was defined by a district identifier and two binary variables, indicating whether all agents in the ward had specialized in general agriculture and whether only one agent was assigned to the ward.

2.4.2 Data

The main data used to evaluate the field experiment was collected from two different sources: administrative government receipts and a survey of service recipients. I designed and conducted the customer survey specifically as part of this project. In addition, I implemented a new procedure of reporting service provision receipts via text message to increase accuracy and usability of the data.

The information provided on official government receipts, and the number of receipts issued, constitutes my provisional outcome data. The information on the receipts details the customer's name, contact number, village, the date of the visit, the total user fee collected and the number of vaccinations applied. After issuing the receipt, agents electronically transmitted the receipt information to a ministry database using a text-message template.[23]

After the end of the vaccination campaign, I administered a survey to service recipients. The survey was conducted over a period of 6 weeks, between March and April 2016, and sampled a randomly selected fraction of 15% of all receipts submitted, stratified by agent. This led to a total sample of 4,516 receipts selected for surveying and verification.[24] The survey was able to contact 3,580 farmers which equates to 80% of receipts sampled.[25] The farmer survey collected detailed information on the service provision and on customer characteristics, thereby verifying that the service was actually provided and collecting verified information on user fees.

I use the information obtained from this section to construct my main outcome measures. In order to arrive at a measure of the total number of farmers served I multiply the number of verifiable receipts per agent with the agent-specific sampling weight of each receipt. I repeat the same procedure

[23]In total, agents issued 31,657 valid receipts, accounting for 702,762 animals vaccinated.

[24]Rounding errors induced by the stratification led to a sample that is slightly smaller than 15% of 31,657.

[25]Among the farmers not reached, enumerators were unable to reach 42% because of incorrect or invalid contact details. In total, phone survey procedures therefore were able to assess the validity of almost 90% of receipts sampled. I treat the remaining receipts as unverifiable and hence incorrect.

for the average price, total revenue collected and the total number of chickens vaccinated.

2.4.3 Estimation and Results

The design of the field experiment allows me to provide evidence that the introduction of a fixed fee affects prices and distort quantities. To this end, I estimate equations of the following form:

$$y_{iwd} = \beta_0 + \beta_1 FixedFee_{iwd} + \beta\mathbf{X} + \gamma_d + \epsilon_{iwd} \tag{2.10}$$

where y_{iwd} is the outcome of interest for participant i in ward w and district d, $FixedFee$ is a binary variable that indicates whether agents' wards were required to pay a fixed fee, and X_{wd} denotes a vector of ward-level stratification variables.[26] The coefficient of interest is β_1. I also include district level fixed-effects (γ_d), as the assignment lottery was stratified by these strata. As the treatment is perfectly correlated within wards, every specification reports robust standard errors clustered at the ward level.

The intervention imposes a fixed fee on agents that is sunk at the time of service delivery. If agents were fully maximizing profits in the status quo, imposing a fixed fee should only induce responses on the extensive participation margin and not affect pricing decisions conditional on participation. Table 2.3 presents the treatment effect estimates from the intervention to tests this hypothesis. Column 1 shows that imposing a fixed fee raises average prices charged by around 11%. Column 2 in table 2.3 shows that this is not a selection effect, as the fixed fee treatment did not affect participation in the vaccination campaign. Inconsistent with a theory based on profit-maximization, agents therefore appear to consider sunk fixed fees when making pricing decisions.

Column 3 in table 2.3 further shows that the rise in prices induces demand effects: the total number of customers served is reduced by around 12% as a result of the fixed fee treatment. As column 4 shows, the fixed fee treatment therefore only reduced collected revenue by 3%, which is statistically insignificant. Using those point estimates, back-of-the-envelope

[26] X_{wd} also contains an indicator for whether a ward was assigned to a cross-cutting treatment ('Price-Cap Treatment). Note that, as I find no evidence of interaction effects between the fixed fee and the cross-cutting treatment, I treat both as separate experiments. Given that the two treatments were assigned as part of a cross-cutting design, treatment effects of the intervention should therefore be interpreted conditional on 50% of the sample being assigned to the cross-cutting treatment.

calculations suggest that as long as the combined travel and vaccination costs per farmer exceed 366 Tanzanian Shillings (or approximately 20% of the average per vaccination price), agents' profits (before deducting the fixed fee) in the treatment group are higher than in the control group. This finding suggests that, instead of optimally setting prices, agents either confuse average with marginal costs or consider a heuristic that is affected by fixed fees when setting prices.[27]

2.5 Evidence from Lab-in-the-Field Experiment 1

The preceding section presented an example of a case where the introduction of a fixed fee affects prices. It is, however, clear that the existing evidence cannot conclusively show that this is, in fact, driven by agents' choices, as the design of the fixed fee treatment could have affected prices through a different channel. For example, agents might have incorrectly assumed limited liability, which relieved them from the responsibility of remitting the fixed fee in case of insufficient revenue collection. This implicit cut-off could then lead to bunching and similar price and quantity effects as observed, yet for completely different reasons. Alternatively, the obligation to pay a fixed fee might have implicitly altered agents' negotiation strategies in ways that cannot be captured through the farmer survey.

In addition, the existing evidence is insufficient to understand the underlying mechanisms and considerations that agents take into account. In particular, agents might choose to set lower prices than optimal in a profit maximizing sense, for example because they are intrinsically motivated. Alternatively, they might use simplified heuristics that only approximate optimal pricing or they might be confused about the appropriate costs to consider. To conclusively understand agents' pricing strategies, more direct evidence is needed. To provide this, I implement a series of lab-in-the-field experiments with a randomly selected subset of the field-experimental sample.

[27]One might be concerned that the fixed fee treatment did not only affect agents' costs but also strengthened their bargaining position and therefore allowed them to extract higher profits from service recipients. To avoid a direct impact on bargaining, the agents were not given any documentation that formally stated the requirement to remit a fixed fee. Agents in the treatment group also were not more likely to mention the need to cover vaccine costs during bargaining with farmers.

2.5.1 Sample and Implementation

I conduct the lab-in-the-field experiments with a total of 311 agents in 14 districts in Tanzania. The sampled districts represent a randomly selected subset of all districts covered in the field experiment. All respondents had also participated in the field experiment. Subjects were individually paired with one enumerator, moved to visually isolated locations for the implementation of the experiment and were given a game sheet to make their choices.[28] In order to ensure independence across participants, subjects did not interact with one another during the experiment and were not informed of other participants' choices.

2.5.2 Design

The lab-experimental design aims to simulate a situation that is similar to the service provision task that agents encounter during the field experiment. However, to incentivize profit maximization the instructions of the lab experiment specifically emphasized that the simulation focuses on a "private" task, therefore framing the exercise as a for-profit interaction with service recipients. Respondents are told that the activity will simulate the delivery of cattle castration services, which, in the Tanzanian context, is a non-public activity typically delivered by livestock extension agents.

As part of the simulation, respondents are informed that they are facing a situation in which they consider castrating the cattle of four different customers. Each customer is associated with an idiosyncratic delivery cost of 1,000, 4,000, 8,000 and 11,000 Experimental Shillings (ES), respectively. Respondents are told that the delivery cost covers fuel for the motorcycle, and that some customers live further away than others, which requires a higher expenditure on fuel. The simulation further assumes that the cost of applying the service once the delivery cost has been paid is zero.

Agents' task is to decide dynamically whether they want to serve each customer and, if so, at what price they are willing to do so. To this end, respondents are told that they should first make up their mind about what the smallest possible price is at which they would prefer to deliver the service to the customer instead of walking away from the business opportunity, starting with the first customer. They then "call" the customer, who offers them a payment. If the price offered is larger than the reservation price, respondents

[28]Figure 2.1 shows a translated version of the game sheet.

deliver the service at the offered price. If it is not, they move on to the next customer.

Given this design, participants' main choice regards their reservation price, i.e. the smallest price for which they would serve a customer instead of walking away from the deal. The profit maximizing reservation price for every customer is the marginal delivery cost, and thus 1,000 ES for the first, 4,000 ES for the second, 8,000 ES for the third, and 11,000 ES for the fourth customer.

The primary objective of this exercise is to elicit agents' perception about the marginal costs associated with delivering services to each customer. Thus, the primary outcome variable obtained from this simulation is the reservation price stated by respondents. I employ a multiple price list (MPL) mechanism to assure that the reservation price stated by agents is equivalent to their true reservation price.[29] As part of the elicitation procedure, respondents are shown ten hypothetical price offers between 1,000 and 10,000 ES for every customer. They are then asked to decide, independently for every offer, whether they would accept the price and pay the delivery cost, or not serve the customer at the offered price.[30] After making the ten choices, a piece of paper representing each choice is put in a bowl and agents draw one offer at random. The choice relating to this offer is then implemented. After the completion of the simulation, participants' earnings are summarized on an information sheet for them. Agents were then paid out 20% of their experimental earnings in cash after the experiment, which further incentivized profit maximization.[31]

This design assures that reporting the true reservation price is incentive compatible. On the one hand, if respondents were to state a higher reservation price, they would risk foregoing service provision opportunities at prices that are acceptable to them. On the other hand, stating a lower price than their true reservation price risks having to serve customers at prices that do not allow respondents to break even. In addition, due to the random nature of the price offer agents are unable to increase payments by stating higher reservation prices. The design of the mechanism thus builds on the same intuition as the

[29]The explanation protocol for the MPL mechanism is available from the author upon request.

[30]Note that, as respondents are only shown potential offers between 1,000 ES and 10,000 ES, it is not optimal to serve the fourth customer at any offered price.

[31]The average profit obtained was 10,271 ES.

widely employed Becker-DeGroot-Marschak mechanism (Becker, DeGroot, and Marschak, 1964).

2.5.3 Sources of Variation

The design of the lab experimental generates two sources of variation that, taken together, provide evidence on the pricing strategies employed by agents. First, I generate exogenous variation in fixed fees by randomly dividing participants in two groups. One group acted as a treatment group and was responsible for paying a fixed fee of 4,000 ES before commencing the simulation. The narrative of the design framed this cost as having to be paid to repair the seller's motorcycle before delivering services. It had to be remitted regardless of whether any customer was served or not, and is thus sunk at the time of decision making. Given this structure, the fixed fee does not affect the profit maximizing choice of reservation prices.

I contrast reservation prices in the treatment group to a control group without a fixed fee requirement. In particular, respondents assigned to the control group received exactly the same instructions as the treatment group but were not required to remit 4,000 ES before commencing the simulation. Every participant played two rounds of this experiment, one with and one without the fixed fee. The order of the two rounds was randomized.

The second source of variation that allows me to identify agents' pricing strategies is generated by the random nature of price offer draws. This generates exogenous variation on the extensive margin, i.e. in whether an agent delivers a service conditional on reservation prices, and on the intensive margin, i.e. in prices paid conditional on delivering a service. Taken together, the random offer draws exogenously vary agents' stock of revenue and average costs at the time of choosing a reservation price for the second, third and fourth customer. The sign of reservation prices' response to increases in price offers can then be combined with the sign of the response to the experimental variation in fixed fees to understand whether agents choose mark-ups over marginal or average costs, or respond to variation in fixed fees through their effect on profits.

2.5.4 Reduced Form Effects

I begin by investigating the reduced form effect of the fixed fee treatment on respondents' choice of indifference points. To this end, I estimate variants of the following model:

$$y_{ir} = \beta_0 + \beta_1 FixedFee_{ir} + \gamma_r + \alpha_i + \epsilon_{ir} \tag{2.11}$$

where y_{ir} denotes the outcome variable of interest for individual i in round r, $FixedFee_{ir}$ is a binary variable indicating whether agents had to remit a fixed fee in a specific round and γ_r denotes round fixed effects. In addition, as each respondent plays the simulation twice I can control for individual level fixed effects, denoted by α_i. The primary outcome variable of interest during the first part of the analysis will be respondents' chosen indifference point for each of the four customers. As randomization was performed at the individual level, I report robust standard errors clustered for every participant. Given the inclusion of round fixed-effects, the model employs within round-variation. As the allocation to the treatment was randomized within rounds, this specification therefore causally estimates its effect on the outcome variables of interest.

A complication arises because respondents did not only have to choose indifference points for every customer, but were initially required to decide whether they would theoretically be willing to serve customers for any possible offer between 1,000 and 10,000 ES. There is thus no data available for possible indifference points exceeding 10,000 ES, as in such cases respondents decided not to serve a given customer. I employ two different methods to address this. First, in most of the specifications I assume that agents who decided not to serve a given customer would have chosen an indifference point of 11,000 ES. As this is the smallest possible indifference point for which respondents would have chosen not to serve a given customer, any findings of a positive treatment effect thus represent a lower bound on the true effects. Second, in a subset of specifications I also estimate effects conditional on agreement to serve a given customer.

The results from the estimation of equation 2.11 directly replicate the effects found in the field experiment. Consider the reduced form treatment effects of the fixed fee treatment on indifference points first. Table 2.4 presents the results, with each column representing one of the four customers.

Note first that average reservation prices in the control group correlate with marginal costs, but fall below marginal costs for the third and fourth customer whose marginal costs are highest. The treatment significantly increases participants' reservation prices for all customers. The incidence of this increase falls primarily on the first and the second customer who have the lowest marginal cost. The indifference point for serving the first and second customer increase by 116% and 23% on average, respectively. Both increases are significant at the 1% level. The treatment also significantly increases reservation prices for the two customers with the highest marginal costs. The increases of 3.3% and 6.3% for customers 3 and 4, respectively, are, however, substantially more modest.

The responses to the fixed fee shown in table 2.4 are inconsistent with profit maximization. Instead, the results suggest that agents indeed base their choice of indifference point on a statistic different to marginal costs. To understand whether agents respond to variation in revenue, I use the variation generated by the random price-offer draws to estimate the effect of revenue collected from previous customers on the indifference point choices.

I begin by showing the reduced form impact of price offer draws on indifference points. Formally, I estimate the following model separately for customers 2, 3 and 4:

$$y_{ir} = \beta_0 + \beta_1 PreviousOffers_{ir} + \gamma_r + \alpha_i + \epsilon_{ir} \tag{2.12}$$

where y_{ir} again denotes the indifference point chosen by individual i in round r and $PreviousOffers_{ir}$ denotes the cumulative value of offers (irrespective of whether service was performed or not) before having reached the respective customer. In addition, γ_r and α_i denote round and respondent fixed effects, respectively. As previous price-offers were assigned randomly, this specification consistently estimates the causal effect of receiving a higher price offer on respondents' reservation price.

Table 2.5 presents the estimation results of equation 2.12. The results show that respondents lower their reservation prices for customers 2 and 3 in response to higher price offers, but do not adjust their indifference points for customer 4. Column 1 shows that increasing the price offered by the first customer by 1,000 ES reduces reservation prices by 3%. Similarly, column 2 shows that increasing the price offered either by the first or the second customer by 1,000 ES reduces reservation prices by 0.8%. Both of those effects are statistically significantly different from zero at the 1% level. In contrast, reservation prices for the fourth customer are not affected by changes to the price offers received. Indeed, column 3 shows that there is no statistically significant association between price offers and indifference points. Taken together, the results presented in table 2.5 provide evidence that agents lower their reservation prices in response to positive shocks to price offers.

2.5.5 Effects of Average Costs and Revenue

The evidence presented so far is insufficient to understand whether agents anchor on average costs or profits, as variation in price offers affects both of those statistics: On the one hand, higher offers increase the likelihood of providing services to a specific customer, which increases average costs. On the other hand, higher offers increase revenue conditional on providing services, which directly affects profits.

To identify the pricing strategies, I turn to estimating my preferred specification which jointly estimates the effect of average costs and revenue on profits. If agents confuse average and marginal costs, changes to revenue should not affect indifference points conditional on average costs. However, the exhibition in section 2.3 shows that if agents target a certain level of profits they will raise reservation prices in response to positive shocks to average costs, conditional on revenue. In addition, respondents will lower reservation prices in response to positive shocks to revenue, conditional on average costs.

Endogeneity concerns make using OLS regression unsuitable to estimate the relationship between indifference points and variation in average costs and revenue.[32] To address this endogeneity concern, I estimate a 2-stage-least-squares model. In the first stage, I project the explanatory variables of interest, average costs and revenue, on an indicator for the fixed fee treatment and a continuous variable for the cumulative price offers received. Formally, the first stage estimates the following two equations:

$$AverageCosts_{icr} = \mu_0 + \mu_1 FixedFee_{ir} + \mu_2 PreviousOffers_{icr} + \gamma_r + \alpha_i + \theta_c + \epsilon_{icr} \tag{2.13}$$

and

$$Revenue_{icr} = \mu_0 + \mu_1 FixedFee_{ir} + \mu_2 PreviousOffers_{icr} + \gamma_r + \alpha_i + \theta_c + \epsilon_{icr} \tag{2.14}$$

$AverageCosts_{icr}$ and $Revenue_{icr}$ denote the average costs and revenue accumulated when respondent i decides whether to serve customer c in round r. In addition, specifications 2.13 and 2.14 control for round, respondent and customer level fixed effects.

[32] A primary concern is reverse causality, as the respondents' choice of indifference points directly influences both average costs and revenue. The sign of the resulting bias will depend on the relationship between indifference point choices and average costs as well as revenue, respectively. As an exogenous increase in indifference points reduces the likelihood of offer acceptance, it causes a reduction in average costs and revenue. The OLS estimate is thus likely to be downward biased.

In the second stage, I employ the predicted values obtained from the first stage to estimate the following model:

$$y_{icr} = \beta_0 + \beta_1 Average\hat{C}osts_{icr} + \beta_2 Reve\hat{n}ue_{icr} + \gamma_r + \alpha_i + \theta_c + \epsilon_{ir} \tag{2.15}$$

where y_{icr} denotes respondent i's reservation price for customer c in round r. The parameters of interest are β_1 and β_2. The null hypothesis for marginal cost pricing versus average cost pricing is $\beta_1 = 0$ and $\beta_2 = 0$, whereas the null hypothesis for average cost pricing against profit targeting is $\beta_2 = 0$.

The fixed fee treatment and the random price offer draws thus act as instrumental variables for average costs and revenue. Given the random assignment of the instrumental variables, the underlying assumption for this specification to produce consistent estimates is that the instrumental variables do not affect respondents' reservation prices through any channel different from average costs and revenue. The structure of the experiment makes this unlikely. First, assignment to the fixed fee treatment directly increases average costs, while leaving all other conditions of the experiment unchanged compared to the control group. Second, random price offer draws affect revenue as well as the likelihood of service provision conditional on indifference points. Through its effect on service provision, random price offer draws thus affect revenue and average costs, but no other factors that might affect respondents' choices of indifference points. The empirical strategy thus produces estimates of the partial causal effect of average cost increases and revenue on prices. The proposed specification pools customers two to four and includes customer fixed-effects. Note that, due to the absence of previous profits, customer 1 is excluded from this analysis.

Table 2.6 presents the results of the 2-stage-least-squares estimation. Columns 1 and 2 show that, in addition to being valid instrumental variables, the indicator for the treatment and the price offer draws are also relevant instrumental variables. Column 1 shows that the fixed fee treatment increases agents' average costs by approximately 48%. Similarly, increasing the price offered by the previous customers by 1,000 ES increases average costs by 4%. At the same time, column 2 shows that the fixed fee treatment, through its effect on average costs, reduces revenue by 17%. Similarly, increasing the price offered by the previous customers by 1,000 ES increases revenue by 16%. Those results are statistically significant at the 1% level, and a test of joint significance of the instruments in the first stages produces a p-value smaller than 0.01. Taken together, those results suggest that the fixed fee treatment and the cumulative offers are relevant instrumental variables.

Turning next to the relationship of interest, column 3 first presents an OLS regression that captures the correlation between respondents' chosen indifference point and average costs as well as revenue. The results show that there is a statistically significant relation between both explanatory variables and reservation prices: a 1,000 ES increase in average costs is associated with a 292 ES, or 5% over the control group mean, increase in respondents' reservation price. Similarly, a 1,000 ES increase in revenue is associated with a 77 ES, or 1.4% decrease in reservation prices. Both of those statistics are significantly different from zero at a 1% level of confidence.

Column 4 then presents the results from the 2-stage-least squares estimation. The estimates show that a 1,000 ES increase in average costs causes agents to adjust their indifference point upwards by 351 ES. This equals an increase of 6% over the control group average. Similarly, a 1,000 ES increase in revenue causes agents to adjust their indifference point downwards by 70 ES, or by 1.3% when compared to the control group average.

Columns 5 and 6 present a robustness check to show that the results are not driven by the imputation of indifference points for customers who agents had decided not to serve. In particular, columns 5 and 6 present the same estimates as in columns 3 and 4, respectively, but exclude observations from customers who respondents had decided not to serve. The results are quantitatively and qualitatively similar to the results that include the imputation.

Two aspects are worth noting about the interpretation of the results. First, the 2SLS estimates are larger than the OLS estimates, which confirms the notion that the OLS estimates are biased downwards. Second, the results present direct evidence that agents indeed use rules-of-thumb when setting prices. Contrary to a model based on profit-maximizing marginal cost pricing, experimental participants adjust their reservation prices in response to changes in average costs, holding marginal costs constant. In addition, the results also present evidence to contradict the notion that agents confuse average and marginal costs. This is because indifference points do not only respond to average costs, but are also adjusted downwards in response to increases in revenue, holding average costs constant. The results presented are instead consistent with a model of profit targeting as presented in section 2.3: agents target an aggregate level of profits and adjust their reservation prices to reach that level. As such, agents will lower their indifference points in response to a positive shock to revenues. At the same time, increases in average costs will lead agents to increase their indifference points.

2.5.6 Effects on Profits

The preceding results have shown that respondents do not maximize profits but instead rely on simplified heuristics. It is possible that in such circumstances increases in costs can increase profits, as they nudge agents to adjust prices towards the profit-maximizing optimum. Table 2.7 presents results that investigate this possibility.

Columns 1 and 2 investigate how increasing average costs by introducing the fixed fee affects agents' propensity to serve the fourth customer. Recall that this customer is associated with a marginal delivery cost of 11,000 ES, which exceeds the highest possible price offer of 10,000 ES. As such, agreeing to serve the fourth customer is not a profitable decision. Yet, as column 1 shows, 76% of respondents in the control group agree to serve customer 4 for one of the prices offered and, as shown in column 2, 34% of respondents receive a price draw that eventually requires them to serve this customer. The results from the experiment provide evidence that, indeed, increasing fixed fees can assist agents in overcoming behavioral biases. In particular, column 1 shows that the fixed fee treatment reduces the likelihood of agreeing to serve the unprofitable fourth customer by 18.9 percentage points. Similarly, the proportion of all agents who end up serving the unprofitable customer is reduced by 11.6 percentage points as a result of the fixed fee treatment. Both results are statistically significant at the 1% level.

While it appears as if increasing fixed fees can nudge participants towards avoiding unprofitable behavior, table 2.7 also shows that the specific treatment employed in the experiment is unable to increase agents' profits. In particular, column 3 shows that increasing fixed fees reduces profits by 20%. Similarly, column 4 illustrates that the absolute value of deviations from the profit maximizing indifference point increases by 48% in response to the fixed fee treatment. Taken together, those results suggest that even though individual biases can be overcome by increasing fixed fees, variation in fixed fees affects average costs which can create further deviations in agents' pricing strategies.

2.6 Evidence from Lab-in-the-Field Experiment 2

To provide further evidence that respondents anchor their supply decisions on the stock of profits instead of choosing prices based on marginal costs, I design and implement an additional lab-in-the-field experiment with the sample described in section 2.5.1. This experiment investigates whether the behavior observed in the previous experiment is unique to prices, or whether agents' behavioral biases also carry through to other areas of economic decision making. To investigate this, I evaluate to what extent production quantity decisions are affected by changes in fixed fees and revenue. To this end, I implement an experiment in which respondents face increasing marginal costs and stochastic returns to production and must decide what aggregate quantity to produce.

2.6.1 Design

The experiment simulates the production decisions of a tailor who produces t-shirts. Production happens upon customer demand, and customers offer either 5,000, 10,000 or 15,000 Experimental Shillings (ES) for a t-shirt. Before accepting to produce an additional t-shirt, respondents are unaware of the actual price offer but only know that each of them can occur with equal probability. As part of this experiment, this is simulated by receiving random price offer draws.

The cost of producing a t-shirt is not constant but increases in the total amount produced. After every customer interaction, respondents decide whether to continue or stop. If they continue, respondents receive an additional price offer draw but are also obliged to remit the cost associated with producing an additional t-shirt. If they decide to stop, the simulation ends and earnings are calculated. The main outcome variable of interest is thus the total quantity of t-shirts produced by every respondent.

Respondents participate in two versions of this simulation and face a different cost schedule in each version. Table 2.8 shows the cost scheduled faced by agents in the first and second version of the simulation. Respondents play two iterations for every version of the game, once without the obligation to pay an initial fixed fee and once with the requirement to remit 10,000 ES before the game commences. Respondents thus play a total of four rounds of the game. Within each version, the order of the iterations is randomized. In addition, I also randomly assigned whether respondents commence the simulation with the first or second version of the cost schedule.

2.6.2 Sources of Variation

Similar to the first experiment, this design allows me to exploit two sources of exogenous variation to understand agents' considerations when making economic decisions. First, the random assignment of the two iterations within each version of the simulation induces exogenous variation in fixed fees. In particular, during the first round of the simulation the group that is assigned the fixed fee acts as a treatment group, whereas the group without a fixed fee acts as a control group. During the second round of the game, the treatment assignment changes, so that the initial control group then turns into the treatment group. Due to the random assignment of the order of iterations, the difference between the groups in every round of the simulation causally identifies the effect of increasing fixed fees on respondents' quantity decisions.

Second, the random price offer draws generate exogenous variation in revenue. For example, agents who draw a higher proportion of offers that equal 15,000 ES have, for a given quantity of production, higher revenue (and profits) than respondents who primarily received price offers of 5,000 ES. As price offers are drawn randomly by the respondent, revenue is orthogonal to respondent characteristics. Variation in revenue caused by price offer draws thus causally identifies the effect of increased revenue on quantity choices. Notice that, in contrast to simulation 1, the variation generated by the price offer draws only affects revenue and does not simultaneously raise respondents' average cost. The effect thus complements the findings from simulation 1 to causally identify the effect of shocks to revenue on quantity choices.

2.6.3 Adjustment to the Conceptual Framework

The notion of the agent targeting a certain level of profits, and thus sticking to rules-of-thumb, can easily be extended to the stopping game played in simulation 2. The task here is to dynamically choose a level of quantities when prices are uncertain. In particular, respondents are given a schedule of (increasing) production costs and are told that customers will offer one of three possible prices for the product, each with equal probability. Respondents' task is then to decide when to stop accepting new price offers and finish production. I will again investigate the three different scenarios presented previously: Rational profit maximization, confusion of average and marginal costs, and reliance on simplified heuristics.

The set-up of the slightly amended model is as follows. Suppose agents faces a cost schedule $C(Q) + K$ with $C'(Q) > 0$ and $C''(Q) > 0$, where Q denotes the aggregate quantity of output produced and K denotes a fixed cost that is sunk at the time of pricing. Agents face a continuum of customers, each of whom demands exactly one unit, and decide dynamically whether to produce another unit or not. The marginal cost of producing one unit of output for customer i is denoted by c_i. Suppose further that the agent is uncertain about the prices offered by potential customers, but has a (rational) expectation about receiving a price of $E(p)$ from customer i. Denote by R_{i-1} the revenue collected before reaching customer i and by Q_i the aggregate quantity produced after having served customer i. Finally, suppose for simplicity that agents are risk neutral. If the agent decides to serve customer i, they receive a payoff of:

$$\pi_i^S = E(p) + R_{i-1} - C(Q_i) - K$$

If, instead, they decide not to serve to serve the customer, they receive:

$$\pi_i^{NS} = R_{i-1} - C(Q_{i-1}) - K$$

Consider first the choices made by an agent who is rationally maximizing profits. It is straightforward to see that such agents serve a given customer as long as the expected payoff exceeds the marginal cost of serving an additional customer:

$$E(p) > C(Q_i) - C(Q_{i-1})$$

The total quantity produced will then be such that:

$$C(Q_i) - C(Q_{i-1}) = c_i = E(p) \tag{2.16}$$

It is straightforward to see that the comparative statics for the perfectly rational agent are then:

$$\frac{\partial Q}{\partial c_i} < 0, \frac{\partial Q}{\partial K} = 0, \frac{\partial Q}{\partial R} = 0 \tag{2.17}$$

Suppose now that agents confuse average and marginal costs. In particular, suppose they perceive marginal costs for customer i to be:

$$\widetilde{c_i} = \omega c_i + (1 - \omega) \frac{C(Q_i) + K}{Q_i}$$

where ω denotes an arbitrary weight between 0 and 1 which measures the extent of agents' confusion between average and marginal costs. The total quantity produced will then be such that:

$$\widetilde{c}_i = \omega c_i + (1-\omega)\frac{C(Q_i)+K}{Q_i} = E(p) \tag{2.18}$$

The comparative statics for the confused agent are then:

$$\frac{\partial Q}{\partial c_i} < 0, \frac{\partial Q}{\partial K} < 0, \frac{\partial Q}{\partial R} = 0 \tag{2.19}$$

Finally, consider an agent who, instead of maximizing profits, chooses quantities based on a simplified heuristic. Suppose the agent's heuristic is, as before, a simple stopping rule which is anchored on profits. In particular, suppose the agent's rule implies to continue producing until profits drop below a customer-specific emergency level of π_i^T. Formally, when deciding whether to serve customer i, agents consider whether their current stock of profits exceeds the emergency amount. They thus serve customer i if:

$$R_{i-1} - C(Q_{i-1}) - K > \pi_i^T$$

Denote by ε_j price-draw deviations from the expected value for customer j and denote their cumulative distribution function by $F(\varepsilon)$. The price received from customer j is $p_j = E(p) + \varepsilon_j$. The ex-ante likelihood of serving customer i is then:

$$\begin{aligned}
Prob(Q_i) &= Prob\left(p_1 - c_1 + R_0 - C(Q_0) - K > \pi_1^T, ..., p_{i-1} - c_{i-1} + R_{i-2} - C(Q_{i-2}) - K > \pi_{i-1}^T\right) \\
&= Prob\left(\varepsilon_1 > \pi_1^T + K + c_1 - E(p), ..., \varepsilon_{i-1} > \pi_{i-1}^T - R_{i-2} + C(Q_{i-2}) + K + c_{i-1} - E(p)\right) \\
&= \prod_{j=1}^{i-1} 1 - F\left(\pi_j^T - R_{j-1} - C(Q_{j-1}) + K + c_j - E(p)\right)
\end{aligned}$$

where R_0 and $C(Q_0)$ equal 0 by definition.

To derive the comparative statics of interest, consider what happens if fixed cost and revenues increase exogenously. First, an exogenous increase in K reduced the probability that any of the $i-1$ draws of ε exceed the target profits which, in turn, reduces the probability of reaching the i^{th} customer. Similarly, an exogenous increase to revenues increases the probability of exceeding the target level ceteris paribus. The comparative statics are thus:

$$\frac{\partial Q}{\partial c_i} < 0, \frac{\partial Q}{\partial K} < 0, \frac{\partial Q}{\partial R} > 0 \tag{2.20}$$

Taken together, this section proposes another test to differentiate between the three possible decision strategies employed by agents. First, if agents rationally maximize profits, their quantity choices are not affected by changes to revenue or fixed costs. In contrast, if agents confuse average and marginal costs, then an exogenous shock to fixed costs decreases quantities, but changes to revenue leave quantities unaffected. Finally, if, instead of solving a maximization problem, agents use rules-of-thumb to determine their quantity choices, quantities respond to increases in fixed costs as well as changes to revenues.

2.6.4 Estimation and Results

I begin the analysis by evaluating how the fixed fee treatment affects respondents' quantity choices. To this end, I estimate the following model separately for version 1 and version 2 of the simulation:

$$Q_{ir} = \beta_0 + \beta_1 FixedFee_{ir} + \gamma_r + \alpha_i + \epsilon_{ir} \tag{2.21}$$

Q_{ir} denotes the aggregate quantity produced by individual i in round r, $FixedFee_{ir}$ is a binary variable indicating whether agents had to remit a fixed fee, γ_r denotes round fixed effects and α_i again denotes individual level fixed effects. As the treatment was randomly assigned, the parameter β_1 causally identifies its effect on respondents' quantity choices.

Table 2.9 presents the results of the estimation of equation 2.21. Column 1 in table 2.9 shows that while agents in the control group produce on average 5.14 units of output in the first version of the game, this figured is reduced by 23% because of the fixed fee treatment. Similarly, in version 2 of the game agents produce on average 4 units of output in the control group, which is reduced by 17% in response to the fixed fee treatment. This again provides evidence of the fact that agents do not make choices on the margin. If they did, changes to fixed fees should not affect their production decisions. Instead, they appear to be anchoring their production decisions on a heuristic that is affected by variation in average costs.

As suggested by the theoretical exposition in section 2.6.3, distinguishing between profit targeting and a confusion between average and marginal costs requires understanding whether agents respond to variation in revenue conditional on average costs or not. A lack of a response to revenue is consistent with confusion between average and marginal costs, whereas a negative response to revenue is consistent with anchoring quantity decisions on target profits. Estimating an OLS regression that regresses quantities on an indicator for the fixed fee treatment and a variable capturing revenue earned is challenging because quantity choices in themselves affect revenue positively, thus causing reverse causality and a possible positive bias of the OLS estimates.

To address this, I employ an instrumental variable strategy that leverages the relationship between random price offer draws and revenue. I first estimate the following first stage regression:

$$Revenue_{ir} = \mu_0 + \mu_1 AverageDraw_{ir} + \gamma_r + \alpha_i + \epsilon_{ir} \tag{2.22}$$

$Revenue_{ir}$ denotes the aggregate revenue collected by individual i in round r, $AverageDraw_{ir}$ is a continuous variable that measures the average price draw received by individual i in round r, and γ_r as well as α_i are round and individual level fixed effects, respectively. I then estimate the following second stage equation:

$$Q_{ir} = \beta_0 + \beta_1 FixedFee_{ir} + \beta_2 Rev\hat{e}nue_{ir} + \gamma_r + \alpha_i + \epsilon_{ir} \tag{2.23}$$

where $Rev\hat{e}nue_{ir}$ denotes the fitted value obtained from the first stage regression.

Given this estimation approach, the value of the average price offer draw obtained by individual i acts as an instrumental variable for revenue. This approach will provide consistent estimates of the causal effect of revenues on quantity choices if the validity condition on the instrumental variable holds. On the one hand, price offers are drawn at random, and any variation in price offers received is thus by definition exogenous to the model. On the other hand, validity also requires that there is no channel other than revenues through which price offer draws affect quantities. Given the design and controlled environment of the lab-in-the-field experiment, this is likely to be satisfied. In particular, price offer draws directly and saliently affect the respondent's revenue but no other relevant statistics.

The average price offer draws are also relevant instrumental variables: The coefficient on the offer draws is statistically significant at the 1% level of significance in the first stage regressions of both versions of the game, leading to the first stage regressions providing significant explanatory power

Columns 3 and 4 in table 2.9 present the results of the instrumental variable estimation. Column 3 shows the results for the first version of the game and column 4 shows the results for the second version. Consistent with the results presented previously, increasing fixed fees by 10,000 ES reduces the aggregate quantity produced by participants by 16% and 14%, respectively. However, in addition to average costs, quantity choices are also affected by revenue. In particular, an increase of 1,000 ES in revenue increases the quantity produced by 0.5% and 0.4%, respectively. This provides direct evidence that instead of confusing average and marginal costs, participants are anchoring their production decisions on profits and thus use rules-of-thumb when making economic choices.

2.7 Conclusion

This paper investigates whether agents treat fixed fees as sunk, and thus evaluates the practical effectiveness of fixed fee contracts. I first present evidence from a field experiment to show that, in contrast to classical economic theory, livestock extension agents' pricing choices respond to the introduction of a fixed fee component in their contracts. I then implement a series of lab-in-the-field experiments to estimate the response of prices to changes in marginal costs, average costs and revenue, which allows me to identify their decisional heuristics. I find that instead of optimally choosing prices as mark-ups over marginal costs, agents rely on rules-of-thumb that anchor their pricing decisions on their aggregate stock of profits. Taken together, the results of this study show that while fixed fee contracts are popular in practice, they do not necessarily maximize incentives.

Beyond its implication for contractual design in the public sector, this paper highlights that, for a fairly generalizable context in a developing country, when faced with complex tasks, economic agents do not make optimal decisions. This finding implies that traditionally non-distortionary policy instruments can become suboptimal. In particular, fixed fees, taxes and transfers are a central component of economic policies aimed at redistributing economic surplus.[33] Such policy instruments are attractive because they allow for redistribution without inducing behavioral distortions if individuals and firms behave optimally. The results of my paper suggest, however, that non-marginal policy instruments can affect choice variables in practice, thus

[33] For example, most income tax systems contain tax credits that are independent of taxpayers' income and tax liability, such as the Child Tax Credit in the United States. The solution to optimal taxation problems as proposed by Mirrlees (1971) typically combines a lump-sum subsidy with a progressive marginal tax rate (Mankiw, Weinzierl, and Yagan, 2009). In addition to redistribute policy contexts, two-part tariffs regularly contain a flat charge in addition to usage-based prices to extract surplus (see, for example, Tirole (1988)).

suggesting that central implications of the theoretical optimal tax literature are not applicable to a real-world context in which agents rely on rules-of-thumb.

More broadly, the findings of this paper suggests that small business growth in developing countries might be constrained by sub-optimal pricing that prevents entrepreneurs from maximizing profits and obtaining the liquidity necessary to invest in growing their business. It therefore provides evidence on managerial capacity constraints in developing countries and suggests potentially high returns to business training programs that improve entrepreneurs pricing choices.[34]

[34] See, for example, Bloom and Van Reenen (2007), Bloom, Eifert, et al. (2013), McKenzie and Woodruff (2015) and Bruhn, Karlan, and Schoar (2017).

Tables

Table 2.1: Summary of Theoretical Predictions (Simulation 1)

	Profit Maximization	Confusion	Rules-of-thumb 1	Rules-of-thumb 2	Rules-of-thumb 3
$\frac{\partial p_i}{\partial \tau}$	> 0	> 0	> 0	$\lesseqgtr 0$	$\lesseqgtr 0$
$\frac{\partial p_i}{\partial C_{i-1}}$	$= 0$	> 0	> 0	> 0	< 0
$\frac{\partial p_i}{\partial R_{i-1}}$	$= 0$	$= 0$	< 0	< 0	> 0

Table 2.2: Treatment Groups (Field Experiment)

Fixed Fee	
No	**Yes**
277 wards (422 agents)	273 wards (410 agents)

Table 2.3: Fixed Fee Treatment Effects on Prices, Quantities and Revenue

Outcome:	**Price**	**Participation**	**# Farmers Served**	**Revenue**
	(1)	(2)	(3)	(4)
Fixed Fee Treatment	7.155**	-0.00405	-5.915*	-2,167
	(3.106)	(0.0267)	(3.551)	(6,526)
Observations	679	832	832	832
District Fixed Effects	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Control Mean	66.16	0.82	50.54	70’701.75
Control St. Dev.	36.03	0.38	56.54	90’531.76

Notes: Standard Errors are clustered at the ward level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Column 1 presents coefficient estimates of a regression of the average price per chicken on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Column 2 presents coefficient estimates of a regression of the total number of farmers served on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Columns 3 presents coefficient estimates of a regression of the revenue collected on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables.

Table 2.4: Reduced Form Effect of Fixed Fee on Indifference Points (Simulation 1)

Indifference Point for:	**Customer 1**	**Customer 2**	**Customer 3**	**Customer 4**
	(1)	(2)	(3)	(4)
Fixed Fee Treatment	2,551***	1,013***	243.9**	508.5***
	(98.38)	(115.4)	(117.9)	(121.9)
Observations	622	622	622	622
Respondent Fixed Effects	Yes	Yes	Yes	Yes
Round Fixed Effects	Yes	Yes	Yes	Yes
Control Mean	2199	4479	7302	8100
Control St. Dev.	1282	1763	2744	3044

Notes: Standard Errors are clustered at the individual level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Columns 1 to 4 present coefficient estimates of a regression of the individual indifference points chosen for customers 1 to 4, respectively, on a treatment indicator and round fixed effects. All specifications control for respondent level fixed effects.

Table 2.5: Reduced Form Effect of Random Price Draws on Indifference Points (Simulation 1)

Indifference Point for:	**Customer 2**	**Customer 3**	**Customer 4**
	(1)	(2)	(3)
Cumulative Previous Offers	-0.145***	-0.0557***	0.000101
	(0.0361)	(0.0204)	(0.0209)
Observations	622	622	622
Respondent Fixed Effects	Yes	Yes	Yes
Round Fixed Effects	Yes	Yes	Yes
Mean of Ind. Variable	4986	7424	8355
St. Dev. of Ind. Variable	2263	2793	3117

Notes: Standard Errors are clustered at the individual level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Columns 1 to 4 present coefficient estimates of a regression of the individual indifference points chosen for customers 1 to 4, respectively, on a continuous variable for the cumulative price draws received before reaching the respective customer and round fixed effects. All specifications control for respondent level fixed effects.

Table 2.6: Effect of Average Costs and Revenue on Indifference Points (Simulation 1)

Outcome:	**Average Costs**	**Revenue**	**Indifference Points**			
	(1)	(2)	(3)	(4)	(5)	(6)
Fixed Fee Treatment	1,373***	-1,229***				
	(34.15)	(156.0)				
Cumulative Offers	0.109***	1.176***				
	(0.00557)	(0.0231)				
Average Costs			0.292***	0.351***	0.184***	0.198***
			(0.0470)	(0.0590)	(0.0478)	(0.0558)
Revenue			-0.0769***	-0.0702***	-0.0618***	-0.0527***
			(0.0106)	(0.0119)	(0.0114)	(0.0121)
Observations	1,866	1,866	1,866	1,866	1,595	1,595
Estimation Method	OLS	OLS	OLS	2SLS	OLS	2SLS
Conditional on Agreement	No	No	No	No	Yes	Yes
Respondent Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Customer Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Round Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Control Mean	2863	7395	5520	5520	5113	5113
Control St. Dev.	1378	6804	3296	3296	3045	3045

Notes: Standard Errors are clustered at the individual level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Columns 1 and 2 present first stage regressions. In particular, column 1 presents coefficient estimates of an OLS regression of the average costs faced by the agent on a treatment indicator for the fixed-cost treatment and a continuous variable for the cumulative price draws received before reaching the respective customer as well as respondent, customer and round fixed effects. Column 2 presents coefficient estimates of an OLS regression of the revenue collected by the agent on a treatment indicator for the fixed-cost treatment and a continuous variable for the cumulative price draws received before reaching the respective customer as well as respondent, customer and round fixed effects. Columns 3 and 5 present OLS estimates of the relationship of interest. In particular, columns 3 and 5 present coefficient estimates of an OLS regression of the respondent's stated indifference point on a continuous variable for the average cost before serving the respective customer and a continuous variable for the cumulative revenue before reaching the respective customer as well as respondent, customer and round fixed effects. Column 3 employs all data, whereas column 5 restricts its attention to responses for which the respondent agreed to serve the customer. Columns 4 and 6 present 2-Stage-Least-Squares estimates of the relationship of interest. In particular, columns 4 and 6 presents coefficient estimates of an 2SLS regression of the respondent's stated indifference point on a continuous variable for the average cost before serving the respective customer and a continuous variable for the cumulative revenue before reaching the respective customer as well as respondent, customer and round fixed effects. Average costs and revenue are instrumented for using a treatment indicator for the fixed-cost treatment and a continuous variable for the cumulative price draws received before reaching the respective customer. Column 4 employs all data, whereas column 6 restricts its attention to responses for which the respondent agreed to serve the customer.

Table 2.7: Effect on Profits (Simulation 1)

Outcome:	**Agreement**	**Served**	**Profit**	$\|mc - p\|$
	(1)	(2)	(3)	(4)
Fixed Fee Treatment	-0.189***	-0.116***	-1,226***	1,069***
	(0.0248)	(0.0302)	(373.9)	(59.12)
Observations	622	622	622	2,488
Customer	4	4	All	All
Respondent Fixed Effects	Yes	Yes	Yes	Yes
Round Fixed Effects	Yes	Yes	Yes	Yes
Customer Fixed Effects	No	No	No	Yes
Control Mean	0.76	0.34	6190	2248
Control St. Dev.	0.43	0.48	4843	2375

Notes: Standard Errors are clustered at the individual level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Column 1 presents coefficient estimates of an OLS regression of a binary variable indicating whether a respondent agreed to serve customer 4 for any of the prices offered on a treatment indicator and round fixed effects. Column 2 presents coefficient estimates of an OLS regression of a binary variable indicating whether customer 4 was actually served in the simulation on a treatment indicator and round fixed effects. Column 3 presents coefficient estimates of an OLS regression of respondents' total profit per round on a treatment indicator and round fixed effects. Column 4 presents coefficient estimates of an OLS regression of the absolute distance between respondents' chosen indifference point and the customer's marginal costs on a treatment indicator as well as customer and round fixed effects. All specifications control for respondent level fixed effects.

Table 2.8: Cost Schedule (Simulation 2)

Quantity	**Total Cost**	
	Version 1	**Version 2**
1	5,000	2,000
2	10,000	6,000
3	15,000	12,000
4	20,000	20,000
5	25,000	30,000
6	40,000	42,000
7	55,000	56,000
8	70,000	72,000
9	90,000	90,000
10	110,000	110,000
11	130,000	132,000
12	150,000	156,000
13	170,000	182,000
14	190,000	210,000
15	210,000	240,000

Table 2.9: Effect of Costs and Revenues on Quantities Produced (Simulation 2)

Outcome:	**Quantity Produced**			
	(1)	(2)	(3)	(4)
Fixed Fee Treatment	-1.163***	-0.665***	-0.814***	-0.558***
	(0.157)	(0.109)	(0.136)	(0.0916)
Revenue x 1000			0.0269***	0.0169**
			(0.00843)	(0.00839)
Observations	622	622	622	622
Simulation Variety	1	2	1	2
Estimation Method	OLS	OLS	2SLS	2SLS
Respondent Fixed Effects	Yes	Yes	Yes	Yes
Round Fixed Effects	Yes	Yes	Yes	Yes
Control Mean	5.14	4.00	5.14	4.00
Control St. Dev.	2.49	1.75	2.49	1.75
First-Stage F-Stat.	N.A.	N.A.	47.24	27.67

Notes: Standard Errors are clustered at the individual level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Columns 1 and 2 present coefficient estimates of an OLS regression of the quantity chosen to produce by the agent on a treatment indicator for the fixed-cost treatment as well as respondent and round fixed effects. Columns 3 and 4 present coefficient estimates of an 2SLS regression of the respondent's chosen quantity on a treatment indicator for the fixed-cost treatment and a continuous variable for the aggregate revenue collected as well as respondent and round fixed effects. Revenue is instrumented for using a continuous variable for the average price offer draw received. The coefficient on revenue is scaled by a factor 1000 and should thus be interpreted as 1000-Shilling increases in Revenue.

Figures

Figure 2.1: Information Sheet for Participants of the Lab-in-the-Field Experiment

Total cost if everyone is served: 24'000

Cumulative Revenue	Cumulative Cost
0	0

Customer Name	Served yes or no?	Price Offered
Mark		
Geoffrey		
Peter		
Karl		

Chapter 3

Priming the Pump - Irrigation, Nutrition and Agricultural Productivity

Attempts to improve agricultural production technology are a common response to undernutrition. In India, groundwater irrigation using tube wells has long been promoted as a means to reduce rainfall-dependence and enhance food security. The merits of adopting tube wells have been debated widely, however, with opponents fearing a deprivation of smaller farmers and impoverishment of rural laborers. To evaluate the causal effects of tube well adoption on nutrition, I employ an instrumental variable framework that exploits variation in land suitability for deep groundwater irrigation caused by differences in hydrogeological structures. I find that groundwater irrigation significantly improves nutrition across the income spectrum: a one standard deviation increase in the proportion of cropped area irrigated with tube wells increases calorie intake by 770 to 915 calories per day. In addition, groundwater irrigation generates positive spillovers on the calorie intake of urban populations and households not employed in agriculture. I present additional evidence which suggests that these effects are driven by increases in agricultural productivity that reduce staple prices and raise wage rates. The findings thus highlight the value of groundwater irrigation in fighting undernutrition and promoting agricultural development.

3.1 Introduction

Undernutrition remains a global challenge. India in particular has been faced with concerns about hunger for decades. Having experienced droughts and a series of severe food crises after independence, undernutrition remains a problem today. For example, 42.5% of children under the age of five were underweight in 2005, compared to only 21% in Africa, with figures even higher in rural areas.[1] Understanding which factors underlie nutritional status, and which policy tools can be used to reduce undernutrition, is thus of utmost importance.

A key response to undernutrition has traditionally been the implementation of policies aimed at increasing food production and raising agricultural productivity (Food and Agricultural Organization of the United Nations, 2003). To overcome the food crises after Indian independence, governments and donors invested in a package of modern agricultural production inputs coupled with investments in research and infrastructure, which fundamentally transformed agriculture (Spielman and Pandya-Lorch, 2009). A central component of this "Green Revolution" was the promotion of modern irrigation technology, which led to a 500% increase in the area irrigated from groundwater between 1960 and 2010 (Garduño and Foster, 2010). The hope was that this technology would reduce the rain-dependence of agricultural production, increase food security and ultimately improve nutrition.

The merits of adopting groundwater irrigation have been debated widely, with opponents fearing a deprivation of smallholders and impoverishment of agricultural laborers when benefits from adoption accrue exclusively to large landowners. With high fixed costs of up to $3000 per well, coupled with the absence of significant public subsidies, small-scale farmers are unlikely to invest in groundwater irrigation (Sekhri (2011), Postel (1999)).[2] It has been feared that this would result in a competitive advantage for richer farmers while depressing output prices and depressing wages for small-scale producers and agricultural laborers (Patniak, 2004). Fears are fuelled by the observation that India appears to suffer from an "agricultural disconnect", where improvements to agricultural productivity are generally

[1] See Deaton and Drèze (2009) for a detailed discussion of how nutrition and calorie intake in India has evolved over the last thirty years. A report by the World Bank (2006) as well as one by the Economist (2015) provides an overview of undernutrition in Indian.

[2] The development of water markets does not imply that small-scale farmers are excluded from groundwater. See, for example, Banerji, Meenakshi, and Khanna (2012)

not accompanied by improved nutritional status of populations (Gillespie, Harris, and Kadiyala, 2012).

Given these countervailing arguments, understanding whether groundwater irrigation indeed improves nutrition across the income spectrum is ultimately an empirical question. Yet, despite a long-lasting debate only limited systematic evidence exists. By focusing on understanding how the adoption of tube wells, the primary type of well used to extract groundwater in India, affects calorie intake, this paper aims to fill this gap.

Studying tube well adoption to investigate potential trade-offs associated with groundwater irrigation is beneficial for two reasons. First, tube wells are the primary type of well used to extract deep groundwater for irrigation purposes in India, with 92% of groundwater extracted through wells used for irrigation (Sekhri, 2014). Second, tube wells penetrate deep underground aquifers, so that their construction is associated with significant drilling and other adoption costs. Third, drilling tube wells is a primarily private activity in India (Sekhri, 2011). Studying tube wells is thus likely to capture the key benefits from groundwater irrigation as well as the potential threats from increased inequality.

One methodological challenge that arises when identifying the causal effect of tube wells on calorie intake is that adoption is not random. Comparing outcomes between adopting and non-adopting regions is therefore unlikely to yield consistent estimates, as adoption decisions are likely to be correlated with other determinants of nutrition, such as agricultural productivity.[3] To overcome this challenge, I propose an empirical strategy that employs variation in hydrogeological structures as a source of exogenous variation in tube well adoption. In particular, I exploit the fact, well documented in the hydrogeological literature on well construction, that the adoption of tube wells in areas that are serviced by aquifers with a hard rock shell is substantially more costly than the adoption in alluvial or sedimentary areas. This heterogeneity in costs generates differential introduction patterns of tube wells between hard and soft rock areas, which I leverage in my empirical strategy. In particular, I focus my analysis on the administrative unit below an Indian state, a district, and exploit the variation generated by the time-varying effect of hard rock aquifers on tube well adoption. This allows me to

[3]Similarly, the efficiency wage literature suggests that nutrition influences productivity and hence wages, which can create a reverse causality between calorie intake and tube well adoption. See Dasgupta and Ray (1987).

use district and time fixed effects that control for level differences in outcomes between districts and years.

My analysis proceeds in three steps. I first show that while there has been a substantial increase in the share of land irrigated using tube wells, the adoption of this technology has primarily occurred in areas where the aquifer was not surrounded by hard rock material: while the proportion of land irrigated using tube wells in soft rock areas had increased by 12.4 percentage points between 1987 and 2004, adoption in hard rock areas only increased by 5 percentage points over the same time period. Moreover, the negative effect of hard rock structures on tube well adoption has increased over time, with hard rock areas having a 7 percentage points smaller share in land irrigated from tube wells in the 1980s, compared to a 16 percentage points difference in the 2000s. This suggests that the time-varying effect of hydrogeological formations does indeed provide a source of variation in tube well adoption.

My main set of results then investigates the effect of tube well adoption on a district's median calorie intake. I find strong evidence that groundwater irrigation improves nutrition: a one standard deviation increase in the fraction of area sown irrigated using tube wells causes an increase in median calorie intake of between 770 and 915 calories per day. This effect is stable across the within-district calorie intake distribution. In particular, I find comparable results when considering the effect of tube well adoption on the 10th to 40th percentile of the within-district calorie intake distribution. While the effect of tube well adoption on calorie intake is largest for households whose main income earner is self-employed in agriculture, I also find that tube well adoption generates spillovers to sectors not directly affected by agricultural productivity improvements: median calorie intake for urban populations and rural households not employed in agriculture rises by 486 and 596 calories, respectively, in response to a one standard deviation increase in the share of area sown irrigated using tube wells. The observed increase in calorie intake in response to tube well adoption is accompanied by a change to expenditure patterns. I find suggestive evidence that groundwater irrigation relaxes calorie constraints and thus allow households to switch to a diet that is richer in more expensive but less nutritious food items. As such, tube well adoption causes households to spend a larger share of their food budget on animal products, while reducing their expenditure on vegetables.

The results presented previously show that improvements to agricultural technology generate inclusive enhancements to the nutritional status of both richer and poorer, as well as urban and rural populations. To understand the mechanisms underlying these results, I investigate how groundwater irrigation affects agricultural productivity as well as food and labor markets. I first show that tube well adoption induces significant direct and indirect improvements to agricultural productivity: rice and wheat yields increase by 0.43 and 1.2 tons per hectare in response to a one standard deviation increase in the share of area sown irrigated using tube wells. The direct productivity enhancement through improved irrigation is complemented by an increased adoption of high-yielding-variety seeds and investments in agricultural machinery.

The improvements to agricultural productivity generate general equilibrium effects on labor and food markets that can account for the observed effects of groundwater irrigation on calorie intake. On the one hand, tube well adoption raises the wage rate of male agricultural workers. The apparent complementarity between groundwater irrigation technology and human labor provides a potential explanation for the observed strong nutrition improvements for households in the agricultural sector. At the same time, spillovers to sectors not directly affected by tube well adoption can occur as groundwater irrigation lowers prices of staple foods: a one standard deviation increase in the share of area sown irrigated using tube wells decreases wheat prices by 21% compared to the sample average. The analysis therefore suggests that a possible driver of the observed increase in calorie intake, and the observed spillovers, is that tube wells increase agricultural productivity. This in turn augments labor input and decreases market prices for staple foods, which makes calories more affordable and leads to improvements in nutritional outcomes.

This paper provides evidence on the relationship between groundwater irrigation and nutrition.[4] By doing this, it complements a small but growing literature that evaluates improvements to irrigation systems.[5] Duflo and

[4] Additional work has also considered the role of water for uses other than nutrition. For example, Rud (2012) exploits that groundwater availability encouraged rural electrification in India. Devoto et al. (2012) show that connecting private dwellings to water mains in Morocco can induce significant quality of life improvements.

[5] In a more general sense, this paper also relates to a literature that investigates the effects of agricultural productivity improvements. For example, Andersen, P. S. Jensen, and Skovsgaard (2016) investigate the effect of adopting the heavy plow on agricultural development in Medieval Europe. Bardhan and Mookherjee (2011) investigate the effect of providing subsidized seeds and fertilizer on farm productivity in West Bengal.

Pande (2007) study the effect of dam construction in India and find that while dams enhance agricultural productivity in downstream districts, they increase poverty in the districts where they are located. Hornbeck and Keskin (2014) investigate the dynamic effect of groundwater irrigation in the southwestern United States and find that while initially irrigation reduced drought sensitivity, it also induced a shift towards more rainfall-sensitive crops.[6] Bardhan, Mookherjee, and Kumar (2012) find that investments in minor irrigation in West Bengal led to growth in farm productivity. The article closest related to this paper is by Sekhri (2014), who shows that groundwater irrigation can reduce rural poverty headcounts.[7] I complement this literature by providing systematic nationally-representative evidence for tube well adoption in India that focuses on nutrition as the key outcome considered when designing agricultural development programs.

More broadly, this paper complements a literature that investigates the determinants of nutrition choices. R. T. Jensen and Miller (2008) provide evidence for Giffen behavior which is driven by subsistence constraints.[8] Atkin (2013) considers how food consumption habits affect the relation between food prices and nutrition. Dubois, Griffith, and Nevo (2014) structurally estimate a demand system for food and find that prices can explain a substantial share of the difference in calorie intake between Europe and the United States. Eli and Li (2017) investigate how calorie requirements affect food consumption choices and calorie intake. The existing literature thus focuses primarily on the determinants of calorie *demand*. I complement this literature by highlighting how factors relating to food *supply* can shape calorie intake and households' food consumption baskets.[9]

The remainder of this article is structured as follows. Section 3.2 provides additional details on groundwater irrigation and tube wells in India. Section

[6]Related, Fishman (2012) finds that irrigation can reduce the impact of rainfall volatility on yields in India.

[7]A more distantly related literature investigates the optimal management of groundwater resources. Sekhri (2011) evaluates whether public or private provision of wells is more effective in conserving groundwater, and finds that public wells can crowd out private wells, which leads to reduced water wastage. See also Banerji, Meenakshi, and Khanna (2012).

[8]Calories are not only needed to meet subsistence constraints but can also enhance labor productivity. Schofield (2014) estimates that an additional calorie intake of 700 calories per day can increase earnings by around 10%. Relatedly, a classical literature has estimated Calorie-Engel curves to understand how calorie intake varies with income. See Subramanian and Deaton (1996).

[9]A related literature investigates policy interventions aimed at enhancing nutrition. Hoynes and Schanzenbach (2009) evaluate the consumption impact of the Food Stamp Program in the United States, and Afridi (2010) evaluates a mandated school meal program in India.

3.3 explains the identification strategy employed in this paper and section 3.4 describes the construction of the data set used in this analysis. Section 3.5 discusses and interprets the results, and section 3.7 concludes.

3.2 Groundwater Irrigation and Tube Wells

Since the 1960s, agriculture in India has undergone an unprecedented transformation. Driven by recurring droughts that caused food crises and famine, a coalition between the Ford Foundation and the Indian government initiated a process that combined the development of improved seed varieties with an expansion of fertilizer use and investments in irrigation infrastructure (International Food Policy Research Institute, 2002). The swift adoption process of new high-yielding-varieties (HYV) seeds and the subsequent rise in agricultural productivity is often referred to as the "Green Revolution". The revolution's effect on agricultural productivity was significant: simulations conducted by the Consultative Group on International Agricultural Research suggest that in the absence of this development food production would have been 20% lower than today's levels (Pingali, 2012).

A central objective of India's agricultural development agenda was to reduce the reliance of agricultural production on monsoon-related rainfall by replacing rainfed agriculture with agricultural production irrigated from groundwater sources (Sekhri, 2011). This was successful: the total area irrigated using groundwater sources increased by approximately 500% between 1960 and 2010 (Garduño and Foster, 2010). Today, around 60% of Indian agriculture is irrigated from groundwater sources (Sekhri, 2014). While gravity irrigation was the main technology used to replace rainfed agriculture until the 1970s, groundwater irrigation is more prevalent today (Sekhri, 2011).[10]

Most groundwater irrigation systems rely on so called tube wells, the number of which has increased from approximately 1 million in 1960 to 19 million in 2000 (International Water Management Institute, 2002). Tube wells consist of a steel tube that is bored into a deep underground aquifer and an electric or diesel-fuelled engine that powers a pump responsible for lifting the water from the aquifers.[11] Compared to other forms of irrigation, especially tank, canal or dug well irrigation, tube wells provide a richer and

[10] Gravity irrigation employs water from tanks or canals.
[11] Aquifers are rock strata which hold and transmit groundwater.

more reliable source of water as the water table levels in the extracted aquifers are higher and less sensitive to weather conditions.

Tube wells were, however, not adopted uniformly across India. As I show in section 3.5, a central factor that determines the installation of groundwater irrigation devices is the rock structure surrounding the aquifer. Adoption costs are lowest in alluvial areas, such as the Indo-Gangetic plain, which account for the largest proportion of area irrigated using tube wells. In contrast, adoption costs are significantly higher in hard rock and mountainous areas. Section 3.3 outlines how I leverage this variation to obtain a causal identification of the effect of tube well adoption on calorie intake.

Despite the seeming success of the "Green Revolution", critics have pointed to its adverse side effects. In particular, it has been argued that the benefits of the increase in irrigation, mechanization and productivity would accrue primarily to richer farmers. The "Green Revolution" has often been linked with putting small producers and agricultural laborers at risk, as lower output and higher input prices reduce earnings from agricultural production, and mechanization reduces wages and employment for agricultural laborers (International Food Policy Research Institute, 2002). Some have even linked the "Green Revolution" to erupting social conflict (Unger, 2014). In this paper I provide systematic evidence to assess the inclusiveness of nutrition improvements induced by the "Green Revolution".

3.3 Identification

Causal identification of the effect of tube well adoption on calorie intake relies on a source of exogenous variation in tube well adoption. This is because OLS estimates that relate tube well adoption to calorie intake might result in spurious correlations, either because of omitted variables or reverse causality. For example, adoption might be more prevalent in areas more dependent on agriculture and less economically developed where calorie intake is lower, leading to a spurious negative correlation between tube well adoption and calorie intake. Similarly, exogenous positive shocks to food supply can reduce incentives to enhance food security through the adoption of tube wells, leading to (negative) reverse causality. In both cases, OLS estimates would understate the true effect of tube well adoption on calorie intake.

In order to causally identify the effect of tube well adoption on calorie intake, I employ differences in hydrogeological structures as a source of exogenous variation. I begin by classifying the rock structure of the aquifers underlying districts in India into two categories: hard rock and soft rock areas. In particular, I define a district as being covered by a hard rock aquifer if the rock structure is predominantly igneous, i.e. consisting of basalt, intrusives, charnokite, granite or banded-gneissic complex (BGC) rocks. In contrast, districts whose aquifers are predominantly metamorphic or sedimentary are defined as soft rock. This classification follows standard definitions in geology.[12]

I then exploit the heterogeneous adoption patterns of tube wells across the two types of aquifers. Since tube wells penetrate deep aquifers, the structure of the rock formation around the aquifer is an important determinant of adoption costs. As such, the adoption of tube wells in areas covered by hard rock aquifers is substantially more costly than the adoption in soft rock areas. This heterogeneity in costs generates differential introduction patterns of tube wells between hard and soft rock areas.

To exploit this variation, the estimation proceeds in two steps. I first use a difference-in-difference approach to predict the effect of hard rock aquifers on tube well adoption patterns over time. In particular, my first stage regression projects the share of total area sown that is irrigated from tube wells in district d and time period t, $TubewellShare_{dt}$, on an indicator variable for whether a district is covered by hard rock aquifers, $HardRock_d$, interacted with an indicator for time period t:

$$TubewellShare_{dt} = \alpha_0 + \sum_{s=2}^{T} \alpha_s \left(HardRock_d \times l_{st} \right) + \gamma_t + \mu_d + \nu_{dt} \tag{3.1}$$

where γ_t and μ_d denote time and district fixed effects and l_{st} denotes a dummy for year s. Under the assumption that tube well adoption had evolved similarly in hard and soft rock areas if they had similar aquifer structures, the coefficient α_t provides an estimate of the causal effect of being located in a hard rock area on tube well adoption in time period t.

This first stage regression provides predicted values of tube well adoption based on heterogeneous adoption caused by differential aquifer structures. In the second stage of the estimation procedure, I employ the projected values from the first stage to estimate the causal effect of tube well adoption on calorie intake and various other outcome measures:

$$y_{dt} = \beta_0 + \beta_1 Tubew\hat{el}lShare_{dt} + \gamma_t + \mu_d + \varepsilon_{dt} \tag{3.2}$$

[12] See, for example, Park (2010).

y_{dt} measures the outcome of interest in district d at time t, $Tubew\hat{e}llShare_{dt}$ denotes the predicted value of the share of area sown irrigated using tube wells obtained through the estimation of equation 3.1.

Similar to a difference-in-difference approach, the identifying assumption necessary to interpret β_1 as measuring the causal effect of tube well adoption is that outcome variables would have evolved similarly in hard and soft rock areas if aquifer structures were similar. While the inclusion of district level fixed effects allows me to control for level differences in outcomes between districts, my identifying assumptions is thus that outcome trends between hard and soft rock areas are comparable with the exception of the effect of aquifers on tube well adoption.

Two caveats to this identification strategy exist. First, the possibility for ground water irrigation may trigger other developments that affect nutrition. For example, Rud (2012) links ground water availability to rural electrification. As I am unable to control for such effects, my estimates should be understood as reduced form in the sense that they combine the direct effect of tube well adoption as well as the indirect effects it triggers. Second, differential pre-trends between hard and soft rock areas that are either exogenous or driven by a spurious third factor would yield inconsistent estimates. In section 3.6 I show that while my point estimates are qualitatively robust to controlling for state specific trends, the associated reduction in identifying variation reduces the precision of my estimates and makes them insignificant. Thus, I cannot conclusively rule out the possibility that trends in outcome variables prior to my data period determine their future development and thus affect my estimates.

The main outcome variables I consider measure median calorie intake and household expenditure patterns per district. In addition, I also investigate how tube well adoption affects agricultural productivity, prices, wages and the distribution of farm sizes. The construction of those variables, as well as the data needed to implement the described analysis, is laid out in the next section.

3.4 Data

To carry out the analysis described in the previous section I combine a number of different data sources. First, detailed micro-level consumption data is required to obtain measures of calorie intake. I obtain this information from the Indian National Sample Survey. Second, identification requires information on hydrogeological structures, which I obtain by digitizing and geoprocessing aquifer maps provided by the Indian Central Ground Water Board. Third, I obtain information on tube well adoption from the Indian Agricultural Statistics. Finally, in order to understand the mechanisms driving the observed effect of tube well adoption on calorie intake, I employ the ICRISAT data set to investigate how tube well adoption affects yields, prices and the farm size distribution.

Merging district level data sets in India can be complicated because of the creation and carving out of new districts, as well as the separation of existing districts. As part of this project, I employ the universe of districts in 2011, as provided by the 2011 census, as a reference list. I then match all districts from other data sets and earlier years to this list that have the same geographic boundaries as the district provided in the 2011 list. Finally, if districts have been merged I combine the information from earlier years to reflect one data point in the 2011 list.

3.4.1 Calorie Intake Data

To measure calorie intake, I employ the Indian National Sample Survey (NSS). The NSS survey is an annually conducted large-scale cross-sectional household expenditure and labor survey. For the present analysis I employ data from 3 rounds of the NSS, spanning the years 1987, 1999, 2004. The data used correspond to three of the quinquennial "thick" rounds of the NSS, namely the 43rd, 55th, and 61st round. As opposed to the annual NSS surveys, the "thick" rounds cover the whole of India, have a substantially larger sample and a more extensive questionnaire. The 43rd, 55th and 61st round have 126,910, 120,309, and 110,521 observations, respectively.[13]

The NSS data contains detailed information on a large number of household expenditure items, including a detailed breakdown of food items. As food purchase quantities are recorded at a detailed level, this data can be converted into a measure of calorie expenditure using the conversion factors provided by Gopalan, Shastri, and Balasubramanian (2004). To convert

[13] Despite the large sample size, the sampling procedures employed by the National Sample Survey Organization do not necessarily guarantee its representativeness at the district level, which is my main unit of analysis. Note, however, that the objective of this paper is not to accurately estimate district-wise trends, but rather to identify the causal relationship between district level outcomes and tube well adoption. As such, as long as differences between the observed district-wise statistics and its true district level counterpart is uncorrelated to differential tube well adoptions patterns caused by aquifer structures, the lack of representativeness at the district level does not threaten my statistical analysis and identification strategy.

the expenditure data into calorie intake data, the main outcome variable considered in this paper, two adjustments are necessary. First, the expenditure data needs to be transformed to take into account the calorie content of meals provided to guests and meals taken outside of the household. As the NSS surveys contain information on the number of meals given to guests and taken outside of the house, as well as the average number of meals per day, I assume a constant calorie content of every meal to adjust the expenditure figures for this. Second, expenditure data is collected at the household level. In order to arrive at a measure of calorie intake at the per-capital level, an assumption about the intra-household allocation of calories is required. I assume that household members receive a fixed share of the total calories purchased and calculate a weighted average within the household to arrive at a per-capita figure.[14]

After converting the expenditure figures into calorie intake data, I collapse the data at the district level. The main outcome variable I consider measures the median calorie intake per district. To trace the effect of tube well adoption for different wealth groups, I also measure the impact of tube well adoption on various district level percentiles of the calorie intake distribution. In addition, I calculate the district level median calorie intake for various sub-populations depending on whether respondents live in urban or rural areas, and their sectoral employment. Collapsing the data at the district level allows me to merge the NSS data set with three other data sets that are crucial for this analysis.

3.4.2 Hydrogeological Data

To causally identify the effect of tube well adoption on calorie intake, I require data on the hydrogeological structure of aquifer systems in India. To this end, I digitize and geoprocess aquifer maps provided by the Indian Central Ground Water Board. In 2012, the Government of India's Ministry of Water Resources published an atlas of the Indian aquifer system. This atlas was produced as part of an on-going effort to improve the management of deep groundwater resources and to prevent conflict over rapidly depleting water sources for irrigation (Central Ground Water Board, 2012). As part of this publication, the Central Ground Water Board published various maps

[14] I assume that children between the ages 0 and 3 consume 15%, children between the ages 4 and 7 consume 33%, children between the ages 8 and 11 consume 50%, and children between the ages 12 and 15 consume 75% of an adult equivalent in calories

that highlight the location of major aquifers and classify the aquifer systems underlying the different parts of India.

The central map I use for this paper shows the distribution of principal aquifer systems across India and is shown in figure 3.1. This map provides color-coded information on 14 principal aquifer systems.[15] To obtain information on the principal aquifer system underlying a given district, I merge the information from the map with a geocoded administrative map of India that contains information on cities and districts based on the 2011 census.[16] I first integrate the map in figure 3.1 into a coordinate system by employing information on city locations as control points to georeference the map. I then digitize the color-coded information on aquifer systems using a supervised maximum likelihood classification routine in which I drill the training sample to the legend provided in figure 3.1. Finally, I collapse the information on hydrogeological structures at the district level by assigning each district to the aquifer systems which is most prevalent in it.

The described routine produces the map output shown in figure 3.2. As can be seen from the map, there is substantial variation in hydrogeological structures across India. Aquifers are predominantly alluvial in northern India's Indo-Gangetic plain. In contrast, basalt aquifers are more prevalent in western and central India, especially in Maharashtra. Using the information in figure 3.2, I can calculate the aforementioned indicator variable for hard rock areas as taking the value 1 if the majority of the district is underlain by basalt, intrusives, charnokite, granite or banded-gneissic complex (BGC) rocks, and 0 otherwise.

3.4.3 Tube Well and Supplementary Data

To measure tube well adoption, I employ two separate data sources. First, I obtain data from the Indian Agricultural Statistics. This data is reported by district administrations and collected annually by the Directorate of Economics and Statistics of the Indian Ministry of Agriculture and Farmers Welfare. It provides annual information on the total cropped area for each district. In addition, the data provides source-wise information on the irrigation area for each district and year as well as explicit information on cultivated area and cropping patterns. The sources of irrigation in the data

[15]These include alluvial, laterite, basalt, sandstone, shale, limestone, granite, schist, quartzite, charnockite, khondalite, gneiss, intrusives and bgc.

[16]The geocoded map on districts is sourced from the Global Administrative Areas data set, developed by Hijmans (2009). The geocoded map on Indian cities is sourced from the Natural Earth Data Set.

are broken down into private and public canals as well as between dug wells and tube wells. Using this data, I calculate the right-hand-side variable of interest as the proportion of the net area sown that is irrigated by tube wells. I use the total area sown in the initial data period as a base to eliminate the endogeneity of the total area sown.

While the advantage of the Indian Agricultural Statistics Data is that it covers all of India, it provides limited information on yields, prices and other outcomes related to agricultural and rural development. To obtain information on this, I also merge my data with the ICRISAT meso data set (International Crops Research Institute for the Semi-Arid Tropics, 2017). This data set provides, among others, information on productivity, irrigation, land use, wages, prices and input use. In addition, it also contains information from the decennial censuses. The data covers 19 Indian states and up to 305 districts as well as the time period between 1966 and today.

Three challenges arise when merging the ICRISAT data with my existing data set. First, the ICRISAT data only covers a subset of Indian districts. In particular, while my original full data set contains complete information for around 329 districts per year, the ICRISAT data only contains such information for 238 districts. To address concerns that the ICRISAT sample of districts might be selected, I provide evidence for my main specifications that the results obtained with the full data set are robust to only focusing on the sample of districts available in the ICRISAT data set. Second, the ICRISAT data covers a longer time series than is available for the NSS data. For the main analysis that employs the ICRISAT data, I opt to use its full time series, spanning the years 1970 to 2009. I provide a robustness check to show that the productivity results are robust to limiting my attention to the years 1987 to 2004 which are available in my constructed sample. Third, while the ICRISAT sample technically spans the years from 1966 until today, the data has gaps. As such, data availability varies between different data series and between districts. As part of this analysis, power considerations lead me to not balance the sample and instead rely on the full sample of district-year pairs that are available in the different series of the ICRISAT data set. This does, however, result in variation in sample size between the different specifications.

Taken together, I use two different data sets and samples for my analysis. First, my main results, in which I relate tube well adoption to calorie intake, are based on the "constructed sample", which combines NSS data with the hydrogeological data and the data from the Indian Agricultural Statistics. This data set covers the years 1987, 1999 and 2004 and the full sample of districts in India. Second, for the agricultural productivity results I employ a merged data set that combines the ICRISAT data with my hydrogeological data. This data set covers the years 1970 to 2009 and a subset of Indian districts. I will refer to this as the "ICRISAT sample". The robustness checks presented in section 3.6 suggest that results from the two samples can be compared.

3.4.4 Summary Statistics

Table 3.1 presents summary statistics of the data sets. Panel A provides information on the constructed sample which combines NSS, aquifer and India Agricultural Statistics data. Panel B provides information on the sample available in the ICRISAT Meso Data. The table shows that while on aggregate 19% of area sown was irrigated using tube wells over the sample period of the constructed sample, this differs substantially between soft and hard rock areas. In particular, 26% of the total area sown was irrigated using tube wells in soft rock areas and only 5% of area sown was irrigated using tube wells in hard rock areas. 30% of the sample is classified as a hard rock area, whereas 70% of districts are underlain by soft rock aquifers. Panel A also provides information on the calorie consumption and food expenditure. The average median per capital calorie intake per district in the sample is 2707 calories per day. In addition, individuals spend 58% of their total expenditure on food.

Panel B shows that the prevalence of tube well irrigation is slightly lower in the ICRISAT sample compared to my constructed sample, with only 13% of area sown irrigated using tube wells in the ICRISAT sample. While this can be partially explained by a higher proportion of hard rock districts in the ICRISAT sample (41% compared to 30% in the constructed sample), the ICRISAT data also goes back to 1970 when tube well adoption was very low in both hard and soft rock areas. Panel B further shows that agricultural productivity is substantially higher in soft than in hard rock areas: Soft rock areas have a higher adoption of high-yielding variety seeds, a larger number of tractors and higher rice as well as wheat yields. This provides a first indication that the aquifer structure affects irrigation, which in turn affects agricultural productivity and nutrition. The next section provides causally identified econometric evidence to further emphasize this point.

3.5 Results

This section presents the estimation results using the empirical methodology outlined in section 3.3. The presentation of the results proceeds in four steps. I first investigate how variation in aquifer structures affects the adoption patterns of tube wells in subsection 3.5.1. Subsections 3.5.2 and 3.5.3 then use the variation generated by aquifer structures to instrument for tube well adoption and investigate its effect on calorie intake and food expenditure patterns. To investigate the mechanisms underlying this result, subsection 3.5.4 investigates through which channels tube well adoption affects agricultural productivity and rural development.

3.5.1 First Stage Results

I begin the evaluation by showing that variation in aquifer structures indeed leads to heterogeneous adoption patterns of tube wells. For illustration purposes, consider figure 3.3 first. The figure plots the coefficients α_1 to α_T from equation 3.1 as well as 95% and 99% confidence intervals, using ICRISAT data. The coefficients measure the differential adoption of tube wells between hard and soft rock areas relative to 1970 for every year from 1971 to 2009, controlling for district fixed effects. Confidence intervals are constructed using robust standard errors clustered at the district level. Two results are worth noting. First, hydrogeological structures do indeed cause heterogeneous adoption patterns between hard and soft rock areas. For example, the share of area sown irrigated using tube wells was approximately 17 percentage points lower in hard rock than in soft rock areas in 2009. Second, hydrogeological structures affect tube well adoption differentially over time. This is crucial, as it provides the identifying variation employed in this paper. As such, the figure shows that tube well adoption is not only lower in hard rock areas than in soft rock areas in 2009, but the difference in the area irrigated using tube wells between hard and soft rock areas has increased by approximately 16 percentage points compared to 1970.

Table 3.2 presents the first stage regression results which employ variation in hydrogeological structures to predict tube well adoption. Panel A shows the estimates of heterogeneous adoption patterns obtained when using data from the three years available in my constructed sample, namely 1987, 1999 and 2004. Panel B uses the full ICRISAT sample which contains district-year observations from 1970 to 2009. Throughout this analysis I report robust standard errors which are clustered at the district level.

Column 1 in table 3.2 first shows that tube well adoption in soft rock areas has increased significantly since 1987. Compared to 1987, the share of area sown irrigated using tube wells in soft rock areas increased by 10 percentage point in 1999 and by 12.4 percentage points in 2004. Compared to 1987, all increases are statistically significant at the 1% level. In contrast, tube well adoption in hard rock areas had only increased by 4.23 percentage points in 1999 and 5.08 percentage points in 2004, which is significantly smaller from a statistical perspective than the increase in soft rock areas over the same time period. The first stage regression presented in column 1 has an F-Statistic of 37.55, highlighting that the variation in rock structures has significant explanatory power for the adoption of tube wells.

Column 2 in table 3.2 shows that this result is robust to limiting the sample of districts to those available in the ICRISAT data set. In particular, while limiting the sample districts to those available in the ICRISAT data set reduces the sample size by approximately 28%, the coefficients from the estimation of equation 3.1 are qualitatively and quantitatively similar to those obtained when employing the constructed sample of districts.

The results presented in panel B in table 3.2 were obtained from a similar analysis as those presented in panel A, but employ longer panel data from the ICRISAT data set. As opposed to the results presented in panel A, which employ information for 1987, 1999 and 2004, results presented in panel B were obtained from annual district level data spanning the years 1970 to 2009. Column 3 in table 3.2 presents coefficient estimates obtained from the following equation:

$$TubewellShare_{det} = \alpha_0 + \sum_{e=2}^{E} \alpha_e \left(HardRock_d \times l_e \right) + \gamma_t + \mu_d + \nu_{det} \tag{3.3}$$

where l_e denote decade dummies. As opposed to equation 3.1, equation 3.3 thus allows the effect of hard rock aquifers on tube wells to vary by decades and not by year of observation.

Column 3 in table 3.2 again shows, using the longer time period available through the ICRISAT data set, that aquifer structures caused heterogeneous adoption patterns of tube wells. In particular, throughout the 1980s the total area sown irrigated using tube wells was 7.13 percentage points lower in hard than in soft rock areas. This difference had increased to 9.8 percentage points and 15.8 percentage points in the 1990s and the 2000s, respectively. With an F-Statistic of 37.90, variation in the effect of aquifer structures explains a statistically significant share of the variance in tube well adoption patterns. Taken together, the results presented in panel A and panel B suggest that the rock structure of the aquifer does indeed cause differential tube well adoption patters which provides identifying variation to estimate the effect of tube well adoption on calorie intake.

3.5.2 Effect of Tube Well Adoption on Calorie Intake

The previous subsection has shown that variation in hydrogeological structures affects tube well adoptions patterns. I now employ this source of variation to investigate how the adoption of tube wells affects calorie intake. Table 3.3 presents the results from the estimation of equation 3.2, using median calorie intake per district as an outcome variable. All specifications are estimated using the constructed sample and thus cover the years 1987, 1999 and 2004. All specifications contain year and district level fixed effects. I report robust standard errors which are clustered at the district level.

Consider first columns 1 and 2 in table 3.3, which show the result from an OLS regression of a district's median calorie intake on the fraction of land irrigated with tube wells. Column 1 employs the full set of districts, whereas column 2 restricts attention to the districts available in the ICRISAT data set. The results show that on average a one standard deviation increase in the fraction of area sown irrigated with tube wells (0.27) is associated with a 197 to 224 calories increase in calorie intake. This is equivalent to a 7% to 8% increase over the sample average. This result is statistically significant at the 1% level and thus documents a correlation between tube well adoption and calorie intake.[17]

Omitted variables and reverse causality make a causal interpretation of the results in columns 1 and 2 challenging. In particular, reverse causality is likely to bias OLS estimates downwards if exogenous positive shocks to food supply reduce the need to enhance food security through the adoption of improved agricultural technology. Furthermore, if tube wells are more likely to be adopted in economically less developed areas, which simultaneously also have lower calorie intake, OLS is again likely to understate the true effect of tube well adoption on calorie intake.

To address this concern, I now employ the variation generated by hydrogeological structures to estimate the causal effect of tube well adoption on calorie intake. Columns 3 and 4 present the results of the corresponding two-stage-least-squares regressions. Column 3 employs the full district sample, whereas column 4 restricts the sample to only those districts available

[17] Table 3.3 also documents what is often referred to as the "India Calorie Puzzle", which states that median calorie intake has declined significantly across India since the 1980s (Deaton and Drèze, 2009).

in the ICRISAT data set. The results show that tube well adoption does indeed cause calorie intake to increase: a one standard deviation increase in the fraction of area sown irrigated using tube wells causes an increase in calorie intake of between 770 and 915 calories. This corresponds to a 28% to 34% increase over the sample average. These results are different from zero at a 1% level of statistical significance. Furthermore, the comparison between column 3 and 4 shows that restricting attention to the sample of districts available in the ICRISAT data set leads to results that are comparable to the output from the constructed sample of districts. The results presented in column 3 plus 4 thus present direct evidence that adopting deep groundwater irrigation technology can improve nutrition. Furthermore, the results are consistent with the notion that OLS estimates are downward biased.

The results presented in table 3.3 show that tube well adoption has increased median calorie intake per district. Tube well adoption does, however, not necessarily affect calorie intake uniformly across the within-district distribution. To understand its distributional effects, and to obtain evidence on the mechanisms through which tube wells increase nutrition, it is thus instructive to investigate the effect of tube well adoption on the calorie intake of various sub-populations. The results presented in table 3.4 therefore break down the effect of tube well adoption on median calorie intake of urban versus rural populations, and between various rural occupational groups.[18] Two results presented in table 3.4 are worth noting. First, tube well adoption has increased calorie intake significantly for both rural and urban populations. On the one hand, the estimates thus show that improvements to agricultural technology generate significant spillovers to urban areas, even though they do not benefit directly from the technology. On the other hand, however, the magnitudes of the point estimates confirm that the benefits of tube well adoption accrue primarily to rural populations who benefit directly from the technology: the comparison between columns 1 and 2 reveals that the effect of tube well adoption is approximately 42%, or 202 calories in response to a one standard deviation increase, larger in rural than in urban areas.

The second result worth noting in table 3.4 is presented in columns 3 to 5. The results show that while tube well adoption generates positive spillovers for occupational groups who do not directly benefit from it, the primary

[18] In order to estimate the effect on urban and rural populations for a comparable sample, some districts that contained only urban or rural populations were dropped from this analysis. This leads to a slightly smaller sample size when compared to table 3.3.

benefits of adoption accrue to independent farmers. In particular, column 3 in table 3.4 shows that a one standard deviation increase in the share of area sown irrigated using tube wells leads to an increase of 596 calories among rural populations not employed in agriculture. In contrast, column 5 shows that the comparable effect on the median calorie intake of those self-employed in agriculture is 731 calories per day. In addition, I find no evidence tube well adoption has impoverished agricultural laborers by substituting labor from workers to machines. In contrast, column 4 in table 3.4 shows that tube well adoption has significantly increased the calorie intake of agricultural laborers to an extent that is comparable to other rural populations. Taken together, tube well adoption thus generated significant spillovers while still primarily improving the nutritional status of rural farmers.

To further investigate the distributional implications of tube well adoptions, table 3.5 investigates whether its effect differs by households' prosperity status. To this end I estimate how tube well adoptions affects different percentiles of the within-district calorie distribution. Column 1 in table 3.5 shows that a one standard deviation increase in the share of area sown irrigated using tube wells increases the calorie intake of households at the 10th percentile of the within-district calorie intake distribution by 688 calories. Columns 2 to 4 in table 3.5 further show that while the point estimates increase slightly for higher percentiles of the calorie intake distribution, consistent with the notion that the tube well adoption has produced larger benefits for wealthier households, the difference between estimates is not statistically significant. I thus find no evidence that tube well adoption has lead to impoverishment and increased nutritional inequality within districts.

Column 5 in table 3.5 shows how tube well adoption affects average calorie intake per district.[19] Consistent with the results presented previously, column 5 shows that a one standard deviation increase in the share of land irrigated using tube wells increases average calorie intake per district by approximately 791 calories.

3.5.3 Effect of Tube Well Adoption on Expenditure Patterns

The results presented previously have shown that tube well adoption increases calorie intake. To further characterize households' behavioral response, I now investigate the associated changes in food expenditure patterns. For example, it is possible that households increase their expenditure share

[19] To reduce noise in the measures of calorie intake, the data series that was used to calculate average calorie intake per district was winsorized at the 95th percentile.

on staples in response to tube well adoption. This could arise in a situation where households are deprived of calories and thus benefit from tube well adoption because it allows them to increase their expenditure on key food items so that they can meet their calorie requirements. Alternatively, tube well adoption could relax calorie constraints and thus allow households to switch to a diet that is richer in more expensive but less nutritious food items. For example, tube wells can reduce manual labor effort and thus negatively affect the need for calorie intake, which allows households to increase their expenditure share on more luxurious food products.

I find supportive evidence for the latter mechanisms. Table 3.6 shows how tube well adoption affects the share of food in total expenditure, and the distribution of food expenditure between various food categories. Column 1 shows that tube well adoption significantly increases food expenditure's share in total expenditure: a one standard deviation increase in the share of area sown that is irrigated using tube wells leads households to increase their food expenditure share by approximately 16 percentage points, which corresponds to a 28% increase when compared to the sample average. The estimate is statistically significant at the 5% level. Columns 2 to 7 show that this increase in expenditure share is accompanied by a shift in food consumption priorities. In particular, column 4 shows that the share of total expenditure on animal products increases significantly. In contrast, households reduce their expenditure share on vegetables and sugar-based food products in response to tube well adoption (see columns 5 and 7). This suggests that tube well adoption causes a change to expenditure patterns by allowing households to shift consumption towards more luxurious food sources, such as animal products.

3.5.4 Mechanism

The previous section has shown that ground water irrigation can have positive effects on nutrition by increasing calorie intake across the income spectrum and allowing households to shift their expenditure towards animal food products. To understand the mechanisms underlying this development, this section investigates how tube well adoption affects agricultural production practices, prices, wages and land usage. To this end, I employ the ICRISAT meso data set, therefore extending the study period to 1970 to 2009, and reducing the district coverage to the subset available through the ICRISAT sample.[20]

Agricultural Production

Consider first table 3.7, which employs ICRISAT data to investigate how tube well adoption affects agricultural productivity and production practices. All specifications use the constructed sample of districts available in the ICRISAT data set and, if available, employ annual data from 1970 to 2009. Every regression controls for year and district fixed effects and standard errors are again clustered at the district level.

Columns 1 and 2 show that tube well adoption significantly increases agricultural productivity, especially for wheat: a one standard deviation in the share of area sown irrigated using tube wells leads to an increase of yields to the magnitude of 0.43 tons per hectare for rice, and 1.2 tons per hectare for wheat. Both results are statistically significant at the 1% level. Table 3.7 further shows that this productivity increase combines the direct effect of tube well adoption through improved irrigation with knock-on effects through changes in agricultural production practices induced by tube well adoption. In particular, columns 3 and 4 highlight that tube well adoption caused the adoption of high-yielding-variety (HYV) seeds. HYV seeds have been developed as part of the Indian Green Revolution since the 1960s and have higher yields, but also higher irrigation and fertilizer requirements compared to normal seeds. The results suggest that a one standard deviation increase in the share of area sown irrigated using tube wells increases the total area sown using HYV seeds by 83,900 hectares for rice and 128,925 hectares for wheat. This exceeds the increases in total area sown induced by tube well adoption as shown in columns 5 and 6 of table 3.7. Tube well adoption has thus induced farmers to shift from traditional seeds to HYV seeds, which in turn increased agricultural productivity.

In addition to HYV seeds, the improvements to irrigation associated with tube well adoption can also act as a complement to improve the productivity of other agricultural investments. In particular, tube well adoption can increase farmers' return on investments in agricultural machinery, thus leading to increased agricultural mechanization. To investigate this possibility, I estimate how tube well adoption affects the total number of tractors and power tillers, two common agricultural mechanization devices,

[20]In the appendix I show that the results are robust to limiting the sample to the years 1987 to 2004.

in a given district. Columns 7 and 8 in table 3.7 show that a one standard deviation increase in the share of area sown irrigated using tube wells leads to a district-wide increase of approximately 12,840 tractors and 1,380 power tillers.[21] This suggests that tube well adoption did indeed increase agricultural mechanization. As tractors and power tillers increase yields independently of tube wells, this further complements their positive direct impact on productivity.

Wages and Prices

The previous section has shown that tube well adoption increases agricultural productivity. To understand the mechanisms underlying the effect of tube well adoption on calorie intake, it is crucial to understand which general equilibrium effects they induce. First, general equilibrium effects induced by tube wells can affect calorie intake through households' income. This effect can be either positive or negative, depending on whether tube wells complement or substitute human labor. Second, tube wells might affect calorie intake through food prices. As shown in the previous section, tube well adoption significantly alters agricultural production practices which can affect prices of staple food products. To evaluate whether possible general equilibrium effects of tube well adoption exists, I now investigate how tube well adoption affects wages and prices.

Consider first the results presented in columns 1 and 2 in table 3.8, which investigate the effect of tube well adoption on wages, separated by gender. The wage data is obtained from the ICRISAT data set and measures hourly field labor wages in nominal Indian rupees. Two results are worth noting. First, tube wells appear to complement human labor input into agricultural production: a one standard deviation increase in the share of area sown irrigated using tube wells increases hourly male and female wages by approximately 27 and 16.2 rupees (approx. \$0.42 and \$0.25), respectively. Second, the wage effect is substantially larger for male wages than for female wages. This is consistent with the notion that tube wells complement labor input if men are more likely to work in agricultural activities that are in direct contact with, or affected by, groundwater irrigation. The wage results thus support the theory that the observed increase in calorie intake is partially

[21]The sample size is smaller for the results on mechanization than for the data on yields and area sown because information on the number of tractors and the number of power tillers in India is only available for an subset of years (and for a subset of districts for some of the years) through the ICRISAT data set. The years for which data is available are 1972, 1982, 1987, 1992 and 2003.

driven by general equilibrium effects, as tube well adoption raises workers' wages and thus enhances their ability to purchase calories. Furthermore, in contrast to arguments brought forward by critics of the "Green Revolution", I find no evidence that improvements to agricultural productivity reduces wages and thus impoverished agricultural laborers.

I now turn to investigating how tube well adoption affects equilibrium prices of food staples. To this end, I employ information on farm harvest prices from the ICRISAT data set.[22] Farm harvest prices measure the price at which farmers sell their output to traders on the village market in units of Indian Rupees per 100 kg of output.[23] As part of this evaluation, I focus on investigating how tube well adoption affects the equilibrium prices of India's two primary staples: rice and wheat. Columns 3 and 4 in table 3.8 show that staple prices decrease in response to tube well adoption. While the point estimate for rice prices is negative, it is not statistically significant. In contrast, wheat prices decrease significantly in response to tube well adoption. This result is consistent with the notion that increased agricultural productivity caused by tube well adoption, reduced output prices. In particular, the results presented in table 3.7 show that wheat yields respond significantly more than rice yields in to tube well adoption. This in turn implies a more negative price response, which is consistent with the results presented here. This analysis therefore suggests that a possible driver of the observed increase in calorie intake is that tube wells increase agricultural productivity and thus decrease market prices for staple foods, which in turn improves the affordability of calories. This effect also provides a potential mechanism to explain the observed spillovers, as calorie demand of populations not directly affected by tube well adoption increases in response to lower prices.

Long-Term Effects

Given the observed strong effects on yields, equilibrium prices and wages, it is possible that tube well adoption has long term effects on the structure of rural economies. For example, changing production practices might affect the farm size distribution if mechanization generates economies of scale that enhance the profitability of larger farms. To investigate this possibility, I estimate the

[22]This data was originally collected by the Indian Ministry of Agriculture and Farmer Welfare and formated by ICRISAT.

[23]Formally, the Indian Ministry of Agriculture and Farmer Welfare defines farm harvest prices as follows: *"Farm Harvest prices of a commodity [...] is defined as the average wholesale price, at which the commodity is disposed of by the producer to the trader at the village site during the specified marketing period after the commencement of harvest."* (Ministry of Agriculture and Farmers Welfare, 2017)

effect of tube well adoption on the number of plots in a given district as well as the distribution of plot sizes. To this end, I employ ICRISAT data on the number of plots of various sizes in a given district.[24]

The results of this exercise are presented in table 3.9. Column 1 highlights that tube well adoption indeed leads to a consolidation of plots: a one standard deviation increase in the share of area sown irrigated using tube wells reduces the number of plots in a given district by approximately 135,657 plots. As columns 4 to 6 highlight, this decrease is driven primarily by a reduction in the number of marginal, small and semi-medium plots. In contrast, the number of large and medium plots has increased significantly in response to tube well adoption: a one standard deviation increase in the share of area sown irrigated using tube wells causes an increase of approximately 19,500 large and medium plots.

The results highlight that improvements to agricultural productivity can have lasting impacts on the structure of rural economies by inducing a consolidation of farms into larger farms at the expense of smaller plots. Table B.1 in the appendix shows that, in addition to a changing farm size distribution, tube well adoption also increases rural populations. While I am unable to evaluate whether the population increase is caused by migration or increased birth rates, the results presented in tables 3.9 and B.1 nevertheless highlight that the adoption of tube wells has induced long-term structural and demographic changes in rural India.

3.6 Robustness

I present four checks to evaluate the robustness of my empirical result. First, the identifying assumption underlying my empirical strategy is that outcome variables, in particular calorie intake, would have evolved in parallel in hard and soft rock districts if aquifer structures were similar. This assumption is violated if trends differ between hard and soft rock areas. While it is impossible to control for diverging trends at the district level, the inclusion of state-year fixed effects or state specific linear trends allow to control for differences in trends between states.

A challenge with controlling for state specific trends is that it substantially reduces statistical power. This is because the identifying variation from my instrumental variable strategy exploits the time varying effect of hard rock

[24]In contrast to other ICRISAT data series, data on land use is only available on a quinquennial instead of an annual basis. The reduces the sample size compared to the previous analysis.

aquifers on tube well adoption. Controlling for state specific trends reduces statistical power as in that case the identifying variation only comes from deviations from the state specific trends, rather than differential trends in themselves.

Columns 2 and 3 in table B.2 show that while qualitatively the point estimates are robust to the inclusion of state specific trends, they are no longer statistically significant. In column 2 of table B.2 I control for state-year fixed effects, whereas column 3 controls for state specific linear trends. Consistent with the notion that the inclusion of state specific trends reduces statistical power, the F-statistics for the first stage of the specifications estimated in columns 2 and 3 suggest that the explanatory power of the first stage regressions is no longer statistically different from zero at the 10% level. However, while the results are not statistically robust to including state specific time trends, the fact that the sign of the point estimates remains robust is indicative that differential pre-trends do not qualitatively affect the subsequent pattern of calorie intake.

Second, I investigate whether the effect of tube well adoption on calorie intake is robust to the inclusion of control variables. In particular, it is possible that systematic differences in hard and soft rock areas exist which simultaneously affect calorie intake, thus leading to a spurious estimated relationship between tube well adoption and calorie intake. To address this, I again estimate specification 3.2, but control for potential heterogeneity between hard and soft rock areas by including control variables for the total area sown, the average price per calorie paid, per capita total expenditure as a proxy for income and district population. The results of this exercise are presented in column 1 of table B.2, which suggests that the estimates are robust to including the aforementioned control variables. In particular, while the original point estimate without controls was 2,851 (presented in column 3 in table 3.3), including controls slightly lowers the point estimate but neither substantially affects its magnitude nor its statistical significance.

Third, I show that the effects of tube well adoption on expenditure patterns presented in table 3.6 are robust to restricting attention to only the districts available in the ICRISAT data set. Table B.3 replicates the analysis presented in table 3.6, but restricts attention to the districts that are represented in the ICRISAT data set. Consistent with the results obtained from the constructed sample, table B.3 shows that tube well adoption causes an increase in the expenditure share on food which is accompanied by an increased expenditure share on animal products and reduced expenditure on vegetables and sugar-based products. These results are thus indicative that the estimates obtained using the ICRISAT sample are also informative about estimates that were obtained using my constructed sample.

Fourth, to further support the notion that results obtained using the ICRISAT data and my constructed sample are comparable, I show that the results of tube well adoption on agricultural productivity are robust to restricting attention to the time period between 1987 and 2004, which is the sample period of my constructed sample. Columns 1 and 2 in table B.4 first show that the estimates of the effect of tube well adoption on yields are qualitatively robust to restricting the analysis period. In particular, similar to the unrestricted results presented in table 3.7, the point estimates for rice and wheat yields have a similar magnitude. While the point estimate for rice yields is no longer significant, I still find a large and statistically significant effect of tube well adoption on wheat yields. Columns 3 and 4 further show that tube well adoption still affects the introduction of HYV seeds, although to a lesser extent than when using the constructed sample. Nevertheless, I still find a positive and statistically significant effect of tube well adoption on the total area sown with HYV seeds for wheat. Taken together, the results thus suggest that the estimates obtained using all ICRISAT time periods are robust to restricting attention to only the time periods available in my constructed sample.

3.7 Conclusion

In this paper, I argue that investment in groundwater irrigation can significantly improve nutrition in India. Exploiting the fact that the adoption of groundwater irrigation devices was slower in areas where the rock structure of aquifers is igneous, I study how tube well adoption affects calorie intake, agricultural productivity, staple prices and wages. I find that groundwater irrigation significantly improves nutrition across the income spectrum and allows households to switch from consuming vegetables to animal based produce. In addition, I find that groundwater irrigation generates positive spillovers on the calorie intake of urban populations and households not employed in agriculture. I present additional evidence which suggests that these effects are driven by increases in agricultural productivity that reduce staple prices and raise wage rates.

By highlighting the potential value of promoting agricultural technology to improve nutritional outcomes, the results presented in this paper have implications for public policy and development programming. In particular, I show that the adoption of groundwater irrigation technology can break India's "agricultural disconnect", raising agricultural productivity and nutrition outcomes at the same time. Furthermore, I find no support for the notion that the productivity improvements associated with the Indian "Green Revolution" left smaller farmers and agricultural laborer impoverished. In contrast, a complementarity between the technology improvements and labor inputs appear to have raised wage rates. This development, together with falling staple prices, has improved nutrition across the income spectrum. Taken together, the results thus highlight that policies which encourage and facilitate groundwater irrigation are likely to effectively combat undernutrition.

While my paper estimates the reduced form relationship between tube well adoption and calorie intake, I am unable to distinguish between the extent to which tube wells directly affect calorie intake, for example by changing rural work practices, and indirectly affect the supply of calories through changes to production patterns. Similarly, my results on the effect of groundwater irrigation on agricultural productivity are unable to distinguish between the direct effect of tube well adoption and the productivity improvements caused by adopting HYV seeds and mechanization devices in response to tube well adoption. Distinguishing those effects is left for future work.

Tables

Table 3.1: Summary Statistics

Panel A: Full Constructed Sample	Total	Hard Rock	Soft Rock
Share of Land Irrigated using Tube Wells	0.19	0.05	0.26
	(0.27)	(0.09)	(0.29)
Median District Calorie Intake	2707	2605	2750
	(344)	(312)	(348)
Food Expenditure Share	0.58	0.57	0.58
	(0.12)	(0.09)	(0.13)
Observations	987	293	694
Panel B: ICRISAT Meso Data Sample			
Share of Land Irrigated using Tube Wells	0.13	0.01	0.20
	(0.22)	(0.04)	(0.26)
Population (in 1,000s)	2188	1964	2373
	(1463)	(1068)	(1703)
Share of Land using HYV Seeds	0.38	0.24	0.48
	(0.35)	(0.18)	(0.40)
Number of Tractors (in 1,000s)	3.81	1.43	5.47
	(9.25)	(2.04)	(11.65)
Rice Yield (tons per hectare)	1.38	1.24	1.48
	(0.86)	(0.76)	(0.91)
Wheat Yield (tons per hectare)	1.57	1.15	1.85
	(0.94)	(0.60)	(1.02)
Observations	6,147	2,490	3,657

Notes: The full sample includes the years 1987, 1999 and 2004. With the exception of the population and the tractor series, the ICRISAT sample contains unbalanced annual district data for the years 1970 to 2009. The population data origins from the decennial census spanning the years 1971 to 2001. The tractor data covers a subset of Indian districts for the years 1972, 1982, 1987, 1992, 2003.

Table 3.2: Differential Adoption of Tube Wells across Hard and Soft Rock Areas

Outcome:	**% of Land irrigated using Tube Wells**		
	(1)	(2)	(3)
Panel A: Years in constructed sample			
1999	0.1000***	0.0961***	
	(0.0110)	(0.0119)	
2004	0.124***	0.123***	
	(0.0117)	(0.0121)	
1999 × Hard Rock	-0.0577***	-0.0513***	
	(0.0133)	(0.0147)	
2004 × Hard Rock	-0.0732***	-0.0696***	
	(0.0148)	(0.0163)	
Panel B: Years in ICRISAT sample			
1980s × Hard Rock			-0.0713***
			(0.00717)
1990s × Hard Rock			-0.0980***
			(0.0144)
2000s × Hard Rock			-0.158***
			(0.0245)
Observations	987	714	6,147
Data Source	Constr.	Constr.	ICRISAT
District Sample	Full	ICRISAT	ICRISAT
Reference Year	1987	1987	1970
District FE	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes
F-Statistic	37.55	32.90	37.90

Notes: Standard Errors are clustered at the district level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Columns 1 and 2 present coefficient estimates of a regression of the proportion of cropped land irrigated using tube wells on district and year fixed effects and an interaction term of years and an indicator variable for areas covered by hard rock aquifers. Column 1 uses all Indian districts, whereas column 2 restricts its attention to the districts available in the ICRISAT data set. Columns 3 presents coefficient estimates of a regression of the proportion of cropped land irrigated using tube wells on district and year fixed effects and an interaction term of the decade that the data was collected in and an indicator variable for areas covered by hard rock aquifers. Column 3 employs the ICRISAT data set, whereas columns 1 and 2 use the NSS/Indian Agricultural Statistics sample.

Table 3.3: Effect of Tube Well Adoption on Median Calorie Intake

	Median Calorie Intake			
	(1)	(2)	(3)	(4)
Tube Well Share	728.4***	827.9***	2,851***	3,390***
	(115.9)	(157.9)	(871.3)	(1,095)
1999	60.33*	66.44*	-112.7	-135.7
	(32.40)	(37.08)	(74.15)	(89.05)
2004	-115.3***	-106.3***	-332.1***	-367.7***
	(30.30)	(36.20)	(88.82)	(109.4)
Observations	987	714	987	714
Data Source	Constr.	Constr.	Constr.	Constr.
District Sample	Full	ICRISAT	Full	ICRISAT
Estimation Method	OLS	OLS	2SLS	2SLS
District FE	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes

Notes: Standard Errors are clustered at the district level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Columns 1 and 2 present coefficient estimates of an OLS regression of daily per capita calorie intake on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Column 1 uses all Indian districts, whereas column 2 only uses data from districts that are available in the ICRISAT data set. Columns 3 and 4 present coefficient estimates of a two-stage-least-squares regression of daily per capita calorie intake on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Column 1 uses all Indian districts, whereas column 2 only uses data from districts that are available in the ICRISAT data set. In columns 3 and 4 the share of cropped land irrigated using tube wells is instrumented for using variation in hydrogeological structures.

Table 3.4: Effect of Tube Well Adoption on Median Calorie Intake, by Population

	Population				
	Urban	Rural			
	Total	Total	Non-ag. self-employed	Ag. laborer	Self-employed in ag.
	(1)	(2)	(3)	(4)	(5)
Tube Well Share	1,800**	2,549***	2,206**	2,482***	2,707***
	(760.2)	(801.1)	(954.9)	(951.9)	(1,041)
Observations	887	887	887	887	887
Data Source	Constr.	Constr.	Constr.	Constr.	Constr.
District Sample	Full	Full	Full	Full	Full
Estimation Method	2SLS	2SLS	2SLS	2SLS	2SLS
District FE	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes

Notes: Standard Errors are clustered at the district level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Column 1 presents coefficient estimates of a two-stage-least-squares regression of the calorie intake of the median urban inhabitant of a district on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Column 2 presents coefficient estimates of a two-stage-least-squares regression of the calorie intake of the median rural inhabitant of a district on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Column 3 presents coefficient estimates of a two-stage-least-squares regression of the calorie intake of the median rural self-employed outside of agriculture inhabitant of a district on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Column 4 presents coefficient estimates of a two-stage-least-squares regression of the calorie intake of the median rural agricultural laborer of a district on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Column 5 presents coefficient estimates of a two-stage-least-squares regression of the calorie intake of the median rural self-employed in agriculture inhabitant of a district on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Estimates are obtained through a two-stage-least-squares regression in which tube well adoption is instrumented for using variation in hydrogeological structures.

Table 3.5: Effect of Tube Well Adoption on Calorie Intake, by Percentile

	Calorie intake of ...				
	10^{th} percentile	20^{th} percentile	30^{th} percentile	40^{th} percentile	Mean
	(1)	(2)	(3)	(4)	(5)
Tube Well Share	2,547***	2,413***	2,572***	2,848***	2,928***
	(949.6)	(833.2)	(790.7)	(831.0)	(949.3)
Observations	987	987	987	987	987
Data Source	Constr.	Constr.	Constr.	Constr.	Constr.
District Sample	Full	Full	Full	Full	Full
Estimation Method	2SLS	2SLS	2SLS	2SLS	2SLS
District FE	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes

Notes: Standard Errors are clustered at the district level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Columns 1 to 5 present coefficient estimates of a regression of the daily per capita calorie intake by an inhabitant in the 10^{th}, 20^{th}, 30^{th} and 40^{th} percentile as well as the average of the calorie intake distribution of a district on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Estimates are obtained through a two-stage-least-squares regression in which tube well adoption is instrumented for using variation in hydrogeological structures.

Table 3.6: Effect of Tube Well Adoption on Food Expenditure

		Share of food expenditure on ...					
	Food exp. share	Staples	Cereals	Animal Prd.	Veg.	Fruit	Sugar
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Tube Well Share	0.592**	0.143	0.220*	0.684***	-0.170**	0.0381	-0.148***
	(0.270)	(0.163)	(0.114)	(0.184)	(0.0743)	(0.0341)	(0.0523)
Observations	987	987	987	987	987	987	987
Data Source	Constr.	Constr.	Constr.	Constr.	Constr.	Constr.	Constr.
District Sample	Full	Full	Full	Full	Full	Full	Full
Reference Year	1987	1987	1987	1987	1987	1987	1987
Estimator	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
District FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Standard Errors are clustered at the district level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Columns 1 to 8 present coefficient estimates of a regression of the average share, as part of monthly household expenditure, of various food items on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Column 1 uses food expenditure as an outcome variable, whereas column 2 uses expenditure on staples, column 3 uses expenditure on non-staple cereals, column 4 uses expenditure on animal products, column 5 uses expenditure on vegetables, column 6 uses expenditure on fruits, and column 7 uses expenditure on sugar and sugar-related food items. Estimates are obtained through a two-stage-least-squares regression in which tube well adoption is instrumented for using variation in hydrogeological structures.

Table 3.7: Effect of Tube Well Adoption on Agricultural Production

	Yield		HYV		Area		Mechanization	
	Rice	Wheat	Rice	Wheat	Rice	Wheat	Tractor	Power Tiller
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Tube Well Share	1.592***	4.446***	310.9***	477.5***	268.3***	378.1***	47.54***	5.127***
	(0.521)	(0.783)	(85.80)	(91.51)	(77.19)	(75.57)	(8.037)	(1.477)
Observations	6,147	6,147	6,147	6,147	6,147	6,147	1,977	1,393
Data Source	ICR.	ICR.	ICR.	ICR.	ICR.	ICR.	ICR.	ICR.
Estimation Method	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
District FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Standard Errors are clustered at the district level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Columns 1 and 2 present coefficient estimates of a two-stage-least-squares regression of rice and wheat yields, respectively, on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Columns 3 and 4 present coefficient estimates of a two-stage-least-squares regression of the total area cropped with HYV rice and wheat seeds, respectively, on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Columns 5 and 6 present coefficient estimates of a two-stage-least-squares regression of the total area cropped with rice and wheat seeds, respectively, on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Columns 7 and 8 present coefficient estimates of a two-stage-least-squares regression of the total number of tractors and power tillers in a given district (in 1000s), respectively, on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. All specifications employ ICRISAT data and cover the years 1970 to 2009. The share of cropped land irrigated using tube wells is instrumented for using variation in hydrogeological structures.

Table 3.8: Effect of Tube Well Adoption on Wages and Prices

	Wages		Prices	
	Male	Female	Rice	Wheat
	(1)	(2)	(3)	(4)
Tube Well Share	100.3**	59.99	-83.79	-315.7***
	(42.30)	(46.35)	(1,272)	(81.85)
Observations	7,778	6,146	2,438	7,217
Data Source	ICRISAT	ICRISAT	ICRISAT	ICRISAT
Estimation Method	2SLS	2SLS	2SLS	2SLS
District FE	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes

Notes: Standard Errors are clustered at the district level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Columns 1 and 2 present coefficient estimates of a two-stage-least-squares regression of male and female wages, respectively, on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Columns 3 and 4 present coefficient estimates of a two-stage-least-squares regression of rice and wheat prices, respectively, on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. All specifications employ ICRISAT data and cover the years 1970 to 2009. The share of cropped land irrigated using tube wells is instrumented for using variation in hydrogeological structures.

Table 3.9: Effect of Tube Well Adoption on Farm Size Distribution

	Number of Plots					
	Total	Large	Medium	Semi-Medium	Small	Marginal
	(1)	(2)	(3)	(4)	(5)	(6)
Tube Well Share	-502.5***	44.83***	27.41**	-156.9***	-329.4***	-88.13
	(136.2)	(8.506)	(12.37)	(30.30)	(58.16)	(111.9)
Observations	2,063	2,063	2,063	2,063	2,063	2,063
Data Source	ICRISAT	ICRISAT	ICRISAT	ICRISAT	ICRISAT	ICRISAT
Estimation Method	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Standard Errors are clustered at the district level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Columns 1 presents coefficient estimates of the total number of plots in a given district on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Columns 2 and 6 present coefficient estimates of a two-stage-least-squares regression of the total number of large, medium, semi-medium, small and marginal plots on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. All specifications employ ICRISAT data. The share of cropped land irrigated using tube wells is instrumented for using variation in hydrogeological structures.

Figures

Figure 3.1: Aquifer Map of India

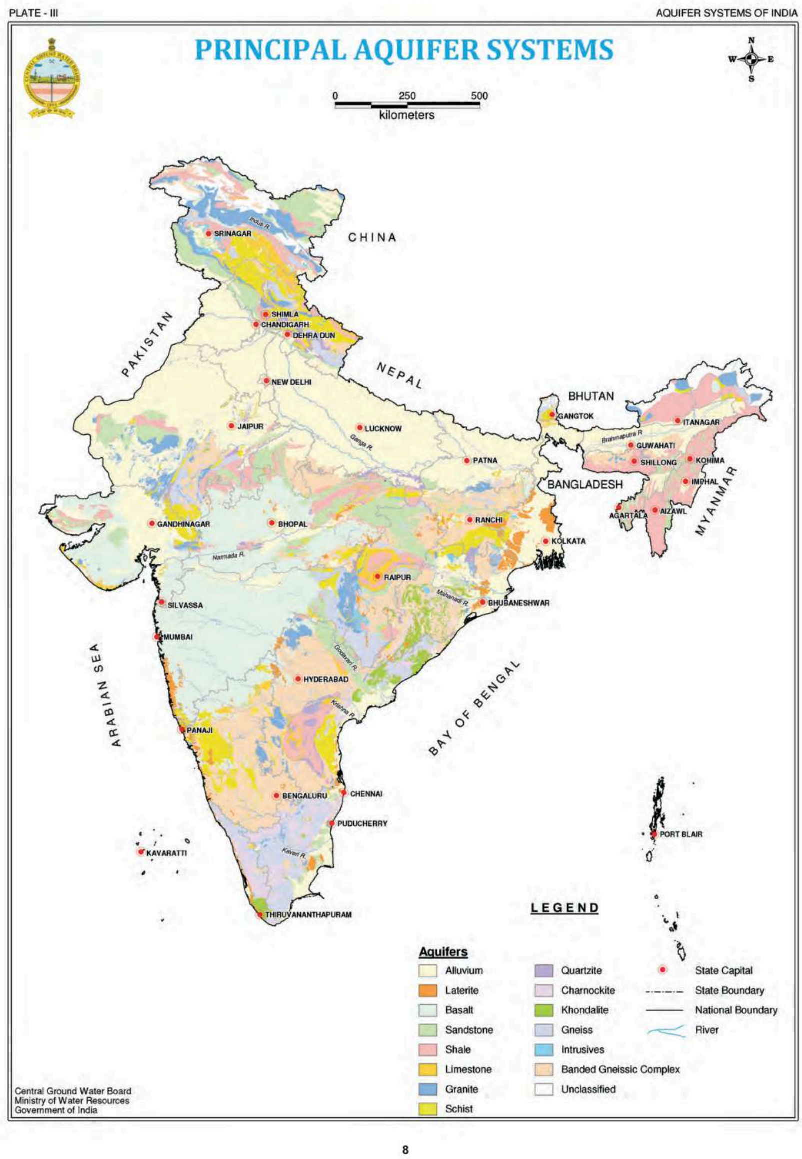

Figure 3.2: District Wise Aquifer Distribution

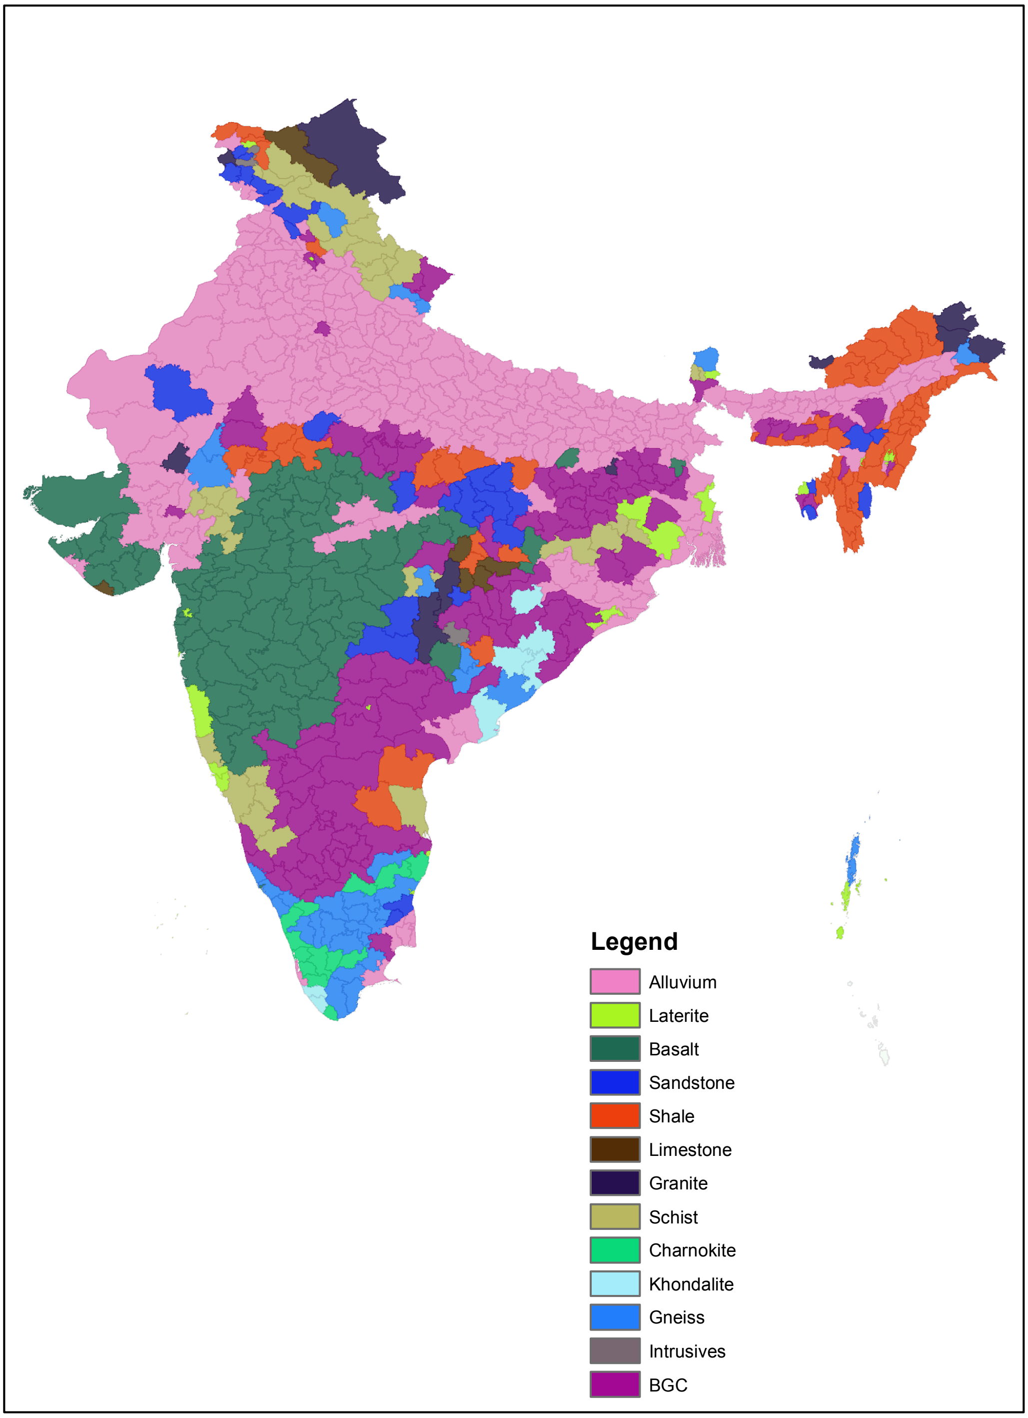

Figure 3.3: Differential Tube Well Adoption in Hard Rock Areas

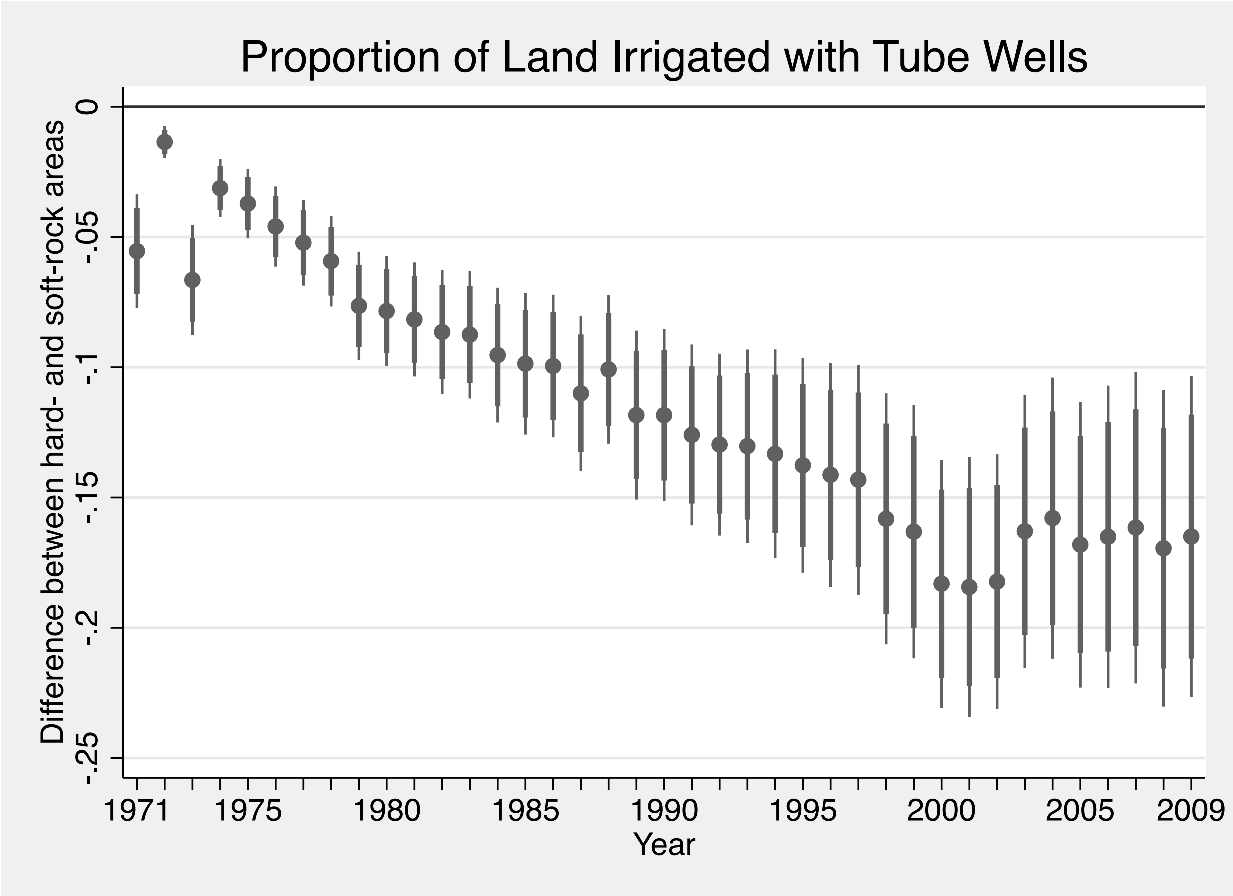

Bibliography

Abel, Jaison R (2000). "The Performance of the State Telecommunications Industry under Price-Cap Regulation: An Assessment of the Empirical Evidence". In: *National Regulatory Research Institute Report*.

Afridi, Farzana (2010). "Child Welfare Programs and Child Nutrition: Evidence from a Mandated School Meal Program in India". In: *Journal of Development Economics* 92, pp. 152–165.

Aghion, Philippe and Jean Tirole (1997). "Formal and Real Authority in Organizations". In: *Journal of Political Economy* 105.1, pp. 1–29.

Alonso, Ricardo and Niko Matouschek (2008). "Optimal Delegation". In: *Review of Economic Studies* 75.1, pp. 259–293.

Altomonte, Carlo, Alessandro Barattieri, and Susanto Basu (2015). "Average-Cost Pricing: Some Evidence and Implications". In: *European Economic Review* 79, pp. 281–296.

Andersen, Thomas Barnebeck, Peter Sandholt Jensen, and Christian Volmar Skovsgaard (2016). "The heavy plow and the agricultural revolution in Medieval Europe". In: *Journal of Development Economics* 118, pp. 133–149.

Anderson, Jock R. and Gershon Feder (2007). "Agricultural Extension". In: *Handbook of Agricultural Economics* 3.06, pp. 2343–2378.

Ashraf, Nava, Oriana Bandiera, and B. Kelsey Jack (2014). "No margin, no mission? A field experiment on incentives for public service delivery". In: *Journal of Public Economics* 120, pp. 1–17.

Ashraf, Nava, Oriana Bandiera, and Scott Lee (2016). "Do-Gooders and Go-Getters: Career Incentives, Selection, and Performance in Public Service Delivery". In: *Mimeo*.

Ashraf, Nava, James Berry, and Jesse M. Shapiro (2010). “Can Higher Prices Stimulate Product Use? Evidence From a Field Experiment in Zambia”. In: *American Economic Review* 100.December, pp. 2383–2413.

Atkin, David (2013). “Trade, Taste and Nutrition in India”. In: *American Economic Review* 103.5, pp. 1629–1663.

Ballinger, T Parker, Michael G Palumbo, and Nathaniel T Wilcox (2003). “Precautionary Saving and Social Learning Across Generations: An Experiment”. In: *Economic Journal* 113.1994, pp. 920–947.

Banerjee, Abhijit, Rema Hanna, Jordan C. Kyle, et al. (2015). “Tangible Information and Citizen Empowerment”. In: *NBER Working Paper* 20923.

Banerjee, Abhijit, Rema Hanna, and Sendhil Mullainathan (2013). “Corruption”. In: *Handbook of Organizational Economics*, pp. 1–72.

Banerji, A., J. V. Meenakshi, and Gauri Khanna (2012). “Social contracts, markets and efficiency: Groundwater irrigation in North India”. In: *Journal of Development Economics* 98.2, pp. 228–237.

Bardhan, Pranab and Dilip Mookherjee (2011). “Subsidized farm input programs and agricultural performance: A farm-level analysis of West Bengal’s green revolution, 1982-1995”. In: *American Economic Journal: Applied Economics* 3.4, pp. 186–214.

Bardhan, Pranab, Dilip Mookherjee, and Neha Kumar (2012). “State-led or market-led green revolution? Role of private irrigation investment vis-a-vis local government programs in West Bengal’s farm productivity growth”. In: *Journal of Development Economics* 99.2, pp. 222–235.

Baumol, William J. and Richard E. Quandt (1964). “Rules of Thumb and Optimally Imperfect Decisions”. In: *American Economic Review* 54.2, pp. 23–46.

Becker, Gordon, Morris DeGroot, and Jacob Marschak (1964). “Measuring Utility by a Single-Response Sequential Method”. In: *Behavioral Science* 9.3, pp. 226–232.

Bloom, Nicholas, Benn Eifert, et al. (2013). “Does managment matter? Evidence from India”. In: *The Quarterly Journal of Economics* 128.1, pp. 1–51.

Bloom, Nicholas and John Van Reenen (2007). “Measuring and Explaining Management Practices across Firms and Countries”. In: *Quarterly Journal of Economics* 72.1, pp. 1–55.

Boven, Michiel van et al. (2008). "Herd immunity to Newcastle disease virus in poultry by vaccination". In: *Avian Pathology* 37.1, pp. 1–5.

Bruhn, Miriam, Dean Karlan, and Antoinette Schoar (2017). "The Impact of Consulting Services on Small and Medium Enterprises: Evidence from a Randomized Trial in Mexico". In: *Journal of Political Economy*, pp. 1–104.

Busso, Matias and Sebastian Galiani (2014). "The Causal Effect of Competition on Prices and Quality: Evidence from a Field Experiment". In: *NBER Working Paper* 20054.

Cary, J W (1998). "Issues in Public and Private Technology Transfer: The Cases of Australia and New Zealand". In: *Privatization of Information and Agricultural Industrialization*.

Central Ground Water Board (2012). *Aquifer Systems of India*. Tech. rep.

Chan, David C. (2015). "Teamwork and Moral Hazard: Evidence from the Emergency Department". In: *Journal of Political Economy* 124.3.

Chaudhury, Nazmul et al. (2006). "Missing in action: teacher and health worker absence in developing countries." In: *Journal of Economic Perspectives* 20.1, pp. 91–116.

Cohen, Jessica and Pascaline Dupas (2010). "Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment". In: *Quarterly Journal of Economics* 75.1, pp. 1–45.

Conlisk, John (1996). "Why Bounded Rationality?" In: *Journal of Economic Literature* 34.2, pp. 669–700.

Crowder, L van et al. (2002). "Uganda: Private Sector Secondment of Government Extension Agents." In: *Contracting for Agricultural Extenison: International Case Studies and Emerging Practices*, pp. 155–162.

Dancey, R J (1993). "The Evolution of Agricultural Extension in England and Wales". In: *Journal of Agricultural Economics* 44.3, pp. 375–393.

Dasgupta, P and D Ray (1987). "Inequality as a Determinant of Malnutrition and Unemployment: Policy". In: *The Economic Journal* 97.385, pp. 177–188.

De Long, Bradford et al. (1990). "Noise Trader Risk in Financial Markets". In: *Journal of Political Economy* 98.4, pp. 703–738.

Deaton, Angus and Jean Drèze (2009). "Food and Nutrition in India: Facts and Interpretations". In: *Economic and Political Weekly* 47.7, pp. 42–65.

Deininger, Klaus and Paul Mpuga (2005). “Does Greater Accountability Improve the Quality of Public Service Delivery? Evidence from Uganda”. In: *World Development* 33.1, pp. 171–191.

Dejene, Alemneh (1989). “The training and visit agricultural extension in rainfed agriculture: Lessons from Ethiopia”. In: *World Development* 17.10, pp. 1647–1659.

Deserranno, Erika (2016). “Financial Incentives as Signals : Experimental Evidence from the Recruitment of Health Workers”. In: *Mimeo*.

Devoto, Florencia et al. (2012). “Happiness on Tap: Piped Water Adoption in Urban Morocco”. In: *American Economic Journal: Economic Policy* 4.4, pp. 68–99.

Dubois, Pierre, Rachel Griffith, and Aviv Nevo (2014). “Do Prices and Attributes Explain International Differences in Food Purchases?” In: *American Economic Review* 104.3, pp. 832–867.

Duflo, Esther (2006). “Poor But Rational?” In: *Understanding Poverty*. Ed. by Abhijit Banerjee. New York: Oxford University Press, pp. 367–368.

Duflo, Esther, Michael Greenstone, et al. (2014). “The Value of Regulatory Discretion”. In: *NBER Working Paper* 20590.

Duflo, Esther and Rohini Pande (2007). “Dams”. In: *Quarterly Journal of Economics* 122.2, pp. 601–646.

Economist (2015). “India ′ s malnourished infants”. In: *The Economist* July 2nd.

Eli, Shari and Nicholas Li (2017). “Caloric Requirements and Food Consumption Patterns of the Poor”. In: *Mimeo: University of Toronto*.

Ellison, Glenn and Drew Fudenberg (1993). “Rules of Thumb for Social Learning”. In: *Journal of Political Economy* 101.4, pp. 612–643.

— (1995). “Word-of-Mouth Communication and Social Learning”. In: *Quarterly Journal of Economics* 110.1, pp. 93–125.

Farrington, John et al. (2002). “Creating a Policy Environment for Pro-Poor Agricultural Extension: The Who? What? And How?” In: *ODI - Natural Resources Perspectives* 80.

Faulhaber, Gerald R. and William J. Baumol (1988). “Economists as Practical Innovators: Products of Theoretical Research”. In: *Journal of Economic Literature* 26.2, pp. 577–600.

Feder, Gershon, Anthony Willett, and Willem Zijp (1999). "Agricultural Extension - Generic Challenges and Some Ingredients for Solutions Extension". In: *Policy Research Working Paper*.

Feldman, Naomi E., Peter Katuscak, and Laura Kawano (2016). "Taxpayer Confusion: Evidence from the Child Tax Credit". In: *American Economic Review* 106.3, pp. 807–835.

Finan, Frederico, Benjamin A. Olken, and Rohini Pande (2015). "The Personnel Economics of the State". In: *NBER Working Paper* 21825, p. 48.

Fisher, Hayden (2014). "Newcastle Disease Control in Africa". In: *Australian Centre for International Agricultural Research Impact Assessment Series* 87.

Fishman, Ram Mukul (2012). "Climate Change, Rainfall Variability , and Adaptation through Irrigation: Evidence from Indian Agriculture". In: *Mimeo*.

Food and Agricultural Organization of the United Nations (2003). "The irrigation challenge: increasing irrigation contribution to food security through higher water productivity from canal irrigation systems". In: *IPTRID Issues Paper No.4*, p. 25.

Garduño, Héctor and Stephen Foster (2010). "Sustainable Groundwater Management: Approaches to Reconciling Demand with Resources". In: *World Bank Strategic Overview Series* 4.

Gautam, Madhur (2000). *Agricultural Extension: The Kenya Experience*.

Gillespie, S., J. Harris, and S. Kadiyala (2012). "The Agriculture-Nutrition Disconnect in India: What Do We Know?" In: *IFPRI Discussion Paper 01187* June.

Gopalan, C., B.V. Rama Shastri, and S.C. Balasubramanian (2004). *Nutritive Value of Indian Foods*. Hyderabad: National Institute of Nutrition, Indian Council of Medical Research.

Grossman, Sanford J and Oliver D Hart (1986). "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration". In: *Journal of Political Economy* 94.4, pp. 691–719.

Haan, Cornelis de et al. (2001). *Livestock Development: Implications for Rurual Poverty, the Environment, and Global Food security*. Washington, D.C.: The World Bank, 75pp.

Hall, R. L. and C. J. Hitch (1939). “Price Theory and Business Behaviour”. In: *Oxford Economic Papers* 2.2, pp. 12–45.

Haltiwanger, John and Michael Waldman (1985). “Rational Expectations and the Limits of Rationality: An Analysis of Heterogeneity”. In: *American Economic Review* 75.3, pp. 326–340.

Hey, John D. and Valentino Dardanoni (1988). “Optimal Consumption Under Uncertainty: An Experimental Investigation”. In: *Economic Journal* 98.390, pp. 105–116.

Hijmans, Robert (2009). *Global Administrative Areas Data Set*.

Hornbeck, Richard and Pinar Keskin (2014). “The Evolving Impact of the Ogallala Aquifer: Agricultural Adaptation to Groundwater and Climate”. In: *American Economic Journal: Applied Economics* 6.1, pp. 190–219.

Howell, John (1986). “Accountability in Extension Work”. In: *Investing in Rural Extension: Strategies and Goals*. Ed. by Gwyn E. Jones.

Hoynes, Hilary W and Diane Whitmore Schanzenbach (2009). “Consumption Responses to In-Kind Transfers: Evidence from the Introduction of the Food Stamp Program”. In: *American Economic Journal: Applied Economics* 1.4, pp. 109–139.

International Crops Research Institute for the Semi-Arid Tropics (2017). *Meso Data Set*.

International Food Policy Research Institute (2002). “Green Revolution: Curse or Blessing?” In: *IFPRI Policy Brief*.

International Water Management Institute (2002). “The Socio-Ecology of Groundwater in India”. In: *Water Policy Briefing* 4.

Ito, Koichiro (2014). “Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing”. In: *American Economic Review* 104.2, pp. 537–563.

Jensen, Robert T. and Nolan H. Miller (2008). “Giffen Behavior and Subsistence Consumption”. In: *American Economic Review* 98.4, pp. 1553–1577.

Kahneman, Daniel (2002). “Maps of Bounded Rationality: a Perspective on Intuitive Judgment and Choice”. In: *Nobel Prize Lecture* December.

Kahneman, Daniel and Shane Frederick (2002). “Representativeness Revisited: Attribute Substitution in Intuitive Judgement”. In: *Heuristics and Biases*. Ed.

by T Gilovich, D Griffin, and Daniel Kahneman. New York: Cambridge University Press, pp. 49–81.

Kessler, JB (2017). “Announcements of Support and Public Good Provision”. In: *Working Paper*.

Khan, Adnan Q., Asim I. Khwaja, and Benjamin A. Olken (2016a). “Making Moves Matter: Experimental Evidence on Incentivising Bureaucrats through Performance-Based Transfers”. In: *Mimeo*.

— (2016b). “Tax Farming Redux: Experimental Evidence on Performance Pay for Tax Collectors”. In: *Quarterly Journal of Economics* 1, pp. 219–271.

Kidd, A. D. et al. (2000). “Privatising Agricultural Extension: Caveat emptor”. In: *Journal of Rural Studies* 16.1, pp. 95–102.

Laffont, Jean-Jacques (2005). *Regulation and Development*. Cambridge: Cambridge University Press.

Laffont, Jean-Jacques and Jean Tirole (1986). “Using Cost Observation to Regulate Firms”. In: *Journal of Political Economy* 94.3, pp. 614–641.

— (1993). *A Theory of Incentives in Procurement and Regulation*. Cambridge: MIT Press.

Lagarde, Mylene and Natasha Palmer (2008). “The Impact of User Fees on Health Service Utilization in Low- and Middle-Income Countries: How strong is the Evidence?” In: *Bulletin of the World Health Organization* 86.11, pp. 839–848.

Liebman, Jeffrey B and Richard J Zeckhauser (2004). “Schmeduling”. In: *Mimeo*.

Mankiw, N. Greogry, Matthew Weinzierl, and Danny Yagan (2009). “Optimal Taxation in Theory and Practice”. In: *The Journal of Economic Perspectives* 23.4, pp. 147–174.

McKenzie, David and Christopher Woodruff (2015). “Business Practices in Small Firms in Developing Countries”. In: *Management Science* August, p. 40.

Ministry of Agriculture and Farmers Welfare (2017). “Farm Harvest Prices of Principal Crops in India”. In: *Government of India Publication*.

Mirrlees, J. A. (1971). “An Exploration in the Theory of Optimum Income Taxation”. In: *Review of Economic Studies* 38.2, pp. 175–208.

Mongin, Philippe (1992). “The "Full-Cost" Controversy of the 1940s and 1950s: A Methodological Assessment”. In: *History of Political Economy* 24.

Msami, Halifa (2007). “Poultry Sector Country Review: Tanzania”. In: *FAO Poultry Sector Country Review*, p. 61.

Muralidharan, Karthik and Venkatesh Sundararaman (2011). “Teacher Performance Pay: Experimental Evidence from India”. In: *Journal of Political Economy* 119.1, pp. 39–77.

Olken, Benjamin A., Junko Onishi, and Susan Wong (2014). “Should Aid Reward performance? Evidence from a Field Experiment on Health and Education in Indonesia”. In: *American Economic Journal: Applied Economics* 6.4, pp. 1–34.

Oprea, Ryan (2014). “Survival Versus Profit Maximization in a Dynamic Stochastic Experiment”. In: *Econometrica* 82.6, pp. 2225–2255.

Oprea, Ryan, Daniel Friedman, and Steven T Anderson (2009). “Learning to Wait: A Laboratory Investigation”. In: *Review of Economic Studies* 76, pp. 1103–1124.

Park, Graham (2010). *Introducing Geology: A Guide to the World of Rocks*. 2nd ed. Edinburgh: Dunedin Academic Press, p. 144.

Patniak, Utsa (2004). “The Republic of Hunger”. In: *Social Scientist* 32.9, pp. 9–35.

Pingali, P. L. (2012). “Green Revolution: Impacts, limits, and the path ahead”. In: *Proceedings of the National Academy of Sciences* 109.31, pp. 12302–12308.

Pingle, Mark and Richard H Day (1996). “Modes of Economizing Behavior: Experimental Evidence”. In: *Journal of Economic Behavior & Organization* 29, pp. 191–209.

Postel, Sandra (1999). *Pillar of Sand: Can the Irrigation Miracle Last?* 1st ed. Worldwatch Environmental Alert Series: W.W. Norto, p. 313.

Rasul, Imran and Daniel Rogger (2016). “Management of Bureaucrats and Public Service Delivery: Evidence from the Nigerian Civil Service”. In: *Economic Journal*.

Reis, Ricardo (2006). “Inattentive Consumers”. In: *Journal of Monetary Economics* 53, pp. 1761–1800.

Rivera, William M. and Daniel J. Gustafson (1991). *Agricultural Extension: Worldwide Institutional Evolution and Forces for Change*. Amsterdam: Elsevier Science Publishers BV.

Rondeau, Daniel and John A. List (2008). "Matching and challenge gifts to charity: Evidence from laboratory and natural field experiments". In: *Experimental Economics* 11.3, pp. 253–267.

Rud, Juan Pablo (2012). "Electricity Provision and Industrial Development: Evidence from India". In: *Journal of Development Economics* 97.2, pp. 352–367.

Sah, Raaj Kumar and Joseph E. Stiglitz (1986). "The Architecture of Economic Systems: Hierarchies and Polyarchies". In: *American Economic Review* 75.4, pp. 716–727.

Sappington, David E M and Dennis L. Weisman (2010). "Price cap regulation: What have we learned from 25 years of experience in the telecommunications industry?" In: *Journal of Regulatory Economics* 38.3, pp. 227–257.

Schofield, Heather (2014). "The Economic Costs of Low Caloric Intake: Evidence from India". In: *Mimeo: University of Pennsylvania*.

Sekhri, Sheetal (2011). "Public provision and protection of natural resources: Groundwater irrigation in rural India". In: *American Economic Journal: Applied Economics* 3.4, pp. 29–55.

— (2014). "Wells, water, and welfare: The impact of access to groundwater on rural poverty and conflict". In: *American Economic Journal: Applied Economics* 6.3, pp. 76–102.

Simon, Herbert A (1978). "Rationality as Process and as Product of Thought". In: *American Economic Review Papers and Proceedings* 68.2, pp. 1–16.

Sims, Christopher A (2003). "Implications of Rational Inattention". In: *Journal of Monetary Economics* 50, pp. 665–690.

Spiegler, Ran (2006). "The Market for Quacks". In: *Review of Economic Studies* 73.4, pp. 1113–1131.

Spielman, David J. and Rajul Pandya-Lorch (2009). *Millions Fed: Proven Successes in Agricultural Development*. Washington, D.C.: International Food Policy Research institute.

Subramanian, Shankar and Angus Deaton (1996). “The Demand for Food and Calories”. In: *Journal of Political Economy* 104.1, pp. 133–162.

Swanson, Burton E, B J Farner, and Ram Bahal (1989). “The Current Status of Agricultural Extension Worldwide”. In: *Global Conference on Agricultural Extension, Rome (Italy), 4-8 Dec 1989*.

Tirole, Jean (1988). *The Theory of Industrial Organization*. Cambridge, Massachusetts: MIT press.

Tversky, Amos and Daniel Kahneman (1974). “Judgment under Uncertainty: Heuristics and Biases”. In: *Science* 185.4157, pp. 1124–1131.

Unger, Corinna R. (2014). “India’ s Green Revolution: Towards a New Historical Perspective”. In: *South Asia Chronicle* 4, pp. 254–270.

Wilson, Mark (1991). “Reducing the Costs of Public Extension Services: Initiatives in Latin America”. In: *Agricultural extension: Worldwide institutional evolution and forces for change.* Pp. 46–69.

World Bank (2006). “What Are the Dimensions of the Undernutrition Problem in India?” In: *India’s Undernourished Children: A Call for Reform and Action*. 1st ed. Washington, D.C. Chap. 1.

Appendices

Appendix A

Appendices to Chapter 1: The Value of Discretion - Price-Caps and Public Service Delivery

A.1 Derivation of price discrimination formula

Recall that social welfare is given by:

$$SWF(\mathbf{p}) = M(c^*(\mathbf{p}))\left[g_{c_1}\mu\int_{p\left(\varepsilon_D^L\right)}^{\infty} v_i - p\left(\varepsilon_D^L\right) dF\left(v_i|\varepsilon_D^L\right)\right]$$
$$+ M(c^*(\mathbf{p}))\left[g_{c_2}(1-\mu)\int_{p\left(\varepsilon_D^H\right)}^{\infty} v_i - p\left(\varepsilon_D^H\right) dF\left(v_i|\varepsilon_D^H\right)\right]$$
$$+ g_a M(c^*(\bar{p}))\left[\mu\left(p\left(\varepsilon_D^H\right)-\tau\right) D\left(p\left(\varepsilon_D^H\right)|\varepsilon_D^H\right)\right]$$
$$+ g_a M(c^*(\bar{p}))\left[(1-\mu)\left(p\left(\varepsilon_D^L\right)-\tau\right) D\left(p\left(\varepsilon_D^L\right)|\varepsilon_D^L\right)\right] - g_a\int_0^{c^*(\mathbf{p})} c_i dM(c_i)$$

Taking derivatives yields:

$$\begin{aligned}
\frac{\partial SWF(\mathbf{p})}{\partial p\left(\varepsilon_D^L\right)} &= \frac{\partial M(c^*(\mathbf{p}))}{\partial p\left(\varepsilon_D^L\right)}\left[g_{c_1}\mu\int_{p\left(\varepsilon_D^L\right)}^{\infty} v_i - p\left(\varepsilon_D^L\right) dF\left(v_i|\varepsilon_D^L\right)\right] \\
&+ \frac{\partial M(c^*(\mathbf{p}))}{\partial p\left(\varepsilon_D^L\right)}\left[g_{c_2}(1-\mu)\int_{p\left(\varepsilon_D^H\right)}^{\infty} v_i - p\left(\varepsilon_D^H\right) dF\left(v_i|\varepsilon_D^H\right)\right] \\
&+ \frac{\partial M(c^*(\mathbf{p}))}{\partial p\left(\varepsilon_D^L\right)} g_a\left[\left(\mu\left(p\left(\varepsilon_D^L\right)-\tau\right) D\left(p\left(\varepsilon_D^L\right)|\varepsilon_D^L\right) + (1-\mu)\left(p\left(\varepsilon_D^H\right)-\tau\right) D\left(p\left(\varepsilon_D^H\right)|\varepsilon_D^H\right)\right)\right] \\
&- M(c^*(\mathbf{p}))\, g_{c_1}\mu\left(1-F\left(p\left(\varepsilon_D^L\right)|\varepsilon_D^L\right)\right) \\
&+ M(c^*(\mathbf{p}))\, g_a\mu\left(1-F\left(p\left(\varepsilon_D^L\right)|p\left(\varepsilon_D^L\right)\right)\right) \\
&+ M(c^*(\mathbf{p}))\left[-g_a\mu\left(p\left(\varepsilon_D^L\right)-\tau\right) f\left(p\left(\varepsilon_D^L\right)|\varepsilon_D^L\right)\right] \\
&- g_a c'^*(\mathbf{p})\, c^*(\mathbf{p})\, m(\mathbf{p})
\end{aligned}$$

The number of low elasticity types served is given by: $N^L(\mathbf{p}) = M(c^*(\mathbf{p}))\, D^L(\mathbf{p})$. Using the definition of $c^*(\mathbf{p})$ as well as the definition of the elasticites and reordering yields the sufficient statistics formula:

$$\frac{\partial SWF(\mathbf{p})}{\partial p\left(\varepsilon_D^L\right)} = \theta N(\mathbf{p})\frac{g_{c_1}\mu CS\left(p|\varepsilon_D^L\right) + g_{c_2}(1-\mu) CS\left(p|\varepsilon_D^H\right)}{p\left(\varepsilon_D^L\right)\left(\mu D^L(\mathbf{p}) + (1-\mu) D^H(\mathbf{p})\right)}$$
$$+ \varepsilon_D^L N^L(\mathbf{p}) g_a\mu\frac{p\left(\varepsilon_D^L\right)-\tau}{p\left(\varepsilon_D^L\right)} - N^L(\mathbf{p})\mu\left(g_{c_1} - g_A\right)$$

A.2 Supplementary Information

A.2.1 Timing

The project proceeded as follows. From July to August 2015, the ministry collected background data on agents' work environment and activities. During this exercise, I conducted a pilot of the experiment. During November 2016 a workshop with senior central and local government officials introduced the experiment, finalized the design and secured political support at all administrative levels.

The intervention was then implemented in January and February 2016 by a mixed team of ministry staff and private enumerators. Both jointly communicated the campaign instructions to participants. To assure data confidentiality, the private enumerators then independently conducted the baseline survey with participants. After the survey, ministry staff was responsible for the distribution of the vaccine and the communication of final technical instructions relating to the correct application and handling of the vaccine. Agents started the vaccination campaign immediately after receiving the vaccines, and were given three weeks from vaccine distribution to complete the task. The last day of vaccination was February 24, 2016.

I then conducted a phone based follow-up survey with service recipients during March and April 2016. Finally, I conducted an in-person follow-up interview with 311 randomly selected experiment participants during May 2016.

A.2.2 Impact on Remote Farmers in Need

I provide an additional piece of evidence on the trade-offs generated by price-caps by focusing on the availability of I-2 substitutes in villages as a dimension of heterogeneity. In my sample, approximately 11% of farmers have access to an imperfect substitute for I-2 provided through private markets. Table 1.5 shows that while the price-cap was effective in reducing prices by roughly 15 Tanzanian Shillings, this effect is driven exclusively by transactions with farmers who don't have access to this substitute. In contrast, the treatment effect for farmers with access to the substitute is positive, small and statistically indistinguishable from zero. This suggests that competition induced by the substitute reduces the surplus available to agents, hence driving prices to a level where the price-cap does not bind.

Price-caps therefore only affect the surplus available to the agent in areas where the substitute is unavailable.

As the substitute is also more likely to be absent in remote areas, this reduction in available surplus resulting from the price-cap substantially reduces agents' incentives to target villages without access to the substitute. To assess this intuition, column 2 in table A.1 investigates how price-caps affect the likelihood that a farmer with no access to the substitute will be served. The results show that the price-cap treatment reduces the likelihood that farmers without access to the substitute are served by 3% to 4%. Taken together, these results further highlight the aforementioned tension: price-caps make services accessible to farmers in need. But when those farmers live far away, price-caps can be counter-productive, as caps reduce the likelihood that agents will travel to the remote farmers.

A.3 Data

A.3.1 Baseline Data

The baseline survey was administered to every participant during the vaccination distribution and was completed before any vaccinations occurred. The survey included detailed questions on agents' demographics, education, work history and alternative income sources. It also collected data on agents' work environment, including information on travel times to villages, transport methods, private providers of veterinary services and agents' interaction with their supervisors.

As part of this survey I also administered two questions aimed at eliciting an incentive-compatible measure of pro-social motivation toward animal health causes. First, I designed a contextualized dictator game. Agents were told that they would receive a lunch allowance of 10'000 Tanzanian Shillings (approximately \$4.50), which they could keep for themselves or donate, in part or in full, to TVLA to purchase subsidized vaccines for the next vaccination campaign. The amount donated is taken as a proxy for the agents' motivation for the cause. The median donation in the dictator game was 1000 Tanzanian Shillings.

Second, agents were given a map with 9 fields, each detailing a possible motivation for why they chose to work as a livestock field officer. Some stated motivations were intrinsic (e.g. *"my job allows me to help farmers when their animals are sick"*) while others reflected extrinsic sources of motivation (e.g. *"my job offers a stable income"*). Enumerators then gave participants 50 maize grains and asked them to distribute the grains between the different fields according to how important each reason was when they were making their career choice. The relative amount of beans allocated to fields that reflect intrinsic motivations then acts as a proxy for the agents' motivation for the cause.

Both measures were designed to increase the likelihood of being rank-preserving in order to assure that measures remain valid even if agents exaggerate their donation or grain allocation because of social pressure.

A.3.2 How accurate is the receipt data?

When assessing the validity of the receipt data it is important to remember that accurate reporting was financially incentivized, as verified receipts attracted a bonus payment of 60 Tanzanian Shillings per vaccination. Crucially, I don't consider receipts that were submitted without a contact phone number for farmers to be complete and therefore don't count them towards the total number of farmers served. Receipts without phone numbers are therefore also ineligible for the bonus payment. Agents were made aware of this rule during the roll-out and were encouraged to identify alternative contact numbers for farmers should they not own a phone, for example by providing the number of their neighbor or of the village leader. While this requirement might have incentivized employees to target farmers more likely to own cellphones, the need to provide phone numbers was present for all treatment groups and is therefore unlikely to challenge the internal validity of the experiment. In addition, identifying farmers' contact numbers does not appear to have been a problem: Less than 4% of receipts were submitted without phone numbers and ministry staff tasked with supervising the campaign did not receive any complaints about challenges with identifying cellphone owners.

A.3.3 Farmer Survey Procedures

For cost reasons, the follow-up survey with farmers was implemented as a phone survey. The phone survey procedures were designed to maximize the likelihood of reaching service recipients. Enumerators were instructed to call each number on three different days, once in the morning and once in the afternoon. After every unsuccessful attempt, enumerators sent a text message to recipients informing them about the objective of the call and asking for an appointment to administer the survey.

A.4 Supplementary Tables

Table A.1: Effect of Treatment on Likelihood of Serving Remote Farmers in Need

Outcome Variable:	**Price**	**No substitute available**
	(1)	(2)
Price-Cap Treatment	3.887	-0.0355**
	(8.379)	(0.0177)
No Substitute	5.702	
	(3.982)	
Price-Cap $\times$ No Substitute	-20.07**	
	(8.520)	
Observations	3,044	3,097
Data Source	Farmer Survey	Farmer Survey
District Fixed Effects	Yes	Yes
Agent Controls	Yes	Yes
Farmer Controls	Yes	Yes
Control Mean	76	0.89
Control St. Dev.	50	0.31

Notes: Standard Errors are clustered at the ward level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Column 1 presents coefficient estimates of the transaction price on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables and an interaction with whether a recipient has access to a private substitute for I-2. Column 2 presents coefficient estimates of a regression of a binary variable indicating whether a recipient has access to a substitute for ND vaccines on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables.

Table A.2: Effects on Participation

Outcome Var.	Training Att.	Participation	Overall Part.	% Verified	% Price Correct
Treatment	(1)	(2)	(3)	(4)	(5)
Price-Cap	0.0130	0.0166	0.0224	0.00476	-0.0358*
	(0.0275)	(0.0265)	(0.0320)	(0.0219)	(0.0188)
Observations	990	832	990	740	675
District FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
Control Mean	0.83	0.82	0.68	0.69	0.83
Control St. Dev.	0.38	0.38	0.47	0.30	0.25

Notes: Standard Errors are clustered at the ward level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. All specifications control for stratification variables and district fixed-effects. The sample for columns 1 and 3 contains the population of agents in all sample districts. The sample for columns 2 are all agents who attended the training and who performed vaccinations, respectively. "Overall Participation" refers to the likelihood of attending training and performing vaccinations after training. % Verified refers to the fraction of transactions reported through receipts that could be verified to have taken place. % Price Correct refers to the fractions of receipts that reported a price that could be verified through follow-ups with farmers.

Table A.3: Robustness - Replication of Price-Cap Treatment Effects excluding Cattle Owners

Outcome: % of Farmers or Transactions with given Characteristic				
Characterstic:	**Owns Cattle**	**Price**	**Main Livelihood is Agriculture**	**Village has no Private Provider**
	(1)	(2)	(3)	(4)
Price-Cap	0.00526	-15.61***	0.0535**	-0.0421**
	(0.0222)	(3.711)	(0.0266)	(0.0205)
Observations	3,098	2,165	2,204	2,204
Observation Level	Transaction	Transaction	Transaction	Transaction
Data Source	Farmer Survey	Farmer Survey	Farmer Survey	Farmer Survey
District Fixed Effects	Yes	Yes	Yes	Yes
Agent Controls	Yes	Yes	Yes	Yes
Farmer Controls	Yes	Yes	Yes	Yes
Control Mean	0.29	79.57	0.76	0.89
Control St. Dev.	0.45	53.69	0.43	0.30

Notes: Standard Errors are clustered at the individual level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. All regressions apart from column 1 exclude cattle owners from the sample. Column 1 presents coefficient estimates of a regression of a binary variable indicating whether a farmer owns cattle or not on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Column 2 presents coefficient estimates of a regression of the transaction price on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Column 3 presents coefficient estimates of a regression of a binary variable indicating whether the recipient's main livelihood is from agriculture on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Column 4 presents coefficient estimates of a regression of a binary variable indicating whether the recipient's village has a private provider of veterinary services on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables.

Table A.4: Robustness - Treatment Effects on Leakage and Transaction Behavior

Outcome Variable	**Vaccine Loss (Proxy)**	**Mentioned Vaccine Cost**
Treatment	(1)	(2)
Price-Cap	-0.0525	-.0047
	(0.0895)	(0.0344)
Observations	819	832
Data Source	Farmer Survey	Farmer Survey
District Fixed Effects	Yes	Yes
Controls	Yes	Yes
Control Mean	0.96	0.65
Control St. Dev.	1.45	0.47

Notes: Standard Errors are clustered at the ward level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Column 1 presents coefficient estimates of the fraction between confirmed vaccinations and the number of vaccine doses collected on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. Column 2 presents coefficient estimates of a binary variable indicating whether agents mentioned the vaccine cost during the service delivery process on an indicator variable for the treatment, district fixed effects as well as ward-level stratification and control variables. All columns employ farmer survey data.

A.5 Supplementary Figures

Figure A.1: Receipt Format

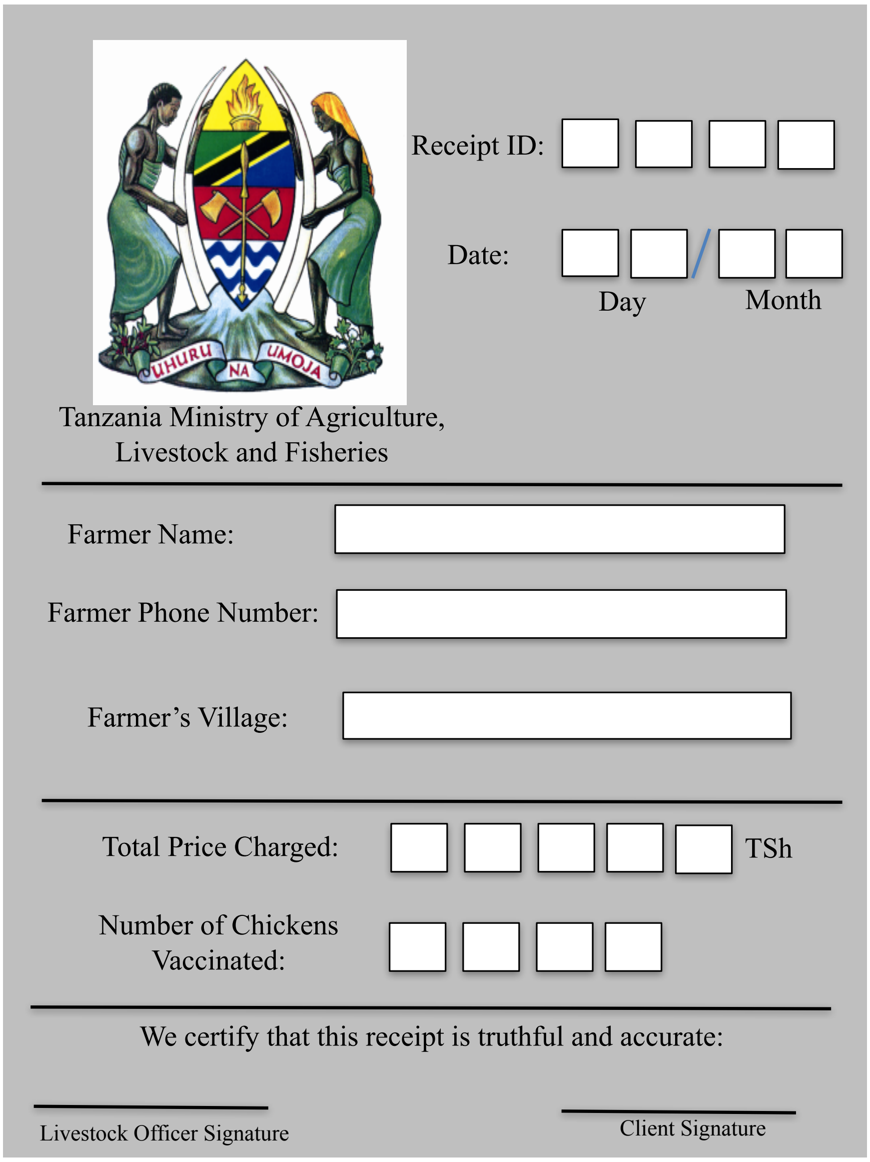

UHURU NA UMOJA

Receipt ID:

Date: /

Day Month

Tanzania Ministry of Agriculture, Livestock and Fisheries

Farmer Name:

Farmer Phone Number:

Farmer's Village:

Total Price Charged: TSh

Number of Chickens Vaccinated:

We certify that this receipt is truthful and accurate:

Livestock Officer Signature

Client Signature

Figure A.2: Distribution of Unverifiable Receipts

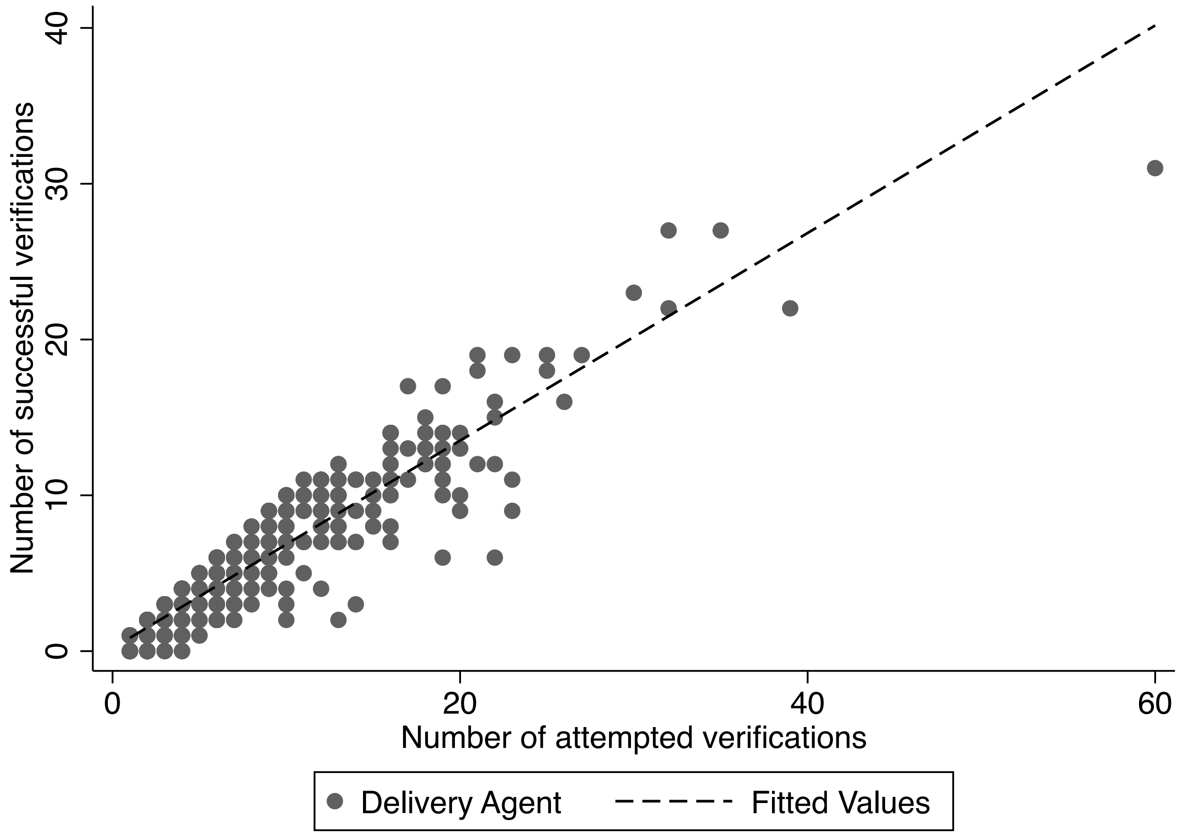

Appendix B

Appendices to Chapter 3: Priming the Pump - Irrigation, Nutrition and Agricultural Productivity

B.1 Supplementary Tables

Table B.1: Effect of Tube Well Adoption on Population

	Population		
	Total	Rural	Urban
	(1)	(2)	(3)
Tube Well Share	1,948*	2,733***	-785.6
	(1,001)	(644.2)	(787.9)
Observations	1,029	1,029	1,029
Data Source	ICRISAT	ICRISAT	ICRISAT
Estimation Method	2SLS	2SLS	2SLS
District FE	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes

Notes: Standard Errors are clustered at the district level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Columns 1 presents coefficient estimates of a two-stage-least-squares regression of the total population in a given district on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Columns 2 presents coefficient estimates of a two-stage-least-squares regression of the rural population in a given district on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Columns 3 presents coefficient estimates of a two-stage-least-squares regression of the urban population in a given district on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. All specifications employ ICRISAT data, which is based on the 1971, 1981, 1991 and 2001 census. The share of cropped land irrigated using tube wells is instrumented for using variation in hydrogeological structures.

Table B.2: Robustness: Effect of Tube Wells on Calorie Intake with Controls and Heterogeneous Time Trends

	Median Calorie Intake		
	(1)	(2)	(3)
Tube Well Share	2,521***	1,002	3,662
	(900.2)	(671.2)	(2,476)
Observations	977	987	987
Data Source	Constr.	Constr.	Constr.
District Sample	Full	Full	Full
Estimation Method	2SLS	2SLS	2SLS
District FE	Yes	Yes	Yes
Year Fixed Effects	Yes	No	Yes
State-Year Fixed Effects	No	Yes	No
State Specific Linear Trend	No	No	Yes
Controls	Yes	No	No

Notes: Standard Errors are clustered at the district level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Columns 1 to 3 present coefficient estimates of a two-stage-least-squares regression of median daily per capita calorie intake on a continuous variable measuring the share of cropped land irrigated using tube wells. Column 1 controls for calorie prices, area sown, total per capita expenditure in the district, district population and district fixed effects. Column 2 controls for state-year fixed effects and district fixed effects. Column 3 controls for state specific linear time trends, district as well as year fixed effects. The share of cropped land irrigated using tube wells is instrumented for using variation in hydrogeological structures.

Table B.3: Robustness: Effect of Tube Well Adoption on Food Expenditure, ICRISAT Districts

	Share of food expenditure on ...						
	Food Exp. Share	Staples	Cereals	Animal Prd.	Veg.	Fruit	Sugar
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Tube Well Share	0.890**	0.222	0.306**	0.721***	-0.232**	0.0357	-0.173***
	(0.344)	(0.176)	(0.150)	(0.231)	(0.0926)	(0.0373)	(0.0653)
Observations	987	987	987	987	987	987	987
Data Source	Constr.	Constr.	Constr.	Constr.	Constr.	Constr.	Constr.
District Sample	ICRISAT	ICRISAT	ICRISAT	ICRISAT	ICRISAT	ICRISAT	ICRISAT
Reference Year	1987	1987	1987	1987	1987	1987	1987
Estimator	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
District FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Standard Errors are clustered at the district level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Columns 1 to 8 present coefficient estimates of a regression of the average share, as part of monthly household expenditure, of various food items on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Column 1 uses food expenditure as an outcome variable, whereas column 2 uses expenditure on staples, column 3 uses expenditure on non-staple cereals, column 4 uses expenditure on animal products, column 5 uses expenditure on vegetables, column 6 uses expenditure on fruits, and column 7 uses expenditure on sugar and sugar-related food items. Estimates are obtained through a two-stage-least-squares regression in which tube well adoption is instrumented for using variation in hydrogeological structures.

Table B.4: Robustness: Effect of Tube Well Adoption on Productivity, 1987 to 2004

	Yield		HYV	
	Rice	Wheat	Rice	Wheat
	(1)	(2)	(3)	(4)
Tube Well Share	1.165	3.096***	170.1	401.2**
	(0.840)	(0.718)	(202.0)	(169.2)
Observations	4,639	4,187	3,121	3,111
Data Source	ICR.	ICR.	ICR.	ICR.
Estimation Method	2SLS	2SLS	2SLS	2SLS
District FE	Yes	Yes	Yes	Yes
Year Fixed Effects	Yes	Yes	Yes	Yes
Base Year	1987	1987	1987	1987

Notes: Standard Errors are clustered at the district level. *** (**) (*) indicates significance at the 1 (5) (10) percent level. Columns 1 and 2 present coefficient estimates of a two-stage-least-squares regression of rice and wheat yields, respectively, on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. Columns 3 and 4 present coefficient estimates of a two-stage-least-squares regression of the total area cropped with HYV rice and wheat seeds, respectively, on a continuous variable measuring the share of cropped land irrigated using tube wells as well as district and year fixed effects. All specifications employ ICRISAT data and cover the years 1987 to 2004. The share of cropped land irrigated using tube wells is instrumented for using variation in hydrogeological structures.