# Three Essays on Energy and Development Economics

by

Faraz Usmani

University Program in Environmental Policy
Duke University

Date: ______________________

Approved:

______________________
Subhrendu Pattanayak, Co-Supervisor

______________________
Marc Jeuland, Co-Supervisor

______________________
Brian Murray

______________________
Robyn Meeks

______________________
Seth Sanders

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in the University Program in Environmental Policy
in the Graduate School of Duke University
2019

# Abstract

## Three Essays on Energy and Development Economics

by

Faraz Usmani

University Program in Environmental Policy
Duke University

Date: ____________________________
Approved:

____________________________________
Subhrendu Pattanayak, Co-Supervisor

____________________________________
Marc Jeuland, Co-Supervisor

____________________________________
Brian Murray

____________________________________
Robyn Meeks

____________________________________
Seth Sanders

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the University Program in Environmental Policy in the Graduate School of Duke University
2019

# Abstract

Global energy-use patterns are characterized by deep inequality. Electricity is indispensable for households, clinics, schools and firms, yet over a billion people live without it. At the same time, nearly three billion rely on traditional stoves and polluting biomass fuels (such as firewood) for their basic energy needs. The resulting household air pollution causes four million deaths annually, a health burden borne disproportionately by women. The international community has hastened to respond to this global energy challenge. This dissertation highlights how—and under what conditions—policies that seek to ensure universal access to modern energy deliver expected environmental and development benefits.

In the first chapter, I ask what drives heterogeneity in the impacts of large-scale rural electrification. Prior evidence on the labor-market impacts of grid electrification is mixed. I hypothesize that variation in local economic conditions—which can complement investments in infrastructure—may help explain why, and combine two natural experiments in India within a regression discontinuity design to test this hypothesis. Most of the world's guar, a crop that yields a potent thickening agent used during hydraulic fracturing ("fracking"), is grown in northwestern India. The rapid rise of fracking in the United States induced a parallel commodity boom in Indian guar production, resulting in a large positive shock to rural economic activity. Leveraging population-based discontinuities in the contemporaneous roll-out of India's massive rural electrification scheme, I show that access to electricity significantly increased non-agricultural employment in villages located in

India's booming guar belt. Where these complementary economic conditions were lacking, electrification had almost no discernible impact. Using a firm-level panel dataset, I then provide suggestive evidence that this growth in non-farm work is partly driven by the rise of electricity-intensive firms that complement agricultural production. In line with the prior literature, I show that electrification alone may not be sufficient to deliver economic benefits, but I also demonstrate that, when combined with complementary economic conditions on the ground, access to electricity can enable individuals, households and firms to take advantage of new opportunities in potentially welfare-enhancing ways.

In the second chapter, I turn to household-level energy use and empirically evaluate the role played by non-governmental organizations (NGOs) in delivering environmental, energy and development interventions in remote, rural settings. I develop a model of household decision-making to evaluate how NGOs address implementation-related challenges and influence intervention effectiveness. To test the model's predictions, I apply quasi-experimental methods to household-survey data from a randomized controlled trial designed to promote clean-cooking solutions in rural India. I uncover a large, positive and statistically significant "NGO effect": prior engagement with the implementing NGO increases the effectiveness of the intervention by at least thirty percent. These findings provide some of the first causal evidence on how NGOs directly influence outcomes, which has implications for the generalizability of experimental research conducted jointly with such local partners. In particular, attempts to scale up findings from such work may prove less successful than anticipated if the role of NGOs is insufficiently understood. Alternatively, policymakers looking to scale up could achieve greater success by fostering partnerships with trusted local institutions.

In the final chapter, I consider how heterogeneity in households' preferences influences demand for energy technologies. I conduct technology-promotion campaigns followed by second-price, sealed-bid ("Vickrey") auctions for two cleaner cooking technologies with over 1,000 households across seventy communities in rural Senegal. I induce exogenous variation in the extent to which these promotion activities cater to heterogeneous preferences by randomly assigning a subset of communities to an auction arm in which both devices are promoted jointly. Consistent with a model in which preferences are constructed—and not simply revealed—as agents make repeated choices, joint promotion lowers willingness to pay for the relatively less familiar alternative compared to settings in which the two devices are promoted exclusively. Rather than simply providing additional choices, implementers looking to enhance uptake of improved technologies must instead devise approaches to help potential end-users think carefully through trade-offs, crystallize and understand their own preferences, and identify solutions that fit their needs.

# Contents

# List of Tables

# List of Figures

# Acknowledgements

I am grateful for the generous support I have received from the United Nations University World Institute for Development Economics Research (UNU-WIDER), the Duke University Energy Initiative, the Duke Global Health Institute (DGHI), the Sustainable Energy Transitions Initiative (SETI), the Nicholas School of the Environment, the Department of Economics, and the Graduate School. The last two chapters of this dissertation were partly made possible by support from the United States Agency for International Development (USAID) and the International Initiative for Impact Evaluation (3ie), respectively.

I am indebted to the members of my committee—Subhrendu Pattanyak, Marc Jeuland, Brian Murray, Robyn Meeks and Seth Sanders—for their guidance. I am especially grateful to my advisors—Subhrendu Pattanayak and Marc Jeuland—for simultaneously serving as mentors, coauthors and friends. I have also benefited immensely from valuable advice, useful critiques and insightful comments from Indermit Gill, Jörg Peters, Alex Pfaff, Brian Pinto, Billy Pizer, Steve Sexton, Smriti Sharma, Saurabh Singhal, Marty Smith, Eswaran Somanathan and Jeff Vincent.

I am incredibly fortunate to count Laura Bellows, Anna Birkenbach, Ipsita Das, Katy Hansen, Esther Heesemann, Farah Hegazi, Rob Fetter, David Kaczan, P. P. Krishnapriya, Luci Lenz, Jennifer Meyer and Josh Rivenbark among my harshest critics *and* my most vocal advocates. This dissertation would be incomplete without them.

Finally, I thank Jawed Usmani and Farah Usmani for always lifting me up, and Sabah Usmani and Daphne Yin for always keeping me grounded.

# Introduction

Electricity is indispensable for households, clinics, schools and firms. Electric lighting enables productive activities outside of daytime hours, electric appliances reduce household drudgery, and electric machinery drives agricultural and industrial productivity (Barnes, 2005). Yet over a billion people (primarily in rural areas of low- and middle-income countries) live without any access to electricity whatsoever, and many more are served by unreliable systems that can support little more than a light bulb (International Energy Agency, 2011). The remote, dispersed nature of rural communities increases the costs associated with extending and maintaining the electrical grid; where the grid has reached, prohibitive upfront connection fees and low expected consumption levels often dissuade rural customers from connecting (Lee et al., 2018; World Bank, 2000). Political interference in electric utilities' pricing and operational guidelines exacerbates these challenges (Mahadevan, 2019).

At the same time, nearly three billion people rely on solid fuels and traditional stoves to meet their daily cooking and heating needs. The resulting household air pollution causes over four million deaths annually, a health burden borne disproportionately by women (Adair-Rohani et al., 2016). Daily collection of solid fuels (such as firewood) restricts opportunities for education and employment; their inefficient combustion generates emissions that intensify climate change. Improved cookstoves (ICS) and modern fuels can help ease this environmental–health–development burden, yet uptake in rural areas—where they are often most needed—remains low (Lewis and Pattanayak, 2012). Lack of awareness about cleaner alternatives, income and liquidity constraints, and the absence of robust supply chains that provide reliable access to modern fuels and technologies pose formidable obstacles.

Inequitable access to modern energy services is not inevitable. In 2016, 87 percent (59 percent) of the global population had access to electricity (cleaner cooking fuels and technologies), up from 71 percent (49 percent) in 1990 (World Bank, 2018b). Large-scale grid- and off-grid electrification programs in countries such as Bangladesh, Kenya and Tanzania have been drivers of these gains (World Bank, 2018a). Similarly, a growing body of evidence demonstrates that policies that reduce monetary costs can increase demand for cleaner cooking fuels and technologies (Bensch et al., 2015; Levine et al., 2018; Pattanayak et al., 2016). Others point to the importance of non-monetary channels, such as social networks (Miller and Mobarak, 2015).

Yet the success of efforts that seek to ensure universal access to energy is also not guaranteed. Indeed, it may be closely tied to specific places, populations and programs, and not easily translatable to new contexts (Ravallion, 2009). Evidence from research looking to rigorously evaluate the impact of access to modern energy services in poor countries is similarly mixed. Differences in contexts across time and space, which can interact with policies in ways that hamper or bolster effectiveness, are likely drivers of this heterogeneity (Pritchett and Sandefur, 2014). Understanding when, where and why policies are most likely to be effective—and not simply whether or not they are—thus requires a closer understanding of the mechanisms through which expected benefits might materialize (Deaton, 2010). In this dissertation, I look to shed light on the underlying mechanisms that connect energy to development through three distinct lenses: (i) how local economic conditions complement large-scale energy infrastructure; (ii) the role local institutional partners play in delivering energy and development interventions in remote, rural settings; and (iii) heterogeneity in beneficiaries' preferences over energy technologies. I study these mechanisms in two key domains: rural electrification and clean cooking.

One channel through which electrification is believed to benefit rural economies is the creation of new economic opportunities and jobs (United Nations, 2018). For instance, access to electricity can greatly increase the productivity of tasks that do not necessarily require electricity (such as tailoring). At the same time, it can give rise to new opportunities that were previously not possible (such as welding). Together, these two pathways can lead to higher wages and large sectoral shifts in employment (e.g., from agricultural to industrial production). Yet the evidence on the labor-market impacts of large-scale grid electrification remains mixed (Bernard and Torero, 2015; Burlig and Preonas, 2016; Dinkelman, 2011; Lenz et al., 2017; Lipscomb et al., 2013). These heterogeneous results may be partly explained by (the lack of) complementary economic conditions. Indeed, household consumption in poor rural economies may not be high enough to generate sufficient demand for a variety of goods and services, which in turn means that increased labor productivity and new opportunities due to electrification may have little to no discernible impact on labor-market outcomes. This is the main finding of my first chapter, which is coauthored with T. Robert Fetter. Specifically, in line with the prior literature, we first show that electrification alone may not be sufficient to deliver expected labor-market benefits. We then demonstrate that access to electricity can drive large shifts out of low productivity agricultural tasks into relatively high productivity non-agricultural employment when combined with economic opportunity on the ground.

We do this by exploiting the interaction of two natural experiments in India. As the hydraulic fracturing ("fracking") boom began in the United States, it induced a parallel commodity boom in northwestern India in the production of a crop called guar, which provides a key input into the fracking process, resulting in a large exogenous shock to rural economies in the region. Almost simultaneously, India began rolling out its massive rural electrification scheme, which aimed to electrify approximately 400,000 villages across 27 states. It prioritized villages for electrification on the basis of a strict population-based threshold, giving rise to discontinuous changes in a village's probability of being electrified. We combine these two natural experiments within a regression discontinuity design to evaluate how the causal effect of electrification on labor-market outcomes varies with exogenous changes in economic contexts. We next dig deeper into mechanisms by providing suggestive evidence that these labor-market dynamics are driven by the rise of complementary non-agricultural opportunities. An increase in guar production necessitates a shift in the labor force towards guar processing, which benefits from upgraded electricity infrastructure. Simultaneously, increased wages and agricultural profits from both the production boom and new processing opportunities can be reinvested in household enterprises, which may also benefit from electricity connections. Consistent with this, we uncover a large increase in (i) the number of workers at firms related to the industrial (electricity-intensive) parts of the guar production chain (such as guar processing); and (ii) home production of income-generating products in electrified guar-growing regions. Finally, we find almost no discernible evidence of any effect of electrification on these labor-market outcomes in villages or regions located in the rest of India, suggesting that complementary economic conditions play a crucial role in driving the impacts of large-scale electrification infrastructure.

In arriving at this conclusion, we demonstrate the importance of explicitly studying drivers of heterogeneity. Differences in local economic conditions and contexts may help explain global variation in the impacts of electrification and other large infrastructure projects. Indeed, there are likely other channels that complement infrastructure investments and drive heterogeneity in similar ways. Insofar as these context-specific characteristics can be identified *ex ante*, they can help improve spatial targeting of resource-intensive infrastructure projects in low- and middle-income countries, where the opportunity costs associated with such investments can be especially high.

The second chapter, coauthored with Marc Jeuland and Subhrendu Pattanayak, turns to household-level energy use and considers the importance of implementing institutions—in particular, non-governmental organizations (NGOs)—and their effectiveness at delivering environmental, energy and development interventions in remote, rural settings. NGOs increasingly play lead roles in implementation on the ground (Werker and Ahmed, 2008). That they are (at least in theory) nimble and efficient has made them attractive partners for international donors. In 2012, for instance, Organisation for Economic Co-operation and Development (OECD) countries channeled over $17 billion of overseas development assistance to—and through—NGOs (Aldashev and Navarra, 2014). Yet outside of a nascent body of work on the association between "implementer identity" and outcomes, little is known about the ways in which NGOs (and the relationships they foster with their beneficiaries) directly impact the effectiveness of the interventions they implement (Bold et al., 2013; Cameron and Shah, 2017; Grossman et al., 2016; Henderson and Lee, 2015; Lewis et al., 2015).

This knowledge gap has serious implications for the growing role of NGOs as partners in research, particularly in the case of randomized controlled trials. In these settings, partnerships with NGOs provide researchers with ready access to target populations, local expertise, human resources and operational infrastructure, all of which lower the costs of doing research. For NGOs, these collaborations create opportunities to conduct rigorous evaluations of the impact of their initiatives. Yet this seemingly symbiotic relationship can mask an underlying "NGO effect"—the direct impact of the NGO–beneficiary relationship on the effectiveness of the intervention—that works to undermine the scalability and generalizability of the solutions deemed effective in applied research (Berge et al., 2012; Peters et al., 2018; Vivalt, 2017).

We provide some of the first rigorous evidence of the magnitude and direction of this effect. To do so, we first use *ex ante* propensity score matching to create a sample of observationally similar villages that are differentiated by prior exposure to a local development NGO. In partnership with this NGO, we then randomly assign nearly 100 geographically distinct hamlets within these villages (covering a sample of almost 1,000 households) to treatment and control groups as part of an experimental ICS-promotion intervention. Our results suggest that the intervention increases adoption: nearly half of all households targeted by the promotion campaign purchased an ICS. However, our study design also allows us to identify the direct impact of the NGO on ICS adoption and energy-use patterns. We uncover a large, positive and statistically significant "NGO effect"—purchase rates are nearly thirteen percentage points (28 percent) higher in treated communities with prior interactions with the NGO. Using a "triple-differences" specification, we find that treated households in NGO communities are also sixteen percentage points more likely to use intervention stoves than treated households in communities without a prior relationship with the NGO, representing a fifty percent increase in the size of the treatment effect. Consistent with these patterns of adoption and use, treated households in NGO communities exhibit significant reductions in the use of solid fuels and in fuel-collection times. In contrast, we find no evidence of similar improvements in energy-use patterns for treated households in non-NGO communities. Our stratified study design, therefore, reveals that we would have considerably overestimated the effectiveness of our intervention had it been a typical randomized evaluation conducted in partnership with the NGO. These findings highlight how attempts to scale up findings from research conducted jointly with NGOs may prove less successful than anticipated if the role of NGOs is insufficiently understood. Alternatively, policymakers looking to scale up could achieve greater success by enlisting trusted local partners.

The third and final chapter, coauthored with Marc Jeuland and Ousmane Ndiaye, maintains a focus on household energy and considers heterogeneity in households' preferences over attributes of energy technologies as a driver of uptake. The context for this chapter is Senegal, where over 95 percent of the rural population uses firewood for its primary energy needs. Against this backdrop, Senegal's government endeavors to lower reliance on polluting traditional fuels by promoting cleaner-burning ICS. The program with the greatest potential reach is FASEN (*Foyers améliorés au Sénégal*), which promotes the *Jambar* stove, a simple, low-cost device that comes in charcoal and firewood versions. Yet FASEN's success has mainly been restricted to urban centers, where charcoal is the primary solid fuel. Making inroads into rural areas, where potential end-users rely almost exclusively on firewood, has proven challenging. One reason for this might be the reliance of ICS promotion efforts on "one size fits all" approaches (where only one type of device is presented to beneficiaries across different settings) that insufficiently cater to households' heterogeneous preferences and unique energy-use needs (Lambe and Atteridge, 2012; Rhodes et al., 2014).

A natural implication of this critique is that the promotion of multiple distinct alternatives should better cater to diverse energy-use needs and enhance intervention effectiveness. To test this hypothesis, we design and implement a technology-promotion campaign followed by second-price, sealed-bid ("Vickrey") auctions featuring two biomass ICS with over 1,000 randomly selected households across seventy rural communities in Senegal. The first of the two devices that feature in our intervention is the *Jambar*, while the second stove is a considerably more efficient device manufactured internationally called the *Jumbo Zama*. By randomly assigning communities to one of three auction arms—two where each stove was promoted and auctioned exclusively, and one where both stoves feature together—we induce exogenous variation in the number of alternatives presented to sample households. Vickrey auctions subsequently allow us to elicit households' willingness to pay (WTP) for one or both of the devices under each of these conditions. In contrast to prior expectations, we find that random allocation to the joint *Jambar–Jumbo Zama* auction lowers WTP for the *Jumbo Zama* while having no distinguishable impact on WTP for the *Jambar*. These results are consistent with a model in which households' preferences over attributes of energy technologies are constructed—and not simply revealed—as repeated choices relating to unfamiliar devices are made (Hoeffler and Ariely, 1999). Therefore, rather than simply providing additional choices, implementers looking to enhance uptake of improved technologies must instead devise approaches to help potential end-users think carefully through trade-offs, crystallize and understand their own preferences, and identify solutions that fit their needs.

While these chapters focus on distinct questions in different geographical contexts, two common threads run through each of them. First, in each chapter, I deploy a diversity of rigorous methodological tools to conduct impact evaluations in the overlapping domains of energy and international development. The recent inclusion of energy-access targets as part of the United Nations' Sustainable Development Goals has thrust energy to the fore of development policy, and governments and international organizations alike are mobilizing considerable resources to achieve access for all (International Energy Agency, 2011; United Nations, 2018). Both electricity and clean cooking are central to this agenda. In bringing to bear rigorous empirical techniques to evaluate the impacts of efforts to enhance access to modern energy services in real-world settings, the chapters in this dissertation help identify and fill crucial, policy-relevant knowledge gaps.

A second common theme that connects these chapters is an aim to rigorously study drivers of heterogeneity directly. For instance, the first chapter points to "average treatment effects" that, in and of themselves, advance our knowledge of the impacts of electrification to a limited degree. Specifically, ignoring underlying spatial variation in complementary economic conditions might lead one to conclude that, on average, access to electricity has little to no discernible impact on labor-market outcomes in thousands of Indian villages. That the policy implications one might draw from such an analysis differ starkly from those that emerge from analyses in which variation in local economic conditions is studied explicitly should give any policy analyst pause when considering overall results that might be masking significant underlying variation. Similarly, in the second chapter, a "naive" analysis that ignores hamlets' prior exposure to the implementing NGO might conclude that the randomized evaluation of our ICS promotion efforts—indeed any randomized evaluation conducted jointly with local implementing partners, more generally—suggests that the intervention was considerably more effective than it truly was. It is precisely this concern for underlying heterogeneity that motivates the auction-based study design presented in the third and final chapter, which focuses on households' heterogeneous preferences as a driver of the effectiveness of interventions that seek to promote uptake of cleaner energy technologies. In contrast to hypotheses implicit in prominent calls for interventions to eschew "one size fits all" approaches to ICS promotion, the results from this chapter suggest that the provision of multiple alternatives alone may not increase uptake. These results are consistent with insights from prior cost-benefit analyses, which suggest that private benefits to households from the use of cleaner cooking technologies may be neither positive nor large (Jeuland and Pattanayak, 2012; Jeuland et al., 2018b; Pinto, 2016).

Taken together, the chapters of this dissertation shed light on when, where and why energy and development policies deliver expected benefits. From a policy perspective, the insights they generate can help guide and improve spatial targeting of policies whose opportunity costs can be especially high in low- and middle-income countries. In addition, these insights also highlight the importance of directly connecting researchers to local policymakers, planners and implementers; researchers provide rigorous evidence of the impacts of public goods, but this evidence is especially valuable when it is informed by local actors' deep contextual knowledge. Ultimately, it is my hope that the insights that emerge from the work in this dissertation contribute to helping low- and middle-income countries meet the ambitious energy-access targets on which they have rightfully set their eyes.

# 1

# Fracking, farmers, and rural electrification in India

*With T. Robert Fetter*

## 1.1 Introduction

Over a billion people worldwide lack access to electricity, and many more are served by unreliable systems capable of supporting little more than a light bulb. The belief that access to reliable electricity catalyzes job creation and economic growth—reflected in the inclusion of energy access targets as part of the United Nations' Sustainable Development Goals—has thrust energy to the fore of development policy (United Nations, 2018). Indeed, governments and international organizations alike are mobilizing considerable resources to ensure access for all. According to the International Energy Agency (2011), over $9 billion was spent in 2009 to extend modern energy services to underserved populations, a figure that it estimates must rise to over $48 billion per year by 2030 in order to achieve universal access. Yet the evidence on the impacts of such efforts remains mixed. Dinkelman (2011) and Lipscomb et al. (2013), for instance, identify large positive effects on employment as a result of rural electrification in South Africa and Brazil, respectively. Burlig and Preonas (2016), on the other hand, find that the effects of rural electrification on labor-market outcomes in India are far more muted. Others have uncovered similarly lackluster impacts in the African context (Bernard and Torero, 2015; Lenz et al., 2017).[1]

This lack of consensus surrounding the benefits of grid expansion highlights both a significant knowledge gap and a critical policy challenge. Indeed, the world's poor are constrained by far more than a lack of access to modern energy services (Banerjee and Duflo, 2007), and there may be profound opportunity costs associated with large-scale investments in energy infrastructure in low- and middle-income settings. India alone is home to nearly 250 million people living without electricity (International Energy Agency, 2015). If electrification by way of resource-intensive grid expansion is foundational in promoting livelihoods among unconnected populations, it represents a necessary first step for development policy. If, on the other hand, expected benefits are highly uncertain—or, worse, illusory—scarce public resources are better targeted elsewhere, and cost-effective approaches that enhance access to only rudimentary energy services (such as basic lighting) may be more appropriate (Grimm et al., 2017).

We exploit the interaction of two natural experiments in India to shed new light on this debate. As the hydraulic fracturing ("fracking") boom began in the United States, it induced a parallel commodity boom in India in the production of an otherwise obscure crop called guar in India. Guar provides a key input into the fracking process and is primarily grown in the semi-arid northwestern tracts of the country by small and marginal farmers (Rai, 2015). Between 2006 and 2011, its price increased by over 1,000 percent, resulting in a large exogeneous shock to rural economies in the region. Almost simultaneously, India began rolling out its massive rural electrification scheme, which aimed to electrify approximately 400,000 villages across 27 states. It prioritized villages for electrification on the basis of a strict population-based threshold, giving rise to discontinuous changes in a village's probability of being electrified. We combine these two natural experiments

[1] In a recent review of the empirical literature, Bonan et al. (2017) note that the current evidence on the impacts of electrification on adults' time allocation and labor activities suggests "mild increases in employment and labor supply, particularly for women, non-agricultural activities and more formal activities" but that the magnitude of such effects "varies significantly across studies and geographical areas."

within a regression discontinuity design to evaluate how the causal effect of electrification on labor-market outcomes varies with exogenous changes in economic contexts.

First, we show that electrification increased non-agricultural employment in villages located in India's booming guar belt by approximately six percentage points (seventy percent). In these same villages, electrification reduced agricultural employment by a corresponding amount, representing a reduction of approximately twenty percent. This is particularly notable given the fact that this region—spread across three states in north-western India—was in the grip of an unprecedented agricultural boom. We next highlight potential mechanisms by providing suggestive evidence that these labor-market dynamics are driven by the rise of complementary non-agricultural opportunities. An increase in guar production necessitates a shift in the labor force towards industrial-scale guar processing, which benefits from upgrades to local electricity infrastructure. Simultaneously, increased wages and agricultural profits from both the production boom and new processing opportunities can be reinvested in household enterprises, which may also benefit from new electricity connections. Consistent with this, we uncover a large increase in (i) the number of workers at firms related to the industrial (electricity-intensive) parts of the guar production chain (such as guar processing); and (ii) home production of income–generating products in electrified guar-growing regions. Finally, we find no discernible evidence of any effect of electrification on these labor-market outcomes in villages or regions located in the rest of India, suggesting that complementary economic conditions play a crucial role in driving the impacts of large-scale electrification infrastructure.

In so doing, we revisit work by Burlig and Preonas (2016), who conduct the first large-scale impact evaluation of India's rural electrification scheme. They show that the program increased electrification rates, but also demonstrate that its impacts on a wide range of socioeconomic outcomes (including those related to the rural labor market) are precisely estimated null results.[2] Our results from non-guar regions of India—using

[2] Results from a randomized controlled trial in Kenya by Lee et al. (2018) echo these findings.

an empirical strategy that follows their own—are consistent with these earlier findings. Using the exogenous shock to economic activity generated by the guar boom, however, also allows us to respond to some of the questions that emerge from this prior body of work and rigorously shed light on important drivers of heterogeneity in the impacts of electrification globally.

Our study, thus, makes three key contributions. First, our results highlight how grid-scale electrification can support potentially welfare-enhancing structural change in the rural economy. Access to electricity alone may not deliver economic and social benefits, as has been demonstrated a number of times in the literature. That electrification significantly enhances non-agricultural employment in boom areas suggests, however, that it can enable individuals, households and firms to fully exploit the opportunities presented by rapidly changing economic contexts.

Second, our findings highlight that the impacts of large-scale investments in grid electrification are crucially tied to local economic conditions. For instance, electricity from the grid may enable local industrial production of certain goods, yet this may make little difference in the short run if complementary factors—such as demand for these locally produced goods, a trained labor force to scale up production to meet that demand, and rural roads that enable access to markets—are not also in place. If they are, however, grid-scale electricity may considerably expand how local actors take advantage of economic opportunities to generate income and enhance welfare. Prior research—which typically estimates the "average treatment effect" of such investments as part of national rural electrification programs—implicitly neglects these context-specific factors.[3] While the particular agricultural boom we study is clearly unique to our setting, it—in combination

[3] This, we contend, is one reason we observe mixed evidence from settings as diverse as Bhutan, Brazil and Vietnam (Khandker et al., 2013; Lipscomb et al., 2013; Litzow et al., 2017). In addition, many national rural electrification schemes are grounded in an obligation—either perceived or real—to ensure universal access to electricity (Tully, 2006). While certainly aligned with broader equity goals, it is not immediately clear that such rights-based approaches are necessarily designed to maximize economic outcomes. That short- or medium-term impact evaluations of such efforts over large spatial scales may yield null results is unsurprising.

with the roll-out of rural electrification—gives us an opportunity to investigate how electrified villages in boom and non-boom areas perform relative to unelectrified villages in the same regions. Insofar as the economic promise or potential of certain areas can be accurately assessed *ex ante*, the insights we generate can be used to inform spatial targeting of resource-intensive infrastructure by allowing policymakers to better gauge cost-benefit trade-offs, and choose appropriate grid-based and off-grid energy solutions for different contexts.[4]

Finally, from a methodological perspective, our study is part of a growing body of work that adopts a rigorous approach to understanding treatment-effect heterogeneity in the real world.[5] That the same intervention can have different impacts in superficially similar settings points to the importance of context-dependence; learning about what these contextual factors are is crucial to learning from these impact evaluations (Vivalt, 2015). Where a sufficiently large number of studies have been conducted, rigorous meta-analyses can shed light on underlying drivers of effectiveness. In most other cases, however, such efforts are typically restricted to relatively crude subgroup analyses, involving interactions of endogenous binary variables representing various subgroups of interest with the main treatment-effect parameter. Our quasi-experimental setting—the combination of an exogenous shock to economic activity with quasi-experimental variation in access to electricity within a regression discontinuity design—provides the first opportunity to study the heterogeneous effects of access to electricity over large spatial scales in a real-world setting.

[4] We emphasize that there may be other channels driving heterogeneity in the benefits generated by large-scale infrastructure projects, such as institutions or access to markets. We believe this is a promising avenue for future research.

[5] In its use of multiple sources of exogenous variation in real-world settings, our study is related to Duque et al. (2018), who examine how early-life exposure to adverse weather shocks (that reduce children's initial skills) in Colombia interacts with the introduction of conditional cash transfers (that promote investments in children's health and education) to influence long-term outcomes. It is also similar to Wysokinska (2017), who studies the determinants of long-run development by similarly examining the interplay between plausibly exogenous variation in institutional and cultural factors in Poland.

This rest of this paper is organized as follows. In Section 1.2, we provide background on our two natural experiments. Section 1.3 highlights our conceptual framework, and discusses our identification strategies. Section 1.4 describes our data. Section 1.5 reports impacts on the first set of outcomes related to the size and composition of the rural labor force. Section 1.6 reports impacts from additional analyses to uncover mechanisms related to the growth of firms. Section 1.7 summarizes results, and discusses policy implications and avenues for future research.

## 1.2 Background

In this section, we first describe India's rural electrification scheme. We then provide a basic overview of hydraulic fracturing ("fracking"). Finally, we discuss guar production in India and, in particular, how it responded to the fracking boom in the United States.

### *1.2.1 Rural electrification*

Rural electrification in India has a checkered past. In 1947, newly independent India had only 1,500 electrified villages, and progress on rural electrification remained slow well into the late 1960s (Banerjee et al., 2014, p. 35). The country's initial electrification efforts focused primarily on urban and peri-urban areas. Severe droughts and food shortages in the early 1960s brought rural electrification into the spotlight, yet subsequent policies prioritized productive uses over household access, and primarily aimed to increase access to electricity for irrigation. Rural household access finally emerged as a key priority area in the late 1970s, and has since featured prominently in India's successive Five-Year Plans. The growing recognition of the role of electrification in rural development—coupled with the existence of multiple national- and state-level electrification agencies with overlapping responsibilities—gave rise to a number of schemes over the decades.[6] The Rajiv Gandhi

[6] For instance, the Kutir Jyoti Yojana was launched in the late 1980s to increase access to electric lighting for households below the poverty line; the Pradhan Mantri Gramodaya Yojana, launched in 2001, extended financing to states to enhance access to public services, including electrification, in rural areas; the Remote

Grameen Vidyutikaran Yojana (RGGVY), launched in 2005, subsumed all existing grid-related rural electrification initiatives.

RGGVY was charged with enhancing access to electricity in over 100,000 unelectrified and 300,000 "partially electrified" villages across 27 Indian states. It aimed to do so primarily by installing and upgrading electricity infrastructure (namely, transmission and distribution lines, and transformers) to support commercial and productive activities in growing rural economies. These included electric irrigation pumps, education and health-care facilities, and small and medium enterprises. In addition to its focus on electricity infrastructure, RGGVY also extended free grid connections to rural households below the poverty line; households above the poverty line could purchase connections.[7] Both groups remained responsible for their own power use as RGGVY did not subsidize electricity consumption.

Although a national program that was largely funded by India's federal government, RGGVY was implemented in practice through decentralized district-level projects overseen by local implementing agencies (such as the State Electricity Board).[8] Electrification under RGGVY proceeded in two steps. First, to qualify for RGGVY funds, the local implementing agency prepared a Detailed Project Report (DPR) for the district in question. The DPR outlined in detail the electrification-related infrastructure needs of the district, the number of households expected to be connected to the grid, and expected project costs. It also identified the set of villages eligible for electrification under RGGVY. These DPRs were

Village Electrification program, launched in 2002, aimed to provide lighting to remote villages using solar photovoltaics and other off-grid energy technologies; and the country's Minimum Needs Program was updated in 2002 to extend financing for rural electrification to states that were seen to be performing especially poorly (Banerjee et al., 2014, p. 37-38).

[7] According to the Ministry of Power (2006), RGGVY's primary mandate included the (i) provision of electricity sub-stations and transmission lines of adequate capacity to establish a "rural electricity distribution backbone;" (ii) electrification of unelectrified villages, including provision of distribution transformers of appropriate capacity; (iii) establishment of decentralized distributed generation and supply in a subset of villages where grid connectivity is infeasible or not cost effective; and (iv) provision of household-level connections for households below the poverty line.

[8] An Indian district is administratively analogous to a county in the United States.

reviewed and approved by India's Rural Electrification Corporation as well as its Ministry of Power before disbursement of funds. Once approved, district-level implementation commenced in line with the village-by-village plan outlined in the DPR.

Districts were allocated to India's Tenth (2002–2007) and Eleventh (2007–2012) Five-Year Plans for funding based on the order in which DPRs were submitted and approved. We refer to these as "RGGVY Phase I" and "RGGVY Phase II" districts, respectively, and identify these districts using state-level five-year-plan progress reports for RGGVY.[9] To keep program costs low, during Phase I, villages containing at least one habitation (a geographically distinct sub-village cluster of households) with a population of 300 or more were eligible to be electrified. Approximately 178,000 villages across 234 Phase I districts in 25 states (as per 2011 administrative boundaries) fit this criterion. Nearly all funds associated with Phase I districts had been disbursed between 2005 and 2008, while funding for Phase II districts—for which the RGGVY eligibility threshold was reduced to 100—was disbursed between 2008 and 2011. In this paper, we specifically focus on Phase I districts (shown in Figure 1.1) as village-level electrification in these districts had been completed well in advance of the release of the 2011 round of the Indian Census, one of our main data sources.[10]

### *1.2.2 Fracking*

Hydraulic fracturing ("fracking") is the process by which fracking fluid (a mixture of mostly water, granular "proppants" such as sand, and chemicals) is injected into crude oil and

[9] For each state, these reports—entitled "Report C-Physical & Financial Progress of RGGVY Projects Under Implementation (Plan-wise)"—list the district name and DPR code, the name of the district-level local implementing agency, details about the financial scope and progress of the project (such as project approval date, total sanctioned amount, and the amount released so far), as well as the scope and progress of electrification (in terms of village- and household-level electrification targets). These reports are available via the website of the Deendayal Updhayaya Gram Jyoti Yojana (DDUGJY)—into which RGGVY was ultimately subsumed—at `http://www.ddugjy.gov.in/`.

[10] Indeed, because enumeration for the 2011 Census began in April 2010, villages electrified as part of RGGVY Phase II would have been only captured inconsistently during Census survey activities. In addition, they would have been electrified for a considerably shorter period of time.

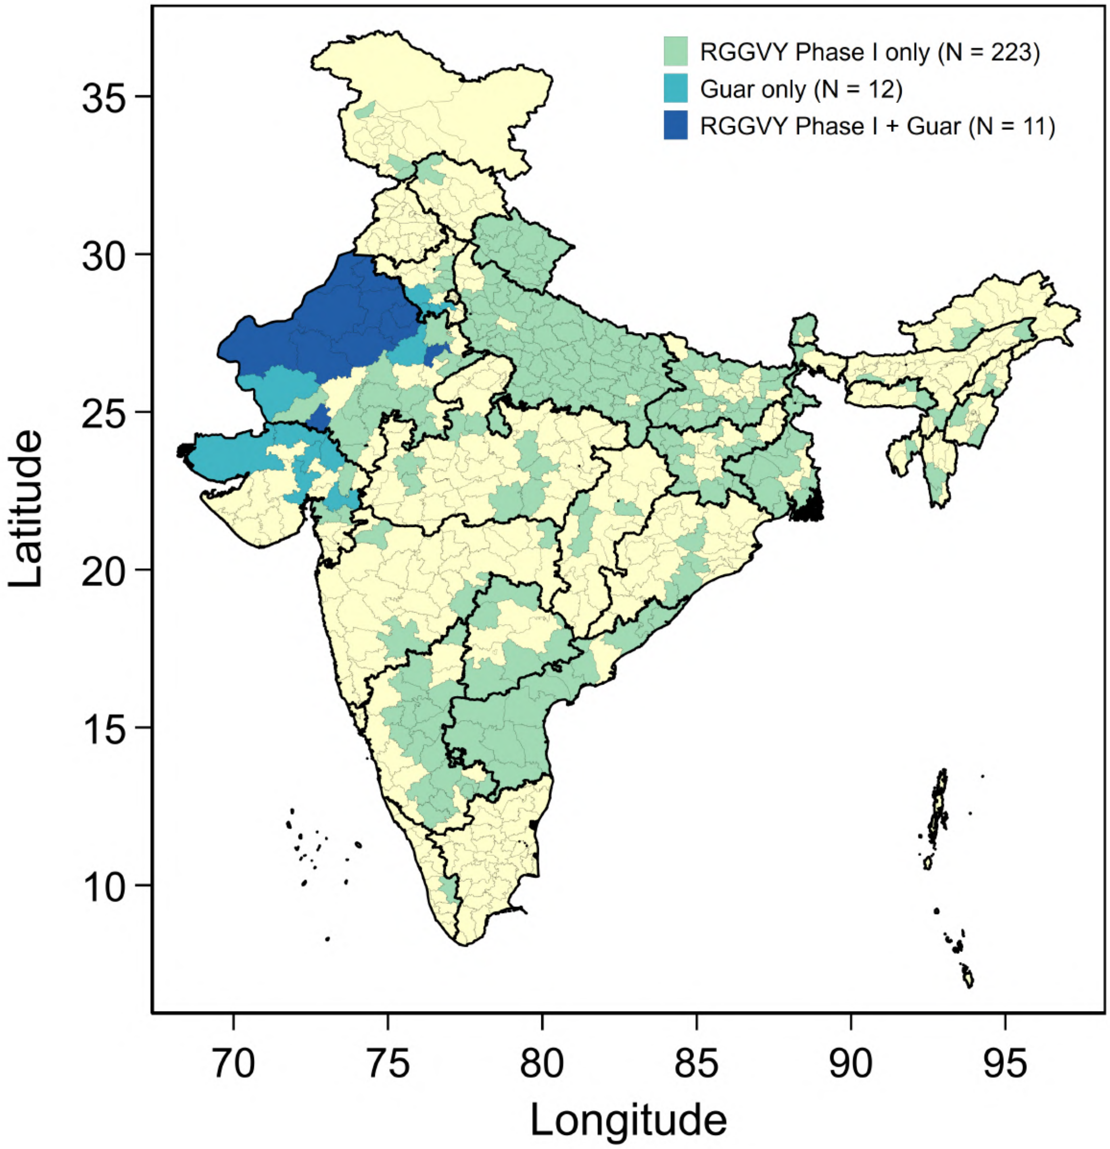


FIGURE 1.1: Districts of India, by guar-production and electrification status. This map shows India's 2011 state (thick lines) and district (thin lines) boundaries. Districts are shaded by their electrification and guar-production status. Unshaded districts were neither approved for the roll-out of electrifcation as part of RGGVY Phase I nor contribute appreciably to guar production in India.

natural gas wellbores at high pressures to create small cracks (fractures) in the underlying rock formation. While not an entirely new approach, recent technological refinements—and, in particular, fracking in combination with horizontal drilling—have considerably increased the effectiveness of the process and transformed the energy landscape in the United States.[11]

Figure 1.2 provides an overview of natural gas (panel *a*) and oil (panel *b*) production from fracked and "conventional" wells in the United States between 2000 and 2015. In 2000, fracked wells produced 3.6 billion cubic feet per day of marketed gas, less than seven percent of the United States' total. Starting in approximately 2005, the industry grew rapidly. By 2015, fracked wells produced around 67 percent of the country's total natural gas. Oil production underwent a similarly momentous shift, albeit slightly later. In 2000, fracked wells yielded less than two percent of the national total. Following a period of growth that began around 2009, approximately half of the United States' total oil output could be traced back to a fracked well in 2015.

A typical "frac job" is preceded by a vertical drill to a depth of around 1,000–5,000 meters, depending on the geophysical characteristics of the shale formation being explored. Upon reaching the desired depth, the well is then drilled horizontally, allowing for greater access to the shale formation. Once drilling is complete, fracking fluid is injected at high pressures into the drill site to induce fractures in the formation. Reductions in pressure following the initial injections cause fluids in the well to return to the surface as "flowback." As production continues, the amount of flowback returning to the surface gradually decreases and the amount of oil or gas increases.

Fracking fluid consists almost entirely of water and proppants; the remaining elements

[11] Compared to conventional (vertical) wells, horizontal wells can typically access greater reserves, and are two to five times more productive (Joshi, 2003). This can lead to considerable cost savings in the long run, despite higher initial drilling costs. Orr (2016) notes that "[a]lthough hydraulic fracturing and horizontal drilling had been used separately to stimulate production at conventional wells since 1947 and 1929, respectively, the combination of these methods has enabled scientists to extract oil and gas trapped in impermeable source rocks such as shale, well-cemented sandstone, and coal bed methane deposits once considered too costly to develop."

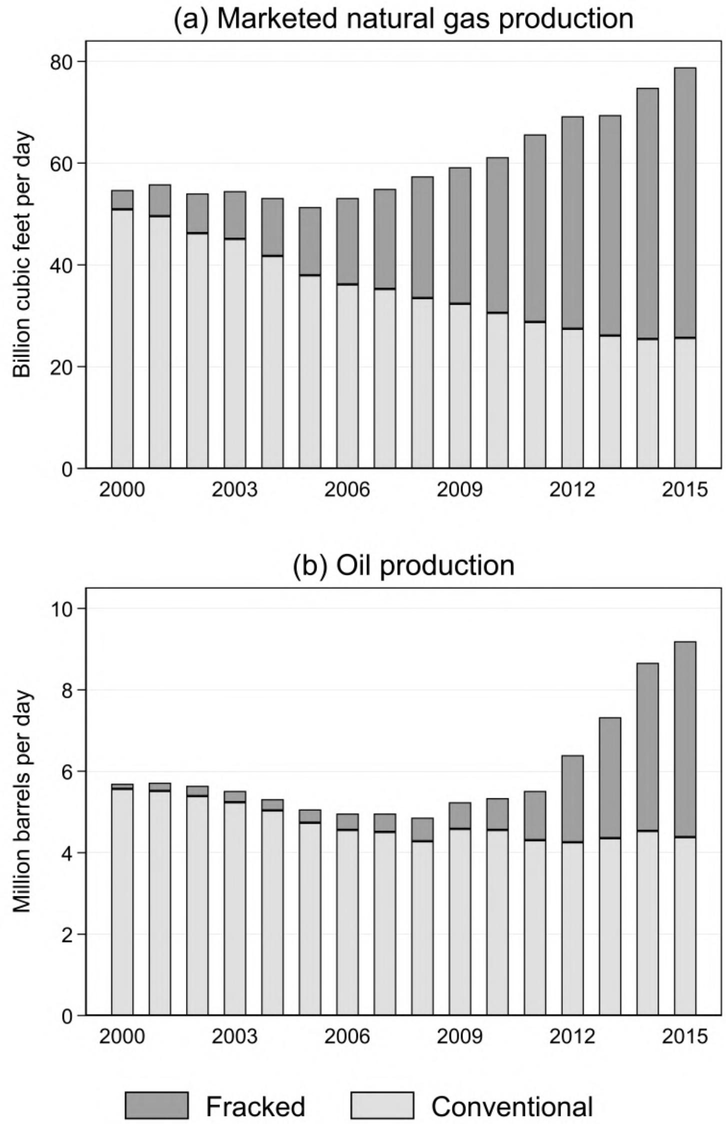


Figure 1.2: Natural gas and oil production in the United States, by source. This figure shows marketed natural gas (panel *a*) and crude oil (panel *b*) produced from fracked and "conventional" wells in the United States between 2000 and 2015. Marketed natural gas production excludes natural gas used for repressuring the well, vented and flared gas, and any nonhydrocarbon gases. Source: United States Energy Information Administration, IHS Global Insight, and DrillingInfo, Inc, as outlined at `https://www.eia.gov/todayinenergy/detail.php?id=26112` and `https://www.eia.gov/todayinenergy/detail.php?id=25372`.

usually include various chemicals that serve, among other things, as gelling agents, corrosion inhibitors, friction reducers, clay controls and biocides (Tollefson, 2013). Of these chemicals, the gelling agent—which increases the viscosity (thickness) of fracking fluid—comprises the largest share. Its use confers two important advantages. First, viscous fluids enable better control of leak-off into the surrounding rock formation, reducing the amount of fracking fluid needed for a given frac job (Barati and Liang, 2014). Second, viscous fluids are more effective at suspending sand and other granular proppants and carrying them deep into the wellbore (Bellarby, 2009). These proppants prevent fractures induced in the rock by high-pressure pumping from closing down completely once the pressure has fallen. These partially open fractures are the passageways through which oil and gas flow out of rocks and into the well.

No particular combination of ingredients is perfect, and operators often face trade-offs.[12] For this reason, experimentation with the specific mix of chemicals used is rife (Fetter, 2018; Fetter et al., 2018). Yet despite operators' readiness to modify the make-up of fracking fluid, guar gum—a powdery substance derived from the bean of the guar plant—is the industry's most widely used gelling agent. Indeed, between 25–50 percent of all fracking operations rely on guar gum, making it "at least two to three times preferred over synthetic [alternatives]" (Elsner and Hoelzer, 2016). This is unsurprising; guar gum is uniquely effective at its job. It can alter the viscosity of fracking fluid by more than two orders of magnitude under certain conditions (Tapscott, 2015). In addition, whereas other natural gums require prolonged cooking, guar gum attains its full viscosity potential in cold water, and is effective even at relatively dilute concentrations (Thombare et al., 2016). Its viscosity potential also remains relatively stable over changes in temperature, and in the acidity or basicity of the solution in which it is mixed (Chudzikowski, 1971). Despite considerable efforts by major chemical companies in recent years, a synthetic alternative

[12] For instance, although more viscous fluids are better able to suspend proppants, they are less "pumpable" and require more energy to be pumped at sufficiently high rates.

that is as effective as guar gum for high-viscosity fracking is yet to be developed (Beckwith, 2012).

### *1.2.3 Guar and guar gum*

Guar (*Cyamopsis tetragonoloba*) is a drought-resistant legume that is primarily cultivated in the semi-arid northwestern tracts of the Indian subcontinent (Kuravadi et al., 2013). It can tolerate relatively high temperatures and requires only sparse but regular rainfall, which makes the rain patterns associated with the monsoon in this region ideal for cultivation (Mudgil et al., 2011). Guar—whose name is derived from the Sanskrit term for "cow food"—has traditionally been cultivated as both fodder and a vegetable crop. It grows well in many different types of soil, and its nitrogen-fixing potential combined with its relatively short planting season also make it an excellent soil-improving crop that fits conveniently within farmers' crop-rotation cycles.[13]

Guar gum (sometimes also called guar flour) is obtained from the endosperm of guar seeds in two distinct energy-intensive steps (Chudzikowski, 1971). Guar seeds are first exposed to a rapid flame treatment, which loosens the hard seed hull (outer shell), which is removed in a scouring or "pearling" operation. The glassy endosperm that this process exposes is then separated from the germ in a milling operation. The resulting guar "splits" can be ground to various levels of fineness to obtain guar gum in powder form. This powder is sometimes further processed and combined with additional chemicals to obtain industry-specific derivatives.

India accounts for approximately eighty percent of global production, making it by far the world's largest producer of guar (National Rainfed Area Authority, 2014).[14] The country occupies a similarly dominant role in the global trade of guar derivatives. Within

[13] Like other legumes, the roots of the guar plant contain nodules inhabited by nitrogen-fixing bacteria, and crop residues—when plowed under—can improve soil fertility and the yield of subsequent crops (Undersander et al., 1991).

[14] Pakistan, the next largest producer, is responsible for approximately fifteen percent of global production.

India, guar is almost exclusively produced in the northwestern part of the country. The state of Rajasthan—which was home to nearly ninety percent of India's total area under guar cultivation and eighty percent of its production in 2013—is the epicenter of this industry. Other important producers include Haryana and Gujarat, which—together with Rajasthan—comprise nearly all of the total area under guar cultivation in India. At the level of the farmer, however, guar cultivation in India is relatively decentralized, and the crop is grown by thousands of small and marginal farmers. While precise data on agricultural practices are unavailable, industry experts also believe most guar cultivation is rainfed, and farmers have typically planted it as a secondary or tertiary crop on small subsistence-level plots of land (Beckwith, 2012).

Nearly all of India's guar is processed domestically, and the country's guar-processing industry dates back to the late 1950s. Indeed, the widespread use of guar gum in the petroleum industry is a relatively recent phenomenon.[15] In addition to its oil and gas applications, guar gum has long been used in a variety of industries, including as a food additive, thickener of cosmetics/toiletries such as toothpaste, and waterproofing agent for explosives (Thombare et al., 2016).

Nevertheless, the unprecedented growth of fracking in the United States in recent years has resulted in an equally unprecedented expansion in guar production in India.[16] Figure 1.3 shows trends in India's global guar gum exports—by total weight and as a share of global trade value—between 2001 and 2015. At the beginning of this period, the value of India's guar gum exports comprised approximately 35 percent of the global

[15] The Department of Agriculture first introduced guar to the United States in 1903 to investigate its potential as a soil-improving legume and as emergency cattle feed. These initial experiments appear to have been disappointing, and the crop fell into relative obscurity until World War II. Spurred on by the sudden unavailability of a thickening agent derived from the locust bean (*Ceratonia siliqua*)—which, until then, had been imported from the Mediterranean region—a search for domestically available alternatives for the paper industry ultimately unveiled the potential of guar (Hymowitz, 1972).

[16] Indeed, as we show in Appendix A using an application of the synthetic control approach to two decades of village-level nighttime luminosity data from India, the start of the fracking boom in the United States led to large increases in economic activity across the guar-growing regions of northwestern India.

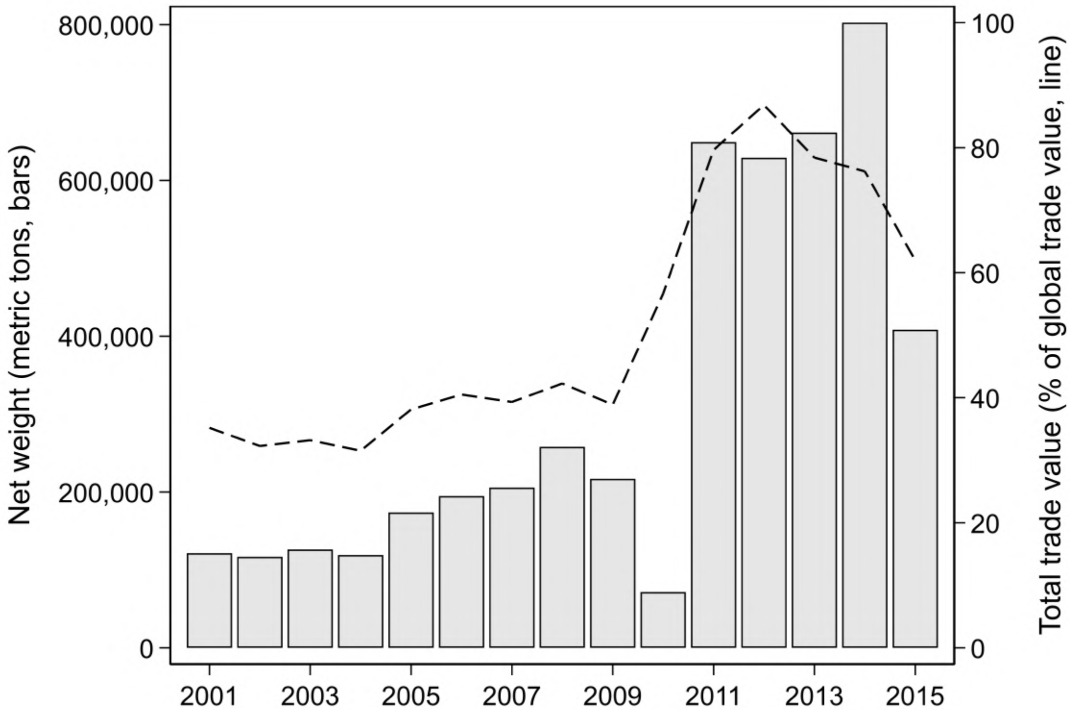


FIGURE 1.3: Weight and share (of global value) of India's guar gum exports. This figure shows the total weight (bar graph; left axis) and share of total global trade value (line graph; right axis) of India's exports of guar gum for each year between 2001 and 2015 based on data for guar gum (product code HS 130232) from the United Nations Comtrade Database (`https://comtrade.un.org/`). Guar cultivation in India exhibited a reduction in 2009-10 on account of drought conditions, resulting in a reduction in the weight of its guar gum exports.

trade in guar gum. This share began to rise starting in 2004-05 as shale-gas exploration became increasingly feasible in the United States. It spiked sharply starting in 2009-10 corresponding to the rise in the use of fracking in oil production. At the height of the boom, nearly ninety percent of the global trade in guar gum (by value) originated in India. The total weight of India's guar gum exports follows a similar pattern, except for a drop in 2009-10 on account of drought conditions in northwestern India (Rai, 2015). Because we rely on data from the 2001 and 2011 rounds of the Indian Census, our main analyses focus on the pre-2011 part of this boom.

## 1.3 Conceptual framework and empirical strategy

In this section, we highlight three main hypotheses that connect access to electricity with household-level labor supply. We use these to develop a simple model of household time allocation. We then describe our regression discontinuity and difference-in-differences empirical strategies, and comment on the identifying assumptions implicit in each.

### *1.3.1 Electrification and labor supply*

There are a number of pathways through which electrification can modify households' labor-supply decisions. One popular argument relates to the time burden imposed by home production activities, such as collecting and preparing traditional fuels for cooking and heating. If electricity can be used for these purposes instead, it frees up household members' time for engaging in market activities.[17] In practice, exclusive reliance on electricity for cooking is relatively uncommon in low- and middle-income countries, and use of traditional fuels such as firewood is widespread, including among electrified households (Barron and Torero, 2017; Pattanayak et al., 2016; Thom, 2000). In India, for instance, 66 percent of all households use biomass-based fuels for cooking (Adair-Rohani et al., 2016). In such settings, access to electricity is unlikely to significantly influence households' time allocation in this way.

Another prominent argument relates to the provision of lighting and its effect on total working hours. If electric lighting can enable households to allocate domestic activities that require good lighting to evening hours, daylight time can be allocated to activities that generate income. Yet this hypothesis also faces a number of limitations. Households in many rural areas have already transitioned away from low-quality kerosene lighting to relatively high-quality electric lamps powered by small-scale batteries (Bensch et al.,

[17] The burden of such household activities in the developing world falls almost entirely on women and girls. For this reason, access to modern energy services (including electricity) is also often promoted as contributing to women's empowerment (O'Dell et al., 2014).

2017). The additional benefits of electric lighting delivered by the grid in such settings are unlikely to be large. More fundamentally, an increase in the total number of well-lit hours may simply lead to an increase in the time households dedicate to leisure activities (Pereira et al., 2011).[18]

A third channel—and one that is the focus of our paper—relates to the productive potential of domestic and income-generating activities that the household can conduct. Specifically, electrification may considerably increase the productivity of domestic or income-generating activities that do not necessarily require electricity, such as water collection or sewing. It may also enable new opportunities to engage in activities that were previously not possible, such as soldering/metalworking or industrial production. Together, these can (i) yield time savings that can be allocated to income-generating activities; and (ii) influence the market wage rate that the household faces, which changes the opportunity cost of not participating in income-generating activities. Depending on the magnitude of these effects, households may reduce the amount of time allocated to leisure, and increase that allocated to home- or market-based activities.[19] Conditional on already being engaged in income generation, households may also reallocate hours to new types of work.

More formally, such changes in individuals' productive potential can be captured in an application of the basic home-production and household time-allocation model (Gronau, 1977). In this framework, the representative individual in household $i$ obtains utility from consumption ($c_i$) and leisure $\left(t_i^l\right)$. Consumption is generated through a home-production function:

$$c_i = c\left(t_i^h, x_i, v_i; \psi_i\right), \tag{1.1}$$

[18] All else constant, an increase in the total number of hours available to households can unambiguously increase time dedicated to leisure. This is because additional hours need not lead to more time allocated to market-based activities (i.e., a "substitution effect") unless also accompanied by a change in the opportunity cost of leisure (e.g., the market wage rate).

[19] We use "leisure" here to mean time not spent engaged in consumption- or income-generating activities at home or in the market.

where $t_i^h$ is the time allocated to home-based work; $x_i$ is a numeraire input to home production that is purchased in the market; and $v_i$ is non-labor income. In addition, $\psi_i$ represents a production productivity parameter. It is, in turn, determined by a productivity production function given by

$$\psi_i = f\left(\eta_i, \epsilon_i, \gamma\right), \tag{1.2}$$

where $\eta_i$ represents the household's electrification status on a continuous scale, thus capturing both basic access and quality. Productivity is also determined by household- and community-level unobserved factors, represented by $\epsilon_i$ and $\gamma$, respectively. For instance, households' stock of education and health can drive the labor productivity of its members. Community-level characteristics—such as weather, institutions, and, in particular, differences in local or regional economic conditions—can play a similar role.

The problem of the household's representative individual is then given by

$$\max_{c_i, t_i^l} u_i = u\left(c_i, t_i^l; \psi_i\right), \tag{1.3}$$

subject to time and budget constraints, given by

$$t_i^m + t_i^h + t_i^l \leqslant T \tag{1.4}$$

and

$$x_i \leqslant w_i t_i^m + v_i, \tag{1.5}$$

where $t_i^m$ is the time allocated to market-based work; $T$ is the total time endowment; and $w_i$ is the market wage. Equations (1.4) and (1.5) together yield the household's full-income constraint:

$$w_i T + v_i = x_i + w_i\left(t_i^h + t_i^l\right). \tag{1.6}$$

The Lagrangian associated with the household's problem is as follows:

$$\max_{c_i, l_i} \mathcal{L} = u\left(c\left(t_i^h, x_i, v_i; \psi_i\right), l_i\right) + \lambda\left(w_i T + v_i - x_i - w_i\left(t_i^h + t_i^l\right)\right). \tag{1.7}$$

As shown in Appendix B, the first-order conditions associated with the household's problem in Equation (1.7) equate the marginal rate of substitution between leisure and consumption with (i) the shadow value of home production; and (ii) the shadow value of market-based activities. Solving this system of equations yields a set of expressions for the household's optimum time allocation:

$$t_i^{j*} = f_j\,(w_i, v_i; \psi) \tag{1.8}$$

for $j = h, l, m$.

We look to investigate how changes in the household's access to electricity ($\eta_i$) interact with community-level factors ($\gamma$) to influence the household's productive potential ($\psi_i$) and ultimately determine the time it allocates to home production, leisure, and market-based activities. Specifically, by exploiting exogenous variation in levels of economic activity across guar- and non-guar-growing regions of India, we aim to shed light on how and why differences in the impacts of access to electricity can emerge.

There are at least two reasons why our model does not offer a clear answer to this question. First, even if we assume that an improvement in the household's access to electricity increases its productivity potential (i.e., $\psi'_{i,\eta} > 0$ and $\psi''_{i,\eta} < 0$), additional assumptions are necessary about the exact shape of the home-production function in Equation (1.1) to predict how changes in productivity as a result of simultaneous changes in electrification and community-level characteristics influence time allocation. Second, even with such assumptions in place, variation in household-level preferences over labor and leisure—the shape of the household utility function—may give rise to counteracting income and substitution effects. Indeed, an increase in its productive potential may ultimately induce a household to allocate *less* time to income-generating activities.

This ambiguity is further compounded by the role household-level characteristics ($\epsilon_i$) can play. The household's opportunity cost of leisure is determined by a variety of factors, such as its stock of education and health, the liquidity or credit constraints it faces, or its "entrepreneurial spirit." Thus, how the impacts of electrification on labor-market outcomes vary with economic conditions is ultimately a question that can be best answered with data. Our study setting allows us a unique opportunity to address this question.

### *1.3.2 Regression discontinuity design*

A comparison of labor-market outcomes in electrified villages located in guar-growing districts before and after electrification is unlikely to yield a causal estimate of the impact of electrification in the presence of high levels of economic opportunity for three reasons.[20] First, this approach lacks a suitable "non-boom" control. Second, it neglects heterogeneity within the set of electrified villages. Among other things, the largest electrified villages are also likely to have better access to schools and health facilities, both of which can directly influence labor-force productivity. Finally, this approach fails to account for changes in other factors over the course of the decade—such as the launch of India's massive rural workfare program in 2006—that can act as confounders. A cross-sectional comparison of guar-growing electrified villages with electrified villages in non-guar-growing regions would yield similarly unreliable estimates. Indeed, most guar-growing districts are located in Rajasthan, which, despite the recent boom, remains one of India's poorest states. A simple *ex post* comparison of guar-growing electrified villages with those in relatively wealthier regions is likely to provide an underestimate of our parameter of interest.

In contrast, we exploit a population-based threshold that guided the roll-out of India's rural electrification scheme as part of a village-level regression discontinuity (RD) design. Villages in districts approved under Phase I of RGGVY were eligible for electrification if they contained a habitation with at least 300 people. Indian villages, however, can contain multiple habitations—typically between one and three—which complicates identification. For instance, a village with a relatively large population may have been ineligible under RGGVY if its population was spread out over multiple habitations; a less populous (but

[20] We describe how we identify India's guar-growing districts in Section 1.4.1.

more concentrated) village may have been electrified. A village's overall population can, thus, be a poor measure of its RGGVY eligibility; comparing villages with overall populations above the RGGVY threshold to villages with populations just below it is unlikely to yield an accurate estimate of the impact of electrification without additional information on sub-village habitation characteristics. To address this concern, we restrict our nationwide sample of villages to single-habitation villages, following the empirical approach developed by Burlig and Preonas (2016). This allows us to similarly estimate the local average treatment effect (LATE) of electrification on labor-market outcomes for villages with overall populations close to RGGVY's eligibility threshold. To highlight the importance of local economic conditions, we pay close attention to differences in the magnitude of the estimated LATE for villages located in guar-growing districts versus those in the rest of India.

We focus on all single-habitation villages in RGGVY Phase I districts with a population within a suitable bandwidth of 300, the RGGVY Phase I threshold for electrification. Within this sample, we look at two overlapping subsets of villages: (i) those that are located in guar-growing districts; and (ii) those with a population greater than or equal to 300 (i.e., those that were electrified as part of RGGVY Phase I). The intersection of these criteria represents our sample of interest: guar-growing villages that were electrified as part of RGGVY Phase I. We compare the impacts of rural electrification in this sample to those in villages that were electrified in non-guar regions of the country.

More formally, we rely on an RD design to estimate

$$\begin{aligned} y_{vds}^{2011} &= \beta_0 + \beta_1 T_{vds} + \beta_2 T_{vds} G_{ds} \\ &\quad + \beta_3 \tilde{P}_{vds}^{2001} + \beta_4 T_{vds} \tilde{P}_{vds}^{2001} + \beta_5 G_{ds} \tilde{P}_{vds}^{2001} + \beta_6 T_{vds} G_{ds} \tilde{P}_{vds}^{2001} \\ &\quad + \beta_7 y_{vds}^{2001} + \gamma_d + \epsilon_{vds} \end{aligned} \tag{1.9}$$

for $-b \leqslant \tilde{P}_{vds}^{2001} \leqslant b$. $y_{vds}^{2011}$ represents an outcome of interest in 2011 for village $v$ located in district $d$ in state $s$, $\tilde{P}_{vds}^{2001} = P_{vds}^{2001} - 300$ (where $P_{vds}^{2001}$ is its population in the 2001 Census round), and $b$ denotes a suitable population bandwidth around the RGGVY's 300-person eligibility threshold. Our preferred specification relies on a narrow bandwidth of fifty people on either side of this cutoff. $T_{vds}$ is a binary variable that equals one if $P_{vds}^{2001} > 300$, i.e., the population of village in $v$ in 2001 is above RGGVY's eligibility threshold. $G_{ds}$ is a binary variable that equals one if village $v$ is located in a guar-growing district. $y_{vds}^{2001}$ is the 2001 value of the outcome variable. $\gamma_d$ represents a district fixed-effect, which allows us to control for all time-invariant district-specific characteristics that make a district more likely to be a guar producer and, thus, independently induce variation in the level of the outcome of interest. $\epsilon_{vds}$ is a village-specific error term. We cluster our standard errors at the district level to allow for correlated unobservables between villages that are located nearby and, in line with RGGVY's implementation structure, electrified and served by the same district-level electrification agency.

In Equation (1.9), $\beta_1$ represents the LATE of electrification on our outcome of interest in villages located in non-guar-growing regions of India. Our parameter of interest is $\beta_2$, which represents the additional effect of electrification in villages affected by the guar boom. If $\hat{\beta}_2$ is statistically different from zero, we conclude that the LATE for electrification in the booming guar-growing regions of India is different from that in the rest of India. Conditional on the inclusion of district fixed-effects, which control for all unobserved spatial differences at the district level, this highlights the degree to which the economic activity generated by the exogenous guar boom augments the impact of electrification.

Identification relies on continuity of potential outcomes in village population (our running variable) at the RGGVY eligibility threshold. This assumption is plausible if (i) villages are not able to manipulate their population levels—either in actuality or in administrative reporting—to influence RGGVY eligibility; and (ii) all observable and unobservable village-level covariates that may be correlated with our outcomes of interest change smoothly at the threshold. The former is unlikely to be a concern in our case.

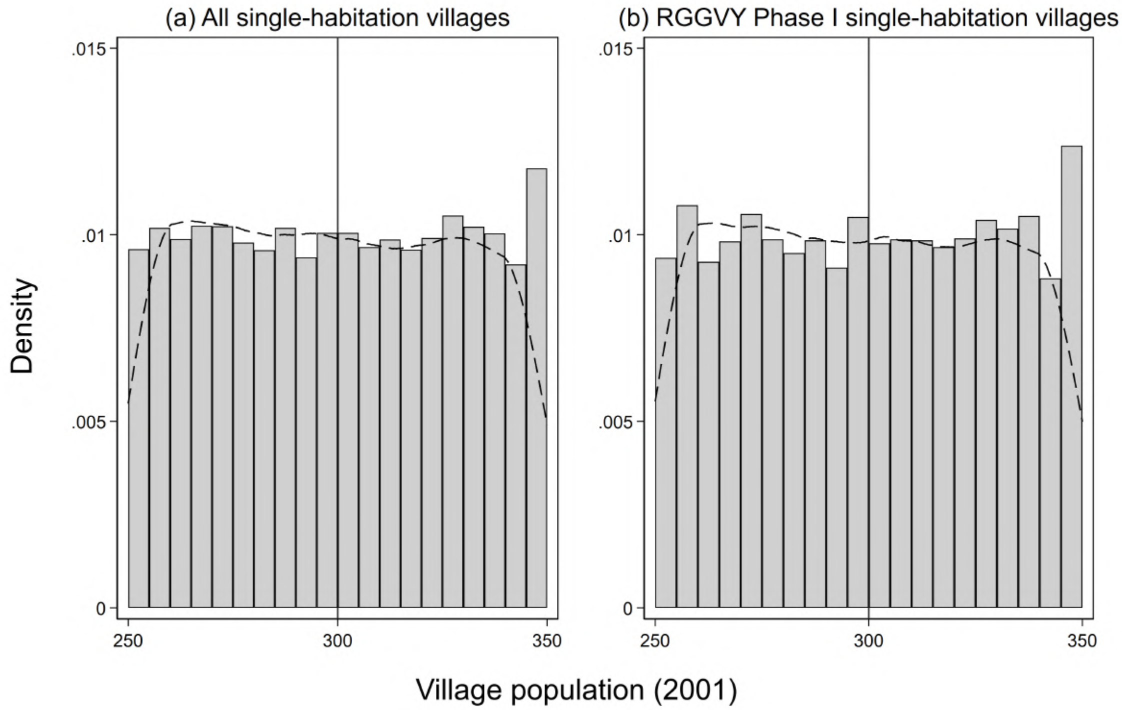


FIGURE 1.4: Village population changes smoothly at RGGVY Phase I eligibility threshold. This figure shows the distribution of village-level population (in five-person bins) for fuzzy-matched single-habitation villages with a 2001 population (as per the Census) that is within a fifty-person bandwidth of RGGVY's 300-person habitation-level eligibility threshold for electrification. Panel (*a*) shows this distribution for all such villages in India ($N$ = 14,668) while panel (*b*) shows it only for villages located within districts approved for the roll-out of rural electrification as part of RGGVY Phase I ($N$ = 7,655).

RGGVY used population figures from the 2001 round of the Indian Census to gauge eligibility (Burlig and Preonas, 2016). These data predate the announcement of RGGVY by at least four years and are thus unlikely to have been manipulated at or near its 300-person eligibility threshold. Nevertheless, following McCrary (2008), in Figure 1.4 we check for bunching at the cutoff—for all single-habitation villages in India that lie within our preferred bandwidth (panel *a*) and for those located in RGGVY Phase I districts (panel *b*)—and find no evidence to suggest that this is the case.

The latter component of this assumption—that all village-level covariates change smoothly at the threshold—is fundamentally untestable. That said, we provide evidence in support of it by examining the pre-RGGVY distribution of key village-level characteristics around the cutoff. We find no evidence to suggest that that these change discontinuously at the 300-person mark prior to the implementation of RGGVY (Table G.1). We are also aware of no other social program in India that uses RGGVY's 300-person habitation-level eligibility criterion.[21]

### 1.3.3 *"Quadruple-differences" estimator*

For certain industry-level outcomes, we use data from all districts of the state of Rajasthan, which is responsible for approximately eighty percent of India's guar cultivation.[22] In these instances, we rely on variation between (i) firms operating within and outside of industries related to guar-gum production and processing; (ii) guar-growing and non-guar districts; (iii) RGGVY Phase I and non-RGGVY Phase I districts; and (iv) the pre- and post–electrification periods to estimate a difference-in-difference-in-difference-in-differences ("quadruple-differences") specification instead.

Consider the following regression:

$$
\begin{aligned}
y_{idt} = {} & \beta_0 + \beta_1 POST_t + \beta_2 INDUSTRY_{id} + \beta_3 \left(INDUSTRY_{id} \times GUAR_d\right) \\
& + \beta_4 \left(INDUSTRY_{id} \times RGGVY_d\right) + \beta_5 \left(INDUSTRY_{id} \times RGGVY_d \times GUAR_d\right) \\
& + \beta_6 \left(INDUSTRY_{id} \times POST_t\right) + \beta_7 \left(GUAR_d \times POST_t\right) + \beta_8 \left(RGGVY_d \times POST_t\right) \\
& + \beta_9 \left(INDUSTRY_{id} \times RGGVY_d \times POST_t\right) + \beta_{10} \left(INDUSTRY_{id} \times GUAR_d \times POST_t\right) \\
& + \beta_{11} \left(GUAR_d \times RGGVY_d \times POST_t\right) + \beta_{12} \left(INDUSTRY_{id} \times GUAR_d \times RGGVY_d \times POST_t\right) \\
& + \gamma_d + \epsilon_{idt},
\end{aligned}
\tag{1.10}
$$

[21] Indeed, to the best of our knowledge, the only other social program that considers habitation-level population data to decide eligibility is the Pradhan Mantri Gram Sadak Yojana (PMGSY), India's rural roads program. PMGSY connected villages containing a habitation with at least 500 people to India's road network, and a growing body of work uses this eligibility cutoff to evaluate the impacts of rural roads on a host of socioeconomic and environmental outcomes (Adukia et al., 2018; Aggarwal, 2018; Asher and Novosad, 2018; Asher et al., 2018). Given our fifty-person bandwidth around RGGVY's 300-person threshold, however, all villages in our analytical sample would have been ineligible for PMGSY.

[22] Our data are described in detail in Section 1.4.

where $y_{idt}$ represents an outcome for interest for industry $i$ in district $d$ in year $t$. $INDUSTRY_{id}$ represents a binary variable that equals one if industry $i$ is related to the production and processing of guar gum, and zero for all other industries. $GUAR_d$ and $RGGVY_d$ represent binary variables that equal one if district $d$ is a guar-growing or a RGGVY Phase I district, respectively, and zero otherwise. $POST_t$ is a binary variable that equals one if year $t$ is in the post-electrification period. $\gamma_d$ represents a district fixed-effect. $\epsilon_{idt}$ represents an industry-year-specific error term.

Our parameter of interest in Equation (1.10) is $\beta_{12}$, the quadruple-differences estimand that sheds light on how industry-level outcomes evolve within the "guar-processing" industry in booming guar-growing districts where rural electrification rolled out. One might be concerned that changes in this specific industry-district group may occur at the expense of other industries or other types of districts. To evaluate the extent to which this might be the case, we compare our estimate for $\beta_{12}$ with our estimates for the other coefficients in Equation (1.10), which highlight changes in other industry-district groups before and after electrification.[23]

## 1.4 Data

We rely on four main sources of data. First, we refer to technical reports published by the governments of both India and the United States to identify India's main guar-growing districts. We complement these data on guar production with information on the roll-out of rural electrification in India to identify those districts that were approved for electrification under RGGVY Phase I. Next, we obtain data on the composition of the village-level labor force from multiple rounds of the Census of India. We complement these with data on individual-level labor-market outcomes and domestic time allocation from multiple rounds of India's National Sample Survey. Finally, we rely on multiple

[23] For example, $\beta_{11}$ highlights changes in the "non-guar-processing" industry in booming guar-growing RGGVY Phase I districts before and after electrification.

rounds of the Economic Census of India to obtain data on the size and sectoral composition of firms in Rajasthan.

### *1.4.1 Guar production*

We review three separate technical reports on guar production in India to identify our sample of guar-producing Indian districts. Two of these—prepared by the Agricultural and Processed Food Products Export Development Authority (2011) and the National Rainfed Area Authority (2014)—represent efforts by the Indian government to systematically quantify and summarize the nationwide production and trade of guar.[24] The third—prepared by the United States Department of Agriculture—signals the growing interest the agency took in guar production as the crop grew to become India's main agricultural export to the United States (Singh, 2014).

For each of these three reports, we systematically create lists of states and districts that they characterize as key producers of guar in India. In particular, we examine changes in state- and district-level rankings along three related metrics: overall production, total cultivated area, and productivity. We then combine each of our generated lists together, and identify the subset of districts that consistently appear on all three. Based on district boundaries at the time of the 2011 Indian Census, we ultimately identify a total of 23 districts: thirteen in the state of Rajasthan, six in Gujarat, and four in Haryana (Figure 1.1). In 2011, these 23 districts were home to nearly 60 million people living over an estimated area of 300,000 km$^2$—roughly equal in terms of both population and size to all of Italy.[25]

To partially validate our selection of these districts, we also estimate their share in total reported production and area under cultivation for guar using national data

[24] The Agricultural and Processed Food Products Export Development Authority (APEDA) is housed within India's Ministry of Commerce & Industry. It is broadly tasked with supporting the development of industries related to products with export potential. The National Rainfed Area Authority (NRAA) is housed within the Ministry of Agriculture & Farmers welfare, where it provides technical advice and monitoring for government schemes operating in rural areas with significant levels of rainfed agriculture.

[25] In this paper, we interchangeably refer to these districts as India's "guar-growing districts" or "guar-growing regions."

from the Ministry of Agriculture on annual district-wise production of the crop between approximately 1999 and 2015.[26] We note that the quality of these data is poor. For instance, districts in the state of Haryana—consistently referred to in the technical reports we use as one of the most important guar-producing states in India after Rajasthan—have non-missing data on guar production only for 2012. At the same time, other districts in regions of India not known for guar production consistently report trivial amounts of production for multiple years in the sample. Nevertheless, we find that the guar-growing districts we identify account for nearly 94 percent of overall guar production in 2012 (the year that contains these statistics for the largest number of districts).

### *1.4.2 Rural electrification*

As mentioned previously, we identify Phase I districts for which DPRs were successfully submitted and approved using state-level five-year-plan progress reports for RGGVY. Identifying villages that were eligible to be electrified within these districts poses additional challenges. RGGVY implementing agencies were directed to determine a village's eligibility for electrification based on the populations of its constituent habitations (geographically distinct sub-village clusters of households). A village was eligible for electrification under RGGVY Phase I if it contained at least one constituent habitation with a population greater than 300. Although a growing number of public-sector interventions are now tracked at the habitation level, to the best of our knowledge, there are only two comprehensive datasets that shed light on habitation-level populations: (i) the census of habitations conducted by the National Rural Drinking Water Program (NRDWP) in 2009; and (ii) the directory of habitation-level populations made available by the Pradhan Mantri Gram Sadak Yojana (PMGSY), India's national rural roads program. In line with the directives for RGGVY implementing agencies, we rely on the former, which contains habitation-level

[26] These data are available via the Ministry of Agriculture's Crop Production Statistics Information System at `https://aps.dac.gov.in/APY/Index.htm`.

population (by caste) for each village.[27] Because the NRDWP data indicates only the name—and not the unique Census code—for each habitation's corresponding village, we adopt a fuzzy matching algorithm originally developed by Asher and Novosad (2018) to match it with a list of Census-designated villages. India's nearly 600,000 villages consist of just over 1.6 million habitations. We are able to successfully match approximately 531,000 (89 percent) of these villages to their constituent habitations. To further validate the quality of these matches, we calculate the discrepancy between the given Census 2011 population for each village and the NRDWP 2009 population estimate that we obtain from summing over the population of all habitations in a village. We drop all villages with a Census-NRDWP population discrepancy of greater than twenty percent; these, we assume, are incorrect fuzzy matches. This leaves us with approximately 370,000 villages.[28]

Our fuzzy-matched dataset consists of village-level identifiers (i.e., state, district, subdistrict and village names, and their corresponding Census codes), village-level count of habitations, village population (obtained by summing over all habitations in a village), population of the largest habitation, and a variable indicating the quality of the match (i.e., distinct groupings based on the extent to which matches across the NRDWP and Census lists of names are exact or fuzzy). The average village in this fuzzy-matched sample contains three habitations; approximately 47 percent of villages contain exactly one habitation.

To obtain the analytical sample with which to estimate Equation (1.9), we restrict our sample of villages in three ways: (i) those located in RGGVY Phase I districts; (ii) those with exactly one habitation; and (iii) those with a Census 2001 population within a narrow fifty-person bandwidth of the 300-person RGGVY Phase I threshold. This yields 7,655 villages located across 22 Indian states; 148 are located in guar-growing districts.

[27] The NRDWP census of habitation was first conducted in 2003, and again in 2009. The 2003 data are no longer publicly available, which is why we rely on the 2009 data, which are available at `https://indiawater.gov.in`.

[28] We describe our fuzzy habitation–village matching procedure in detail in Appendix C.

### *1.4.3 Rural labor-market outcomes*

Our data on the make-up of the rural labor force come from the 2001 and 2011 rounds of the Indian Census. Specifically, in addition to data on population for each of India's approximately 600,000 villages, the Primary Census Abstract (PCA) data tables in the Census report information by gender on three distinct village-level subgroups: (i) "main workers," who engage in any economically productive activity for at least six months a year; (ii) "marginal workers," who do so for less than six months a year; and (iii) "non-workers," who do not engage in any economically productive activity. Within the first two subgroups, workers are further categorized as cultivators, agricultural laborers, household-industry workers, or "other." A person is classified as a cultivator if they are engaged in cultivation of land that they own or lease, implying that they bear the risks associated with cultivation. In contrast, a person is classified an agricultural laborer if they work on another person's land for payment. In rural areas, a household industry is defined as "production, processing, servicing, repairing, or making and selling (but not merely selling) of goods" that is done by one or more members of a household within the confines of the village. Finally, "other" workers include all professions not captured by the other three categories, such as government employees, teachers and traders.[29]

For each village–year in our Census panel, we combine cultivators and agricultural laborers (both main and marginal) to calculate the population of agricultural workers, overall and by gender. We similarly combine household-industry and other workers to obtain corresponding figures for the village-level population of non-agricultural workers. These data—together with information on village population as well as the breakdown of that population into workers and non-workers—allow us to evaluate impacts along two dimensions: (i) the extensive margin, i.e., the net change in the overall labor force as a percentage of the village population; and (ii) the sectoral composition of the labor force,

[29] Additional information about these definitions is available in the 2011 Census' meta data documentation at `http://www.censusindia.gov.in/2011census/HLO/Metadata_Census_2011.pdf`.

namely, the relative shares of agricultural versus non-agricultural workers.

We complement our village-level data on the composition of the labor force with individual-level data on labor-force outcomes and domestic time allocation from the Employment-Unemployment surveys conducted as part of the 2004 ($60^{\text{th}}$) and 2011-12 ($68^{\text{th}}$) rounds of India's National Sample Survey (NSS).[30] These quinquennial surveys are representative at the level of the NSS region, a non-administrative sampling unit below the state but above the district. NSS regions typically consist of two or more contiguous districts and do not cross state boundaries. We combine our data on district-level guar production and roll-out of rural electrification with these NSS regions to create a region-level repeated cross-section covering over 400,000 people across rural India. In particular, we focus on (i) respondents' "usual principal activity" (the activity an individual contributed the bulk of their time to over the past year); (ii) for those in the labor force, the industry to which they belong; and (iii) for those not in the labor force, the extent to which they engage in home production in addition to domestic duties.

#### *1.4.4 Firm-level data*

Our data on the universe of firms and establishments employing more than ten people in the state of Rajasthan come from the Economic Census (EC) of India. Specifically, we rely on the "Directory of Establishments" associated with the 2005 (Fifth) and 2013-14 (Sixth) rounds of the EC.[31] This directory reports information on basic firm characteristics, including name, address, number of employees, and the sector/industry to which the firm belongs, as indicated by a National Industrial Classification (NIC) code.

In 2005, this directory listed a total of 20,715 firms in Rajasthan. By 2013, this number had increased to 27,803. We combine these two rounds of the EC in a district-level panel

[30] Information on how NSS data can be purchased from the Ministry of Statistics and Program Implementation is available at `http://mospi.gov.in/sample-surveys`.

[31] The Directories of Establishments for the Fifth as well as the Sixth round of the EC are available from the Ministry of Statistics and Program Implementation at `http://www.mospi.gov.in/economic-census-3`.

dataset with which to study changes in the nature and composition of firms in response to guar boom and the roll-out of rural electrification in Rajasthan. We focus, in particular, on firms in industries related to the guar production chain (such as industrial guar-processing units).

## 1.5 Size and sectoral composition of the rural labor force

In this section, we estimate how rural electrification affects the size and composition of the labor force across guar- and non-guar-growing regions of India. We measure this using data on population and employment at the village- and region levels from the Indian Census and National Sample Survey (NSS), respectively. We find no evidence to suggest that electrification has a net effect on the overall size of the labor force in electrified villages located in guar-growing regions of India. We show next, however, that access to electricity substantially reduces (increases) the share of agricultural (non-agricultural) workers in these villages. In electrified villages located in non-guar-growing districts across the rest of India, in contrast, we find no evidence that access to electricity has any discernible effect on the labor-market outcomes that we study.

### *1.5.1 Size of the rural labor force*

We begin by studying the impacts of electrification on the size of the overall labor force (agricultural and non-agricultural workers together) as a share of the village population. We obtain data on the total number of workers in each village from the Indian Census, and apply the RD strategy outlined in Equation (1.9) to identify the effects of electrification on labor-force size separately in guar- and non-guar-growing regions.[32]

Figure 1.5 plots the share of total workers—overall and by gender—just above and below the RGGVY threshold separately for villages located in guar- and non-guar-growing

[32] Total workers includes both "main" and "marginal" cultivators, agricultural laborers, household industry workers, and "other" workers. In 2011, workers comprised approximately 44 percent of the total population of India's villages.

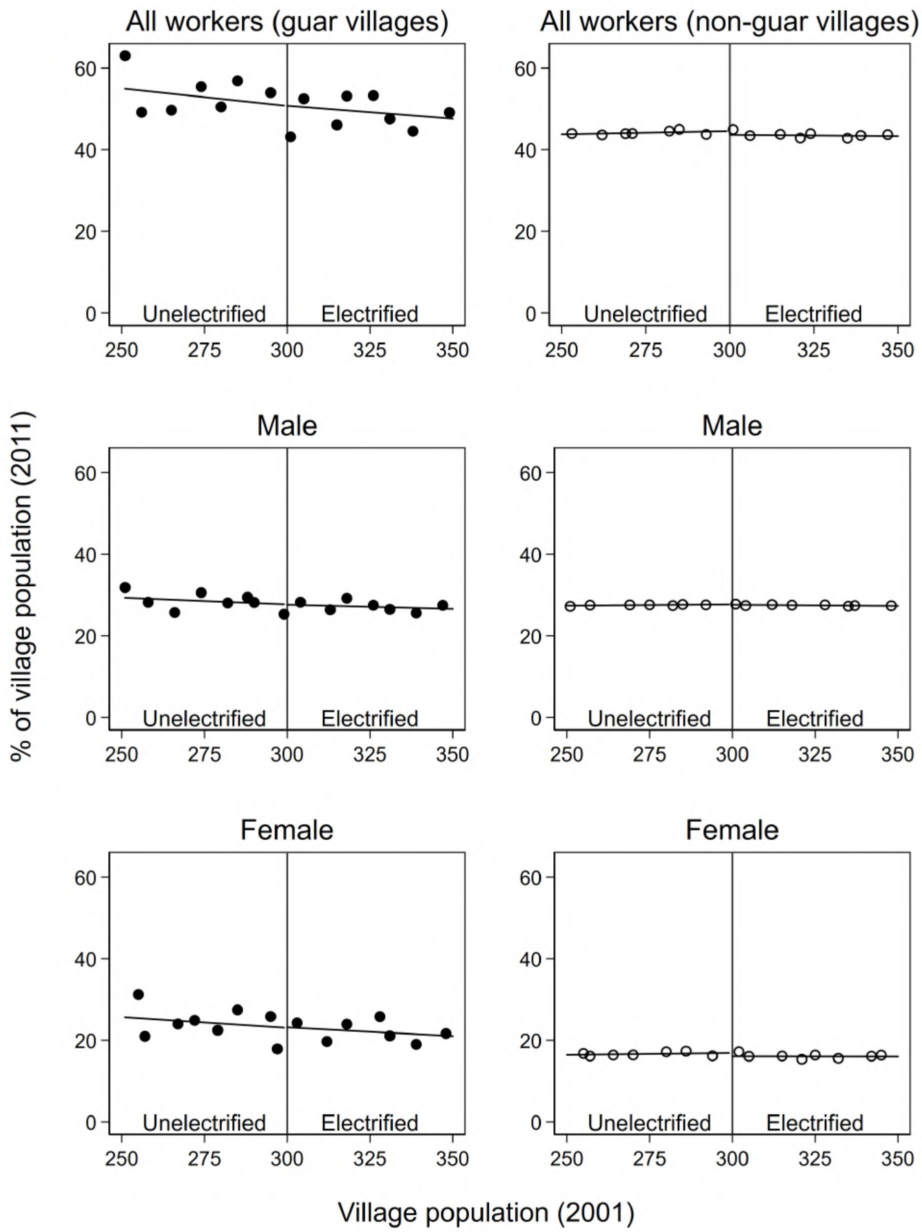


FIGURE 1.5: RD results of impact of electrification on size of labor force. This figure shows the results from estimating the regression specification outlined in Equation (1.9). The left panels show the results for villages located in guar-growing districts; the right panels show the corresponding results for villages located in non-guar-growing districts. Table 1.1 reports associated numerical estimates. Best-fit lines are estimated using the predicted values from the regression. Each solid (hollow) dot represents the mean predicted values for approximately ten (500) villages in fifteen-person bins.

Table 1.1: RD estimates of impact of electrification on size of labor force

| | | (1) | (2) | (3) |
|---|---|---|---|---|
| | | All workers (% of 2011 population) | | |
| | | All | Male | Female |
| $\hat{\beta}_1$ | $\mathbb{1}$ (Village pop. (2001) > 300) | -0.78 | -0.13 | -0.62 |
| | | (0.55) | (0.20) | (0.46) |
| $\hat{\beta}_2$ | $\mathbb{1}$ (Village pop. (2001) > 300) × $\mathbb{1}$ (Village in guar-growing district) | 0.14 | -0.13 | 0.07 |
| | | (2.57) | (1.37) | (1.43) |
| District FEs | | Yes | Yes | Yes |
| Census (2001) controls | | Yes | Yes | Yes |
| $N$ | | 7649 | 7649 | 7649 |
| Adjusted $R^2$ | | 0.39 | 0.38 | 0.39 |
| Mean of outcome | | 43.98 | 27.51 | 16.47 |

This table shows results from estimating Equation (1.9). These results correspond to those presented graphically in Figure (1.5). The outcome variable for each regression comes from the Primary Census Abstract tables of the 2011 round of the Indian Census. Each regression includes all single-habitation villages in RGGVY Phase I districts with a 2001 population within a fifty-person bandwidth of RGGVY's 300-person eligibility threshold. Estimates associated with the population running variable $\left(\tilde{P}^{2001}_{vds}\right)$ are omitted. Following Correia (2015), six singleton observations are excluded. Standard errors—in parentheses—are clustered at the district level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

regions. This figure graphically depicts the results from our RD specification. It shows that electrification has no discernible effect on the size of the labor force in villages in either guar or non-guar-growing regions. Examining these labor-market dynamics separately for male and female workers yields strikingly similar results.

The regression results presented in Table 1.1 support these findings and attach a magnitude to the effects. The estimate in the first row of this table represents the effect of electrification in non-guar-growing districts; as the indicator variable suggests, these villages are located just above RGGVY's eligibility threshold. The estimate in the second row, the interaction of the preceding parameter with the indicator variable for if the electrified village is located in a guar-growing district, thus represents the degree to which the impact of electrification is augmented by the guar boom in villages in India's guar belt. Column (1) reports the main RD estimates for these two parameters for the overall working population. The magnitude of the estimates is small. In non-guar-growing villages, for instance, the results point to a reduction in the overall size of the workforce by 0.8 percentage points (s.e. 0.6), an imprecisely estimated decrease of less than two percent. The estimated coefficient for the additional effect in electrified guar-growing villages in the second row is similarly small. Importantly, neither of these result are statistically significant at conventional levels, and we are unable to reject the hypothesis that access to electricity had no effect on the overall size of the labor force in these two settings.

Columns (2) and (3) report the same specification estimated separately for the share of male and female workers, respectively. The estimates are similar: electrification has no discernible effect on the share of male or female workers in both guar and non-guar villages.

Taken together, these results suggest that, on net, households do not respond to electrification by adjusting their labor choices along the extensive margin.[33] Although we cannot rule out that large-scale entry and exit of workers in response to electrification may be taking place, these findings stand in contrast to those from earlier work (e.g., Dinkelman, 2011) that finds that access to electricity can increase net labor-force participation (especially for women).

### *1.5.2 Sectoral composition of the rural labor force*

*Village-level RD* To shed more light on underlying labor-market dynamics, we turn next to impacts of electrification on the sectoral composition of the rural labor force (agricultural

[33] We also examine the extent to which electrified villages experience large-scale in-migration. We do this by testing for discontinuous changes in the 2011 population of these villages. We find at the RGGVY threshold, electrified villages exhibit a discontinuous increase in the population of the village, driven entirely by an increase in the male population. However, as shown in Table G.2, the magnitude of this change is small (an average increase of approximately three people or less than one percent relative to the sample mean). We, thus, rule out that electrified villages are on the receiving end of large-scale in-migration due to electrification.

and non-agricultural workers separately). Our first set of analyses once again use data from the Indian Census, this time on the village-level population of ("main" and "marginal") cultivators, agricultural laborers, household-industry workers and "other" workers, and non-workers.[34] We combine all workers belonging to the first two of these occupational sub-categories—cultivators and agricultural laborers—to calculate the population and share of agricultural workers for each of the villages in our sample. We similarly combine the last two of these sub-categories—household-industry and "other" workers—to obtain corresponding figures for non-agricultural workers. We use these as our outcome variables to study how the sizes of the agricultural and non-agricultural sector change relative to the size of the non-working population in response to rural electrification.

We find that, in guar-growing regions, electrification substantially reduces the size of the agricultural labor force and increases the size of the non-agricultural labor force. In addition, we find no differential impact of electrification on the share of the non-working population across villages in guar- and non-guar-growing regions of India. Table 1.2 provides numerical results from estimating Equation (1.9) separately for each of these three subgroups. Having electricity reduces the share of agricultural workers in the population of non-guar villages by 1.2 percentage points (s.e. 0.6) relative to a sample mean of approximately 36 percent (column 1). Guar-growing villages, in contrast, exhibit an additional reduction in this share of over six percentage points (s.e. 1.7). The guar boom, thus, leads to an approximately fivefold augmentation in the impact of electrification on the share of the agricultural labor force. Comparing the estimates in the third row for male (column 2) and female (column 3) agricultural workers suggests that the magnitude of this effect is especially large for women. The guar boom augments the reduction in the share of male agricultural workers due to electrification by 2.9 percentage points (thirteen percent) and that of female agricultural workers by 3.3 percentage points (24 percent).

Columns (4)–(6) report corresponding estimates for the non-agricultural labor force.

[34] We describe each of these categories in detail in Section 1.4.3.

Table 1.2: RD estimates of impact of electrification on share of agricultural and non-agricultural workers, and non-workers

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
|---|---|---|---|---|---|---|---|---|---|
| | Agricultural workers (% of 2011 population) | | | Non-agricultural workers (% of 2011 population) | | | Non-workers (% of 2011 population) | | |
| | All | Male | Female | All | Male | Female | All | Male | Female |
| $\hat{\beta}_1$ $\mathbb{1}$ (Village pop. (2001) > 300) | -1.17** | -0.27 | -0.91** | 0.53 | 0.17 | 0.27 | 0.78 | 0.26 | 0.50 |
| | (0.59) | (0.29) | (0.39) | (0.40) | (0.24) | (0.24) | (0.55) | (0.20) | (0.44) |
| $\hat{\beta}_2$ $\mathbb{1}$ (Village pop. (2001) > 300) × | -6.39*** | -2.85*** | -3.25** | 5.60*** | 2.30** | 3.22*** | -0.14 | 1.66 | -1.65 |
| $\mathbb{1}$ (Village in guar-growing district) | (1.71) | (0.97) | (1.34) | (1.19) | (1.12) | (1.23) | (2.57) | (1.55) | (1.50) |
| District FEs | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Census (2001) controls | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| $N$ | 7649 | 7649 | 7649 | 7649 | 7649 | 7649 | 7649 | 7649 | 7649 |
| Adjusted $R^2$ | 0.37 | 0.36 | 0.38 | 0.16 | 0.25 | 0.07 | 0.39 | 0.40 | 0.34 |
| Mean of outcome | 35.96 | 22.27 | 13.68 | 8.02 | 5.23 | 2.79 | 56.02 | 23.66 | 32.36 |

This table shows results from estimating Equation (1.9). These results correspond to those presented graphically in Figures 1.6 (columns 1–3), 1.7 (columns 4–6) and 1.8 (columns 7–9). Outcome variables for regressions reported in columns (1)–(6) are constructed using data from the Primary Census Abstract tables of the 2011 round of the Indian Census. Specifically, "agricultural workers" represents a village-level sum of main and marginal cultivators and agricultural laborers, while "non-agricultural workers" represents a village-level sum of main and marginal household-industry and workers. Each regression includes all single-habitation villages in RGGVY Phase I districts with a 2001 population within a fifty-person bandwidth of RGGVY's 300-person eligibility threshold. Estimates associated with the population running variable $\left(\tilde{P}^{2001}_{vds}\right)$ are omitted. Following Correia (2015), six singleton observations are excluded. Standard errors—in parentheses—are clustered at the district level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The first row of these columns shows that electrification appears to have no discernible impact on the share of the non-agricultural workforce in villages in non-guar districts. Column (4) shows that electrification in guar-growing regions leads to a simultaneous growth in the size of the non-agricultural labor force, which increases by an additional 5.5 percentage points (s.e. 1.2), representing a seventy percent increase relative to the sample mean. This increase is nearly identical to the reduction in the share of agricultural workers in column (1). The second row of columns (5) and (6) shows that this effect is, once again, driven especially by the female workforce. As shown in column (6), the guar boom augments the increase in the share of female non-agricultural workers in guar-growing villages by over three percentage points (s.e. 1.2), approximately 115 percent of the sample mean. For male agricultural workers, the 2.3 percentage point (s.e. 1.1) additional increase represents an increase of just under 45 percent.

Columns (7)–(9) show that these labor-market dynamics in guar villages do not appear to be accompanied by an increase in the relative size of the non-working population. More broadly, comparing the results for guar- and non-guar-growing electrified villages in Table 1.2 shows that while the estimated coefficients for the effect of electrification in non-guar-growing villages generally have the same signs as those for guar-growing villages, the former are considerably smaller in magnitude and largely indistinguishable from zero. In other words, on average, electrification appears to have no discernible effect on the relative size of the agricultural, non-agricultural and non-working population in villages in the RGGVY Phase I districts in the rest of India. Figures 1.6, 1.7 and 1.8 graphically represent the results from our RD specification for agricultural workers, non-agricultural workers and non-workers, respectively, and visually highlight the large differences in impact across the two settings for the first two of these subgroups.

*Does the guar boom drive these results?* Districts that are home to guar production could differ from the rest of India along a variety of metrics, such as agro-ecological conditions, income,

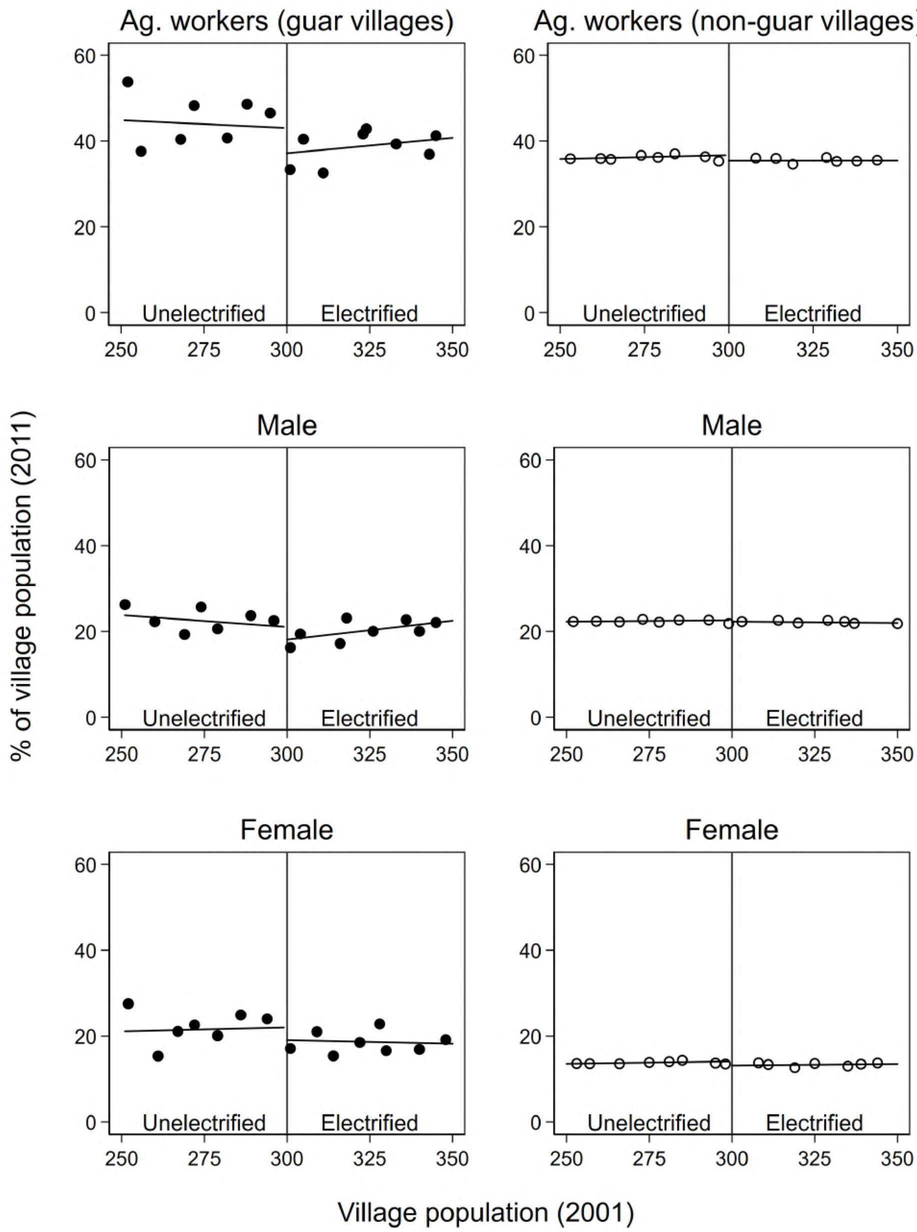


Figure 1.6: RD results of impact of electrification on share of agricultural labor force. This figure shows the results from estimating the regression specification outlined in Equation (1.9). The left panels show the results for villages located in guar-growing districts; the right panels show the corresponding results for villages located in non-guar-growing districts. Columns (1)–(3) of Table 1.2 report associated numerical estimates. Best-fit lines are estimated using the predicted values from the regression. Each solid (hollow) dot represents the mean predicted values for approximately ten (500) villages in fifteen-person bins.

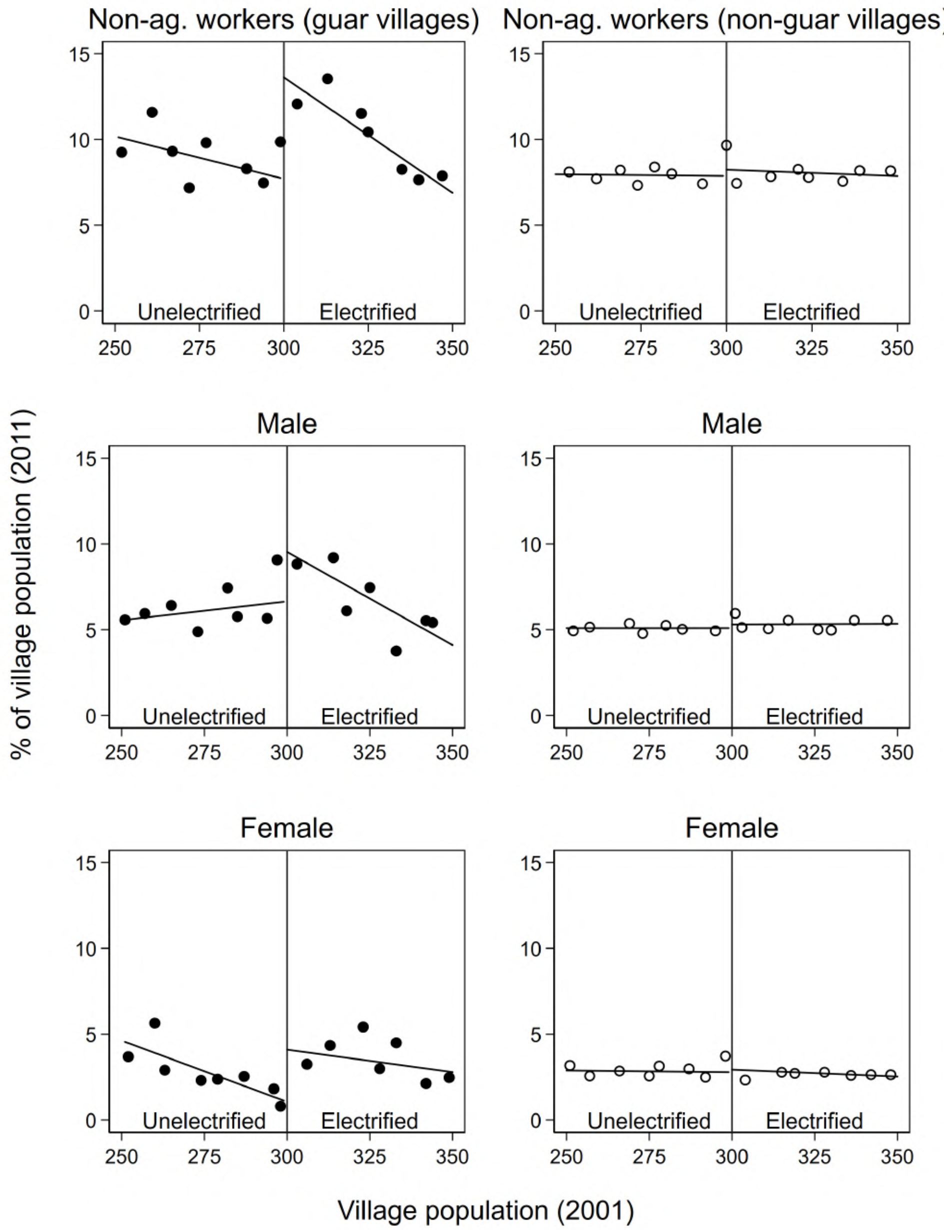


FIGURE 1.7: Regression discontinuity results of impact of electrification on share of non-agricultural labor force. This figure shows the results from estimating the regression specification outlined in Equation (1.9). The left panels show the results for villages located in guar-growing districts; the right panels show the corresponding results for villages located in non-guar-growing districts. Columns (4)–(6) of Table 1.2 report associated numerical estimates. Best-fit lines are estimated using the predicted values from the regression. Each solid (hollow) dot represents the mean predicted values for approximately ten (500) villages in fifteen-person bins.

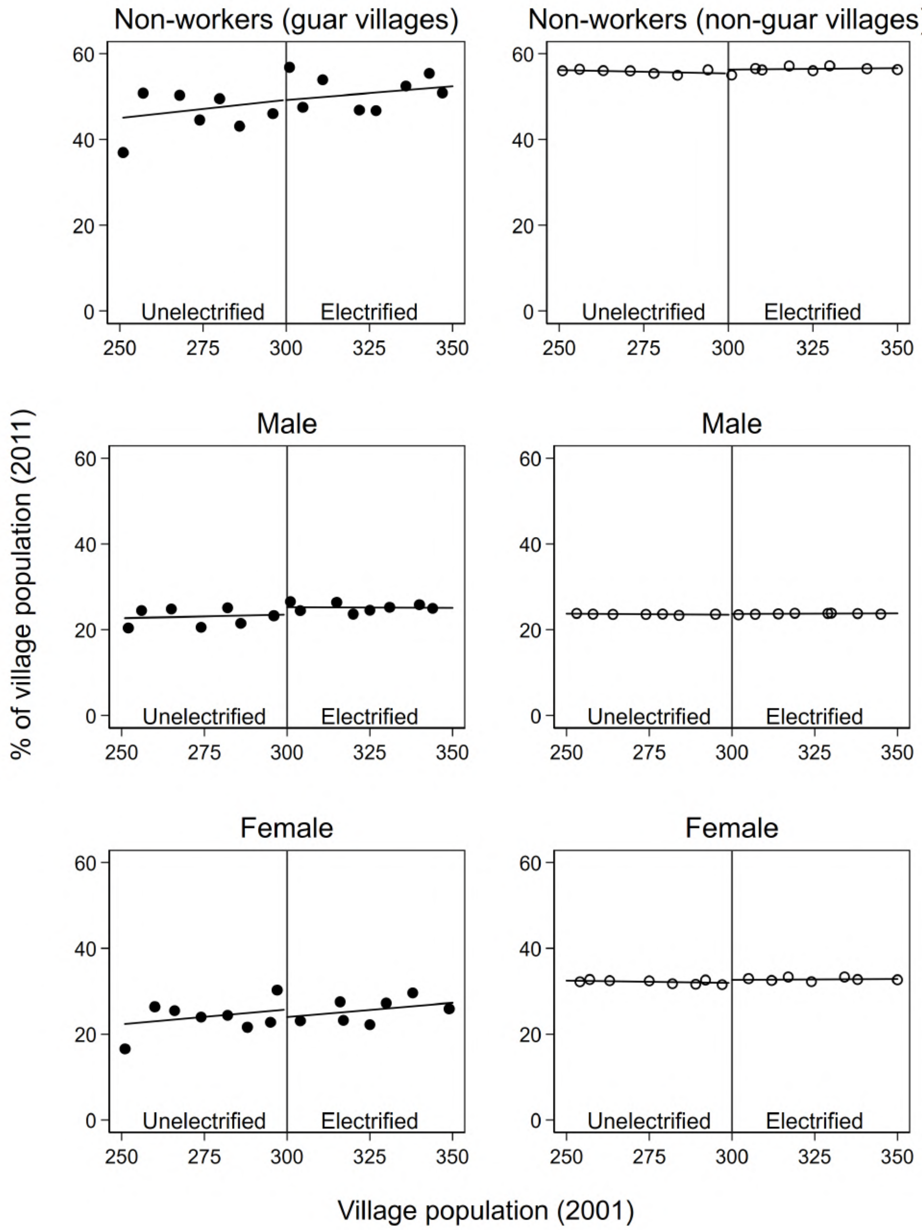


FIGURE 1.8: RD results of impact of electrification on share of non-working population. This figure shows the results from estimating the regression specification outlined in Equation (1.9). The left panels show the results for villages located in guar-growing districts; the right panels show the corresponding results for villages located in non-guar-growing districts. Columns (7)–(9) of Table 1.2 report associated numerical estimates. Best-fit lines are estimated using the predicted values from the regression. Each solid (hollow) dot represents the mean predicted values for approximately ten (500) villages in fifteen-person bins.

and demographics. These could drive the results reported in Table 1.2 independently of the roll-out of rural electrification as part of RGGVY Phase I. The inclusion of district fixed-effects in the RD specification outlined in Equation (1.9), however, soaks up all such unobserved spatial differences. In addition, conditional on state fixed-effects, we find that villages in guar- and non-guar-growing districts are statistically indistinguishable in 2001—before the guar boom or rural electrification—along a host of key socioeconomic indicators (Table G.3).

Nevertheless, if there is considerable district-level heterogeneity in the impacts of rural electrification across India, any random subset of RGGVY Phase I districts can potentially exhibit the differential impacts that we identify in Table 1.2. In other words, the interaction of the guar boom with the roll-out of rural electrification in the eleven RGGVY Phase I districts that are guar growers need not be driving our results; it could be the case that we observe the results that we do simply by chance. To test this, we turn to a randomization-based inference procedure (Athey and Imbens, 2017).

Our approach relies on randomly assigning placebo guar-growing districts and re-estimating Equation (1.9). We repeat this process 1,000 times for the share of agricultural and non-agricultural workers (overall and by gender) to obtain a distribution of placebo estimates for $\hat{\beta}_2$ for each of these outcomes. Figure 1.9 shows these distributions and highlights their 90 and 95 percent confidence intervals. If the differential effect of electrification on the share of agricultural and non-agricultural workers that we observe in guar-growing districts was due to chance, we would expect to observe our actual estimated values for this parameter from Tables 1.1 and 1.2—indicated by the dashed lines in Figure 1.9—near the middle of these distributions. Instead, we find that our estimates of $\hat{\beta}_2$ are extreme values outside the 90 or 95 percent confidence intervals of these distributions in all cases; any other configuration of RGGVY Phase I districts is highly unlikely to yield estimates that are as large in magnitude. Taken together, this strongly suggests that it is

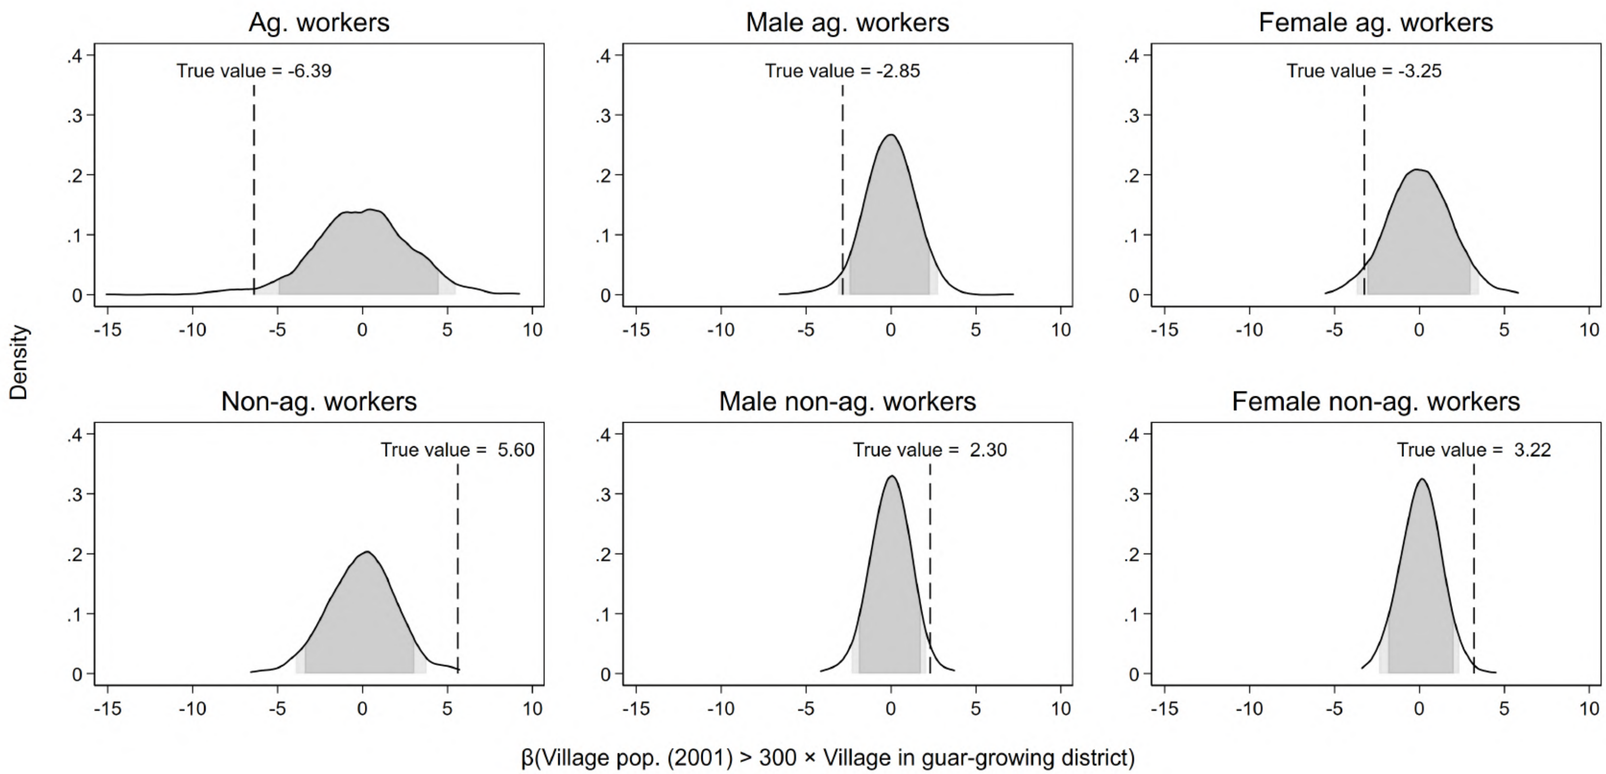


FIGURE 1.9: Evaluating differential impact of electrification in guar-/non-guar-growing districts using randomization inference. Each panel of this figure plots the distribution of 1,000 estimated values of $\hat{\beta}_2$ from a randomization-based inferential procedure (Athey and Imbens, 2017). In each iteration, we randomly assign eleven RGGVY Phase I districts to placebo guar- and non-guar-growing groups, and re-estimate Equation (1.9) to obtain a $\hat{\beta}_2$ placebo value for the degree to which the guar boom augments the impact of electrification. The dashed vertical line indicates the magnitude of the true value of $\hat{\beta}_2$, as reported in Table 1.2. Dark (light) shading represents the 90 (95) percent confidence interval of each distribution.

indeed the advent of the guar boom and its interaction with the simultaneous roll-out of rural electrification as part of RGGVY Phase I that drives the results we observe.

*Validating with region-level difference-in-differences* As an additional test of this difference between the effect of electrification in guar- and non-guar-growing districts of India, we turn to data from the 2004 ($60^{\text{th}}$) and 2011-12 ($68^{\text{th}}$) rounds of the Employment-Unemployment surveys carried out as part of India's National Sample Survey (NSS). These quinquennial surveys report data on individual-level labor-force outcomes that are representative at the level of the NSS region, a non-administrative sampling unit below the state but above the district.[35] Because region boundaries have evolved over the years, we first generate 67 custom region groupings that remain unchanged between 2004 and 2011. We combine these to create a region-level repeated cross-section (before and after the roll-out of rural electrification) containing basic employment and demographic data for nearly 500,000 individuals across rural India. We assume that a region is home to guar production if it encompasses at least one of the 23 guar-growing districts shown in Figure 1.1. Similarly, a region is assumed to be part of the roll-out of rural electrification in India if at least one of its constituent districts was approved for RGGVY Phase I. We restrict our analyses to these latter regions in order to evaluate differences between guar- and non-guar-growing RGGVY Phase I regions.[36]

Central to the Employment-Unemployment survey's categorization of the labor-force status of respondents is their "usual principal activity," the economic or non-economic activity to which the respondent has dedicated the bulk of their time over the past year. For individuals in the labor force, this categorization includes seven mutually exclusive

[35] NSS regions typically consists of two or more neighboring districts that broadly share geophysical characteristics. Larger states are separated into multiple regions. Smaller states and union territories are often encompassed by a single region. Regions do not cross state boundaries. The 2011-12 NSS round divided India into 88 regions, while the 2004 round had 78 regions.

[36] Per this definition, there are 45 RGGVY Phase I regions across the two NSS rounds; three of these are also guar-growing regions. Together, they contain labor-market data for 406,935 rural inviduals.

types of work-related activities.[37] We generate a binary variable that equals one if the respondent reported having engaged in any one of these seven (that is, they were in the labor force), and use this to estimate the following difference-in-differences specification:

$$y_{irst} = \beta_0 + \beta_1 \left(GUAR_{rs} \times POST_t\right) + X_{irst} + \gamma_r + \gamma_{st} + \epsilon_{irst}, \tag{1.11}$$

where $y_{irst}$ is a binary variable that equals one if respondent $i$ in region $r$ in state $s$ in year $t$ is in the labor force and zero otherwise. $GUAR_{rs}$ is a binary variable that equals one if region $r$ contains a guar-growing district. $POST_t$ is a binary variable that equals one for the post-electrification period. $X_{irst}$ represents a control for the age of respondent $i$. $\gamma_r$ is a region fixed-effect, $\gamma_{st}$ is a state–year fixed-effect, and $\epsilon_{rst}$ is a region–year-specific error term.

Table 1.3 presents our results. Column (1) shows that the labor-force participation rate is no different between guar- and non-guar-growing NSS regions that saw the roll-out of rural electrification. This is consistent with the RD results reported in Table 1.1, which show that the effect of electrification on the share of total workers in the village population is broadly indistinguishable across electrified villages located in guar- and non-guar-growing districts.

To shed light on sectoral changes, we use the National Industrial Classification (NIC) code that the NSS provides for each respondent in the labor force to identify the industry to which they belong. To ensure consistency with our RD results, we create a new variable that indicates whether a particular working individual is engaged in agricultural or in non-agricultural work.[38] We use this to estimate Equation (1.11) separately for each of

[37] These are: (i) worked in a household enterprise as an own account worker (self-employed); (ii) worked in a household enterprise as an employer; (iii) worked in a household enterprise as an unpaid family worker; (iv) worked as salaried/wage employee; (v) worked as casual wage labor in public works; (vi) worked as casual wage labor in other types of work; and (vii) did not work but seeking/available for work.

[38] The 2011-12 ($68^{\text{th}}$) round of the NSS lists a five-digit NIC code as per the 2008 NIC system. The 2004 ($60^{\text{th}}$) round relies on the older NIC 1998 system. To match accurately across these two systems, we rely on the first two digits of each NIC system to identify the "division" each individual is employed within; this is the broadest categorization within India's NIC system and largely concordant across the two different NIC

Table 1.3: Differential impact on labor-force participation in guar-growing electrified regions

| Outcome variable | (1) $GUAR \times POST$ Coef. | (2) Std. Err. | (3) $N$ | (4) Adj. $R^2$ | (5) Mean of outcome |
|---|---|---|---|---|---|
| $\mathbb{1}$ (In the labor force) | 0.017 | (0.031) | 406,935 | 0.19 | 0.37 |
| Male | 0.013 | (0.016) | 209,546 | 0.37 | 0.54 |
| Female | 0.016 | (0.046) | 197,389 | 0.16 | 0.19 |

This table shows results from estimating Equation (1.11) on a repeated cross-section of individual-level data from 45 RGGVY Phase I National Sample Survey (NSS) regions. Each row represents a separate regression for all individuals in the sample, overall and by gender. An NSS region consists of two or more contiguous districts within a state, and does not cross state boundaries. An NSS region is defined as a RGGVY Phase I region if it contains at least one RGGVY Phase I district; a RGGVY Phase I region is assumed to also be a guar-growing region if it contains at least one guar-growing district (as shown in Figure 1.1). The underlying data cover a total of 406,935 rural individuals sampled in 2004 ($POST = 0$) and 2011-12 ($POST = 1$). The NSS regions in the dataset are formulated to ensure consistency in regions across those used in 2004 ($60^{\text{th}}$) and 2011-12 ($68^{\text{th}}$) rounds of the NSS. All models include region fixed-effect and state-by-year fixed-effects, and control for the age of the respondent. Standard errors—in column (2)—are clustered at the NSS region level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

the seven subcomponents of labor-force participation.

Table 1.4 reports results. Column (1) shows large reductions in the share of workers employed in agricultural household industries. The share of workers reporting self-employment (the single largest subcomponent of labor-force participation) in agricultural fields, for instance, falls by nearly ten percentage points (fifteen percent relative to the mean) in electrified guar-growing regions relative compared to electrified non-guar-growing regions. This is consistent with the RD results pointing to a reduction in the size of the agricultural labor force in electrified guar-growing villages in Table 1.2.

Finally, we turn to home production, the central component of the conceptual framework we outline in Section 1.3. For the subset of respondents not in the labor force, the NSS indicates whether their "usual principal status" is related to domestic responsibilities

systems. We assume an individual is engaged in agricultural work if their employment type falls within "Division 01: Crop and animal production, hunting and related service activities."

Table 1.4: Differential impact on farm-related labor force in guar-growing electrified regions

| Outcome variable: 𝟙 (Agricultural worker) | (1) *GUAR* × *POST* Coeff. | (2) Std. Err. | (3) $N$ | (4) Adj. $R^2$ | (5) Mean of outcome |
|---|---|---|---|---|---|
| *Labor-force subcategory:* | | | | | |
| 𝟙 (Household enterprise: Own account worker) | -0.095*** | (0.016) | 57,257 | 0.14 | 0.62 |
| 𝟙 (Household enterprise: Employer) | -0.65*** | (0.058) | 1,391 | 0.19 | 0.67 |
| 𝟙 (Household enterprise: Unpaid family worker) | -0.024** | (0.0093) | 34,079 | 0.081 | 0.84 |
| 𝟙 (Salaried/wage employee) | 0.0098 | (0.018) | 18,336 | 0.13 | 0.04 |
| 𝟙 (Casual wage labor: Public works) | 0.0068 | (0.0053) | 1,350 | 0.18 | 0.05 |
| 𝟙 (Casual wage labour: Other types of work) | -0.038 | (0.074) | 34,126 | 0.29 | 0.57 |
| 𝟙 (Seeking/available for work) | – | (–) | 5,195 | – | 0.00 |

This table shows results from estimating Equation (1.11) on a repeated cross-section of individual-level data from up to 45 RGGVY Phase I National Sample Survey (NSS) regions. Each row represents a separate regression for distinct subgroups of the rural labor force; in each of these regressions, the outcome variable is a binary variable that equals one if the respondent is in a farming-related industry (Division 01 "Crop and animal production, hunting and related service activities," as per the 2008 National Industrial Classification [NIC] system). An NSS region consists of two or more contiguous districts within a state, and does not cross state boundaries. An NSS region is defined as a RGGVY Phase I region if it contains at least one RGGVY Phase I district; a RGGVY Phase I region is assumed to also be a guar-growing region if it contains at least one guar-growing district (as shown in Figure 1.1). The underlying data cover a total of 406,935 rural individuals sampled in 2004 ($POST = 0$) and 2011-12 ($POST = 1$). The NSS regions in the dataset are formulated to ensure consistency in regions across those used in 2004 ($60^{th}$) and 2011-12 ($68^{th}$) rounds of the NSS. All models include region fixed-effect and state-by-year fixed-effects, and control for the age of the respondent. Singleton observations are dropped. Standard errors—in column (2)—are clustered at the NSS region level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

only or to domestic responsibilities in combination with home production.[39] We use these to estimate a linear probability model in line with Equation (1.11) with the sample of the rural population that is not in the labor force. Table 1.5 presents results. Column (1) shows a six percentage point (sixty percent) reduction in the share of the population engaged in only domestic duties in electrified guar-growing regions. At the same time, the share of the population engaged in domestic duties in combination with home production in these regions increases by a nearly identical amount, albeit one that is less precisely estimated. Column (1) also shows that this change is almost entirely driven by a shift in the share of women engaged in home production in addition to domestic duties. These results,

[39] Specifically, the NSS contains two mutually exclusive categories for a subset of the population not in the labor force: (i) "domestic duties only," which includes all activities that constitute the care economy, such as looking after the young, the sick and the elderly as well as other healthy household members, cooking, cleaning and provisioning for the household; and (ii) what we refer to as "domestic duties and home production," which also includes being engaged in free collection of goods (vegetables, roots, firewood, cattle feed), sewing, tailoring, weaving, etc. for household use.

Table 1.5: Differential impact on home production in guar-growing electrified regions

| Outcome variable | (1) $GUAR \times POST$ Coeff. | (2) Std. Err. | (3) $N$ | (4) Adj. $R^2$ | (5) Mean of outcome |
|---|---|---|---|---|---|
| $\mathbb{1}$ (Domestic duties only) | -0.061*** | (0.0079) | 406,935 | 0.042 | 0.10 |
| Male | -0.00058 | (0.0011) | 209,546 | 0.0043 | 0.003 |
| Female | -0.13*** | (0.019) | 197,389 | 0.095 | 0.21 |
| $\mathbb{1}$ (Domestic duties and home production) | 0.052* | (0.026) | 406,935 | 0.060 | 0.10 |
| Male | -0.0084** | (0.0033) | 209,546 | 0.0025 | 0.003 |
| Female | 0.12** | (0.052) | 197,389 | 0.14 | 0.20 |

*Notes.* This table shows results from estimating Equation (1.11) on a repeated cross-section of individual-level data from 45 RGGVY Phase I National Sample Survey (NSS) regions. Each row represents a separate regression for distinct subgroups of individuals, overall and by gender. "Domestic duties" includes all activities that constitute the care economy, such as looking after the young, the sick and the elderly as well as other healthy household members, cooking, cleaning and provisioning for the household, while "home production" includes being engaged in free collection of goods (vegetables, roots, firewood, cattle feed), sewing, tailoring, weaving, etc. for household use. An NSS region consists of two or more contiguous districts within a state, and does not cross state boundaries. An NSS region is defined as a RGGVY Phase I region if it contains at least one RGGVY Phase I district; a RGGVY Phase I region is assumed to also be a guar-growing region if it contains at least one guar-growing district (as shown in Figure 1.1). The underlying data cover a total of 406,935 rural individuals sampled in 2004 ($POST = 0$) and 2011-12 ($POST = 1$). The NSS regions in the dataset are formulated to ensure consistency in regions across those used in 2004 ($60^{th}$) and 2011-12 ($68^{th}$) rounds of the NSS. All models include region fixed-effect and state-by-year fixed-effects, and control for the age of the respondent. Standard errors—in column (2)—are clustered at the NSS region level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

too, are broadly consistent with the impacts of electrification on household's labor and time-allocation decisions being driven by local economic contexts.

*RD robustness checks* We first use RGGVY Phase II districts in India to conduct a large-scale placebo test. Specifically, using only those districts of India that were not approved for rural electrification as part of RGGVY Phase I, we estimate Equation (1.9) for the overall share of workers, agricultural workers, non-agricultural workers, and non-workers in the village population.[40] As large-scale roll out of rural electrification did not occur in these districts over the period covered by our data, we should not expect to see an impact of a village's 2001 population being above RGGVY's eligibility threshold in either guar-growing or non-guar districts. Table G.4 confirms this intuition.

We turn next to our analytical sample of villages. In constructing the sample of single-

[40] As shown in Figure 1.1, twelve non-RGGVY Phase I districts are also guar-growing districts.

habitation villages for our main RD analyses, we made two key choices: (i) during our village-habitation fuzzy matching procedure, we discarded any village with a discrepancy of greater than twenty percent between its total Census 2011 population and its total NRDWP 2009 population (calculated by combining the population in each of its matched habitations); and (ii) we restricted our sample to villages within a narrow fifty-person bandwidth of RGGVY's 300-person eligibility threshold. We test the sensitivity of our main results to each of these choices.

We first estimate Equation (1.9) allowing for increasingly greater levels of population discrepancy in our sample but keeping our preferred fifty-person RD bandwidth fixed. Figure H.1 shows how $\hat{\beta}_2$, our parameter of interest, evolves as we relax our definition of what we consider a successful match, thereby increasing the size of the underlying analytical sample. As the sample expands to contain an increasing number of villages that are unlikely to have been good matches, the magnitude of $\hat{\beta}_2$ generally attenuates gradually as expected. In particular, we do not observe erratic changes in the magnitude of this estimated parameter.

Next, we fix the sample population discrepancy rate at our preferred level of twenty percent and vary the size of the RD bandwidth around RGVVY's 300-person eligibility threshold. Figure H.2 shows how $\hat{\beta}_2$ evolves as the RD bandwidth widens. Once again, as the analytical sample expands to contain an increasingly dissimilar number of villages on either side of the RGGVY eligibility threshold, the magnitude of $\hat{\beta}_2$ attenuates smoothly.

Finally, we adjust our inference to account for multiple hypothesis testing using the free step-down resampling methodology of Westfall and Young (1993). This bootstrap-based procedure controls the family-wise error rate (the probability of a type I error when testing a "family" of hypotheses).[41] We combine all regressions reported in Tables 1.1 and 1.2 into a family of hypotheses and use this approach to control the family-wise error

[41] See Jones et al. (2018) for a detailed description of how this is implemented.

rate associated with $\hat{\beta}_2$. Table G.5 reports that our main result—that electrified villages in guar-growing districts see a large reduction in the share of agricultural workers and a corresponding increase in the share of non-agricultural workers relative to electrified villages in non-guar districts—is robust to this adjustment.

## 1.6 Growth of firms

Many scholars contend that firms' location decisions (the extensive margin) are informed by differences in comparative advantage that arise due to spatial heterogeneity driven by agglomeration, market size, and production or transport costs (Amiti and Javorcik, 2008; Carlton, 1983; Wheeler and Mody, 1992). In the short run, a shock that differentially impacts some locations may also induce changes in firm size (the intensive margin) (Adhvaryu et al., 2013).

In our setting, such firm-level location and size decisions are naturally linked to the rural labor market. Industry and agriculture exist side-by-side in rural India (Srivastava and Srivastava, 2010). In addition, firms' choices are are often influenced by the availability of infrastructure (Martin and Rogers, 1995) and closely related to the employment effects associated with commodity-price shocks (Lederman and Porto, 2015). These factors are the focus of our analyses. Changes in firm-level characteristics at the extensive and intensive margin are, therefore, a useful way to uncover potential mechanisms.

In this section, we estimate how the impact of rural electrification on firm creation and size differed across guar- and non-guar-growing districts. To measure impacts on firm proliferation, we rely on a "triple-differences" specification applied to district-level panel data on the universe of firms in the state of Rajasthan. To examine firm size, we further exploit variation between firms in two broad industrial groups (firms within and outside industries related to the production and processing of guar gum), which allows us to employ a "quadruple-differences" approach. Along the extensive margin, we do not find an increase in the number of firms in "guar-related" industries in electrified guar-growing districts of Rajasthan relative to its unelectrified and/or non-guar districts. However, we uncover large increases in the size of firms in industries related to guar production in terms of the number of workers. Specifically, we show that "guar-related" firms located in guar-growing districts where rural electrification rolled out grew in size, broadly consistent with the increase in non-agricultural employment demonstrated in Section 1.5.2. Importantly, this growth does not appear to happen at the expense of firms operating in other industries, non-guar districts or non-RGGVY Phase I districts.

### *1.6.1 Proliferation of guar-processing firms*

We look first at differential impacts of electrification on the proliferation of firms across guar and non-guar districts of Rajasthan. Our data come from the "Directory of Establishments" associated with the Fifth (2005) and Sixth (2013-14) rounds of the Economic Census (EC) of India. This directory reports information on basic firm characteristics, including name, number of employees (within a range), and the industry to which the firm belongs, as indicated by a National Industrial Classification (NIC) code. Because guar processing does not have its own NIC code, we use the 2013 EC's directory (which lists a total of 30,000 establishments in Rajasthan containing at least ten employees) to identify the set of NIC codes that can be assigned to guar processors. We start by finding guar-processing units in the directory that are easily identifiable as such (based on their use of "guar," "guar gum" or some variant thereof in their names) and record the NIC codes assigned to them. We complement this step with a review of the detailed breakdown of NIC codes prepared by India's Central Statistical Organization to identify additional codes that can contain guar-processing units.[42] Ultimately, we identify five three-digit NIC codes that can contain industrial units most directly related to guar processing.[43] Together, these

[42] The 2013 EC uses a 2008 update of the NIC system. This document is available at `http://mospi.nic.in/classification/national-industrial-classification`.

[43] As per the 2008 NIC system, these are: (i) Support activities to agriculture and post-harvest crop activities (016); (ii) Manufacture of basic chemicals, fertilizer and nitrogen compounds, plastics and synthetic rubber in primary forms (201); (iii) Manufacture of prepared animal feeds (108); (iv) Manufacture of non-metallic

represent approximately ten percent of all listed establishments in Rajasthan.

For each district-year in the EC, we calculate the number of firms that belong to one of these industries as a percentage of the total number of firms to create a district-level panel. We use this to estimate a triple-differences version of the specification outlined in Equation (1.10).[44] We find no evidence that the share of firms belonging to guar-related industries in guar-growing RGGVY Phase I districts changes at a different rate relative to other types of districts between 2004 and 2013. Column (1) of Table 1.6 reports our numerical estimates. The coefficient for the triple-interaction term (representing the triple-differences estimate of the additional effect of electrification in guar-growing districts) is statistically indistinguishable from zero. This suggests that an increase along the extensive margin (that is, the establishment of new guar-processing units) is not the main channel through which firms respond.

### *1.6.2 Growth in the size of guar-processing firms*

We turn next to the relative sizes of guar-related firms to shed light on firm-level responses along the intensive margin. The Directory of Establishments in both rounds of the EC categorizes each firm into one of three groups: those with 10-100 employees, with 101-500 employees, or with greater than 500 employees. For each district–year, we calculate the share of guar-related and non-guar firms that belong in each group and use this to estimate Equation (1.10).

Figure 1.6.2 plots these shares for the first two firm-size groups—which account for nearly all firms in our sample—and conveys the essence of our quadruple-differences approach. This graph shows that, in guar-growing RGGVY Phase I districts, the share of guar-related firms with 10-100 employees fell by over four percentage points between 2005

mineral products (239); and (v) Wholesale of agricultural raw materials and live animals (462). We use the concordance tables prepared by India's Central Statistical Organization to map these codes to the 2004 NIC system, which is used in the 2005 EC.

[44] Specifically, because our outcome variable is the share of firms in guar-related industries, we do not exploit inter-industry variation.

Table 1.6: Impact of electrification on firm type and size in Rajasthan

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | Guar-related firms (% of all firms) | Number of employees (% of all firms) | | |
| | | 10-100 | 101-500 | >500 |
| *POST* | -0.11 | 0.58 | -0.72 | 0.14 |
| | (0.96) | (0.71) | (0.64) | (0.24) |
| *GUAR* × *POST* | 0.84 | -1.03 | 1.52 | -0.49 |
| | (1.81) | (1.41) | (1.15) | (0.36) |
| *RGGVY* × *POST* | 1.80 | 0.41 | -0.23 | -0.18 |
| | (2.25) | (0.89) | (0.81) | (0.26) |
| *GUAR* × *RGGVY* × *POST* | -3.18 | 1.40 | -1.65 | 0.25 |
| | (2.82) | (1.54) | (1.29) | (0.38) |
| *INDUSTRY* | | -3.65 | 3.87 | -0.23 |
| | | (6.59) | (6.59) | (0.62) |
| *INDUSTRY* × *GUAR* | | -1.99 | 2.32 | -0.33 |
| | | (8.11) | (8.36) | (0.75) |
| *INDUSTRY* × *RGGVY* | | 2.91 | -3.07 | 0.16 |
| | | (6.73) | (6.72) | (0.68) |
| *INDUSTRY* × *RGGVY* × *GUAR* | | -0.72 | 0.48 | 0.24 |
| | | (8.36) | (8.61) | (0.82) |
| *INDUSTRY* × *POST* | | -0.66 | -0.89 | 1.55 |
| | | (3.29) | (3.00) | (1.02) |
| *INDUSTRY* × *RGGVY* × *POST* | | 0.05 | 1.02 | -1.07 |
| | | (3.44) | (3.13) | (1.16) |
| *INDUSTRY* × *GUAR* × *POST* | | 3.30 | -2.46 | -0.84 |
| | | (3.31) | (3.15) | (1.15) |
| *INDUSTRY* × *GUAR* × *RGGVY* × *POST* | | -12.24** | 11.39** | 0.85 |
| | | (5.30) | (5.26) | (1.29) |
| District FEs | Yes | Yes | Yes | Yes |
| Number of districts | 32 | 33 | 33 | 33 |
| Mean of outcome | 10.36 | 94.60 | 4.86 | 0.54 |
| *N* | 64 | 129 | 129 | 129 |
| Adj. $R^2$ | 0.81 | 0.32 | 0.32 | 0.23 |

This table shows results from estimating Equation (1.10). The results reported in columns (2)–(4) are related to those presented graphically in Figure (1.6.2). The outcome variable for each regression is calculated using data from the "Directory of Establishment" of the 2005 ($POST = 0$) and 2013-14 ($POST = 1$) rounds of the Economic Census of India. Standard errors—in parentheses—are clustered at the district level. A firm is assumed to belong to a "guar-related" industry if its 2008 National Industrial Classification (NIC) code is one of the following: (i) Support activities to agriculture and post-harvest crop activities (016); (ii) Manufacture of basic chemicals, fertilizer and nitrogen compounds, plastics and synthetic rubber in primary forms (201); (iii) Manufacture of prepared animal feeds (108); (iv) Manufacture of non-metallic mineral products (239); and (v) Wholesale of agricultural raw materials and live animals (462). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

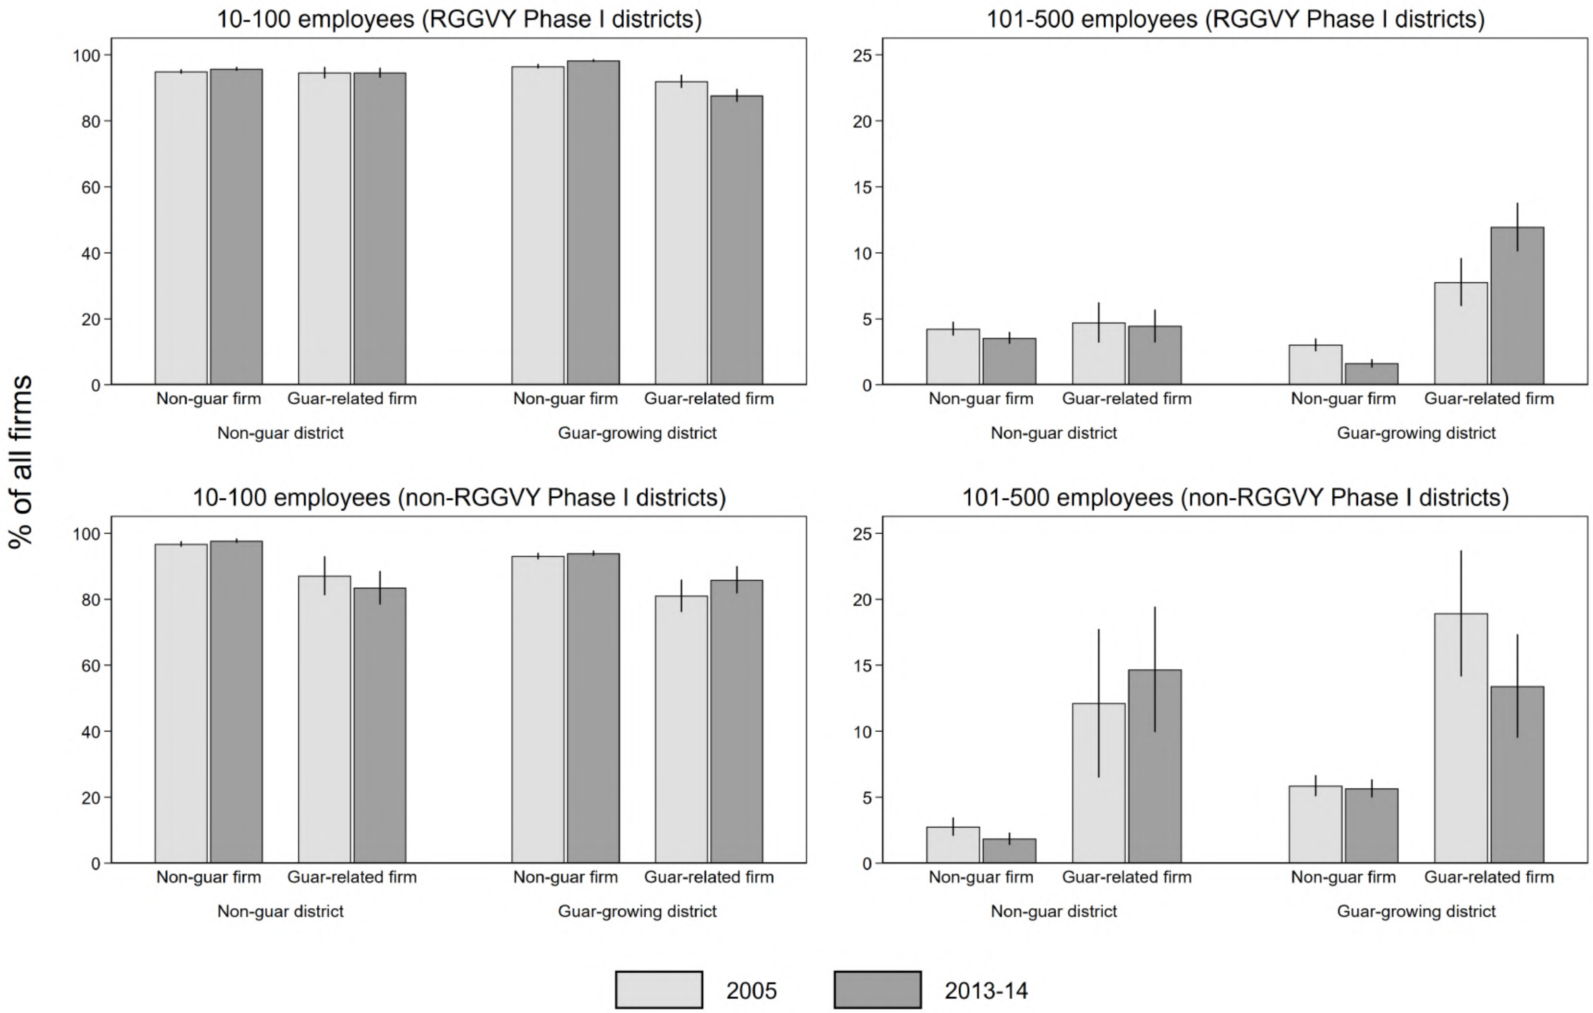


FIGURE 1.10: Distribution of guar-related firms in RGGVY Phase I districts of Rajasthan. This figure shows the distribution, by firm size, of the firms listed in the "Directory of Establishments" associated with the 2005 ($N$ = 13,816) and 2013-14 ($N$ = 18,336) rounds of the Economic Census of India. A firm is assumed to be in a guar-related industry if its 2008 National Industrial Classification (NIC) code is one of the following: (i) Support activities to agriculture and post-harvest crop activities (016); (ii) Manufacture of basic chemicals, fertilizer and nitrogen compounds, plastics and synthetic rubber in primary forms (201); (iii) Manufacture of prepared animal feeds (108); (iv) Manufacture of non-metallic mineral products (239); and (v) Wholesale of agricultural raw materials and live animals (462). All firms in the sample are located in one of Rajasthan's RGGVY Phase I districts. Error bars indicate 95 percent confidence intervals for the means.

and 2013; over the same period, those with 101-500 employees increased by nearly the same amount. This suggests that firms in industries related to guar processing expanded their operations in guar-growing districts that were approved for rural electrification as part of RGGVY relative to those that were not. This graph also shows that these expansions do not appear to be accompanied by contractions in non-guar RGGVY Phase I districts.

The regressions in columns (2)–(4) of Table 1.6 confirm this finding. Column (2) reports the quadruple-differences estimate for the share of guar-related firms with 10-100 employees. In guar-growing RGGVY Phase I districts, this share falls by approximately twelve percentage points (s.e. 5.3) between 2005 and 2013. Over the same period, the share of guar-related firms with 101-500 employees increases by a nearly identical amount, as shown in column (2); relative to the sample mean of approximately five percent, this represents a more than doubling of the share of guar-related firms that employ between 101-500 people. In addition, we find no evidence to suggest that this growth in firm size occurs at the expense of firms in other industries or firms in non-guar/non-RGGVY Phase I districts; the estimates in all other rows of columns (2)–(3) are relatively small in magnitude and statistically indistinguishable from zero.

Broadly, these findings are consistent with reports of substantial increases in projected installed capacity by guar-gum manufacturers in northwestern India as the fracking boom began in the United States (Rai, 2015). We add to this largely observational evidence by demonstrating that firms in industries benefiting from the boom are cognizant of local infrastructural contexts, and restrict their expansions primarily to electrified areas.

## 1.7 Conclusion

In this paper, we combine two natural experiments—an exogenous fracking-induced boom in the production of a crop called guar in northwestern India, and population-based discontinuities in the contemporaneous roll-out of India's massive rural electrification scheme—within a regression discontinuity design to evaluate how the causal effect of rural electrification on labor-market outcomes changes with exogenous variation in economic conditions and contexts. We assemble a variety of evidence from multiple large administrative datasets to reach three main conclusions. Our first finding is that, in villages located in India's guar-growing regions, access to electricity led to a large increase in non-agricultural employment relative to agricultural employment, especially among women. Our second finding is that these labor-market dynamics appear to be driven by an increase in employment by electricity-intensive industrial firms that complement guar production (such as guar-processing units) near these communities. It is also related to a proliferation of household-level enterprises and home production in these areas. Finally, our third finding is that, on average, access to electricity appears to have no discernible impact on these labor-market outcomes in villages located in the rest of India.

The main implication of these findings is that complementary economic conditions and contexts are crucial for the ultimate impacts of large-scale electrification. Proponents have long claimed that reliable electricity delivered by the grid is foundational for the structural transformation of rural economies. Its potential to drive job creation and employment growth is often central to this argument.[45] Yet the evidence base on this point remains thin. In particular, impact evaluations are typically unable to rigorously shed light on drivers of spatial and temporal heterogeneity. We show that access to electricity from the grid led to large-scale structural transformation of the rural economy in large swathes of northwestern India, which saw the rise of complementary economic opportunities. In the rest of India, where these complementary conditions were lacking on average, the impacts of grid-scale electrification on rural labor-market outcomes were largely negligible.

These results highlight the role electrification—and large-scale infrastructure, more

[45] In an evaluation of one of its grid expansion projects in Namibia, for instance, the Swedish International Development Corporation Agency notes that "the most important effect of electricity expansion on the poor is …through the effect it can have on the general economic development by providing power to new investments in industry and small businesses" (Goppers, 2006). Other international donors echo these sentiments (Independent Evaluation Group, 2008; United Nations, 2018).

broadly—can play in low- and middle-income countries. Alone, such investments may be insufficient, yet built in anticipation of (and to support) other policies and changes, large-scale infrastructure can provide a foundation for sustained economic growth and development. In our setting, access to grid-scale electricity allowed individuals, households, and firms to respond to rapidly changing economic contexts in ways that potentially deliver economic benefits and improve welfare. We believe that rigorously identifying other potential drivers of the success of large-scale infrastructure is a promising avenue for future research.

# 2

# NGOs and the effectiveness of interventions

*With Marc Jeuland and Subhrendu K. Pattanayak*

## 2.1 Introduction

Implementation challenges stymie policies and programs in virtually every sector. Transaction costs borne to identify beneficiaries and service providers, gauge their trustworthiness, bargain to reach consensus, and monitor agreements to ensure they are fulfilled can undermine the effectiveness of interventions—particularly in low-income settings in which such frictions are especially important (Holloway et al., 2000; Ito, 2009; Jack and Suri, 2014; Schaner, 2016). Where public-service delivery has proven difficult, a variety of non-state, non-market institutions—most notably nongovernmental organizations (NGOs)—have emerged (Werker and Ahmed, 2008). Indeed, NGOs increasingly play lead roles in implementation on the ground; by some estimates, India alone is home to over three million of them—a figure that outnumbers its public hospitals, schools and police force (Anand, 2015).[1] That they are (at least in theory) nimble and efficient has made them attractive partners for international donors. In 2012, for instance, Organisation for Economic Co-operation and Development (OECD) countries channeled over $17 billion of

[1] In contrast, the total number of international NGOs is more conservatively estimated to be closer to 30,000 (Aldashev and Navarra, 2014).

overseas development assistance to—and through—NGOs (Aldashev and Navarra, 2014). Private charitable giving to international development NGOs may be considerably higher (Micklewright and Wright, 2003). Yet little is known about the ways in which NGOs directly impact the effectiveness of the interventions they implement.

In this study, we develop a model of household decision-making motivated by transaction costs to evaluate how NGOs influence intervention outcomes. We then use a novel matched-experimental study design to test our model's main theoretical claim: that an NGO lowers transaction costs where it has been more active, which in turn increases intervention effectiveness. We test this claim in the context of an intervention designed to promote improved cookstoves (ICS) in rural India. Nearly three billion people globally rely on traditional stoves and solid fuels for their primary energy needs. ICS, which are designed to increase the efficiency of solid-fuel combustion, have long been seen as a potential solution to the environmental-health-development burden these energy-use patterns impose (Jeuland and Pattanayak, 2012).[2] Yet widespread uptake has proven challenging (e.g., Mobarak et al., 2012) and—as in other arenas—the role played by local institutions in driving adoption remains poorly understood despite considerable scholarship on what encourages uptake (Lewis and Pattanayak, 2012).

We first use *ex ante* propensity score matching to create a sample of observationally similar villages that are differentiated by prior exposure to a local development NGO. In partnership with this NGO, we then randomly assign nearly 100 geographically distinct hamlets within these villages (covering a sample of almost 1,000 households) to treatment and control groups as part of an experimental ICS-promotion intervention. Our results suggest that the intervention increases adoption: nearly half of all households targeted by the promotion campaign purchased an ICS. However, our study design also allows

[2] Reliance on traditional fuels and stoves imposes a tremendous burden on household health (via exposure to household air pollution); local forests (via unsustainable extraction of fuelwood); and the global climate (due to widespread biomass burning). Bailis et al. (2015), Jeuland and Pattanayak (2012), Jeuland et al. (2015b) and Venkataraman (2005) provide an overview of the global nature and magnitude of these environmental, health and welfare challenges.

us to identify the direct impact of the NGO on ICS adoption and energy-use patterns. We uncover a large, positive and statistically significant “NGO effect”—purchase rates are nearly thirteen percentage points (28 percent) higher in treated communities with prior interactions with the NGO. This finding is robust to multiple ways of defining the scope of the NGO’s prior relationships with communities. Placebo tests inspired by randomization-based inferential procedures also reveal that it is highly unlikely to have been the result of the chance selection of the set of villages in our sample. Using a “triple-differences” specification—which further relaxes our identifying assumptions—we find that treated households in NGO communities are also sixteen percentage points more likely to use intervention stoves than treated households in communities without a prior relationship with the NGO, representing a fifty percent increase in the size of the treatment effect. Consistent with these patterns of adoption and use, treated households in NGO communities exhibit significant reductions in the use of solid fuels and in fuel-collection times. In contrast, we find no evidence of similar improvements in energy-use patterns for treated households in non-NGO communities. Our stratified study design, therefore, reveals that we would have considerably overestimated the effectiveness of our intervention had it been a typical randomized evaluation conducted in partnership with the NGO.

These findings thus highlight some of the implications of the increasing roles NGOs play as partners in research, especially in recent years in response to growing calls for more “evidence-based” policies (Banerjee et al., 2007, 2017). Charity evaluators (such as GiveWell and Giving What We Can) routinely release lists of top NGOs whose effectiveness is rigorously evaluated for would-be donors in search of causes. In the fiercely competitive world of fundraising, such signals do not go unnoticed, and NGOs are often keen to work with researchers and have the impacts of their programs assessed. For applied researchers, partnerships with NGOs provide ready access to target populations, local expertise, and human resources and operational infrastructure. These lower the costs of doing research and create opportunities for researchers to test new theories directly. Randomized controlled trials (RCTs) in developing countries have particularly benefited from these trends. Indeed, proponents argue that "many of the best RCTs have come from long-term partnerships between researchers and NGOs or other local partners" (Glennerster, 2015). NGO–researcher collaborations undoubtedly yield valuable insights about the effectiveness of environmental, health and development interventions (e.g., Banerjee et al., 2015; Brooks et al., 2016; Jayachandran et al., 2017; Miguel and Kremer, 2004). That said, we show that this seemingly symbiotic relationship may also mask the unique roles NGOs often play in remote, rural settings—with serious implications for the generalizability and scalability of solutions deemed effective in applied research (Berge et al., 2012; Peters et al., 2018; Vivalt, 2017).

More broadly, our study is motivated by calls for research that sheds light on how and why certain interventions are effective, not simply whether or not they are (Deaton, 2010; Pattanayak et al., 2017; Rosenthal et al., 2017). In particular, we demonstrate how solutions from applied research that is tied to the places, populations and programs associated with specific NGOs may have very different impacts when implemented in alternative settings (Ravallion, 2009).[3] Our study is also related to a growing literature on the role of implementer identity, which finds that NGO-led interventions are generally more effective than comparable efforts by other actors (Bold et al., 2013; Cameron and Shah, 2017; Grossman et al., 2016; Henderson and Lee, 2015). Yet such comparisons implicitly assume that there is something inherently different about NGOs that leads to implementation effectiveness. As we show, this is not necessarily the case, and even heretofore "effective" NGOs struggle to overcome implementation-related barriers when forced to operate in

[3] This is often referred to as "site selection bias" (Allcott, 2015) and can arise in at least two different ways. First, as Lin et al. (2012) show in the context of forestry interventions, the spatial distribution of social-welfare programs is non-random. The presence of such programs for the purposes of evaluation is correlated with past levels of NGO activity, which is itself strategically determined by NGOs (Brass, 2012; Fruttero and Gauri, 2005; MacLean et al., 2015). Second, NGOs capable of managing complex RCTs (designed to provide evidence of program effectiveness) are also likely to implement programs more effectively than the average implementer.

new settings in which their stock of social capital is low. This is because effective local NGOs often spend years fostering trust in the communities in which they operate. These typically unobserved context- and NGO-specific characteristics interact with aspects of the intervention, influence transaction costs associated with implementation, and ultimately help determine final outcomes.[4]

Nevertheless, we recognize that ours is not the only study grappling with these questions, and that there are tangible opportunities for us to incorporate lessons from this broader literature on the importance of implementer identity. To evaluate the relative strength of the evidence we uncover, we use estimates from Bold et al. (2013) and Cameron and Shah (2017)—who compare the effectiveness of an NGO-led intervention to a government-led one—to formally specify a distribution for our prior understanding of the effects NGOs might have on final outcomes. How does our study contribute to this prior knowledge? To answer this question, we use this prior distribution as a core component of a multilevel Bayesian analysis, with which we revisit our main results. We show that a synthesis of existing insights and our results in this way yields posterior distributions of the magnitude of the "NGO effect" that are considerably more precise. Importantly, we demonstrate that our results overwhelmingly point to the direction of this effect being positive.

This paper proceeds as follows: in Section 2.2, we develop our theoretical model of household decision-making in the presence of NGOs and transaction costs; Section 2.3 provides an overview of our data and sample-selection methods; Section 2.4 outlines our empirical framework and identification strategy; Section 2.5 presents results; Section 2.6 demonstrates how insights from the related literature inform our results within a Bayesian framework; and Section 2.7 concludes.

[4] For instance, if beneficiaries are unsure about the benefits and costs associated with unfamiliar welfare-improving technologies (such as ICS), trusted NGOs can leverage these relationships to reduce households' perceptions of the risks associated with technology adoption.

## 2.2 NGOs, transaction costs and household decision-making

Transaction costs are inherent in the adoption of new technologies—especially in remote, rural settings characterized by low information (Bernard et al., 2017; Foster and Rosenzweig, 2010; Suri, 2011). For example, while the material costs of new technologies are usually borne immediately, benefits are often uncertain and realized far in the future; resources are thus required to learn about the valuable attributes of the technology and gauge the full extent of its costs and benefits. In addition, the search and exchange process that technology adoption entails can impose significant costs. These are often related to the size of the market. Relatively large markets—such as major urban areas—feature a multitude of retailers and suppliers, competing along technology price, quality and differentiability criteria. In contrast, relatively small markets—such as remote, rural settings—are characterized by weak supply chains, a paucity of suppliers, and limited options. These contextual characteristics influence households' decisions about technology adoption, particularly in smaller markets, in which kinship ties, reciprocal exchange and repeated dealings (in other words, social capital) are more salient (Kranton, 1996).

These insights motivate our model of household decision-making. Drawing on Jeuland et al. (2015b)—who in turn build on more fundamental work in environmental and health economics (Grossman, 1972; Pattanayak and Pfaff, 2009)—we develop a model in which households decide whether to invest in technologies that avert environmental health risks (such as ICS). These decisions necessarily involve a trade-off with consumption of other goods and leisure: households must maximize utility by allocating limited resources to environmental and health investments, consumption, and leisure. Accordingly, household utility ($u$) is a function of consumption ($c$), leisure ($l$), technology adoption ($a$), time spent sick ($s$) and household environmental quality ($e$). Sickness is determined by household environmental quality, which is itself determined by the household's adoption decisions. The household's utility function is assumed to be twice differentiable, continuous and concave.

The household faces a full-income constraint, whereby its exogenous income ($y$) must be allocated to consumption ($c$) as well as the materials ($m$), time ($t$) and knowledge ($k$)—with prices $p$, $w$ and $r$, respectively—required for technology adoption. Similarly, total available time ($T$) must be allocated to leisure, time spent sick and time allocated to the technology-adoption decision.

The Lagrangian associated with the household's optimization problem is as follows:

$$\begin{aligned} \max_{a,l,c,t,m,k} \mathcal{L} = u\left[c, l, a, s\left(e\right), e\left(a\right)\right] + \lambda\left[y - c - p\left(\xi\right) \cdot m - r\left(\xi\right) \cdot k + w\left(T - s\left(e\right) - l - t\left(\xi\right)\right)\right] \\ + \mu\left[T - l - s\left(e\right) - t\left(\xi\right)\right]. \end{aligned} \tag{2.1}$$

Note that the material and time costs associated with the technology adoption are presented in Equation (2.1) as a function of the activities of an NGO, denoted by $\xi$.[5] In addition, while NGO activity is relatively less likely to influence the price of time (the local market wage rate), it can influence the amount of time a household needs to allocate to the technology-adoption decision; accordingly, $t$ is also a function of $\xi$.

The first-order conditions associated with the maximization problem outlined in

[5] Most fundamentally, differences in NGO activity arise because NGOs are present in certain locations and not in others. However, such differences may also arise due to the amount of time an NGO has been active in a particular community, or due to the nature and intensity of the programs it chooses to implement there.

Equation (2.1) are as follows:

$$\mathcal{L}_a = u_a + u_s s_e e_a + u_e e_a - \lambda \left[p(\xi) \cdot a_m + r(\xi) \cdot a_k + w\left(a_{t(\xi)} + s_e e_a\right)\right] - \mu \left[a_{t(\xi)} + s_e e_a\right] = 0 \tag{2.2}$$

$$\mathcal{L}_l = u_l - \lambda \cdot w - \mu = 0 \tag{2.3}$$

$$\mathcal{L}_c = u_c - \lambda = 0 \tag{2.4}$$

$$\mathcal{L}_t = u_a a_{t(\xi)} + u_s s_e e_a a_{t(\xi)} + u_e e_a a_{t(\xi)} - \lambda w \left(1 + s_e e_a a_{t(\xi)}\right) - \mu = 0 \tag{2.5}$$

$$\mathcal{L}_m = u_a a_m + u_s s_e e_a a_m + u_e e_a a_m - \lambda\left(p(\xi) + w s_e e_a a_m\right) - \mu = 0 \tag{2.6}$$

$$\mathcal{L}_k = u_a a_k + u_s s_e e_a a_k + u_e e_a a_k - \lambda\left(r(\xi) + w s_e e_a a_k\right) - \mu = 0 \tag{2.7}$$

$$\mathcal{L}_\lambda = y - c - pm - rk + w\left(T - s(e) - l - t(\xi)\right) = 0 \tag{2.8}$$

$$\mathcal{L}_\mu = T - l - s(e) - t \geq 0\,; \mu\left(T - l - s(e) - t(\xi)\right) = 0. \tag{2.9}$$

Assuming that all individuals work some non-zero hours, it follows from Equation (2.9) that $\mu = 0$. Then, from Equation (2.2), the optimal level of technology adoption ($a^*$) must fulfill

$$\underbrace{\frac{u_a + u_s s_e e_a + u_e e_a}{\lambda} - w s_e e_a}_{\text{Marginal benefit, } MB(a)} = \underbrace{p(\xi) \cdot a_m + r(\xi) \cdot a_k + w a_{t(\xi)}}_{\text{Marginal cost, } MC(a,\xi)}. \tag{2.10}$$

The first term on the left-hand side of Equation (2.10) represents the monetary value of the change in utility arising from marginal investments in welfare-improving technologies; the second term is the opportunity cost of sickness valued at the price of time (the wage rate, $w$). The right-hand side represents the marginal cost of investing in an additional unit of $a$, disaggregated by the costs associated with materials, knowledge and time commitment necessary for technology adoption.

Let $\pi(a, \xi)$ represent household welfare (or net benefit). Then:

$$\left.\frac{\partial \pi(a, \xi)}{\partial a}\right|_{a=a^*} = MB(a) - MC(a, \xi) = 0 \tag{2.11}$$

$$\frac{\partial^2 \pi(a, \xi)}{\partial a^2} = \frac{\partial MB(a)}{\partial a} - \frac{\partial MC(a, \xi)}{\partial a} < 0, \tag{2.12}$$

where Equation (2.12) follows from the concavity of the household's utility function.

We further assume that:

$$p'(\xi) < 0, p''(\xi) < 0 \tag{2.13}$$

$$r'(\xi) < 0, r''(\xi) < 0 \tag{2.14}$$

$$t'(\xi) < 0, t''(\xi) < 0. \tag{2.15}$$

For the intuition behind the identities outlined in Equations (2.13)–(2.15), consider the roles local NGOs might play in facilitating technology adoption in remote, rural settings. Effective NGOs are often intimately familiar with local infrastructure limitations (such as access to electricity) and work to identify the technologies most suited to local contexts on behalf of their beneficiaries. This reduces households' time costs. Their status as trusted sources of relevant information may similarly reduce costs associated with the acquisition of knowledge, reducing the need for households to independently attempt to verify claims: believing that an NGO is reliable and trustworthy—the NGO's existing stock of social capital within its community—lowers the risks of adopting technologies recommended by it. Finally, NGOs may also provide beneficiaries with subsidies or discounts, lowering costs associated with the acquisition of materials directly. These subsidies or discounts need not be externally funded. For instance, an NGO's relatively larger size may allow it to exploit economies of scale associated with acquiring technologies in bulk, resulting in cost savings that it could pass onto its beneficiaries.[6]

Given these conditions, we can evaluate how NGO activity influences technology

[6] It is worth noting settings in which these assumptions might not hold. For instance, while an NGO's positive reputation may increase the willingness of beneficiaries to adopt solutions suggested by it, its assessment of its own "reputation risks" (Herman et al., 2003) could lead it to restrict the set of interventions it chooses to implement. Indeed, our local implementation partner expressed reservations about implementing a potentially welfare-improving ICS-promotion intervention in the study setting, highlighting in particular the reputation risks the NGO faced should ICS-promotion activities related to our study (described in Section 2.4.1) be received poorly by its beneficiaries. In addition, while our implementing partner lobbied during the design phase for reductions in the baseline price of ICS that its beneficiaries would ultimately face, it is not difficult to imagine a rent-seeking NGO behaving very differently. Needless to say, there are a variety of organizations active in remote, rural settings, and we do not intend to insinuate that our model provides general principles for how these actors will behave under all conditions.

adoption directly using the implicit-function theorem as follows:

$$\frac{da}{d\xi} = -\frac{\partial^2 \pi \big/ \partial a \partial \xi}{\partial^2 \pi \big/ \partial a^2} \tag{2.16}$$

$$= -\frac{-\,\partial MC\,(a,\xi) \big/ \partial \xi}{\partial^2 \pi \big/ \partial a^2} \tag{2.17}$$

$$> 0. \tag{2.18}$$

The numerator in Equation (2.17) is positive. This follows from Equations (2.13)–(2.15), which imply that the marginal cost of technology adoption is decreasing in $\xi$. The negative sign of the denominator is established in Equation (2.12), implying that the sign of the expression in Equation (2.17) is positive. In other words, the welfare-maximizing level of technology adoption increases with an increase in NGO activity, all else being constant.[7]

Our model thus yields a tractable definition of the transaction costs associated with technology adoption: the opportunity costs of allocating time, materials and knowledge to the adoption process. In structuring risks and variability, fostering collective action, and influencing existing practices, NGOs influence the set of costs associated with investments in welfare-improving technologies—and, indeed, in participating in environmental, health and development interventions more broadly. Recent work has looked at the roles NGOs play in enhancing monitoring and enforcement (Aldashev et al., 2015; Grant and Grooms, 2017); improving public-service delivery (Devarajan et al., 2013); and working with het-

[7] The implicit-function theorem cannot be used to evaluate the impact of discrete changes in NGO activity. One may be interested, for instance, in how an NGO's decision to begin operating in a community in which it had not previously done so affects households' avoidance behavior. To see the effects of such a change, note that an increase in $\xi$ from $\xi'$ to $\xi^*$ causes a discrete reduction in marginal cost of avoidance behavior, analogous to the continuous case. Under this condition, household welfare exhibits increasing marginal returns to $\xi$, as demonstrated below:

$$\begin{aligned}\frac{\partial \pi\,(a,\xi^*)}{\partial a} - \frac{\partial \pi\,(a,\xi')}{\partial a} &= (MB\,(a) - MC\,(a,\xi^*)) - \left(MB\,(a) - MC\left(a,\xi'\right)\right)\\ &= -MC\,(a,\xi^*) + MC\left(a,\xi'\right)\\ &> 0.\end{aligned}$$

Let $\arg\max \pi\,(a,\xi') = a'$ and $\arg\max \pi\,(a,\xi^*) = a^*$. Then $a^* > a'$ by the strict-monotonicity theorem (Edlin and Shannon, 1998).

erogeneous beneficiaries (Bengtsson, 2013). Outside of a nascent stream of research on the strategic nature of NGOs' location decisions, however, little is known about the roles that they might play in directly determining the outcomes of applied interventions (Brass, 2012; Fruttero and Gauri, 2005; MacLean et al., 2015). Indeed, while some have noted the presence of differential impacts across communities with and without NGO activity (e.g., Niehaus and Sukhtankar, 2013; Sharma et al., 2015), to the best of our knowledge no study has rigorously examined how the presence of an effective local NGO and the specific institutional context that represents might directly influence household decision-making. In contrast, our model provides a clear, empirically testable hypothesis: where NGO activity has been higher, intervention effectiveness—in our case, in the form of uptake of ICS technologies—will also be higher. The remainder of this paper turns to an empirical evaluation of this claim.

## 2.3 Data and descriptive statistics

Our sampling frame consists of over 1,000 households across 97 geographically distinct hamlets (*toks*) located in 38 Census-delineated villages (*gram panchayats*) in the northern Indian state of Uttarakhand. In this section we (i) describe the creation of observationally equivalent groups of these 38 NGO and non-NGO villages using *ex ante* propensity score matching; (ii) outline the random selection of households in hamlets located within these villages; and (iii) present descriptive statistics for our full sample of households.

### *2.3.1 Creation of the sample using propensity score matching*

Given our interest in the influence of community-level institutional factors—as well as the fact that interventions are often initiated at the community level—we developed a sampling strategy designed to minimize the risk associated with community-level confounders influencing our analysis. Specifically, our strategy for the pre-baseline selection of communities (and distribution of households across these communities) relied on a matched-experimental design that sought to ensure sufficient variation to create the contextual strata of interest for the study, namely previous engagement with our known institutional partner (an Uttarakhand-based environmental NGO).[8]

We first conducted an exhaustive enumeration of all villages in two districts of Uttarakhand—Bageshwar and Nainital—that had previously been targeted in a program implemented by a local development NGO.[9] In total, we identified 148 distinct villages lying in our NGO stratum, which we refer to as "NGO villages." Using data on community-level characteristics for these villages from the 2001 round of the Indian Census, we next applied propensity score matching to identify observationally similar "non-NGO villages" in which the NGO had not previously intervened (Rosenbaum and Rubin, 1984).[10] The matching approach allowed us to purposely select for variation within our institutional stratum, while maintaining similarity across the strata in terms of a large number of observable Census characteristics. In the first stage of this approach, we excluded any village with fewer than forty households from our sampling frame, and then estimated logistic regression models for the outcome of selection into a previous NGO intervention as a function of community-level Census variables.[11]

[8] Pattanayak et al. (2009) present a step-by-step process of creating such a sample with an application to a similar environmental health program in rural India. King et al. (2007) describe a large-scale application of a similar approach in Mexico to evaluate the impacts of universal health insurance.

[9] The NGO leads activities related to agriculture and forestry (promotion of sustainable agricultural practices, sustainable fodder cultivation and promotion of culinary herbs), health (local hospitals/clinics), education (local schools), village-level groups (self-help groups, youth groups and vocational cooperatives), and water management (watershed renewal and spring-water recharge).

[10] This does not mean that no NGO had ever worked in these villages, only that our local implementation partner had not. Our analyses, thus, focus on how relationships with particular implementing NGOs influence intervention effectiveness.

[11] In total, we estimated three distinct specifications of these logit models. In the first specification of this model, we included sub-district (block) fixed effects, which helped restrict the set of controls to very local villages, and omitted variables that were frequently missing in the Census data (e.g., access to bus services, tap water and/or electricity availability characteristics). The second specification dropped these sub-district fixed effects, whereas the third specification was similar to the second except that it included the additional controls (with missing values in the Census assumed to be zero, or omitted in the case of tap-water availability). We eliminated the second specification because it was clearly less robust to the distance restrictions for our sample, i.e., fewer matches were preserved when dropping the sub-districts far away from our base of operations. We also found that the quality of matches on a number of variables was

In the first stage we find that, in general, NGO villages are slightly smaller and have proportionally more Scheduled Caste members than the average village (Table G.6).[12] NGO villages are also more likely to have infrastructure or village-level institutions (schools, credit societies and bus facilities), but are more remote (further away from large towns, and with less access to paved roads and telephones) than non-NGO villages.

In the second stage, following the logit estimation, we estimated the probability of receiving a previous NGO intervention for all villages in our eligible districts. These predicted probabilities constitute the propensity score for each village. We matched NGO villages to non-NGO villages with the most similar propensity scores (allowing for replacement, and limiting matches to the support region with the greatest overlap in density of NGO and non-NGO villages). In this way, we ensured that the communities selected for our sampling frame were as similar as possible. For each model, we developed a matching routine that restricted our potential sample villages on the basis of size (greater than forty households) and distance from the base of operations for our survey (working in sub-districts that could be reached within one day). We next eliminated the worst ten percent of matches on the basis of propensity score distance—a process known as "trimming"—to ensure that pairs that are poor matches are not selected simply due to the inclusion of villages that do not happen to have good matches (Crump et al., 2009).

Finally, to draw our precise sample of matched pairs, we studied each of the individual pairs remaining after trimming in detail. Here, we paid particular attention to match quality and overall balance between the NGO and non-NGO strata with regards to key contextual factors (such as population, distance from nearby towns, and the presence of village-level groups and societies). Our final sample consisted of 38 villages (nineteen matched pairs of NGO and non-NGO villages). At the conclusion of our matching exercise,

greatly reduced with inclusion of the sub-district-level controls used in the first specification. Thus, our final analysis is based on the third specification (see Figure H.3 and Table G.6 for additional details).

[12] The Scheduled Castes and Scheduled Tribes are various historically disadvantaged or indigenous groups in India that have received official recognition as such from the Indian government.

Table 2.1: Post-match balance on village-level characteristics across NGO and non-NGO villages

| Village-level characteristic | (1) NGO villages Mean | (2) NGO villages Std dev. | (3) Non-NGO villages Mean | (4) Non-NGO villages Std dev. | (5) Normalized difference |
|---|---|---|---|---|---|
| Area (km$^2$) | 146.7 | 94.5 | 175.0 | 268.6 | −0.10 |
| Total population | 386.8 | 136.6 | 376.7 | 130.0 | 0.05 |
| Scheduled Caste$^\dagger$ population (proportion) | 0.26 | 0.29 | 0.28 | 0.30 | −0.03 |
| Scheduled Tribe$^\dagger$ population (proportion) | 0.0065 | 0.028 | 0.00036 | 0.0016 | 0.21 |
| Number of primary schools | 1.11 | 0.46 | 1.05 | 0.23 | 0.10 |
| Number of middle schools | 0.37 | 0.50 | 0.32 | 0.48 | 0.08 |
| Number of health centres | 0 | 0 | 0 | 0 | – |
| Number of primary health centres | 0 | 0 | 0 | 0 | – |
| Number of telephone connections | 0.26 | 0.56 | 0.42 | 0.51 | −0.20 |
| $\mathbb{1}$ (Bus services) | 0.11 | 0.32 | 0.053 | 0.23 | 0.13 |
| $\mathbb{1}$ (Credit societies) | 0 | 0 | 0 | 0 | – |
| $\mathbb{1}$ (Approach to village: paved road) | 0.21 | 0.42 | 0.16 | 0.37 | 0.09 |
| Distance from nearest town (km) | 25.1 | 16.1 | 19.9 | 11.9 | 0.25 |
| Forest area (hectares) | 34.1 | 63.9 | 26.0 | 51.8 | 0.10 |
| $\mathbb{1}$ (Tap water) | 0.89 | 0.32 | 1 | 0 | −0.32 |
| $\mathbb{1}$ (Electricity for all purposes) | 0.053 | 0.23 | 0 | 0 | 0.22 |
| Observations | 19 | | 19 | | 38 |

Columns (1) and (3) show means for NGO and non-NGO villages, respectively, for each of the village-level variables from the 2001 round of the Indian Census used for our final propensity score matching exercise. Columns (2) and (4) show the respective standard errors for these means. Column (5) shows the normalized difference between the two means. $^\dagger$The Scheduled Castes and Scheduled Tribes are various historically disadvantaged or indigenous groups in India that have received official recognition as such from the Indian government. $^{*}$ $p < 0.10$, $^{**}$ $p < 0.05$, $^{***}$ $p < 0.01$.

our matched pairs of NGO and non-NGO villages were balanced on all community-level characteristics used during our propensity-score estimation (Table 2.1).[13]

### 2.3.2 *Sub-cluster survey samples and household surveys*

Uttarakhand is one of India's least densely populated states, with terrain that gives rise to "sub-clusters" (geographically distinct hamlets known as *toks*) within villages.[14] These

[13] Table G.7 presents additional tests of cross-sectional balance using more detailed community-level data from the 2011 round of the Indian Census. We did not use this Census round for our matching exercise as it had only been released provisionally at the time. Nevertheless, we find that NGO and non-NGO villages are balanced across multiple dimensions using these data as well, particularly once we make adjustments to account for multiple hypothesis testing.

[14] Topographically, Uttarakhand is characterized by "hilly terrain, rugged and rocky mountains, deep valleys, high peaks, swift streams and rivulets, rapid soil erosion, frequent landslides and widely scattered

hamlets typically vary in terms of cultural and socioeconomic characteristics. To maximize sample variation along these socioeconomic lines we randomly selected between two and four hamlets within each matched village for our final sample. Specifically, we determined that our survey teams would work within at least two hamlets in small villages, at least three in medium villages, and at least four in large villages (owing to population differences across these three groups of villages). Our final sample consisted of 97 hamlets.

Households were randomly selected to participate in survey activities; baseline surveys occurred during the summer of 2012.[15] If household members were unavailable during the entire day of survey activities—or if they refused to participate—neighboring households were randomly selected as replacements. Field supervisors performed household introductions and obtained informed consent, recorded GPS coordinates and elevation data, and oversaw quality-control checks in each village. A random subsample of households was also selected for detailed weighing of daily solid-fuel use.[16]

The ICS-promotion intervention (described in detail in Section 2.4.1) began in August 2013, with midline and endline surveys taking place around November 2013 and November 2014—the three- and fifteen-month marks, respectively.

### *2.3.3 Descriptive statistics*

Table 2.2 presents an overview of our main sample of 943 households, disaggregated by NGO and non-NGO villages.[17] The average household in our study consists of just under

habitation" (Maurya, 2014).

[15] Highly variable village structures and geographic constraints created variation in the number of hamlets and households sampled in each village. A minimum of twenty surveys were completed in small villages, thirty in medium ones, and forty in large ones. If a village was divided into distinct geographical sub-units (e.g., half the village was to the north of the main road, while the other half was to the south), the target number of surveys was split equally among these groups.

[16] This process involves asking households to collect an amount of firewood and other solid fuels that is slightly more than what they anticipate using over the next day. This amount is weighed by the field team, which returns approximately 24 hours later to weigh the remaining amount.

[17] We interviewed 1,063 households at baseline. We restrict our main analytical sample to the 943 households that were also located and interviewed during the midline (three-month) and endline (fifteen-month) survey rounds.

Table 2.2: Overview of sample households

| Household-level characteristic | (1) NGO villages Mean | (2) NGO villages Std dev. | (3) Non-NGO villages Mean | (4) Non-NGO villages Std dev. | (5) Normalized difference |
|---|---|---|---|---|---|
| *Demographics* | | | | | |
| Household size | 5.23 | 2.12 | 4.52 | 1.89 | 0.24*** |
| Mean number of children aged five and under | 0.55 | 0.86 | 0.38 | 0.72 | 0.16*** |
| Age of household head (years) | 53.8 | 13.7 | 54.2 | 14.0 | −0.023 |
| $\mathbb{1}$ (Female-headed household) | 0.25 | 0.44 | 0.27 | 0.44 | −0.026 |
| *Socioeconomic characteristics* | | | | | |
| Education level of household head (years) | 6.23 | 4.51 | 6.14 | 4.62 | 0.014 |
| Education level of primary cook (years) | 4.66 | 4.32 | 4.55 | 4.62 | 0.018 |
| $\mathbb{1}$ (Below poverty line) | 0.58 | 0.49 | 0.58 | 0.49 | −0.0074 |
| *Stove- and fuel-use characteristics* | | | | | |
| $\mathbb{1}$ (Owns traditional stove) | 0.99 | 0.12 | 0.97 | 0.16 | 0.061 |
| $\mathbb{1}$ (Owns improved stove) | 0.31 | 0.46 | 0.28 | 0.45 | 0.038 |
| Traditional-stove use (minutes per day) | 300.3 | 142.3 | 284.5 | 137.6 | 0.080 |
| $\mathbb{1}$ (Used an improved stove in past week) | 0.29 | 0.46 | 0.28 | 0.45 | 0.028 |
| $\mathbb{1}$ (Uses traditional fuels) | 0.99 | 0.11 | 0.97 | 0.16 | 0.073 |
| $\mathbb{1}$ (Uses a clean fuel daily) | 0.28 | 0.45 | 0.26 | 0.44 | 0.034 |
| Traditional-fuel collection (minutes per day) | 129.6 | 101.5 | 101.8 | 87.5 | 0.20** |
| *Beliefs and perceptions* | | | | | |
| $\mathbb{1}$ (Heard of stoves that produce less smoke) | 0.30 | 0.46 | 0.22 | 0.41 | 0.13** |
| $\mathbb{1}$ (Heard of fuels that produce less smoke) | 0.37 | 0.48 | 0.26 | 0.44 | 0.16*** |
| $\mathbb{1}$ (Thinks cookstove emissions are unsafe) | 0.50 | 0.50 | 0.50 | 0.50 | −0.0060 |
| *Health status* | | | | | |
| $\mathbb{1}$ (At least one case of cough/cold in past week) | 0.25 | 0.43 | 0.20 | 0.40 | 0.086 |
| Observations | 469 | | 474 | | 943 |

This table presents baseline (pre-intervention) summary statistics for 943 households that are part of the final study sample. "Traditional stove" includes traditional braziers (*angithi*), clay stoves (*mitti ka chulha*), coal/fuelwood heaters (*sagarh*), pan-shaped coal stoves and three-stone fires. "Improved stove" includes stoves fuelled by biogas, electricity, LPG and kerosene, and commercially available efficient biomass cookstoves. "Traditional fuel" includes crop residue, dung, fuelwood, leaves and household waste (trash). "Clean fuel" includes biogas, electricity (for cooking), kerosene and LPG. Variables for time spent using traditional stoves per day and for time spent collecting traditional fuels per day are winsorized at the 97.5 percentile level. Missing values in variables for knowledge of stoves or fuels that produce less smoke are replaced with zero. $p$ values associated with differences in means—reported using asterisks in column (5)—are obtained from a regression model of the form: $Y_{ij} = \beta_0 + \beta_1 \cdot \mathbb{1}\,(\text{NGO village}) + v_{ij}$, where $Y_{ij}$ represents a household-level characteristic for household $i$ in hamlet $j$; $\mathbb{1}$ (NGO village) represents an indicator that equals one if hamlet $j$ is located in an NGO village, and $v_{ij}$ represents a normally distributed error term. Standard errors are clustered at the hamlet level, and $p$ values are adjusted using the free step-down resampling methodology of Westfall and Young (1993), as operationalized by Jones et al. (2018). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

five members. The average household head is 54 years old and has had approximately six years of formal education. Only about one-quarter of surveyed households are headed by women, and more than half fall below the Indian poverty line.

As in many other parts of rural India, reliance on traditional stoves and fuels is practically universal.[18] In contrast, only about one-third of households own any type of improved stove; this is almost exclusively limited to liquefied petroleum gas (LPG) stoves, which are typically owned by relatively wealthy households. Awareness of modern alternatives to traditional cooking technologies is low: only about one-quarter to one-third of households profess an awareness of the existence of stoves or fuels that produce less smoke. This is not necessarily due to a lack of awareness about the harms associated with exposure to household air pollution. Indeed, half of surveyed households believe that the smoke generated by their primary stove is unsafe. Households also report spending up to two hours per day on average collecting traditional fuels for household use. In addition, around one in five households report that at least one family member suffered from a case of cough or cold in the past two weeks. Together, this suggests that the welfare burden imposed by widespread reliance on traditional energy sources is substantial.

Balance tests in Table 2.2 reveal that randomly selected households in NGO villages are broadly similar to their counterparts in non-NGO villages in terms of demographic and socioeconomic characteristics as well as in terms of stove ownership and use patterns. We note that NGO-village households are somewhat larger, report spending approximately thirty more minutes per day collecting traditional fuels, and report higher awareness of the existence of cleaner stoves and fuels. We control for these differences in all our analyses explicitly or via the inclusion of household fixed-effects.

[18] In our setting, "traditional stoves" include traditional braziers (*angithi*), clay stoves (*mitti ka chulha*), coal/firewood heaters (*sagarh*), pan-shaped coal stoves and three-stone fires. "Improved stoves" include stoves fuelled by biogas, electricity, LPG and kerosene, and commercially available efficient biomass cookstoves. Similarly, "traditional fuels" include crop residue, dung, fuelwood, leaves and household waste (trash), while "clean fuels" include biogas, electricity (for cooking), kerosene and LPG.

## 2.4 Empirical framework and identification strategy

Our empirical framework combines *ex ante* community-level matching (described in detail in Section 2.3) with a randomized intervention design and a quasi-experimental difference-in-difference-in-differences estimation approach for identification. In this section we (i) provide an overview of our experimental intervention; and (ii) outline our estimation and identification strategy.

### *2.4.1 Randomized ICS-promotion intervention design*

Figure 2.1 presents an overview of our intervention design and implementation timeline. The intervention was randomized at the level of the hamlet; roughly seventy percent of hamlets—and, by implication, seventy percent of households—stratified across NGO and non-NGO villages were randomly assigned to the ICS-promotion treatment group prior to the start of the intervention in August 2013. As part of ICS-promotion efforts, treated households were visited by trained enumerators who identified themselves as affiliated with our partner NGO. Treated households received a personalized demonstration of two distinct ICS technologies: the Greenway biomass cookstove and the G-Coil electric stove. At the end of the demonstration, survey teams presented these households with an offer to purchase one or both of the stoves. This offer consisted of a financial plan (the opportunity to make payments in three installments) combined with one of three randomized level of rebate (high, medium, or low—representing a reduction in the cost of each stove by about two, twenty and thirty percent, respectively).[19] These rebates were randomized at the household-level and delivered as a discount counting against the final installment payment if a household was found to be still using the stove by that time. Households in control hamlets did not receive the intervention; survey teams visited control households

[19] The market price of the electric stove was approximately INR 1,000 (USD 20 in 2012) while that of the biomass stove was INR 1,400 (USD 28). The highest rebate amount, therefore, was around USD 6–8, depending on stove type.

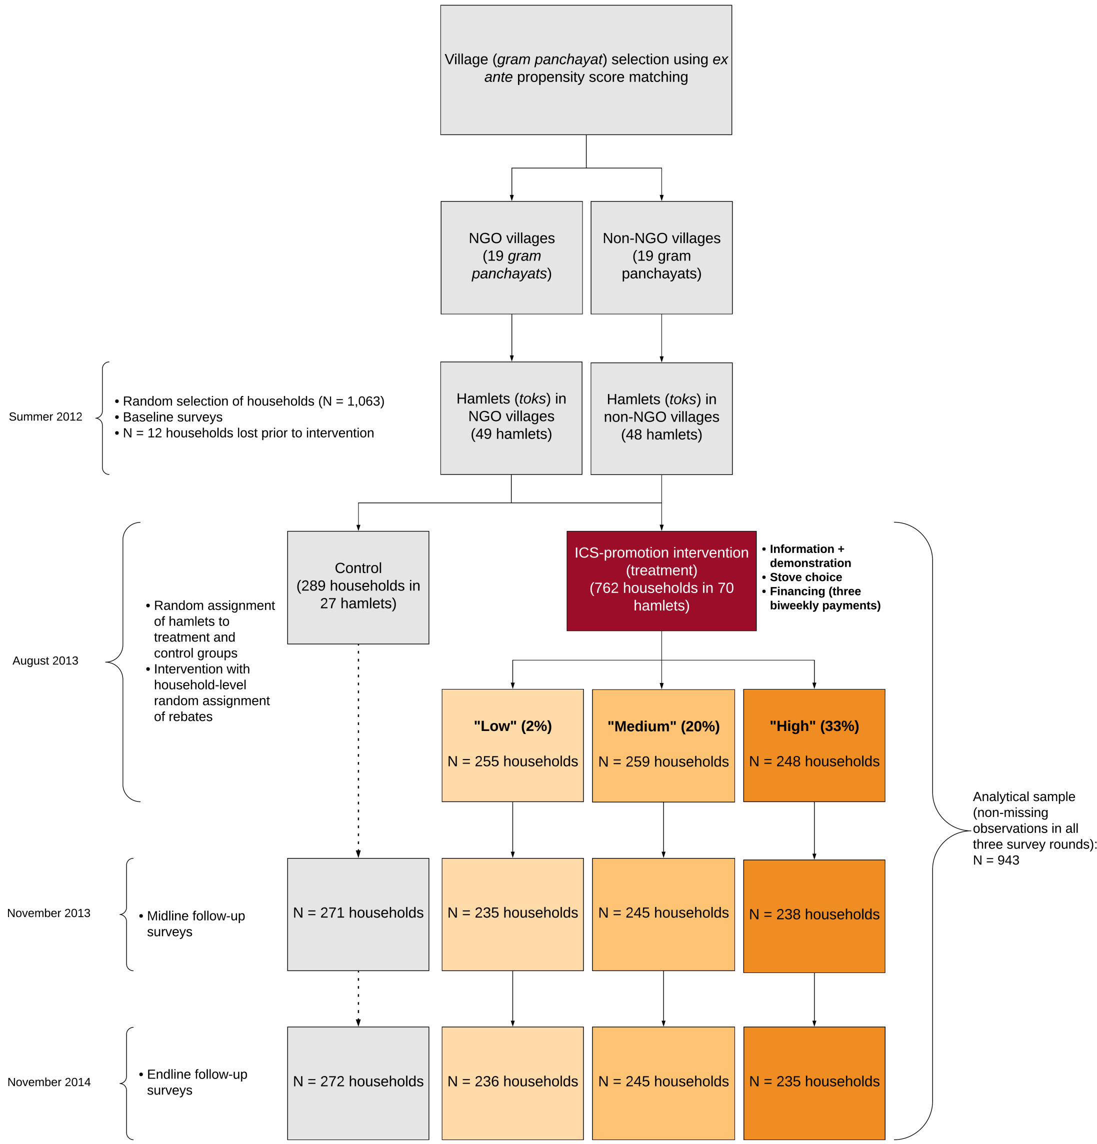


FIGURE 2.1: Study design and timeline. This figure presents an overview of our intervention design and implementation timeline. Random household-level rebates were provided as a percentage of the price of the stove; the market price of the electric stove was approximately INR 1,000, while that of the biomass stove was INR 1,400. USD 1 ≈ INR 50 in 2012.

at the same time as treated households to conduct follow-up surveys.

The stratified study design shown in Figure 2.1 enables us to compare the differential impact of the same randomized intervention delivered by the same field team professing to be affiliated with the same NGO across two institutionally distinct settings, namely communities with which the NGO had a preexisting relationship and those with which it did not. This allows us to isolate how NGOs—and the trust and social capital they foster in their local communities—influences the outcomes of interventions directly. Since our intervention is designed principally to increase uptake of cleaner cooking technologies, our main outcome of interest is the purchase rate of intervention ICS. To investigate heterogeneity in this rate across treated households located in NGO and non-NGO villages separately, we estimate the following specification:

$$Y_{ij} = \beta_0 + \beta_1 \left(TREATMENT_j\right) + \beta_2 \left(NGO_j\right) + \beta_3 \left(TREATMENT_j \times NGO_j\right) + \sum_n \beta_n X_{i,n} + \epsilon_{ij}, \tag{2.19}$$

where $Y_{ij}$ is a binary variable that equals one if household $i$ in hamlet $j$ purchased at least one of the two intervention ICS offered during intervention activities and zero if it did not; $TREATMENT_j$ is a binary variable that equals one if hamlet $j$ is randomly assigned to the treatment group and zero if it is assigned to the control group; $NGO_j$ is a binary variable that equals one if hamlet $j$ is located in an NGO village and zero if it is in a non-NGO village; $X_{i,n}$ represents a set of household-level controls; and $\epsilon_{ij}$ is a normally distributed error term. Our coefficient of interest is $\beta_3$, which sheds light on the additional impact of the randomized ICS-promotion intervention on purchase rates in the NGO stratum of villages.

### 2.4.2 *Difference-in-difference-in-differences specification*

Although the villages across the NGO and non-NGO strata are matched on sixteen different community-level characteristics (Table 2.1), one may still be concerned that unobservable community-level differences drive either the selection of NGOs into certain villages, the selection of households into villages in the NGO stratum, or both. Identification may be threatened, for instance, by an NGO-stratum-specific factor that affects households' responsiveness to the intervention. To address this concern, we leverage the multiple rounds of our survey in a difference-in-difference-in-differences ("triple-differences") specification. Specifically, we estimate the following model:

$$
\begin{aligned}
Y_{ijt} = {} & \beta_4 + \beta_5 (POST_1) + \beta_6 (POST_2) \\
& + \beta_7 \left(TREATMENT_j \times POST_1\right) + \beta_8 \left(TREATMENT_j \times POST_2\right) \\
& + \beta_9 \left(NGO_j \times POST_1\right) + \beta_{10} \left(NGO_j \times POST_2\right) \\
& + \beta_{11} \left(TREATMENT_j \times NGO_j \times POST_1\right) + \beta_{12} \left(TREATMENT_j \times NGO_j \times POST_2\right) \\
& + \gamma_i + \epsilon_{ijt},
\end{aligned}
\tag{2.20}
$$

where $Y_{ijt}$ represents ICS-related adoption, use, or impact for household $i$ in hamlet $j$ in survey round $t$. In Equation (2.20), $POST_1$ and $POST_2$ represent binary variables that are equal to one if data for the relevant observation were collected during the first and second follow-up survey rounds, respectively, and zero otherwise; these variables capture time trends over our multiple survey rounds. We also include household fixed-effects (represented by $\gamma_i$) to control for unobserved household-level differences.[20] Our coefficients of interest are now $\beta_{11}$ and $\beta_{12}$, which represent the additional impact of our intervention for treated households located in NGO villages relative to treated households located in non-NGO villages during the midline and endline survey rounds, respectively.

It is worth noting that the fully interacted triple-differences specification outlined in Equation (2.20) considerably relaxes our identifying assumptions. Identification would only be threatened by a confluence of factors—say, if hamlets in the treatment arm were located closer to urban areas; if NGO villages exhibited a greater degree of rural-to-urban

[20] Collinearity of the $TREATMENT_j$ and the $NGO_j$ binary variables—not included separately in Equation (2.20)—with the household-specific binary variables implies that $\gamma_i$ also captures any differences that may exist across households in treated and untreated hamlets, and in NGO and non-NGO villages.

migration (unobserved by us); *and* if the ICS-promotion intervention spanned a period that entailed a seasonal return of said (relatively cash-rich) urban migrants back to their homes, resulting in a time-varying shock specific to treated NGO-stratum hamlets that positively influenced households' purchase of ICS technologies. While certainly possible, we contend that this is unlikely in practice. Random allocation of hamlets to the intervention and control arms should preclude community-level characteristics in treated and untreated hamlets from differing significantly. In addition, our matching approach controls for a host of observed community-level differences between NGO- and non-NGO villages; the inclusion (and interaction) of survey-round time trends accounts for changes over time in hamlets in the intervention arm as a whole as well as in the NGO stratum specifically; and household fixed-effects soak up time-invariant unobservable differences.[21]

## 2.5 Results and discussion

We now turn to a discussion of the results of our empirical analyses. Our main results focus on heterogeneity in the effectiveness of the ICS-promotion intervention across matched NGO and non-NGO settings. Although our intervention is not designed to evaluate the impact of ICS in real-world settings, as part of secondary analyses we also investigate heterogeneity in impacts of ICS promotion on households' energy-use and time-allocation patterns across NGO and non-NGO communities.

### *2.5.1 Effectiveness of the intervention across NGO and non-NGO villages*

*Household-level purchase of intervention ICS technologies* Our primary outcome of interest is household-level purchase of ICS technologies promoted during the ICS-promotion

[21] Because we evaluate the impact of our ICS-promotion intervention using primary data collected during one baseline (pre-intervention) and two follow-up (post-intervention) survey rounds, we are unable to verify whether pre-intervention trends for our ICS-related adoption, use and impact outcomes of interest are parallel across NGO and non-NGO villages. Instead, we test for differences in pre-trends across NGO and non-NGO villages for a host of community-level characteristics that are likely to be correlated with our outcomes of interest using the 2001 and 2011 rounds of the Indian Census. We find no evidence of differences in pre-trends across NGO and non-NGO villages over this period (Table G.8).

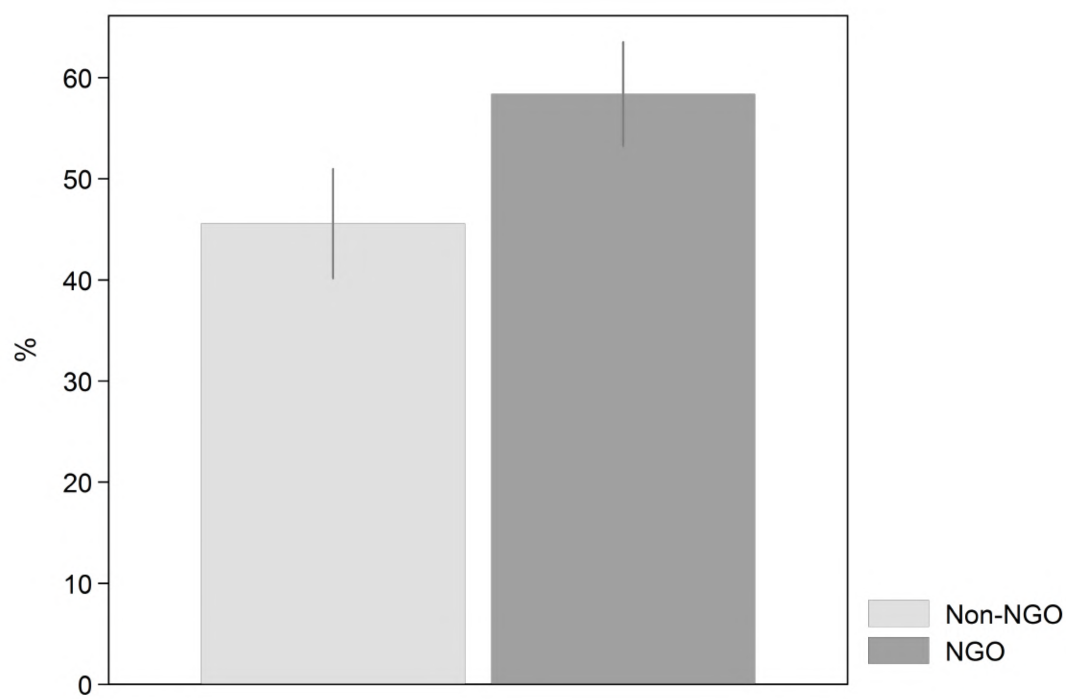


FIGURE 2.2: Mean intervention ICS purchase rates in treated hamlets in NGO and non-NGO villages. This figure plots the share of households in NGO and non-NGO villages that purchased at least one intervention ICS in response to the ICS-promotion intervention as a percentage of all treated households in the respective stratum. Error bars represent 95 percent confidence intervals for the means.

intervention. Figure 2.2 highlights the mean purchase rate for households located in treated hamlets in NGO and non-NGO villages. On average, nearly sixty percent of treated households in NGO villages purchased at least one of the two intervention ICS. In non-NGO villages, the corresponding figure is approximately 45 percent. We next evaluate this difference more rigorously in a linear regression framework. Table 2.3 presents our results. We find that the promotion campaign is extremely effective at encouraging the uptake of the intervention stoves; as shown in column (1), over half of targeted households purchase at least one of the two promoted ICS technologies. However, when we disaggregate our results by NGO and non-NGO villages following Equation (2.19) in column (2), we find that the purchase rate is nearly thirteen percentage points

Table 2.3: Effect of promotion on intervention ICS purchase in matched NGO/non-NGO villages

| | (1) | (2) | (3) |
|---|---|---|---|
| | $\mathbb{1}$ (Purchased intervention ICS) | | |
| $TREATMENT_j$ | 0.52*** | 0.46*** | 0.45*** |
| | (0.029) | (0.049) | (0.048) |
| $NGO_j$ | | 0.00*** | −0.017* |
| | | (0.00) | (0.010) |
| $TREATMENT_j \times NGO_j$ | | 0.13** | 0.13** |
| | | (0.058) | (0.060) |
| Constant | −0.00*** | −0.00*** | −0.079* |
| | (0.00) | (0.00) | (0.043) |
| Mean dep. (control) | 0.00 | 0.00 | 0.00 |
| Observations | 943 | 943 | 943 |
| Household-level controls | No | No | Yes |
| Adjusted $R^2$ | 0.23 | 0.24 | 0.24 |

The outcome variable is an indicator that equals 1 if household $i$ in hamlet $j$ purchased at least one of the two ICS promoted during the intervention. Column (1) presents aggregated results; results are disaggregated by NGO and non-NGO villages (as shown in Equation 2.19) in column (2). Baseline household-level controls for household size, number of children under five, awareness of existence of cleaner stoves and fuels, and total traditional-fuel collection time per day are included in column (3). 'Traditional fuel' includes crop residue, dung, fuelwood, leaves, and household waste (trash); missing observations for total time spent collecting traditional fuel for 32 households are replaced with the sample mean value. Standard errors (in parentheses) are clustered at the hamlet level. $^{*}$ $p < 0.10$, $^{**}$ $p < 0.05$, $^{***}$ $p < 0.01$.

(28 percent) higher in treated hamlets located in NGO villages—a statistically significant and positive "NGO effect."[22] These results are robust to the inclusion of household-level

[22] In Figure H.4 we present results from the application of an approach inspired by randomization-based inferential procedures (Athey and Imbens, 2017) to village-level NGO stratum allocation. Our approach relies on randomly assigning villages to placebo NGO and non-NGO strata, and re-estimating Equation (2.19); this process is repeated 1,000 times to obtain a distribution of placebo "NGO effect" estimates. If the effect we observe was due to the chance selection of the villages in our NGO and non-NGO strata, we would expect to observe our actual estimate located near the middle of this distribution. Instead, we find that only three percent of these placebo estimates are greater in magnitude than our actual estimates. In addition, in Appendix D, rather than characterizing NGO and non-NGO villages using a binary variable, we

controls that were found to be unbalanced at baseline in Table 2.2, as shown in column (3).

*Ownership and use of intervention ICS* This apparent "NGO effect" is not limited to the initial purchase decision. Figure 2.3 highlights that ownership (panel *a*) and reported use (panel *b*) of intervention ICS remains higher in NGO villages over multiple survey rounds. To rigorously evaluate differences in these trends, we separately estimate the triple-differences specification outlined in Equation (2.20) for ownership and use of intervention ICS during the midline (three-month) and endline (fifteen-month) follow-up surveys. Our results are shown in Table 2.4, which also includes results from estimating a conventional difference-in-differences specification (without NGO-specific interactions) for comparison. While approximately forty percent of treated households report owning an intervention ICS and thirty percent report having used it recently (columns 1 and 3, respectively), reported ownership and use by treated households in NGO villages during the first follow-up (in columns 2 and 4, respectively) are sixteen percentage points higher. This represents an increase in the size of the treatment effect of between 50–80 percent relative to ownership and use by households in treated non-NGO hamlets. By the endline follow-up survey (conducted approximately fifteen months after the start of the intervention), the difference in intervention ICS ownership rates across treated NGO and non-NGO hamlets remains positive, but is no longer statistically significant.[23] Similarly, although the difference in reported intervention ICS use rates between the two treated NGO and non-NGO groups remains positive by the endline survey, it is no longer statistically significant.

characterize the village-specific "intensity" of NGO activity based on two different measures: (1) the number of projects/initiatives the NGO has implemented in a particular village; and (2) the number of years it has been active in a particular village. We use these two measures to separately re-estimate Equation (2.19) for heterogeneity in rates of ICS purchase and find that these analyses provide evidence that is consistent with our main results.

[23] That said, we are unable to reject that the difference between the two coefficients—for ownership of intervention ICS by households in treated NGO hamlets at midline and at endline—is statistically zero.

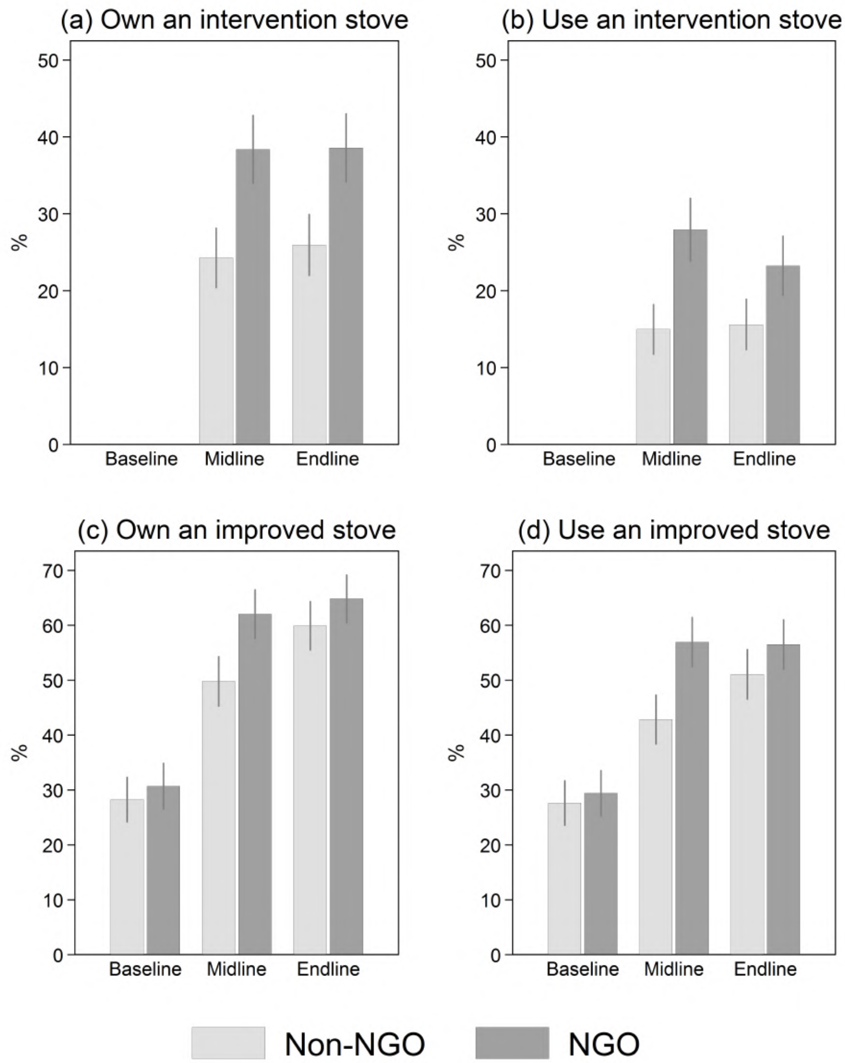


FIGURE 2.3: ICS ownership and use in NGO and non-NGO villages. Panels (a) and (c) plot the share of households that own at least one of the two intervention ICS or an improved ICS, respectively, as a percentage of all (treated and untreated) sample households in the indicated stratum of villages. Panels (b) and (d) plot the subset of these households that report having used these devices in the week prior to the survey as a percentage of all (treated and untreated) sample households in the indicated stratum of villages. "Improved stove" includes stoves fuelled by biogas, electricity, LPG, kerosene, and commercially available efficient biomass cookstoves; we also include the two ICS promoted as part of the promotion intervention in this definition. Error bars represent 95 percent confidence intervals for the means. Baseline survey activities occurred approximately one year before the intervention; midline and endline surveys occurred approximately three and fifteen months, respectively, after the intervention (see Figure 2.1).

Table 2.4: Effect of promotion on intervention ICS adoption in matched NGO/non-NGO villages

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | $\mathbb{1}$ (Owns intervention ICS) | | $\mathbb{1}$ (Uses intervention ICS) | |
| $POST_1$ | 0.012* | 0.014 | 0.0039 | 0.0068 |
| | (0.0064) | (0.0091) | (0.0038) | (0.0066) |
| $POST_2$ | 0.039** | 0.027* | 0.027** | 0.020* |
| | (0.016) | (0.016) | (0.011) | (0.011) |
| $TREATMENT_j \times POST_1$ | 0.41*** | 0.33*** | 0.29*** | 0.21*** |
| | (0.030) | (0.047) | (0.026) | (0.034) |
| $TREATMENT_j \times POST_2$ | 0.39*** | 0.34*** | 0.23*** | 0.20*** |
| | (0.034) | (0.053) | (0.027) | (0.042) |
| $NGO_j \times POST_1$ | | −0.0046 | | −0.0068 |
| | | (0.013) | | (0.0066) |
| $NGO_j \times POST_2$ | | 0.027 | | 0.016 |
| | | (0.034) | | (0.025) |
| $TREATMENT_j \times NGO_j \times POST_1$ | | 0.16*** | | 0.16*** |
| | | (0.058) | | (0.048) |
| $TREATMENT_j \times NGO_j \times POST_2$ | | 0.098 | | 0.061 |
| | | (0.070) | | (0.056) |
| Mean dep. (baseline non-NGO control) | 0.00 | 0.00 | 0.00 | 0.00 |
| Observations | 2,829 | 2,829 | 2,829 | 2,829 |
| Adjusted $R^2$ | 0.56 | 0.56 | 0.34 | 0.35 |
| Household fixed-effects | Yes | Yes | Yes | Yes |

The outcomes variables are ownership (columns 1–2) and use (columns 3–4) of at least one of the two intervention ICS as measured during baseline, and midline and endline follow-up surveys (conducted approximately three and fifteen months after the start of the intervention, respectively). The outcome variable for ownership is an indicator that equals one if household $i$ in hamlet $j$ reports owning at least one of the two ICS promoted during the intervention in survey round $t$; for use, it is an indicator that equals one if the household—conditional on ownership—reports having used at least one of these two ICS in the past week. The even-numbered columns present the results of estimating the triple-differences specification outlined in Equation (2.20). The accompanying odd-numbered columns present results from estimating a double-differences specification without NGO-specific interactions for comparison. Standard errors (in parentheses) are clustered at the hamlet level. $^{*}$ $p < 0.10$, $^{**}$ $p < 0.05$, $^{***}$ $p < 0.01$.

*Ownership and use of improved stoves* We next separately estimate Equation (2.20) for reported ownership and use of all improved stoves—and not only the two ICS promoted during the intervention. Here, we aim to evaluate the effectiveness of our intervention more broadly.[24] Indeed, although we limited promotion activities to two specific ICS technologies, the primary goal of such an intervention—were it to be implemented as part of national or regional policies—would arguably be to increase reliance on cleaner, more efficient cooking technologies more generally. Understanding how institutional actors influence intervention effectiveness at this higher level is thus crucial.

As shown in Table 2.5, we find that the differential adoption and use patterns associated with uptake of the intervention ICS in treated NGO hamlets (highlighted in Tables 2.3 and 2.4) do not appear to be related to adoption and use of improved stoves more broadly. Specifically, while the intervention is effective at increasing ownership and use of improved stoves in treated hamlets (columns 1 and 3, respectively), we detect no differential effect of the intervention between treated hamlets in NGO and non-NGO villages (columns 2 and 4). In fact, the negative coefficient for both of our triple-differences estimates hints at an underlying "technology substitution" effect—households that choose to purchase intervention ICS in treated NGO hamlets are those that might have adopted an improved stove in the absence of the intervention anyway.[25]

At the same time, our results also highlight other differential dynamics. Recall that the first and second follow-up surveys occurred approximately three and fifteen months after the start of ICS-promotion activities, respectively. Given the concerted efforts of the Indian government in recent years to enhance access to cleaner cooking fuels and technologies—

[24] Recall that in addition to the two intervention ICS, "improved stoves" include kerosene burners, LPG and biogas stoves, as well as other electric or efficient biomass alternatives. In our study area, LPG stoves were the main improved alternative, owned by approximately one-third of households at baseline.

[25] Ownership and reported use of improved stoves at baseline (pre-intervention) are balanced across NGO and non-NGO hamlets (Figure 2.3, panels *c* and *d*), and baseline ownership of improved stoves does not predict subsequent purchase of an intervention ICS (Table G.9), suggesting that intervention ICS are typically being purchased by households that did not already have an improved stove.

Table 2.5: Effect of promotion on improved-stove adoption in matched NGO/non-NGO villages

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | $\mathbb{1}$ (Owns improved stove) | | $\mathbb{1}$ (Uses improved stove) | |
| $POST_1$ | 0.016 | −0.048 | 0.0039 | −0.054 |
| | (0.036) | (0.042) | (0.034) | (0.041) |
| $POST_2$ | 0.14*** | 0.12*** | 0.12*** | 0.095* |
| | (0.039) | (0.044) | (0.040) | (0.048) |
| $TREATMENT_j \times POST_1$ | 0.34*** | 0.38*** | 0.29*** | 0.30*** |
| | (0.047) | (0.066) | (0.042) | (0.056) |
| $TREATMENT_j \times POST_2$ | 0.26*** | 0.28*** | 0.18*** | 0.20*** |
| | (0.048) | (0.065) | (0.047) | (0.061) |
| $NGO_j \times POST_1$ | | 0.15** | | 0.14** |
| | | (0.068) | | (0.063) |
| $NGO_j \times POST_2$ | | 0.040 | | 0.067 |
| | | (0.083) | | (0.085) |
| $TREATMENT_j \times NGO_j \times POST_1$ | | −0.10 | | −0.045 |
| | | (0.092) | | (0.080) |
| $TREATMENT_j \times NGO_j \times POST_2$ | | −0.047 | | −0.059 |
| | | (0.10) | | (0.098) |
| Mean dep. (baseline non-NGO control) | 0.36 | 0.36 | 0.35 | 0.35 |
| Observations | 2,829 | 2,829 | 2,829 | 2,829 |
| Adjusted $R^2$ | 0.56 | 0.57 | 0.52 | 0.52 |
| Household fixed-effects | Yes | Yes | Yes | Yes |

The outcome variables are ownership (columns 1–2) and use (columns 3–4) of all improved stoves as measured during baseline, and midline and endline follow-up surveys (conducted approximately three and fifteen months after the start of the intervention, respectively). The outcome variable for ownership is an indicator that equals one if household $i$ in hamlet $j$ reports owning at least one improved stove in survey round $t$; for use, it is an indicator that equals one if the household—conditional on ownership—reports having used such a device in the past week. The even-numbered columns present the results of estimating the triple-differences specification outlined in Equation (2.20). The accompanying odd-numbered columns present results from estimating a double-differences specification without NGO-specific interactions for comparison. "Improved stove" includes stoves fuelled by biogas, electricity, LPG, kerosene, and commercially available efficient biomass cookstoves; we also include the two ICS promoted as part of the promotion intervention in this definition. Standard errors (in parentheses) are clustered at the hamlet level. $^{*}$ $p < 0.10$, $^{**}$ $p < 0.05$, $^{***}$ $p < 0.01$.

and LPG, in particular (Barnwal, 2017; Kumar et al., 2016)—we would expect ownership of improved stoves to naturally increase over this period. In Equation (2.20), we control for these time trends by including dummy variables for both follow-up survey rounds. As part of our triple-differences specification, we also interact these variables with the NGO-village dummy, which allows us to separately control for NGO-village-specific time trends. With these controls, we find that only three months after the intervention, ownership and use rates of improved stoves in control hamlets in NGO villages (represented by the coefficient for $NGO_j \times POST_1$) are nearly fifteen percentage points higher.

To shed light on these two phenomena, we separately analyze trends in ownership of the three most important types of improved stoves in our setting (LPG, electric and efficient biomass stoves) across treatment and control, and NGO and non-NGO communities.[26] Figure H.5 suggests that changes in ownership of improved stoves in both NGO and non-NGO control hamlets are almost entirely driven by changes in ownership of LPG stoves. This trend is confirmed when we estimate Equation (2.20) for only LPG stoves, indicated by the near equality in magnitude of the coefficients on the $NGO_j \times POST_1$ interaction term in columns (1) and (2) of Table G.10. In addition, we find that our intervention indeed appears to have *lowered* ownership of LPG stoves in treated NGO hamlets initially relative to treated non-NGO hamlets—as indicated by the negative coefficient on the $TREATMENT_j \times NGO_j \times POST_1$ triple interaction term—in line with our "technology substitution" hypothesis.

The mechanisms that lead to households in control NGO communities adopting LPG stoves at these elevated rates are unclear. It may be the case that opportunities for households living in different hamlets to interact are higher in NGO villages. This could lead to information spillovers across treated and control hamlets within NGO villages that induce households in the latter to adopt relatively readily available LPG stoves. It

[26] Recall that one of our two intervention ICS is an electric stove while the other is an efficient biomass stove.

could also be the case that our NGO partner—having gained experience implementing improved-stove interventions—began to enable households in its control communities to avail of the Indian government's LPG-promotion schemes. Indeed, it is only one year later at the second follow-up that we see non-NGO control communities (represented by the identical coefficient on $POST_2$ in columns 1 and 2 of Table G.10) beginning to catch up in terms of ownership of LPG stoves, consistent with the expected spread of improved LPG technologies over time. While we are unable to speak definitively about these underlying dynamics, our stratified study design does allow us to see how our intervention may have crowded out improved energy-technology acquisition in treated hamlets with prior interaction with the NGO. Because the socioeconomic and environmental benefits associated with an intervention like ours ultimately depend on community-level uptake of improved stoves broadly, this NGO-specific crowding out also has implications for intervention effectiveness.

### *2.5.2 Heterogeneity in impacts across NGO and non-NGO villages*

Finally, we turn to an evaluation of the socioeconomic and environmental impacts of the ICS promotion with an eye to investigating heterogeneity in impacts across NGO and non-NGO villages.[27] Specifically, we separately estimate Equation (2.20) for an objective measure of daily fuelwood use and for reported time spent collecting fuels for household use per day.

Columns (1) and (2) of Table 2.6 present our results for an objective measure of fuelwood use, obtained from a subsample of households that were randomly selected for weight-based measurements of their solid-fuel use over a 24-hour period.[28] As shown

[27] We note that our intervention was designed to evaluate the effectiveness of tools to promote the sale, adoption and use of ICS, and was not intended to evaluate ICS impacts in real-world settings. That said, our stratified design does allow us to investigate impact heterogeneity. Bensch and Peters (2015), Bensch et al. (2015), Beyene et al. (2015), Brooks et al. (2016), Hanna et al. (2016), Lewis et al. (2016), Meeks et al. (2016) and Somanathan and Bluffstone (2015) are some recent examples of evaluations of the economic, environmental and health impacts of various non-traditional cooking technologies.

[28] Households were instructed to set aside an amount of fuelwood they expected to use over the next

Table 2.6: Impact of promotion on fuel collection and use in matched NGO/non-NGO villages

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | Fuelwood use (kilograms per day) | | Fuel-collection time (minutes per day) | |
| $POST_1$ | 4.42*** | 3.38*** | 3.95 | −19.0 |
| | (0.68) | (0.57) | (18.4) | (18.5) |
| $POST_2$ | 1.50** | 1.44 | −9.83 | −30.7** |
| | (0.75) | (1.18) | (15.5) | (12.5) |
| $TREATMENT_j \times POST_1$ | −2.01** | −0.61 | −26.0 | 19.2 |
| | (0.93) | (0.91) | (21.9) | (25.4) |
| $TREATMENT_j \times POST_2$ | 0.65 | 0.47 | −14.4 | 12.8 |
| | (1.08) | (1.53) | (17.0) | (15.9) |
| $NGO_j \times POST_1$ | | 2.24** | | 53.3 |
| | | (1.12) | | (33.2) |
| $NGO_j \times POST_2$ | | 0.12 | | 48.5* |
| | | (1.46) | | (29.0) |
| $TREATMENT_j \times NGO_j \times POST_1$ | | −2.93* | | −95.9** |
| | | (1.67) | | (40.4) |
| $TREATMENT_j \times NGO_j \times POST_2$ | | 0.34 | | −60.8* |
| | | (2.12) | | (32.1) |
| Mean dep. (baseline non-NGO control) | 8.90 | 8.90 | 113.6 | 113.6 |
| Observations | 1,143 | 1,143 | 2,829 | 2,829 |
| Adjusted $R^2$ | 0.19 | 0.20 | 0.011 | 0.029 |
| Household fixed-effects | Yes | Yes | Yes | Yes |

The outcome variables are daily fuelwood use (columns 1–2) and total time spent collecting fuels (columns 3–4) as measured during baseline, and midline and endline follow-up surveys (conducted approximately three and fifteen months after the start of the intervention, respectively). Daily fuelwood use is derived from a 24-hour before–after household-level fuel-weighing test; we restrict the analysis in columns (1) and (2) to the subsample of households that participated in such tests in at least two of our three survey rounds ($N = 388$). The outcome variable for fuel-collection time is derived from self-reported data on time spent (per day, week or month) collecting fuelwood, crop residue, leaves, dung, biomass pellets, kerosene, LPG, biogas and—if relevant—any other fuel used by the household; missing observations for time spent collecting fuel for up to six households in each survey round are replaced with the survey-round-specific sample mean value. The even-numbered columns present the results of estimating the triple-differences specification outlined in Equation (2.20). The accompanying odd-numbered columns present results from estimating a double-differences specification without NGO-specific interactions for comparison. Standard errors (in parentheses) are clustered at the hamlet level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

in column (1), the intervention appears to reduce fuelwood use by approximately two kilograms per day in the relatively short run (as indicated by the negative estimated coefficients for $TREATMENT_j \times POST_1$). However, our triple-differences estimate for the midline in column (2) reveals that this effect is almost entirely driven by reductions in fuelwood use by households in treated NGO hamlets. These household appear to use nearly three kilograms less fuelwood per day—evidence of the "NGO effect" that is consistent with reported use of purchased intervention stoves. By the time of the endline survey one year later (that is, the coefficient on $TREATMENT_j \times NGO_j \times POST_2$), this effect appears to attenuate somewhat, and we no longer detect a significant difference between fuelwood use in NGO and non-NGO communities. That said, it is worth noting that we are unable to reject that the difference between the triple-differences estimates for the midline and endline survey rounds is zero.

Consistent with these fuelwood use patterns, in column (3) we find that while there is no detectable effect of the intervention on reported daily fuel-collection times for treatment hamlets in general, households in treated hamlets in NGO villages reported significant reductions in time spent collecting fuels for household use at the time of the midline and endline follow-up surveys (column 4). This effect appears to be driven primarily by reductions in time spent collecting firewood and other traditional fuels (Table G.11). Once again, the two triple-differences estimates are statistically indistinguishable from each other.

## 2.6 Bayesian synthesis of the evidence

How do our findings contribute to the limited evidence on the roles NGOs play in shaping outcomes of interventions? To answer this question, our final set of analyses combines the insights from the implementer identity literature—which finds that NGO-led interventions

24-hour period. This amount was weighed by the field team, which returned the next day to reweigh the remaining amount. We underscore that the relatively involved nature of the fuel-weight test limited the available sample size and, consequently, our ability to detect meaningful impacts for this outcome variable.

are often more effective than comparable efforts by other actors—with the results of our matched-experimental study design in a simple Bayesian regression framework.

Suppose $\beta_3$ is a random variable that represents the true effect of the NGO on the outcomes of the ICS promotion intervention in our setting—the same term we use to indicate the coefficient on the $TREATMENT_j \times NGO_j$ interaction term in Equation (2.19). Our first step is to suitably characterize the prior evidence base for this parameter. For this, we focus on two existing randomized evaluations. First, we turn to Bold et al. (2013), who evaluate a nationwide education reform in Kenya that expanded funding for hiring "contract teachers." Such teachers are hired directly by schools—typically at wages that are below those offered to tenured public school teachers—to address teacher shortages; they are also not accorded the same tenure protections available to their civil-service colleagues. There are a number of pathways through which contract teachers might improve educational outcomes.[29] Yet, as in many other cases, the evidence on the efficacy of such programs is from relatively small-scale interventions. As part of their large-scale evaluation, Bold et al. (2013) randomly assign a nationwide sample of public schools to a control group and one of two treatment groups that differ only in the identity of the implementer—in one, the implementation is led by an NGO while in the other it is led by the government. They find that "an additional contract teacher in a school where the program is managed by the NGO increased test scores by roughly 0.18 standard deviations." In contrast, the treatment effect is both smaller and statistically insignificant in the government-implementation schools. Digging deeper into mechanisms, they also find that contract teachers were actually hired and in place for at least twelve percent more months in NGO-implementation schools over the course of the intervention.

We turn next to Cameron and Shah (2017), who evaluate the scale-up of a community-

[29] Specifically, Bold et al. (2013) highlight three pathways: (1) contract teachers are typically hired from a waiting list of candidates for civil-service teaching appointments, and are thus similarly skilled and less expensive to hire; (2) a "selection effect," whereby only relatively good teachers are retained over time as poor-quality teachers do not have their contracts extended; and (3) an "incentive effect," whereby the lack of permanent contracts introduces dynamic incentives to increase teaching effort.

led total sanitation (CLTS) intervention in the Indonesian province of East Java. Nearly one billion people engage in open defecation, partly due to a lack of access to improved sanitation facilities. CLTS relies primarily on social pressures—rather than subsidies or grants—to encourage adoption of latrines and induce sustained behavior-change.[30] In addition to evaluating the impact of CLTS relative to a control group, Cameron and Shah (2017) investigate heterogeneity in impacts across treatment communities in which implementation was carried out by different actors. Specifically, treated communities were nearly evenly split; implementation was led by government staff in one half and by local non-governmental "resource agencies" in the other.[31] They find that households are over five percentage points more likely to build a latrine in communities in which NGOs triggered the CLTS intervention—a 42 per cent increase in the rate of latrine ownership relative to the mean in the control communities. Once again, the treatment effect is both smaller and statistically insignificant in the government-implementation communities.

While not perfectly analogous to our setting, *a priori*, the results in these two studies offer insights about the possible range of $\beta_3$—the "NGO effect." A 42 percent increase in the size of the treatment effect—as found by Cameron and Shah (2017)—would translate in our setting into an increase in the ICS purchase rate by households in treated NGO hamlets of nearly twenty percentage points relative to the purchase rate across treated non-NGO hamlets. In contrast, the more conservative twelve percent estimate from Bold et al. (2013) would imply a difference in the purchase rate between treated NGO and non-NGO hamlets of closer to five percentage points. The midpoint of these two estimates is 12.5 percentage points. This prior information suggests that we use a prior distribution $f(\beta_3)$ that assigns most of its probability to the interval $(0.05, 0.20)$, and that the expected

[30] Dickinson et al. (2015), for instance, describe a village-level CLTS intervention in India that entailed a "walk of shame," defecation mapping and fecal weighing—all designed to invoke an emotional response about the ubiquity of defecation sites in and around rural communities.

[31] While not technically randomized across implementers, the authors contend that no systematic process guided the selection of implementing teams. They also note that baseline village- and household-level characteristics are balanced across implementer arms.

value of $\beta_3$ under $f(\beta_3)$ be close to 0.125. We, therefore, represent our prior information about $\beta_3$ as follows:

$$\beta_3 \sim \text{Beta}(2, 14)\,. \tag{2.21}$$

The density of this prior distribution—a beta distribution with shape parameters $\alpha = 2$ and $\beta = 14$—is represented by the dashed line in panel (a) of Figure 2.4. The expected value of $\beta_3$ under this prior is 0.125, and the most probable value is approximately 0.07, corresponding to the peak in the density function. Just under two-thirds of the area under the curve lies between 0.05 and 0.20. Importantly, this distribution only has positive support over the interval $[0, 1]$. Together, these characteristics implicitly capture the prior information that the "NGO effect" is strictly positive but not excessively high.[32]

With this prior specified, we fit a Bayesian multilevel mixed-effects model broadly corresponding to the specification outlined in Equation (2.19) to investigate heterogeneity in purchase rates across treated NGO and non-NGO hamlets (described in detail in Appendix E).[33] Our results are shown in panel (a) of Figure 2.4, where the solid line approximates the posterior distribution we obtain for $\beta_3$. This distribution has less density in its tails and is more peaked, reflecting our updated beliefs given our specified prior and our data. The estimated posterior mean for $\beta_3$ is 0.124. The ninety percent "credible interval" (the range that has ninety percent posterior probability to contain the true effect) is between 0.034 and 0.24. This represents a 7–50 percent increase approximately in the size of the

[32] Admittedly, the choice of the specific beta prior outlined in Equation (2.21) is somewhat subjective as any number of alternative distributions would satisfy our outlined mean- and interval-related criteria. We note that the relative dearth of related evidence that we might draw upon to examine the distribution of estimates in more detail contributes to this subjectivity. Given this constraint, we believe our specified prior distribution does a reasonable job of characterizing the existing evidence. In addition, we also test the robustness of our results to prior specification using both "weak" and "strong" prior distributional assumptions.

[33] Multilevel models account for hierarchical structures within the data. In our case, households are located within hamlets that are part of villages, which themselves are in distinct NGO and non-NGO strata. These level-specific effects also vary randomly across levels based on specified prior distributions. To facilitate comparability with our main analyses—where we cluster our estimated standard errors at the level of the hamlet—we restrict the nested structure to the level of the hamlet. In addition, we specify diffuse ("uninformative") priors for all random model parameters besides our parameter of interest.

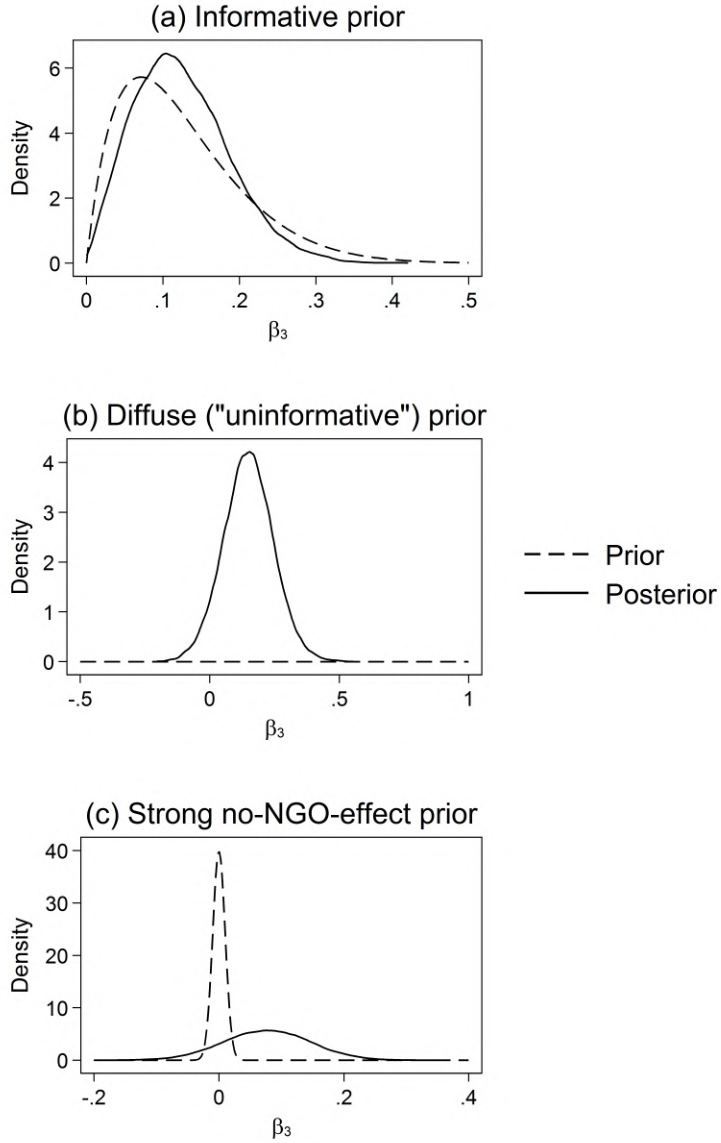


FIGURE 2.4: Bayesian analysis of heterogeneity in ICS purchase across NGO and non-NGO hamlets. This figure plots prior and posterior distributions of the $\beta_3$ parameter, which represents the additional effect of the ICS promotion intervention on purchase rates in treated NGO hamlets relative to treated non-NGO hamlets. Posterior distributions are estimated via Markov Chain Monte Carlo (MCMC) simulation. Specifically, we ran 50,000 MCMC samples after a burn-in period of 10,000 iterations, with thinning every fifth iteration. In panel (a), $f(\beta_3) = \text{Beta}(2, 14)$; in panel (b), $f(\beta_3) = \mathcal{N}(0, 10000)$; and in panel (c), $f(\beta_3) = \mathcal{N}(0, 0.01)$. Diffuse priors are used for all other parameters in the underlying model for each panel (see Appendix E).

treatment effect, respectively.

To test the robustness of this result to prior specification, we repeat our analysis two additional times with alternatively specified prior distributions. First, we assume that:

$$\beta_3 \sim \mathcal{N}\left(\hat{\beta}_3^{\mathrm{OLS}}, 10000\right). \tag{2.22}$$

This normal distribution (centerd at our estimated coefficient for $\beta_3$ from Table 2.3) is highly diffuse ("uninformative") and has positive support over the real line; these characteristics reflect the *a priori* beliefs of someone who only weakly suspects that the "NGO effect" may be positive. Panel (b) of Figure 2.4 presents our results, in which the solid line represents the posterior distribution we obtain. The resulting posterior mean for $\beta_3$ is 0.15. The ninety percent credible interval is $(-0.01, 0.31)$. In addition, the posterior probability that the "NGO effect" is positive is 0.94.

Next, we assume that:

$$\beta_3 \sim \mathcal{N}(0, 0.01). \tag{2.23}$$

This distribution, in contrast, reflects relatively strong *a priori* beliefs that NGOs do *not* influence the outcomes of the interventions they implement.[34] Panel (c) of Figure 2.4 presents our results. The resulting posterior mean for $\beta_3$ of 0.08 is expectedly diminished but remains positive. Indeed, the posterior probability that the "NGO effect" is positive is 0.87—relatively unaffected by the considerably stronger prior centerd on zero (compared to the highly diffuse prior in panel *b*). The ninety percent credible interval is $(-0.04, 0.19)$.

Our analyses, thus, demonstrate that the evidence for a large and positive effect of the NGO on the effectiveness of the ICS-promotion intervention is relatively robust—even under strong prior distributional assumptions about the lack of such an effect. We also show how evidence from related research can be used to inform these distributional assumptions and guide causal inference.

[34] Arguably, this belief is implicit in the act of conducting applied research in partnership with NGOs without having in place a study design similar to ours, which explicitly attempts to identify NGO-related heterogeneity.

## 2.7 Conclusion

Using data from an experimental intervention covering nearly 1,000 households across 97 geographically distinct hamlets in rural Uttarakhand, India, we highlight how NGOs influence the outcomes of applied interventions. We first develop a model of household decision-making grounded in transaction costs. We posit that NGOs lower transaction costs, and thus enhance the effectiveness of the interventions they implement. To empirically test our model's prediction, we use *ex ante* propensity score matching to create a sample of observationally similar rural communities that are differentiated by prior exposure to a local development NGO. In partnership with this NGO, we then stratify an experimental intervention designed to promote ICS on this institutional variable to identify heterogeneity in adoption, use and impacts.

We uncover a large, positive and statistically significant "NGO effect"—prior exposure to the NGO increases the effectiveness of the intervention by nearly thirty percent. Specifically, in line with our model's predictions, ICS purchase rates for households in treated hamlets located in "NGO villages" are thirteen percentage points (28 percent) higher than for households in treated hamlets located in matched "non-NGO villages." Using a difference-in-difference-in-differences ("triple-differences") specification, we find an even larger NGO effect on ICS use in these communities: households in such communities are up to sixteen percentage points more likely to have used an ICS, representing a fifty percent increase in the size of the treatment effect. Consistent with these patterns of ownership and use, households in villages with prior exposure to the NGO also exhibit reductions in daily fuelwood use and fuel-collection time; in contrast, we find no evidence of any impact of the intervention on energy-use patterns for treated households in non-NGO villages.

Although previous work has noted the presence of differential impacts across communities with and without NGO activity, to the best of our knowledge, we are the first to rigorously examine how the presence of an effective local NGO—and the specific institutional context that represents—directly influences household decision-making and ultimately determines the effectiveness of interventions. As such, our study begins to address a knowledge gap that has significant implications for the policy relevance of experimental research conducted in partnership with NGOs and other civil society organizations. Effective local organizations may be crucial for the implementation of environmental, health and development interventions in remote, rural settings. Subsequent attempts to scale-up findings deemed effective in applied research conducted in partnership with such institutions into national or regional policies may prove much less successful than anticipated if their roles and contributions are insufficiently accounted for. Alternatively, promoters of scaled-up interventions could achieve greater success if they enlist the assistance of trusted local partners.

# 3

# Preference heterogeneity and willingness to pay for energy technologies

*With Marc Jeuland and Ousmane Ndiaye*

## 3.1 Introduction

Three billion people use solid fuels and traditional stoves to meet their daily cooking and heating needs. The resulting household air pollution causes over four million deaths annually, a health burden borne disproportionately by women (Adair-Rohani et al., 2016). Daily collection of solid fuels (such as firewood) restricts opportunities for education and employment, and exacerbates pressures on local forests; their inefficient combustion generates emissions that intensify global warming. Improved cookstoves (ICS) can help ease this environment–health–development burden, yet uptake in rural areas—where they are most needed—remains low (Lewis and Pattanayak, 2012). Some scholars have demonstrated that policies that reduce monetary costs are effective at increasing demand for ICS (Bensch et al., 2015; Levine et al., 2018; Pattanayak et al., 2016). Others point to the importance of non-monetary channels, such as social networks (Miller and Mobarak, 2015). To date, however, little work has considered beneficiaries' heterogeneous preferences, namely, whether efforts to promote "improved" devices often fail because they prioritize technical factors over end-users' unique energy needs and cultural contexts.

In this paper, we first develop a simple conceptual framework to highlight the relationship between heterogeneous household preferences and willingness to pay (WTP) for ICS. We then present results from a field-based, empirical test of the main hypotheses that emerge from this framework. Specifically, we conduct a technology-promotion campaign followed by second-price, sealed-bid auctions featuring two biomass ICS with over 1,000 randomly selected participants across seventy rural communities in Senegal. The first device is a simple, locally made stove called the *Jambar*, while the second is a considerably more efficient device manufactured internationally called the *Jumbo Zama*. By randomly assigning communities to one of three auction arms—two where each stove is promoted and auctioned exclusively, and one where both stoves feature together—we induce exogenous variation in the number of alternatives presented to sample households and, by extension, in the degree to which stove promotion caters to households' heterogeneous preferences. We find that random allocation to the joint *Jambar–Jumbo Zama* auction lowers WTP for the *Jumbo Zama* relative to a single-stove scenario while having no distinguishable impact on WTP for the *Jambar*. This main result is robust to the use of different estimation strategies and the inclusion of a host of household- and village-level controls.

These results are consistent with a model in which preferences over attributes are constructed—and not simply revealed—as agents make repeated choices in unfamiliar domains. Researchers and policymakers alike routinely stress the importance of catering to households' preferences and energy-use needs as a way to enhance the effectiveness of ICS promotion (e.g., Lambe and Atteridge, 2012; Rhodes et al., 2014). "One size fits all" ICS promotion (where only one type of device is presented to beneficiaries across different settings) is criticized for failing to account for contextual factors that determine rural energy-use patterns and are thus reflected in stove design. These include demographic characteristics (e.g., household size), environmental factors (e.g., weather, availability of biomass fuels) and cultural idiosyncrasies (e.g., diets, cuisine). These critiques may well be legitimate but, to the best of our no knowledge, no study has directly tested how the inclusion of additional alternatives that cater to households' heterogeneous preferences influences demand for ICS. We demonstrate that the presence of additional choices can, in fact, dampen demand for cleaner alternatives. Insofar as robust demand is necessary to foster self-sustaining rural markets for energy technologies, our results have implications for the design of policies that aim to introduce cleaner cooking solutions in remote areas. Rather than simply providing additional choices, implementers looking to enhance uptake of improved technologies must devise approaches to help potential end-users think carefully through trade-offs, crystallize and understand their own preferences, and identify solutions that fit their needs.

This paper proceeds as follows: in Section 3.2, we present a conceptual framework that connects household preferences to WTP for ICS; Section 3.3 provides contextual background and an overview of our study design; Section 3.4 describes our analytical sample and empirical specifications, and presents results; Section 3.5 discusses our findings; and Section 3.6 concludes.

## 3.2 Conceptual framework

How might heterogeneity in preferences influence WTP for cooking technologies? Let some energy technology $\tau$ be represented by a vector, $\Lambda^\tau = (\lambda_1^\tau, \lambda_2^\tau, \ldots, \lambda_n^\tau)'$, of $n$ attributes (Rosen, 1974). Each $\lambda \in \Lambda$ may be thought of as denoting a specific device attribute, such as fuel efficiency, emissions or maximum temperature. The maximum price ($p$) that household $i$ is willing to pay for device $\tau$ is determined by a household response function, which takes as inputs these attributes.[1] Specifically:

$$p(\tau) = p_i(\lambda_1^\tau, \lambda_2^\tau, \ldots, \lambda_n^\tau; X_i), \tag{3.1}$$

[1] We assume that each individual household is a price-taker in the market for these devices.

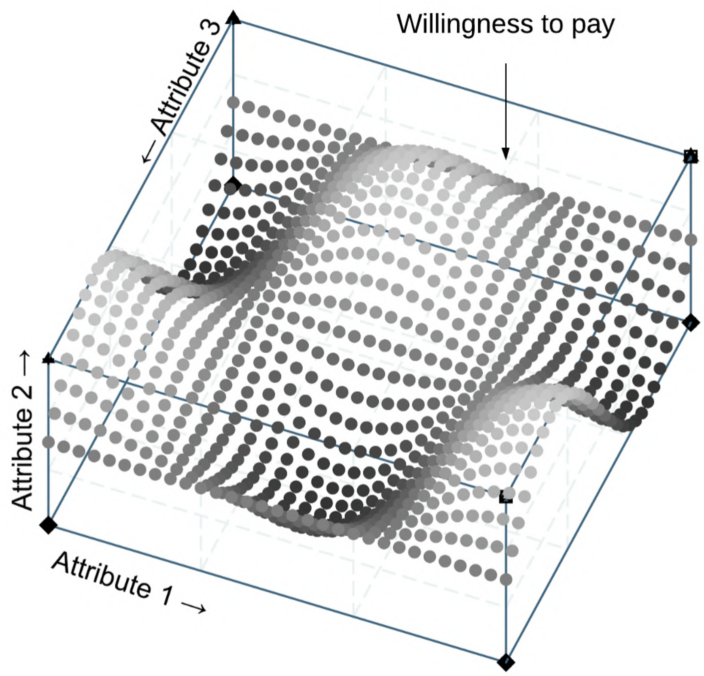


FIGURE 3.1: Shape of the household's technology WTP response surface. This figure presents a household's WTP for a hypothetical technology as a function of three continuous attributes. It demonstrates how underlying household preferences over attributes might give rise to WTP responses that need not be "well behaved" (e.g., non-monotonic changes in WTP in response to changes in particular attributes).

where $X_i$ represents a vector of characteristics for household $i$.

What does this underlying relationship between preferences and WTP mean for efforts to promote adoption for cleaner cooking technologies in remote, rural settings? To develop a conceptual framework to engage with this question, we draw on Nadel and Pritchett (2016), whose work on the limited external validity of individual randomized controlled trials (RCTs) serves as a valuable starting point. Consider the response function shown in Equation (3.1). Without additional assumptions about the shape of this function, it is neither possible to know a household's WTP for a particular stove in advance, nor is it possible to evaluate whether a household values one device more or less than another. Indeed, there is no reason to expect *ex ante* that households' WTP is "well behaved."

Figure 3.1 illustrates this challenge using a case where a hypothetical household's WTP is fully determined by three attributes. The three-dimensional WTP response surface shown in this figure emerges from distinct combinations of these attributes. WTP changes in unpredictable ways (e.g., non-monontonically in response to changes in particular attributes). Heterogeneity in preferences implies each household's response surface is distinct. Aggregating households over larger spatial scales (such as the village) can, in principle, yield a community-level response surface, yet the shape of this surface is also unknown to the analyst or the policymaker. Indeed, each intervention that sheds light on demand for a particular device reveals only a single point on the surface; it does little to map the response space, which is ultimately what is needed for the design of successful policies.[2]

If—consistent with neoclassical assumptions in economics regarding the stability of agents' preferences—each household's response surface is fixed, increasing the number of stove choices that are made available to beneficiaries as part of an intervention should weakly increase adoption but have no influence on WTP. After all, different devices appeal to different households, and a diverse portfolio of products to choose from should induce adoption of cleaner stoves by a larger share of households that identify options in line with their preferences; it should not influence households' WTP for the devices relative to a situation where each alternative is promoted exclusively since these underlying valuations are fixed. Prior evidence in the domain of clean cooking suggests that this might be the case. For instance, Jeuland et al. (2015a) use stated preferences techniques to find that households' willingness to pay for distinct attributes of a hypothetical cooking device—such as the smoke emissions it reduces—varies systematically with observable

[2] Nadel and Pritchett (2016) apply their framework more broadly to the success of interventions. They argue that the world of policy design is often characterized by (i) high-dimensional design spaces (the number of design elements that can be varied in a particular intervention or policy); (ii) "rugged" outcome-response surfaces over those design elements (whereby small design changes can lead to unpredictable changes in outcomes); (iii) context-specific outcome-response surfaces. In such a setting, they conclude, learning about the effectiveness of policies through the use of RCTs is sub-optimal compared to "crawling the design space" through an iterative approach.

household characteristics. Households' preference orderings over these attributes are then found to be effective at predicting the kinds of stoves they ultimately adopt (Jeuland et al., 2018a). This work builds on insights from earlier, pilot-scale efforts that suggest that the provision of multiple devices is associated with higher uptake rates (Lewis et al., 2015). While these are promising results, this prior work lacks an appropriate control group to rigorously draw inferences about how introducing additional choices drives outcomes.

More fundamentally, response surfaces may themselves not be fixed, further compounding the uncertainty faced by policymakers. In particular, the availability of additional choices—and the ability to more directly compare substitutes—might directly influence the shape of the underlying response surface, in line with a "constructive processing approach" to preference generation, whereby preferences are "constructed—not merely revealed—in the generation of a response to a judgment or choice task" (Payne et al., 1992). This may occur if, for instance, the very process of considering the advantages and disadvantages of different alternatives serves to highlight the idiosyncratic benefits or deficiencies of particular devices.

Hoeffler and Ariely (1999), who study the process by which preferences are learned and developed over time, present a prominent early attempt to bridge these two traditions. They contend that reality lies between the neoclassical approach (in which preferences are fixed) and the constructive processing approach (in which preferences are constructed based on the factors available at the time of preference elicitation). Specifically, using a series of lab-based choice experiments, they demonstrate that prior experience (or lack thereof) in a particular domain influences preference development, and repeated choices within that domain lead to preference stability. In other words, a household's preferences are malleable when familiarity is low but stabilize as the household gains more experience.

Without additional assumptions, theory alone is insufficient at highlighting how the availability of additional alternatives and choices influences WTP for cleaner energy technologies in remote, rural settings. If preferences over stove attributes are fixed and stable, we should expect to see no effect of the inclusion of additional alternatives on WTP. If, on the other hand, additional choices between substitutes highlight the idiosyncratic benefits of particular devices (especially unfamiliar ones), they might increase observed WTP. If, on the other hand, additional choices make the shortcomings of these devices more salient, WTP may be reduced. Two related research questions emerge:

1. Are preferences over stove attributes fixed and stable, as perceived from observed WTP?
2. If not, what is the direction and magnitude of the impact that the inclusion of additional choices has on WTP for cleaner energy technologies?

Accordingly, the remainder of this paper focuses on an empirical approach to evaluate how and to what extent the availability of additional choices influences households' WTP for cleaner cooking technologies through second-price, sealed-bid ("Vickrey") auctions in rural Senegal.

## 3.3 Background and study design

In this section, we provide a brief overview of the rural energy challenge in Senegal (including a description of the two stoves we use in our study) before turning to a detailed description of our study design.

### *3.3.1 Rural energy in Senegal*

Senegal faces an exceptionally heavy energy-related burden: over 95 percent of its rural population uses firewood for its primary energy needs. The consequences of these energy-use patterns for the local environment and people's living conditions are massive. Environmental concerns in particular are heightened by deforestation and its role in desertification in the Sahel (Brandt et al., 2014). Most of Senegal (except southern Casamance) is arid; firewood is scarce, and fuel collection in rural areas often exceeds ten hours per week—a workload that is mostly borne by women (Bensch and Peters, 2015). In addition, the World Health Organization (2009) estimates that overall 6,300 people die every year in the country from diseases related to household air pollution generated by energy use.

Against this backdrop, Senegal's government endeavors to lower reliance on polluting fuels. The program with the greatest potential reach is FASEN (*Foyers améliorés au Sénégal*), which promotes the *Jambar* stove, a simple, low-cost device (comprising a metal cylinder and a clay inlay) that comes in charcoal and firewood versions. The firewood *Jambar* is the first of two stoves that we use in our study. FASEN fosters local production of the *Jambar* by training ceramic and metalworking artisans. It also sponsors regular radio and television advertising to ensure that messages about clean cooking reach households. Ultimately, it looks to establish self-sustaining markets, in which potential customers are aware of the multidimensional set of benefits the Jambar delivers and are thus willing to pay cost-covering prices.[3] Yet FASEN's successes to date have largely been in and around urban centers, where charcoal is the primary solid fuel. Making inroads into rural areas, where potential end-users rely almost exclusively on firewood, has proven challenging.

The second stove that we use in our study is the *Jumbo Zama*, which, in contrast, is manufactured in South Africa and is generally unavailable in Senegal. It is made of heat-resistant steel with an outer cage that remains cool to the touch even after prolonged use, which makes it both durable and transportable. It was originally designed as a larger version of the *Zama Zama* stove, a wood-burning device whose design aims to maximize combustive efficiency without electric ventilation. In addition to being visually distinct, therefore, the *Jumbo Zama* stove can potentially deliver higher firewood savings and air-quality improvements than the *Jambar* stove.

[3] Specifically, the *Jambar* has been shown to lower consumption of firewood and "decrease early indicators for respiratory diseases and eye infections" in real-world settings (Bensch and Peters, 2015). That said, we note that a growing body of work suggests that simple, biomass-fueled stoves such as the *Jambar* may not be efficient enough to deliver expected improvements in firewood use, air quality or health (Aung et al., 2016; Hanna et al., 2016; Mortimer et al., 2017). For this reason, there are increasing calls for research on the potential of electric- and natural-gas-based cooking technologies (Smith, 2014, 2017).

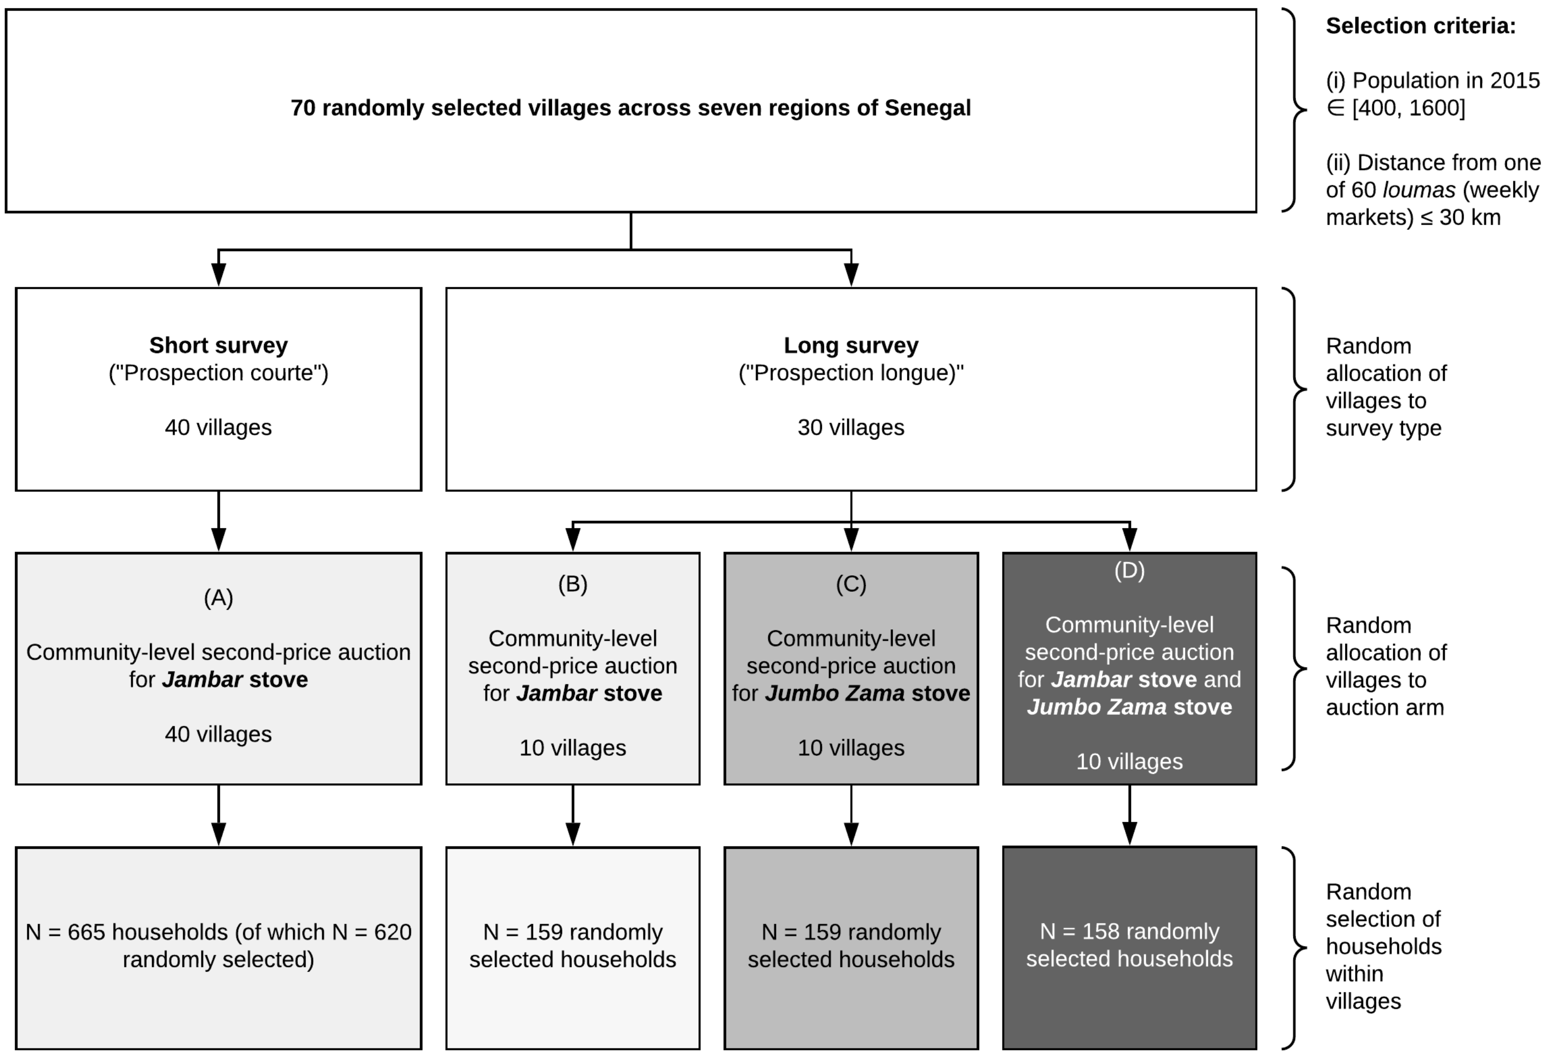


FIGURE 3.2: Randomized allocation of villages to auction arms

### 3.3.2 *Study design*

Our study design relies on (i) random allocation of villages into one of four different auction arms; and (ii) random selection of households within villages to serve as auction participants. Specifically, as shown in Figure 3.2, we start with a sample of seventy randomly selected villages from across seven regions of Senegal.[4] Each of these villages is then randomly assigned to either a short survey (*prospection courte*) or long survey (*prospection longue*) arm. With these survey arms, villages are next assigned to auction arms in which we conduct second-price, sealed-bid auctions following the mechanism first described by Vickrey (1961).[5]

There are two primary differences between survey arms. First, although households in both survey arms complete a survey to capture households' socioeconomic and demographic characteristics, perceptions related to environmental health risks and energy-use patterns, field teams use a more detailed survey instrument in the long survey arm. Second, households in the long survey arm cast their auction bids by entering bid amounts on paper auction sheets as part of household surveys. Field teams collect these auction sheets and host a public event at a prespecified location in the village at the end of the day, at which point the bids on each auction sheet are revealed and the results of the auction are announced. In contrast, in the short survey arm, both the casting of bids on paper auction sheets as well as the announcement of the auction results take place sequentially at a public event in the village.

---

[4] These villages were randomly selected from the universe of all Senegalese villages (i) whose population (in 2015) was between 400 and 1,600 people; and (ii) that are located within thirty kilometers of a set of sixty pre-identified *loumas* (weekly peri-urban/rural markets). The region-wise distribution of this sample of villages is as follows: Diourbel (22 villages); Fatick (14); Kaffrine (8); Kaolack (10); Louga (5); Saint-Louis (3); and Thiès (8).

[5] In this type of auction, bidders submit secret bids without knowing the bids of other auction participants. The highest bidder wins the auction and pays a price equal to the second-highest bid. This mechanism is incentive compatible; the optimal strategy of each bidder is to bid—and thus reveal—her true valuation. In combination with the mechanism's relative simplicity, which makes it easy to explain in low-literacy settings, the mechanism's incentive compatibility has made it a popular tool to elicit valuations in field-based setting (e.g., Demont et al., 2013).

All villages in the short survey arm are assigned to Auction Arm "A," which consists of a community-level second-price, sealed-bid auction for the *Jambar* stove. Each village in the long survey arm is randomly assigned to one of three auction arms:

- Auction Arm "B," which also consists of a second-price, sealed-bid auction for the *Jambar* stove;
- Auction Arm "C," which consists of a second-price, sealed-bid auction for the *Jumbo Zama* stove; and
- Auction Arm "D," where both the *Jambar* and the *Jumbo Zama* stoves are auctioned using a second-price, sealed-bid mechanism.

Finally, in each village, approximately fifteen households are randomly selected from household lists and invited to participate in auction and survey activities.[6] If household members are initially unavailable or engaged in other tasks, field teams return later in the day; if a household does not wish to participate, field teams select replacements randomly from household lists. Once a household consents to participating, field teams share information about how to use the stove(s) being auctioned, associated benefits (such as the device's potential to lower fuel consumption), and answer questions. Appropriate sample stoves are available during this time for household members to examine. Field teams then carefully explain the second-price auction mechanisms. In particular, field teams highlight that a bid that corresponds to the maximum amount one is willing to pay is ideal, given the rules of the second-price auction mechanism.[7]

[6] Because both the casting of bids as well as the announcement of auction results occurs publicly in the short survey arm, members of households that are not randomly selected and invited to participate are also present and, in some cases, participate in auction and survey activities. Specifically, as shown in Figure 3.2, 45 out of the 665 households (approximately seven percent) in the short survey arm were not randomly selected to participate. While we do not restrict the participation of these individuals in the field, we check the robustness of our results to the presence of this small group of "self-selected" auction participants by excluding them entirely from our analytical sample and find that our results remain unaffected.

[7] The results of these field-based auctions are not binding in any legal sense. Indeed, in villages where

The highest bidder is deemed the winner of the auction at the public event to announce auction results, and invited to purchase the device in question at a price equal to the second-highest bid. If the individual refuses, field teams make note of the refusal, and invite the second-highest bidder to purchase the device at a price equal to the third-highest bid.[8]

This study design—in particular, our four auction arms—allow us to evaluate how the presence of additional choices influences WTP. Specifically, results from auction arms "A" and "B" shed light on baseline WTP for the *Jambar* stove, those from auction arm "C" do the same for the *Jumbo Zama* stove, and those from auction arm "D" highlight how the presence of a distinct alternative that caters to the preferences of a different group of households influences households' WTP for each device.

## 3.4 Empirical specifications and results

In this section, we first provide an overview of our analytical sample and test for balance across our auction arms. We then highlight and estimate our main empirical specifications, and present results. We find that jointly auctioning the *Jambar* and *Jumbo Zama* stoves reduces WTP for the latter by at least 25 percent; WTP for the *Jambar*, on the other hand, is unaffected. This result is robust to the use of a variety of different empirical approaches

both stoves are auctioned jointly, field teams inform households that should they end up winning both auctions, they have the option to purchase only one should they be so inclined. This ensures households bid their true WTP for each device, unencumbered by concerns related to budget constraints under dual-win scenarios. In practice, very few households (between three and six percent of the auction participants, depending on the device) do not follow through with a purchase upon winning the auction. The script used by field teams to explain the second-price auction mechanism in villages assigned to the long survey arm is shown in Appendix F. Field teams used a similar script to explain the auction mechanism in villages in the short survey arm.

[8] In case multiple bidders have cast the highest bid, a winner is selected randomly from the set of highest bidders. If this individual declines the purchase offer, field teams first randomly select and approach one of the other individuals with the highest bid before moving on to individuals with lower bids. Similarly, if multiple individuals have cast the second-highest bid and the highest bidder declines to purchase offer, field teams randomly select one of the second-highest bidders, who is invited to purchase the device at a price equal to the third-highest bid. The auction concludes when an invitation made to purchase the device in this way is accepted.

as well as to inclusion of a host of household- and community-level controls.

### *3.4.1 Descriptive statistics and balance tests*

A total of 1,141 households from across seventy villages participated in stove auctions across all four auction arms. Of these, 1,096 (96 percent) were randomly selected by field teams and invited to participate in auction and survey activities; 45 households self-selected into the auctions conducted in Auction Arm "A," where both the casting of sealed bids as well as the announcement of results occurred during a public event. Table 3.1 presents descriptive statistics and tests for balance across auction arms for the sample of randomly selected households and villages.

Panel (a) of Table 3.1 presents data on key household-level socioeconomic and demographic characteristics.[9] As shown in column (1), the average household in Auction Arm "A" is large, made up of approximately thirteen people. This is consistent with patterns of marriage and household formation in Senegal, where polygyny (whereby a man has more than one wife) is widely practiced (Lardoux and de Walle, 2003).[10] The average household head is approximately fifty years old and overwhelmingly male, and just over half of household heads can read. Only fourteen percent of households have a household member with a bank account, suggesting that households' access to formal banking services is low. The average households sees itself as approximately in the middle of the village-level wealth distribution.

Panel (b) of Table 3.1 presents key village-level characteristics, reported by the village chief (or another key informant) during community surveys, which were conducted prior to beginning field activities within each village. As shown in column (1), the average village in Auction Arm "A" is reported to have just under 1,000 people, consistent with our *ex ante* selection of villages with Census populations lying between 400 and 1,600

[9] For consistency, we restrict comparisons in Table 3.1 to key variables present on both versions of the survey used during field activities.

[10] We define a "household" as a group of individuals who typically prepare and consume meals together.

Table 3.1: Descriptive statistics and tests for balance across auction arms

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) |
|---|---|---|---|---|---|---|---|---|---|---|
| | Auction Arm "A" | | Difference between Auction Arm "A" and: | | | | | | | |
| | Mean | Std Dev | Auction Arm "B" | | Auction Arm "C" | | Auction Arm "D" | | $N$ | Adj $R^2$ |
| | | | Coeff | Std Err | Coeff | Std Err | Coeff | Std Err | | |
| (a) *Household-level characteristics* | | | | | | | | | | |
| Household size | 12.9 | 6.91 | -0.54 | (0.86) | 1.61* | (0.92) | 0.51 | (0.78) | 1096 | 0.005 |
| Age of household head | 51.0 | 14.06 | 2.28 | (1.78) | 0.26 | (2.07) | 3.55* | (1.96) | 1077 | 0.021 |
| $\mathbb{1}$ (Female household head) | 0.05 | 0.23 | 0.0062 | (0.021) | -0.028* | (0.015) | -0.03* | (0.018) | 1096 | -0.001 |
| $\mathbb{1}$ (Household head can read) | 0.55 | 0.50 | -0.021 | (0.080) | 0.014 | (0.064) | -0.037 | (0.077) | 1096 | 0.000 |
| $\mathbb{1}$ (Bank account) | 0.14 | 0.35 | -0.0027 | (0.041) | 0.075 | (0.047) | 0.025 | (0.048) | 1059 | 0.031 |
| Relative wealth perception (out of 5) | 2.55 | 1.28 | -0.14 | (0.17) | 0.16 | (0.14) | 0.13 | (0.13) | 1076 | 0.015 |
| $\mathbb{1}$ (Household owns *Jambar*) | 0.02 | 0.15 | 0.00035 | (0.013) | 0.0088 | (0.017) | -0.00028 | (0.013) | 1096 | -0.004 |
| (b) *Village-level characteristics* | | | | | | | | | | |
| Population | 945.7 | 725.6 | -174.2 | (204.8) | 566.0 | (360.7) | 221.9 | (241.9) | 56 | 0.008 |
| $\mathbb{1}$ (Transport facilities) (e.g., bus stop) | 0.03 | 0.16 | 0.21 | (0.14) | 0.13 | (0.10) | 0.17 | (0.12) | 70 | 0.168 |
| $\mathbb{1}$ (Educational facilities) (e.g., school) | 0.75 | 0.44 | 0.098 | (0.15) | 0.018 | (0.16) | 0.094 | (0.17) | 70 | -0.086 |
| $\mathbb{1}$ (Health facilities) (e.g., community health center) | 0.48 | 0.51 | -0.20 | (0.17) | -0.018 | (0.20) | -0.057 | (0.18) | 70 | -0.089 |
| $\mathbb{1}$ (Bank/microfinance facilities) | 0.05 | 0.22 | 0.051 | (0.12) | -0.038 | (0.037) | 0.048 | (0.11) | 70 | -0.057 |
| $\mathbb{1}$ (Women's groups) | 1 | 0 | -0.074 | (0.069) | -0.0086 | (0.0096) | 0.0094 | (0.022) | 70 | 0.121 |
| $\mathbb{1}$ (Youth groups) | 0.98 | 0.16 | -0.34** | (0.16) | -0.0052 | (0.043) | -0.16 | (0.13) | 70 | 0.173 |
| $\mathbb{1}$ (Religious groups) | 1 | 0 | 0 | – | 0 | – | 0 | – | 70 | . |

This table shows difference in selected household- and village-level characteristics across the four auction arms indicated in Figure 3.2. Columns (1) and (2) show the mean and standard deviation, respectively, for the indicated characteristic for households/villages in Auction Arm "A." Columns (3), (5) and (7) of panel (a) report coefficients $\hat{\beta}_1$, $\hat{\beta}_2$ and $\hat{\beta}_3$, respectively, from a linear regression of the form: $Y^A_{ijr} = \beta_0 + \beta_1 B_j + \beta_2 C_j + \beta_3 D_j + \gamma_r + \epsilon_{ijr}$, where $Y^A_{ijr}$ represents the respective household-level characteristic for household $i$ in village $j$ in region $r$; $B_j$, $C_j$ and $D_j$ are binary variables that equal one if village $j$ is randomly assigned to Auction Arms "B," "C" and "D," respectively; $\gamma_r$ is a region fixed-effect; and $\epsilon_{ijr}$ is household-specific error term. The associated standard errors reported in columns (4), (6) and (8) of panel (a) are clustered at the village level. Columns (3), (5) and (7) of panel (b) report coefficients $\hat{\beta}_1$, $\hat{\beta}_2$ and $\hat{\beta}_3$, respectively, from a linear regression of the form: $Y^A_{jr} = \beta_0 + \beta_1 B_j + \beta_2 C_j + \beta_3 D_j + \gamma_r + \epsilon_{jr}$, where $Y^A_{jr}$ represents the respective village-level characteristics for village $j$ in region $r$; $\epsilon_{jr}$ is village-specific error term; and other parameters are defined as before. Heteroscedasticity-robust standard errors are reported in columns (4), (6) and (8) of panel (b). All analyses are restricted to households that were selected randomly by field teams and invited to participate in auction and survey activities. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

people. Only about three percent of villages have formal transportation facilities (such as a bus stop) within their boundaries. Approximately three-quarter, however, have a school and approximately half have a health clinic. Only five percent of villages have formal banking facilities, consistent with the low rates of households with access to a bank account. Finally, nearly all villages have active village-level women's groups, youth groups and religious groups.

Column (2)–(8) of Table 3.1 report estimated coefficients and associated standard errors from a linear regression of the respective household- or village-level characteristics on binary variables that equal one for Auction Arms "B," "C" and "D," and region fixed-effects.[11] These, therefore, shed light on any systematic differences between villages and households across our auction arms. Consistent with the random allocation of villages across auction arms and random selection of households within villages, we find that our sample is generally well balanced across all four auction arms.

### *3.4.2 Joint auctions and willingness to pay*

Figure 3.3 shows the distribution of bids (including median and upper/lower quartile values) for the *Jambar* and *Jumbo Zama* stoves in each of our four auction arms. Panel (a) of Figure 3.3 suggests that random allocation to the long survey arm increases median WTP for the *Jambar* relative to WTP in the short survey arm. It also suggests that jointly auctioning a *Jambar* stove along with a *Jumbo Zama* stove does not appear to have any additional effect. In contrast, as shown in panel (b), random allocation to the joint auction arm appears to substantially reduce WTP for the *Jumbo Zama* stove relative to when it is auctioned alone.

To more rigorously evaluate the impact of having an alternative choice on WTP for these devices, we separately estimate the following specifications using ordinary least

[11] See Table 3.1 for additional details.

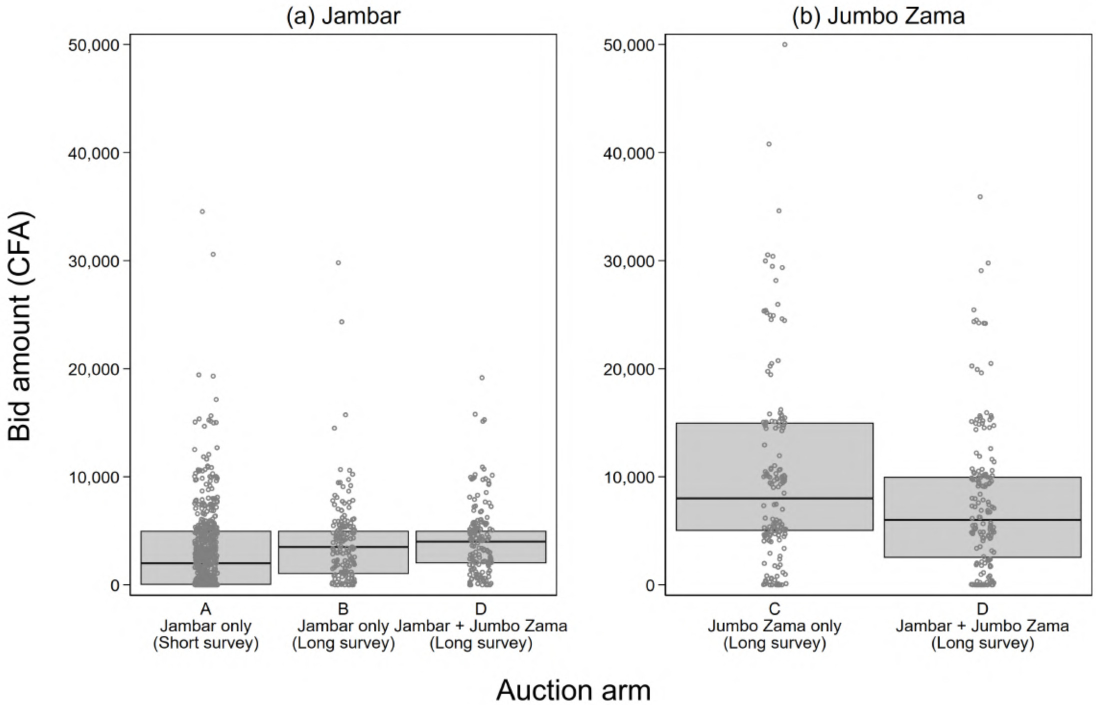


Figure 3.3: Distribution of bids across auction arms. This figure shows the distribution of bids cast in second-price auctions across the four auction arms indicated in Figure 3.2. Panel (a) shows the distribution of bids for *Jambar* stoves while panel (b) shows the distribution of bids for the *Jumbo Zama* stove. Markers indicate bid values; boxes indicate median and upper/lower quartiles for bids within each auction arms. Bid values have been jittered to visually highlight their distribution.

squares:

$$Y_{ijr}^{Jambar} = \beta_0 + \beta_1 Auction_j + \beta_2 Survey_j + \gamma_r + \epsilon_{ijr} \tag{3.2}$$

$$Y_{ijr}^{Zama} = \beta_3 + \beta_4 Auction_j + \gamma_r + \epsilon_{ijr}. \tag{3.3}$$

$Y_{ij}^{Jambar}$ and $Y_{ij}^{Zama}$ represent the bid cast for the *Jambar* and/or *Jumbo Zama* stoves, respectively, by household $i$ in village $j$ in region $r$. $Auction_j$ is a binary variable that equals one if village $j$ is randomly allocated to Auction Arm "D," in which both stoves are auctioned together. $\beta_1$ and $\beta_4$, our coefficients of interest, therefore represent the impact of the availability of an alternate choice on WTP for the *Jambar* and *Jumbo Zama* stoves, respectively. $Survey_j$ in Equation (3.2) is a binary variable that equals one if

Table 3.2: Impact of joint stove auction on willingness to pay

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | *Jambar* bid amount (CFA) | | | *Jumbo Zama* bid amount (CFA) | | |
| $\mathbb{1}$ (Joint stove auction) | -228.6 | -67.07 | -339.8 | -2260.5*** | -3288.7*** | -3494.1*** |
| | (641.2) | (679.8) | (644.0) | (772.2) | (897.1) | (797.1) |
| $\mathbb{1}$ (Long survey arm) | 1456.3*** | 973.7* | 1292.5** | | | |
| | (537.2) | (541.8) | (528.7) | | | |
| Region FEs | Yes | Yes | Yes | Yes | Yes | Yes |
| "True valuation" FEs | No | Yes | Yes | No | Yes | Yes |
| Household- and village-level controls | No | No | Yes | No | No | Yes |
| $N$ | 982 | 982 | 982 | 317 | 317 | 317 |
| Adjusted $R^2$ | 0.038 | 0.264 | 0.292 | 0.067 | 0.199 | 0.240 |
| Mean of outcome | 3272.3 | 3272.3 | 3272.3 | 8873.8 | 8873.8 | 8873.8 |

This table shows results from estimating Equations (3.2) and (3.3) using ordinary least squares on the full sample of bid data from *Jambar* and *Jumbo Zama* sealed-bid, second-price auctions in columns (1) and (4), respectively. All models include region fixed-effects. Columns (2) and (4) also include fixed-effects for households that won an auction but declined to purchase the respective stove at the next-highest price. Columns (3) and (6) include controls for all household- and village-level characteristics shown in Table 3.1. Missing values in control variables for (i) reported village population are replaced with region-level means; (ii) relative wealth perception, household size, and age of household are replaced with village-level means; and (iii) households' bank account access, ownership of the *Jambar* stove, and household head's ability to read are replaced with zeros Additional binary variables that equal one for any household for which missing values are replaced in this way are also included in the estimation. Standard errors—in parentheses—are clustered at the village level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

village $j$ is randomly allocated to the long survey arm. The coefficient $\beta_2$, thus, represents the impact of (i) receiving a personalized household-level demonstration of the *Jambar* stove; (ii) being able to cast a bid for it in a household (as opposed to a community) setting; and (iii) participating in a longer household survey on WTP for the *Jambar* stove.[12] $\gamma_r$ represents a region fixed-effect, which controls for all unobserved region-level characteristics that might independently induce variation in bid values. $\epsilon_{ijr}$ represents a household-specific error term. To account for unobserved intra-village correlation in bids, we cluster standard errors at the village level.

Column (1) of Table 3.2 presents results from estimating Equation (3.2) using the full sample of *Jambar* second-price auctions. The first row shows the estimated coefficient for the impact of joint stove auctions on WTP for the *Jambar*. Although it is negative,

[12] Because all villages in the short survey arm participate in *Jambar*-only auctions, we include $Survey_j$ only in the specification shown in Equation (3.2).

this coefficient is small in magnitude (approximately seven percent relative to the mean *Jambar* bid) and statistically indistinguishable from zero. We are thus unable to reject that hypothesis that joint *Jambar–Jumbo Zama* auctions have no impact on WTP for the *Jambar*. The second row of column (1) displays the estimated coefficient for the impact of being in the long survey arm on WTP for the *Jambar* stove. We find that random assignment to the long survey arm results in *Jambar* bids that are on average over CFA 1,450 (s.e. 537) higher, a statistically significant increase in WTP by approximately 45 percent relative to the mean *Jambar* bid.

Column (4) of Table 3.2 presents corresponding results from estimating Equation (3.3) using the full sample of *Jumbo Zama* second-price auctions. In contrast to the null result shown in column (1), we find that joint *Jambar–Jumbo Zama* auctions lead to a large reduction in WTP for the *Jumbo Zama*. Specifically, the average *Jumbo Zama* bid is nearly CFA 2,300 (s.e. 772) lower in villages randomly assigned to the joint stove auction, representing a 25 percent reduction in WTP for the device.

We next look to control for bias introduced by bidders not bidding their true valuation. Recall that although field teams carefully explain the rules of the second-price auction to all auction participants and highlight the optimal bidding strategy (Appendix F), approximately six percent of households that cast bids in a *Jambar* auction and three percent that did so in a *Jumbo Zama* auction decline to purchase the respective device upon winning. To account for these bidders, we construct a binary variable that equals one for any household that is deemed an auction winner, invited to purchase a device and refuses, and estimate Equations (3.2) and (3.3) with this variable included. This "true valuation" fixed-effect controls for all unobserved household-level characteristics that would have led to certain households casting bids higher than their true valuation. Columns (2) and (5) of Table 3.2 present estimated coefficients. Our results are qualitatively unchanged: random assignment to the joint *Jambar–Jumbo Zama* auction considerably lowers households' WTP for the *Jumbo Zama* stove and has no statistically distinguishable impact on WTP for the *Jambar* stove.

Finally, in addition to the "true valuation" fixed-effect, we include all household- and village-level controls indicated in Table 3.1 and estimate Equations (3.2) and (3.3). Estimated coefficients are shown in columns (3) and (6) of Table 3.2. Once again, we find that random allocation to the joint auction lowers WTP for the *Jumbo Zama* stove and has no statistically distinguishable impact on WTP for the *Jambar* stove.

### *3.4.3 Additional analyses and robustness checks*

We cluster our standard errors at the village level in all analyses presented in Table 3.2 to account for unobserved intra-village correlation in bids. Cluster-robust standard errors, however, "can produce misleading inferences when the number of clusters … is small, even if the model is consistent and there are many observations in each cluster" (Esarey and Menger, 2018). Specifically, too few clusters can substantially increase false positive rates. This may be a concern in our setting, especially in the case of *Jumbo Zama* auctions, which occur in two auction arms covering a total of only twenty villages (Figure 3.2).[13] We address this potential limitation by using the wild cluster bootstrap-$t$ procedure developed by Cameron et al. (2008) to adjust the standard errors for all our main analyses. These adjustments (shown in Table G.12) do not change our main results.

Recall that, in Auction Arm "A," we do not stop households that had not been randomly selected by field teams from participating in auction and survey activities, and a total of 45 household self-select into participation in this way. Therefore, we next repeat our main analyses after restricting our sample to only those households that had been randomly selected as part of field activities. Table G.13 shows that our results are unaffected by this sample-selection choice.

Finally, in our setting, a household's bid cannot be lower than zero even if its true

[13] Because *Jambar* auctions also include villages randomly allocated to the short survey arm, a second-price, sealed-bid auction for the *Jambar* stove occurs in a total of sixty villages.

valuation for either or both of the stoves included in our study may be negative. Ordinary least squares estimation with data censored in this way can lead to inconsistent estimates. We, therefore, estimate all specifications presented in Table 3.2 using the Tobit regression model, which is designed to estimate linear relationships with a dependent variable that is censored from above or below (Tobin, 1958). Table G.14 presents these estimates, which show that our results are qualitatively unchanged.

## 3.5 Discussion

We return now to the research questions that emerged from the conceptual framework outlined in Section 3.2. Had we observed no difference in households' WTP for the *Jambar* and the *Jumbo Zama* stoves between the single- and joint-stove auctions, we would have concluded that households' underlying WTP response surfaces (that is, their preferences) are heterogeneous but fixed and unaffected by the number of alternatives in a given choice set.[14] In a scenario with consumers behaving "rationally" in this way, the provision of additional choices that cater to the needs of an increasingly larger proportion of beneficiaries as part of ICS promotion activities would have weakly increased adoption rates.

However, we document an observed *reduction* in demand for at least one of the two stoves in the joint-auction scenario. A simple explanation for this result might be that households cast lower bids for the *Jumbo Zama* during the joint *Jambar–Jumbo Zama* auction due to budgetary reasons. Specifically, if households are concerned that they may not be able to afford both devices in case they won both auctions by bidding their true valuations, they may systematically cast lower bids for one or both alternatives as a way

[14] This is partly related to the independence of irrelevant alternatives (IIA) axiom in decision theory, which states that if, within the choice set $\{A, B\}$, $A$ is preferred to $B$, then that preference ordering is maintained in the expanded choice set $\{A, B, C\}$. We stress, however, that while we see a significant reduction in WTP for the *Jumbo Zama* stove when the *Jambar* stove is included in the choice set, we do not observe changes in households' preference between the two devices. Indeed, mean WTP for the *Jumbo Zama* is always larger than that for the *Jambar*.

to minimize this risk. Recall, however, that field teams specifically inform all bidders in the joint-auction arm that should this situation occur, the winner has the choice to purchase both or only one of the devices. In addition, we separately compare differences in the rates of households that won an auction for the*Jambar* or *Jumbo Zama* stove but ultimately decline to purchase the respective device across our four auction arms.[15] If strategic concerns are driving households to systematically lower their bids in the joint-auction arm, we would expect to observe lower rates of households bidding above their "true valuations" for one or both of the stoves in this way. We find no evidence to suggest that these rates are different across auction arms (results available upon request).

Taken together, our results thus suggest that household decision-making related to technology adoption may not be fully "rational." Jointly auctioning the *Jambar* and *Jumbo Zama* appears to change households' WTP response surfaces and their preferences over stove attributes, consistent with a process of consolidation and stabilization over preferences given repeated choices and gained experience. Indeed, in a recent survey of cooking practices and preferences over stove attributes in a rural Senegalese community, Hooper et al. (2018) find that although ICS are viewed as more desirable than traditional alternatives, "first-hand experience with these stoves [is] limited." Even in our study sample, only about two percent of households own a *Jambar* stove prior to the beginning of field activites (Table 3.1). The *Jumbo Zama* is largely unavailable in Senegal. The ability to make choices (in particular, repeated choices during the joint auction) and consider difficult trade-offs in a new domain may have helped participants crystallize their heretofore malleable preferences and arrive at a bid value that more accurately reflected their true valuation.

The implications for policymakers and practitioners looking to enhance widespread adoption of cleaner cooking solutions in rural areas are clear. The provision of additional choices alone is unlikely to be sufficient. A one-shot increase in the number of available

[15] Recall that this comprises between three and six percent of our sample of households.

alternatives in a scenario characterized by unfamiliarity may even have an adverse effect on demand. Rather, implementers must "find ways to help consumers think deeply about the trade-offs in their product domain so that they understand better their own preferences and purchase products that better fit their needs" (Hoeffler and Ariely, 1999).

## 3.6 Conclusion

Researchers and policymakers alike routinely stress the importance of catering to local preferences and needs when designing interventions that seek to increase adoption of cleaner cooking technologies in remote, rural settings. We develop a simple conceptual framework that connects heterogeneity in households' preferences over technology attributes to demand for these devices. The predictions of our framework hinge on the nature of underlying preferences. If households' preferences are heterogeneous but fixed, the simultaneous provision of different types of energy technologies that cater to distinct energy-use patterns and needs should weakly increase adoption of improved alternatives without having any impact on willingness to pay. If, on the other hands, preferences are malleable and emerge through the very act of making choices in unfamiliar domains, the impact of simultaneously providing multiple alternatives is unclear.

To test these hypotheses, we conduct technology-promotion campaigns followed by second-price, sealed-bid ("Vickrey") auctions for two cleaner cooking technologies—a locally made stove, the *Jambar*, and an imported, more efficient alternative, the *Jumbo Zama*—with over 1,000 households across seventy communities in rural Senegal. We induce exogenous variation in the extent to which these promotion activities cater to heterogeneous preferences by randomly selecting a subset of communities where both devices are promoted jointly. Consistent with a model in which preferences are constructed—and not simply revealed—through repeated choices, joint promotion lowers willingness to pay for the relatively less familiar *Jumbo Zama* stove without having any impact on demand for the *Jambar*.

These results suggest that the provision of additional choices alone may not be sufficient to drive widespread adoption of cleaner cooking solutions. Rather, implementers and policymakers should devise approaches to help potential end-users think carefully through the trade-offs inherent in the choice between unfamiliar alternatives, crystallize and understand their own preferences, and identify solutions that meet their distinct energy-use needs.

# Conclusion

Electricity is indispensable for households, clinics, schools and firms, yet over a billion people live without it. At the same time, nearly three billion rely on traditional stoves and dirty biomass fuels (such as firewood) for their basic energy needs. The resulting household air pollution causes four million deaths annually, a health burden borne disproportionately by women. The international community has hastened to respond to this global energy challenge. This dissertation highlights when, where and why policies that seek to ensure universal access to modern energy deliver expected environmental and development benefits.

In the first chapter, I ask what drives heterogeneity in the impacts of large-scale rural electrification. Prior evidence on the labor-market impacts of grid electrification is mixed. I hypothesize that variation in local economic conditions—which can complement investments in infrastructure—may help explain why, and combine two natural experiments in India within a regression discontinuity design to test this hypothesis. Most of the world's guar, a crop that yields a potent thickening agent used during hydraulic fracturing ("fracking"), is grown in northwestern India. The rapid rise of fracking in the United States induced a parallel commodity boom in Indian guar production, resulting in a large positive shock to rural economic activity. Leveraging population-based discontinuities in the contemporaneous roll-out of India's massive rural electrification scheme, I show that access to electricity significantly increased non-agricultural employment in villages located in India's booming guar belt. Where these complementary economic conditions were lacking, electrification had almost no discernible impact. Using a firm-level panel dataset, I then provide suggestive evidence that this growth in non-farm work is partly driven by the rise of electricity-intensive firms that complement agricultural production. In line with the prior literature, I show that electrification alone may not be sufficient to deliver economic benefits, but I also demonstrate that, when combined with complementary economic conditions on the ground, access to electricity can enable individuals, households and firms to take advantage of new opportunities in potentially welfare-enhancing ways.

In the second chapter, I turn to household-level energy use and empirically evaluate the role played by non-governmental organizations (NGOs) in delivering environmental, energy and development interventions in remote, rural settings. I develop a model of household decision-making to evaluate how NGOs address implementation-related challenges and influence intervention effectiveness. To test the model's predictions, I apply quasi-experimental methods to household-survey data from a randomized controlled trial designed to promote clean-cooking solutions in rural India. I uncover a large, positive and statistically significant "NGO effect": prior engagement with the implementing NGO increases the effectiveness of the intervention by at least thirty percent. These findings provide some of the first causal evidence on how NGOs directly influence outcomes, which has implications for the generalizability of experimental research conducted jointly with such local partners. In particular, attempts to scale up findings from such work may prove less successful than anticipated if the role of NGOs is insufficiently understood. Alternatively, policymakers looking to scale up could achieve greater success by fostering partnerships with trusted local institutions.

In the final chapter, I consider how heterogeneity in households' preferences influences demand for energy technologies. I conduct technology-promotion campaigns followed by second-price, sealed-bid ("Vickrey") auctions for two cleaner cooking technologies with over 1,000 households across seventy communities in rural Senegal. I induce exogenous variation in the extent to which these promotion activities cater to heterogeneous preferences by randomly assigning a subset of communities to an auction arm in which both devices are promoted jointly. Consistent with a model in which preferences are constructed—and not simply revealed—as agents make repeated choices, joint promotion lowers willingness to pay for the relatively less familiar alternative compared to settings in which the two devices are promoted exclusively. Rather than simply providing additional choices, implementers looking to enhance uptake of improved technologies must instead devise approaches to help potential end-users think carefully through trade-offs, crystallize and understand their own preferences, and identify solutions that fit their needs.

Two prominent themes connect these three chapters. First, each deploys a diversity of rigorous methodological tools to conduct impact evaluations in the overlapping domains of energy and international development. In so doing, the chapters in this dissertation help identify and fill crucial, policy-relevant knowledge gaps related to global energy access. Second, each of these chapters present efforts to rigorously study drivers of heterogeneity directly. These highlight mechanisms through which investments in energy infrastructure and technologies yield expected benefits, and the ways in which such energy investments interact with context-specific characteristics and influence final outcomes. They also suggest that a variety of local factors may be driving heterogeneity in effectiveness. Rigorously identifying these drivers of success is a promising avenue for future research as this knowledge base can help guide and improve spatial targeting of energy and development policies globally. Ultimately, it is my hope that the insights that emerge from the work in this dissertation contribute to helping low- and middle-income countries meet the ambitious energy-access targets on which they have rightfully set their eyes.

# Appendix A

## Using nighttime luminosity to evaluate the impact of the fracking-induced guar boom on economic activity

Did the fracking-induced guar boom in northwestern India have a meaningful impact on economic activity? To answer this question, we rely on the synthetic control methodology (SCM) applied to two decades of nighttime luminosity data covering nearly all of India's approximately 600,000 villages. We find that guar-growing districts shine brighter at night as a result of the start of the guar boom compared to a synthetic "counterfactual." As nighttime luminosity is a widely accepted proxy for regional economic activity, these results point to a large increase in economic activity in India's guar-growing regions due to the start of the United States' fracking boom.

### A.1 Synthetic control methodology

Like the conventional difference-in-differences estimator, the SCM relies on differences between "treated" and "untreated" units before and after an event of interest (Abadie and Gardeazabal, 2003; Abadie et al., 2010). However, SCM does not give equal weight to all untreated units. Instead, it hinges on using a linear combination of untreated units to generate a weighted average whose pre-treatment outcome trends closely match those of the treated unit. This synthetic "counterfactual" unit is then projected into the post-treatment period and compared with the treated unit to gauge the direction and magnitude of impacts.

This feature makes it particularly attractive for estimating treatment effects in small-sample settings such as our own, in which only 23 mostly contiguous districts in nothwestern India are assumed to be "treated" by the fracking boom. Indeed, many applications have featured only one treated unit that is compared with multiple untreated units over time (e.g., Coffman and Noy, 2011; Singhal and Nilakantan, 2016).

Formally, let $T_0$ represent the number of pre-treatment periods (out of $T$ total periods) and $J$ represent the number of untreated units. Let $\mathbf{W} = \left(w_1, \ldots, w_J\right)$ be a $(J \times 1)$ vector of non-negative weights such that $\sum_{j=1}^{J} w_j = 1$. Each $w_j \in \mathbf{W}$ represents the weight of the $j^{\text{th}}$ untreated unit. Let $\mathrm{Y}_1$ be a $(T_0 \times 1)$ vector of outcome measures in the treated unit for each pre-treatment period $t$. Similarly, let $\mathrm{Y}_0$ be a $(T_0 \times J)$ matrix that contains the same outcome measures for each untreated unit $j$ in pre-treatment period $t$. Broadly, the aim of the SCM is to pick $\mathbf{W}^*$ such that:

$$\mathbf{Y}_1 = \mathbf{Y}_0\mathbf{W}^*. \tag{A.1}$$

Applications of the SCM typically specify a set of $k$ pre-treatment characteristics $\mathbf{X}$ as predictors, where $\mathbf{X}$ includes observed covariates $\mathbf{Z}$ that are unaffected by the treatment as well as linear combinations of the pre-treatment outcomes $\mathbf{Y}$. Given $\mathbf{Y}$ and $\mathbf{X}$, $\mathbf{W}$ is picked so as to minimize the root-mean-squared prediction error (RMSPE) of the predictors:

$$\mathbf{W}^* = \arg\min_{\mathbf{W}} \left\{ \sqrt{(\mathbf{X}_1 - \mathbf{X}_0\mathbf{W})'\, \mathbf{V}\, (\mathbf{X}_1 - \mathbf{X}_0\mathbf{W})} \right\}, \tag{A.2}$$

where the subscripts denote treated and untreated units as in Equation (A.1), and $\mathbf{V}$ represents a $(k \times k)$ matrix that specifies the relative importance of the predictors.[1] Placebo

[1] Abadie and Gardeazabal (2003) choose $\mathbf{V}$ so as to minimize the RMSPE of the outcome variable in the

tests determine the statistical significance of the effects observed in the post-treatment period. Specifically, the treated unit is excluded from the sample, and the analysis is repeated for each untreated unit, which is now assumed to have been treated instead. The presence of many large effects in the resulting distribution of post-treatment placebo effects suggests that the original estimated effect may have been the result of chance.[2]

## A.2 Nighttime luminosity

Nighttime luminosity measures are increasingly used by economists to investigate changes in regional economic activity over time (Doll et al., 2006; Henderson et al., 2012). Recent applications also demonstrate that they serve as useful proxies for information on socioeconomic outcomes in low-income settings, where high-quality statistical data are often missing (Chen and Nordhaus, 2011; Pinkovskiy and Sala-i-Martin, 2016). This work typically uses data generated as part of the Defense Meteorological Satellite Program (DMSP) led by the National Oceanic and Atmospheric Administration (NOAA). DMSP satellites take pictures of the Earth every night. NOAA processes and cleans these nightly images to remove irregularities (such as cloud cover or solar glare), averages them across years, and makes the annual composite images publicly available.[3] Each pixel of these annual images—representing 30 arc seconds or approximately 1 km$^2$ at the equator—is assigned a number on a relative brightness scale ranging from 0 to 63.

Most prior research has relied on these annual composites. While annual averages certainly provide useful information, they smooth away substantial variation in brightness over the calendar year and are, therefore, less precise (Min et al., 2017). We use a considerably richer dataset of monthly village-level nighttime luminosity measures developed

pre-treatment period.

[2] Given the geographical spread of the guar shock across many districts in northwestern India, our analysis relies on an extension to this basic approach developed by Cavallo et al. (2013), who generalize the application of SCM to multiple treated units possibly at different time periods.

[3] NOAA's annual composite images are available at `https://ngdc.noaa.gov/eog/dmsp/downloadV4composites.html`.

by Gaba et al. (2016), who revisit the complete archive of raw visible band (VIS) imagery captured during every night in India between 1993 and 2013 to generate each observation. Because the DMSP includes multiple satellites, this archive consists of approximately 30,000 high-resolution image strips. Brightness values are extracted from these images for each date from each pixel corresponding to the latitude and longitude of each of India's approximately 600,000 villages. These values are processed in line with NOAA recommendations to remove irregularities, and the resulting 4.4 billion observations are aggregated to the village-month level by taking the median measurement for each village over the course of a month. In addition, because the 0–63 relative brightness levels in the raw data are not directly comparable over time, additional image processing and background noise reduction procedures are applied to generate statistically recalibrated visible band (SR-VIS) measures, which enable more reliable comparisons both cross-sectionally and across time.[4]

We use these data to evaluate differential impacts of the fracking-induced guar boom on nighttime luminosity—and, by proxy, economic activity—across guar- and non-guar-growing regions of India. Because we identify guar-growing regions of India at the district level, in our analysis we rely on district-month measures of nighttime brightness.[5]

## A.3 Results

We specify a parsimonious predictive model of nighttime luminosity, namely, one in which nighttime luminosity in district $d$ in year $t$ is a function of nighttime luminosity in year $t-1$.[6] Figure A.1 presents our main results. The solid line highlights the trend

[4] Min et al. (2017)—who use SR-VIS data to study power-supply irregularity across rural India—describe these image-processing procedures in more detail. The data are available at `http://api.nightlights.io/`.

[5] Gaba et al. (2016) determine these by identifying the median village light output within each district boundary for each month.

[6] Prior applications of the SCM have often used contemporaneous or lagged values of the outcome variable for all units $j'$ as the sole predictor in estimation of treatment effects for unit $j$ (e.g., Acemoglu et al., 2016). The justification for this approach is that the outcome variable fully characterizes all observed and

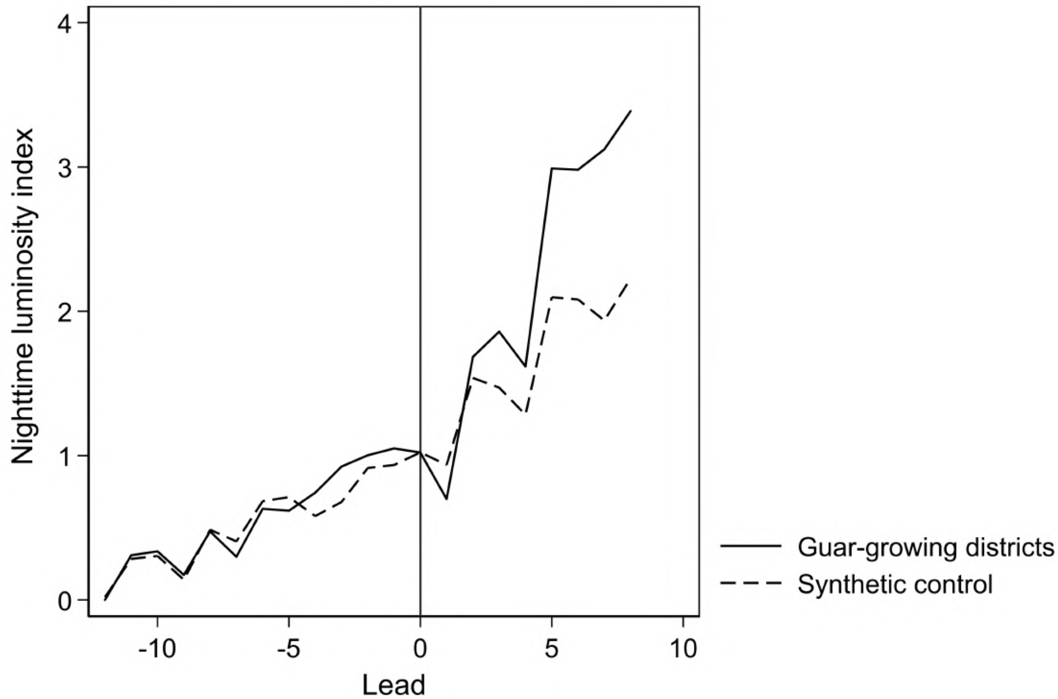


FIGURE A.1: Pre-/post-guar-boom trends in nighttime luminosity in guar-growing districts. This figure presents results from a synthetic control approach to evaluate the impact of the start of the fracking-induced guar boom in India on nighttime luminosity in India's guar-growing districts (as shown in Figure 1.1). The outcome variable is an index of nighttime luminosity, aggregated to the district-year level from the village-month level. The fracking-induced guar boom is assumed to begin in 2006, indicated by the vertical line. Other years (covering the period 1993-2013) are presented as leads and lags relative to 2006.

in mean monthly nighttime brightness for India's guar-growing districts. The vertical line represents the start of the fracking boom in the United States (assumed to be 2006). The dashed line represents mean monthly nighttime brightness for a "counterfactual" set of guar-growing districts (unaffected by the fracking-induced guar boom). As described earlier, this is generated by estimating a set of weights for monthly nighttime brightness data for all other Indian districts over the pre-fracking-boom period (1993-2005) that are used to most closely track pre-boom—and predict post-boom—nighttime brightness

unobserved determinants.

Table A.1: Impact of fracking-induced guar boom on nighttime luminosity in Rajasthan

| (1) Year | (2) Estimated coefficient | (3) $p$-value |
|---|---|---|
| 2007 | -0.24*** | 0.0006 |
| 2008 | 0.15 | 0.58 |
| 2009 | 0.39** | 0.04 |
| 2010 | 0.33** | 0.01 |
| 2011 | 0.89 | 0.14 |
| 2012 | 0.90** | 0.03 |
| 2013 | 1.19*** | 0.004 |

This table presents the estimated effect of the fracking-induced guar boom on nighttime luminosity in India's guar-growing districts (relative to a synthetically generated set of guar-growing districts) for each post-boom year (column 2). Column (3) presents $p$-values associated with each estimated coefficient, obtained by adjusting the observed effect sized by the pre-treatment match quality as outlined by Cavallo et al. (2013). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

trends in the guar-growing areas. The divergence in the two lines in the post-boom period is stark, and suggests that the start of fracking-induced boom resulted in sizable increases in nighttime brightness—and, by extension, economic activity in India's guar-growing regions. Indeed, $p$-values estimated year-by-year using placebo tests for each post-boom year indicate that by 2011, the probability of this increased economic activity being detectable from space in this way by chance is extremely low (Table A.1).

# Appendix B

# Home production and labor supply

The Lagrangian associated with the household's problem described in Section 1.3.1 is as follows:

$$\max_{c_i,l_i} \mathcal{L} = u\left(c\left(t_i^h, x_i, v_i; \psi_i\right), t_i^l\right) + \lambda\left(w_i T + v_i - x_i - w_i\left(t_i^h + t_i^l\right)\right). \tag{B.1}$$

Ignoring the $i$ subscripts, this yields the following first-order conditions for an interior solution:

$$\mathcal{L}_{t^l} = u_{t^l} - \lambda w = 0 \tag{B.2}$$

$$\mathcal{L}_{t^h} = u_c c_{t^h} - \lambda w = 0 \tag{B.3}$$

$$\mathcal{L}_x = u_c c_x - \lambda = 0 \tag{B.4}$$

$$\mathcal{L}_\lambda = wT + v - x - w\left(t^h + t^l\right) = 0. \tag{B.5}$$

These first-order conditions indicate that household's time allocations are chosen to equate the marginal rate of substitution between leisure and consumption with (i) the shadow value of home production; and (ii) the shadow value of market-based activities. Specifically, from Equations (B.2), (B.3) and (B.4):

$$\frac{u_{t^l}}{u_c} = c_{t^h} = c_x w. \tag{B.6}$$

From this, the general form of the household's optimal time allocation to home production is obtained:

$$t^{h*} = f_{t^h}\left(w, v; \psi\right). \tag{B.7}$$

Equations (B.2), (B.4) and (B.5) can be solved jointly to obtain the household's optimal time allocation to leisure and its demand for the market-purchased home-production input:

$$t^{l*} = f_{t^h}\left(w, v; \psi\right) \tag{B.8}$$

$$x^{*} = f_{x}\left(w, v; \psi\right). \tag{B.9}$$

Equation (B.9) and Equation (B.7) combined with the household's consumption production function yield the household's optimal consumption:

$$c^{*} = c\left(t^{h*}, x^{*}, v; \psi_i\right). \tag{B.10}$$

Finally, combining the household's time constraint with Equations (B.7) and (B.8) yields the household's time allocation to market-based activities:

$$t^{m*} = T - t^{h*} - t^{l*} = f_{t^m}\left(w, v; \psi\right). \tag{B.11}$$

# Appendix C

## Habitation-Village matching procedure

We use a multi-step matching procedure to identify villages eligible for electrification under RGGVY Phase I based on the populations of their constituent habitations, and identify corresponding village names from the 2001 and 2011 Census to those in the 2009 census of habitations conducted by the National Rural Drinking Water Program (NRDWP). The NRDWP habitation census covers 1.65 million habitations in 574,259 villages.[1] Because the NRDWP survey indicates only the name of each village (and not its unique Census code), matching on names is necessary; however, not all village names match exactly between the names used in NRDWP and those used in the Census, even conditional on an exact match for state and district. Accordingly, our matching process incorporates a combination of exact and fuzzy name matches, prioritizing exact matches where possible.

We treat the 2001 Primary Census Abstract (PCA) villages as the master dataset.

[1] This includes five of the seven Union Territories—Chandigarh, Dadra and Nagar Haveli, Daman and Diu, Lakshadweep, and Puducherry—and Goa. However, we exclude these from the merge process because Goa and all seven Union Territories were fully electrified prior to 2005, so were excluded in RGGVY (Ministry of Power, 2012). Excluding the seven Union Territories and Goa, the 2009 survey covers 1.65 million habitations in 573,702 villages.

As a first step for matching village names with the 2009 NRDWP habitations data, we standardize state, district, block, and village names to correct minor differences in spelling between the names in use by the NRDWP and the Census. We also account for districts that were renamed between 2001 and 2009. Our procedure for standardizing state and district names is sufficiently comprehensive to achieve a 100 percent match among state and district names between the NRDWP and Census, except for a handful of cases where districts are split or combined (not just renamed) between 2001 and 2009.[2]

We use information from the state, district, block, and village level, and prioritize exact matches. Where exact name matches are not possible, we employ a fuzzy match, using the "masalafied Levenshtein" distance and "Masala merge" code in Stata and Python (Asher and Novosad, 2018). This is a modification of the standard Levenshtein string distance metric, one that lowers the cost of certain substitutions that are common in Indian languages.[3] We thus create a five-tier matching hierarchy:

1. Exact match on state, district, block, and village name;
2. Exact match on state, district, and village name, with a fuzzy match on block name;
3. Exact match on state and district name, with a fuzzy match on block and village name;
4. Exact match on state, district, and village name (without regard to block name); and

[2] One approach to match villages in split or combined districts would be to geolocate all villages from the old district(s) into the new district(s). We take a somewhat less intensive approach and look for name-based village matches in a proper subset of the old or new district area—specifically, an area of known overlap between old and new. For instance, Tiruppur district in Tamil Nadu was formed in 2009 from parts of Coimbatore and Erode. Among villages in the NRDWP belonging to Tiruppur district, we look for matching Census village names within Erode district, but not within Coimbatore district. We also flag any matches associated with split or combined districts. We have run our matching algorithm excluding these flagged matches and, after completing all five steps of the multi-step procedure, achieved virtually identical results.

[3] Additional information about Masala merge (including its underlying code) is available at `http://www.dartmouth.edu/~novosad/code.html`.

5. Exact match on state and district name, with a fuzzy match on village name (without regard to block name).

Of the 563,338 villages in the 2001 Census, we match 531,325 to villages in the NRDWP habitation census (94.3 percent). This includes 400,457 exact matches (71 percent), of which 271,774 (48 percent) are exact matches on state, district, block, and village name.[4] Further, our algorithm achieves a 90 percent or greater match rate across every state with the exception of Tripura (36 percent), Tamil Nadu (76 percent), Jammu and Kashmir (78 percent), Nagaland (82 percent), and Assam (83 percent). We also match at least 96 percent of villages in each of the three northwestern states where guar is produced (98 percent in Rajasthan and Gujarat, and 96 percent in Haryana).

As a further verification step, we compare the village population recorded by the NRDWP in 2009 to the village population recorded by the 2011 PCA. For any village name match in which these figures diverge by more than 20 percent, we exclude the village from the matched set.[5] Using this matched sample, we identify single-habitation villages, and use the population of each of these in the 2001 round of the Census to gauge its eligibility for electrification under RGGVY Phase I.

[4] Our match rate is comparable to others in the literature. For instance, Burlig and Preonas (2016) report matching 86 percent of villages from the 2003 and 2009 NRDWP habitation surveys to corresponding Census villages. While Asher and Novosad (2018) do not report a village-level match rate, they do indicate they matched over 85 percent of habitations listed in the PMGSY to corresponding Census villages. Aggarwal (2018), who also evaluates the impact of India's rural roads program, reports a match rate of 80 percent.

[5] We have also run our analysis using thresholds other than 20 percent and find substantially similar results (Figure H.1).

# Appendix D

## Redefining “NGO village”

There is often considerable spatial heterogeneity in the scale and scope of an NGO’s activities. The same NGO can be deeply invested in the welfare of one particular community while at the same time only superficially involved with another. If this is the case in our setting, our characterization of NGO and non-NGO villages using a binary variable may be too crude. We, therefore, turn to two additional ways of defining the level of an NGO’s involvement with each of our study villages: (1) the number of activities it is leading in a particular community; and (2) the number of years since it first began operating there.

We obtain information on spatial variation in the NGO’s portfolio of activities through reviews of its annual reports, newsletters, and other promotional materials.[1] Specifically, we rely on these materials to identify which activities occur in which villages, and when the NGO’s operations first began there. This presents challenges as these promotional documents are typically not sufficiently detailed to allow us to comprehensively construct both measures. Recall that our sample contains 38 villages, of which half are NGO villages.

[1] Recall that the NGO leads activities related to agriculture and forestry (promotion of sustainable agricultural practices, sustainable fodder cultivation, and promotion of culinary herbs), health (local hospitals/clinics), education (local schools), village-level groups (self-help groups, youth groups, and vocational cooperatives), and water management (watershed renewal and spring-water recharge).

In total, we are able to construct a detailed breakdown of the NGO's activities in 12 of its 19 villages; we are able to identify the NGO's commencement year in an equal number of villages. A count of projects among these villages with non-missing implementation data reveals that the NGO leads just over four initiatives (with a minimum of one and maximum of seven) in each of its villages. Similarly, immediately prior to the start of the intervention, the NGO has been operating for just over 15 years in the average village, ranging from four years in the newest village to 25 years in the oldest one.

Using these additional measures (i.e. a count of the number of the NGO's active projects, and the overall age of its engagement in each village) we can investigate heterogeneity in purchase of intervention ICS—now across villages with relatively different "intensities" of NGO activity. We account for uncertainty introduced by missing data through a simulation-based bootstrap. Specifically, let $\mathbf{NGO} = \mathbf{NGO}_{\text{obs}} \cup \mathbf{NGO}_{\text{miss}}$, where $\mathbf{NGO}$ represents our data (on the village-specific count of NGO projects or the age of its engagement there) and $\mathbf{NGO}_{\text{obs}}$ and $\mathbf{NGO}_{\text{miss}}$ represent non-overlapping observed and missing components of it, respectively. For each bootstrap simulation $n \in N$, we then proceed as follows:

1. Randomly generate:

   (a) $\mathbf{NGO}^{\text{count},n}_{\text{miss}} \overset{\text{iid}}{\sim} \text{unif}\,(0, 7)$

   (b) $\mathbf{NGO}^{\text{age},n}_{\text{miss}} \overset{\text{iid}}{\sim} \text{unif}\,(0, 25)$

2. Construct:

   (a) $\mathbf{NGO}^{\text{count},n} = \mathbf{NGO}_{\text{obs}} \cup \mathbf{NGO}^{\text{count},n}_{\text{miss}}$

   (b) $\mathbf{NGO}^{\text{age},n} = \mathbf{NGO}_{\text{obs}} \cup \mathbf{NGO}^{\text{age},n}_{\text{miss}}$

3. Randomly sample hamlets (with replacement) and estimate the specification outlined in Equation (2.19):

(a)

$$Y_{ij}^{\text{count},n} = \beta_0^{\text{count},n} + \beta_1^{\text{count},n}\left(TREATMENT_j\right) + \beta_2^{\text{count},n}\left(\mathbf{NGO}_j^{\text{count},n}\right) + \\ + \beta_3^{\text{count},n}\left(TREATMENT_j \times \mathbf{NGO}_j^{\text{count},n}\right) + v_{ij}$$

(b)

$$Y_{ij}^{\text{age},n} = \beta_0^{\text{age},n} + \beta_1^{\text{age},n}\left(TREATMENT_j\right) + \beta_2^{\text{age},n}\left(\mathbf{NGO}_j^{\text{age},n}\right) + \\ + \beta_3^{\text{age},n}\left(TREATMENT_j \times \mathbf{NGO}_j^{\text{age},n}\right) + v_{ij}$$

4. Save the estimated regression coefficients:

   (a) $\hat{\beta}_3^{\text{count},n}$

   (b) $\hat{\beta}_3^{\text{age},n}$

In other words, for villages lacking data on the count of projects, we replace missing count observations with random draws from a uniform distribution over the interval lying between the minimum (i.e. zero in non-NGO villages) and maximum (seven) number of projects observed in each village in our data. For villages lacking data on the age of NGO engagement, we similarly replace missing observations with random draws from a uniform distribution over the interval lying between the minimum (zero in non-NGO villages) and maximum (25) ages in our data. Having replaced these missing observations, we sample hamlets (with replacement) from our study sample to generate a bootstrapped sample, and separately estimate the specification outlined in Equation (2.19) by replacing our NGO binary variable with the count of NGO activities or the age of NGO engagement. For each specification, we repeat this process 10,000 times to obtain a distribution of the estimated coefficient for the $TREATMENT_j \times NGO_j$ interaction term. We note that this is a relatively conservative approach to dealing with the uncertainty surrounding missing observations. It is almost certainly the case that the NGO has active projects in each of its villages, and has operated in them for at least a few years. Nevertheless, the uniform distributions we use to replace missing observations have positive support over zero (the value these variables are assigned for non-NGO villages). In addition, our randomly generated replacements never exceed the maximum value observed in the (non-missing) data.

Figure D.1 presents our results. As shown in panel (a), we find that purchase rates by households in treated NGO hamlets are, on average, approximately 2 percentage points higher for every additional project that the NGO leads in that particular village. This result is statistically significant at the 10 per cent level, as measured by the 90 per cent confidence interval of the distribution of our simulated regression coefficients. Similarly, panel (b) shows that every additional year of the NGO's presence in a village resulted in an increase in rates of ICS purchase by households in treated NGO hamlets by just under 1 percentage point. Broadly, these results serve as a robustness check for our main result (Table 2.3). They are also consistent with our model and show that NGO activity—defined in a variety of ways—influences the effectiveness of interventions.

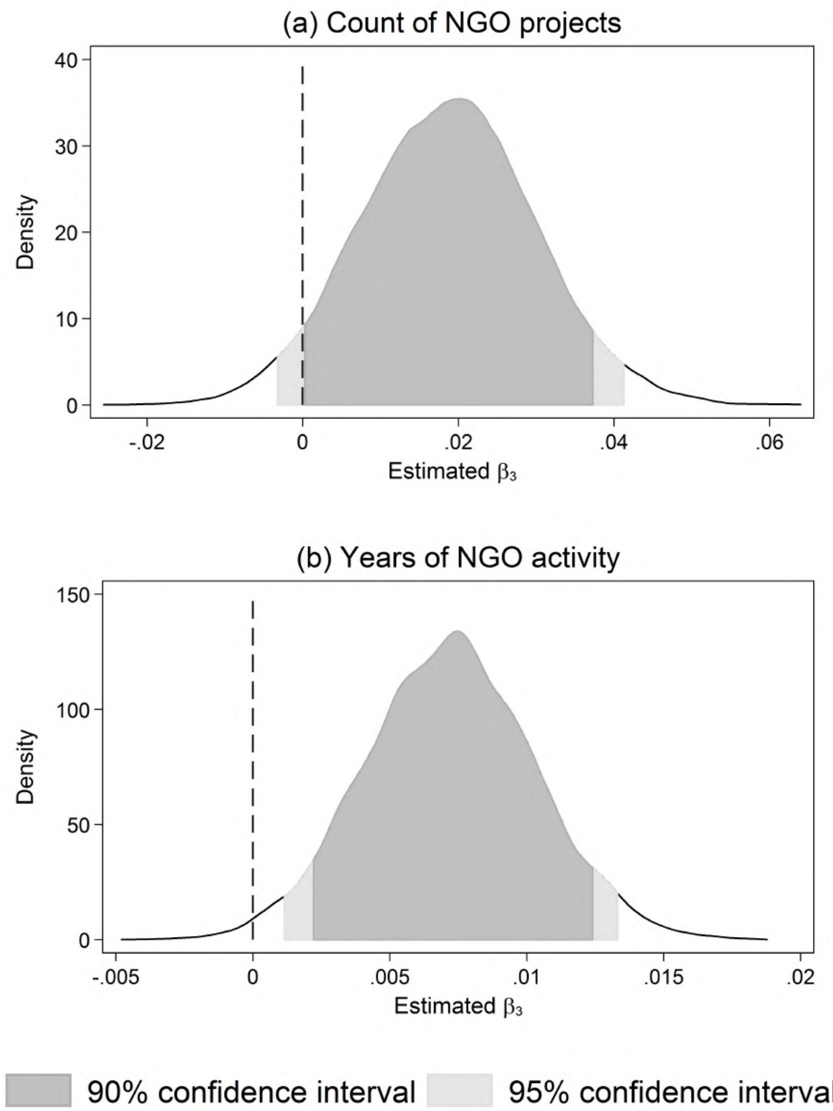


Figure D.1: Heterogeneity in ICS purchase rates based on alternative definitions of "NGO village." This figure plots the distribution of 10,000 $\beta_3$ coefficients obtained from estimating Equation (2.19) using a simulation-based bootstrap approach. This coefficient represents the additional impact of the ICS promotion intervention on ICS purchase rates in treated hamlets located in villages with relatively higher levels of NGO activity. In panel (a), the NGO activity variable is a count of the number of active projects being implemented by the NGO in each village; in non-NGO villages, this variable equals zero. In panel (b), the NGO activity variable is the number of years since the NGO first began operating in each village; once again, in non-NGO villages, this variable equals zero. The outcome variable is an indicator that equals one if household $i$ in hamlet $j$ purchased at least one of the two ICS promoted during the intervention.

# Appendix E

## Bayesian analysis of "NGO effect"

We complement our main analyses using Bayesian methods. Standard frequentist approaches assume that underlying statistical parameters are fixed. Conditional on these fixed parameters, the data are one realization of infinitely many samples, and can be combined with assumptions about large-sample approximations (e.g., asymptotic normality) for inference. In contrast, Bayesian techniques assume that the true values of parameters are random variables, and assign distributions to these parameters based on a-priori information. Conditional on the observed data and the specified prior distribution, a posterior distribution for each parameter can be estimated and used for inference.

*E.0.1 Likelihood*

We specify our likelihood function for the data as follows:

$$Y_{ij} \sim \mathcal{N}\left(\beta_0 + \beta_1\left(TREATMENT_j\right) + \beta_2\left(NGO_j\right) + \beta_3\left(TREATMENT_j \times NGO_j\right) + \sum_{j=1}^{97} u_j \gamma_j, \sigma_v^2\right), \tag{E.1}$$

where $Y_{ij}$ is a binary variable that equals one if household $i$ in hamlet $j$ purchased at least one of the two intervention ICS offered during intervention activities and zero if it did not;

$TREATMENT_j$ is a binary variable that equals one if hamlet $j$ is randomly assigned to the treatment group and zero if it is assigned to the control group; and $NGO_j$ is a binary variable that equals 1 if hamlet $j$ is located in an NGO village and zero if it is in a non-NGO village. In addition, we include hamlet-specific random effects, represented by each of the $\gamma_j$ terms with coefficients $u_j$.

### *E.0.2 Priors*

*Model I: Informative prior for $\beta_3$ and diffuse priors for other parameters*

$$\beta_0 \sim \mathcal{N}(0, 10000)$$

$$\beta_1 \sim \mathcal{N}\left(\hat{\beta}_1^{\text{OLS}}, 10000\right)$$

$$\beta_2 \sim \mathcal{N}\left(\hat{\beta}_2^{\text{OLS}}, 10000\right)$$

$$\beta_3 \sim \text{Beta}(2, 14)$$

$$u_j \sim \mathcal{N}\left(0, \sigma_\gamma^2\right)$$

$$\frac{1}{\sigma_\nu^2} \sim \Gamma(0.01, 0.01)$$

$$\frac{1}{\sigma_\gamma^2} \sim \Gamma(0.01, 0.01)$$

*Model II: Diffuse priors for all parameters*

$$\beta_0 \sim \mathcal{N}\left(0, 10000\right)$$

$$\beta_1 \sim \mathcal{N}\left(\hat{\beta}_1^{\text{OLS}}, 10000\right)$$

$$\beta_2 \sim \mathcal{N}\left(\hat{\beta}_2^{\text{OLS}}, 10000\right)$$

$$\beta_3 \sim \mathcal{N}\left(\hat{\beta}_3^{\text{OLS}}, 10000\right)$$

$$u_j \sim \mathcal{N}\left(0, \sigma_\gamma^2\right)$$

$$\frac{1}{\sigma_\nu^2} \sim \Gamma\left(0.01, 0.01\right)$$

$$\frac{1}{\sigma_\gamma^2} \sim \Gamma\left(0.01, 0.01\right)$$

*Model III: Strong no-NGO-effect prior for $\beta_3$ and diffuse priors for other parameters*

$$\beta_0 \sim \mathcal{N}\left(0, 10000\right)$$

$$\beta_1 \sim \mathcal{N}\left(\hat{\beta}_1^{\text{OLS}}, 10000\right)$$

$$\beta_2 \sim \mathcal{N}\left(\hat{\beta}_2^{\text{OLS}}, 10000\right)$$

$$\beta_3 \sim \mathcal{N}\left(0, 0.01\right)$$

$$u_j \sim \mathcal{N}\left(0, \sigma_\gamma^2\right)$$

$$\frac{1}{\sigma_\nu^2} \sim \Gamma\left(0.01, 0.01\right)$$

$$\frac{1}{\sigma_\gamma^2} \sim \Gamma\left(0.01, 0.01\right)$$

### *E.0.3 Results*

Posterior distributions of the parameters in the two models are estimated via Markov Chain Monte Carlo (MCMC) simulation. Specifically, we ran 50,000 MCMC samples after

Table E.1: Markov Chain Monte Carlo results

| | (1) Posterior mean | (2) Posterior standard deviation | (3) 95% credible interval |
|---|---|---|---|
| **(a) Model I** | | | |
| $TREATMENT_j$ | 0.46 | 0.055 | [0.34, 0.56] |
| $NGO_j$ | 0.018 | 0.063 | [−0.11, 0.14] |
| $TREATMENT_j \times NGO_j$ | 0.12 | 0.062 | [0.02, 0.26] |
| Constant | −0.007 | 0.048 | [−0.10, 0.09] |
| **(b) Model II** | | | |
| $TREATMENT_j$ | 0.44 | 0.064 | [0.32, 0.57] |
| $NGO_j$ | 0.0004 | 0.082 | [−0.17, 0.16] |
| $TREATMENT_j \times NGO_j$ | 0.15 | 0.097 | [−0.04, 0.34] |
| Constant | 0.0002 | 0.053 | [−0.10, 0.11] |
| **(c) Model III** | | | |
| $TREATMENT_j$ | 0.48 | 0.056 | [0.37, 0.59] |
| $NGO_j$ | 0.05 | 0.066 | [−0.08, 0.19] |
| $TREATMENT_j \times NGO_j$ | 0.08 | 0.069 | [−0.06, 0.21] |
| Constant | −0.02 | 0.048 | [−0.12, 0.07] |

The outcome variable is an indicator that equals one if household $i$ in hamlet $j$ purchased at least one of the two ICS promoted during the intervention. Columns (1) and (2) present the mean and standard deviations, respectively, for the MCMC sample. Column (3) presents the 95 percent credible interval for the MCMC sample. Estimates for 97 hamlet-specific random effects not reported for brevity. As in Table 2.3, $N = 943$ households.

a burn-in period of 10,000 iterations, with thinning every fifth iteration. Table E.1 presents our results.

# Appendix F

# Second-price, sealed-bid (“Vickrey”) auction script

This appendix includes English translations of the script used by field teams to explain the mechanisms of the second-price, sealed-bid (“Vickrey”) auctions for the *Jambar* and *Jumbo Zama* stoves in villages in the long survey arm. Field teams used similar scripts to explain the auction mechanism in villages in the short survey arm.

## F.1 Auction Arms “A,” “B” and “C”

I will now offer you the opportunity to participate in an auction for a [**NAME OF STOVE**]. First, I will give you a little demonstration of the stove, after that I will explain the rules of the sale.

[**CONDUCT STOVE DEMONSTRATION**]

Now, for the sale of the [**NAME OF STOVE**] I am going to ask you to write your bid—the maximum price you would pay to buy the [**NAME OF STOVE**]—on this auction sheet, which will contain your name and your proposed bid. All households surveyed in your village will be invited to submit their bids as well. All bids will be put in this box. Once all the bids are together, we will open the box and make a public reading of the proposals. The household that writes the highest bid will win the auction. The winning household will be able to buy the [**NAME OF STOVE**] at a price equal to the second-highest proposed in the box. For example, suppose you offer CFA 50,000 for the stove and it is the highest proposal in the village. Assume the second best proposal is CFA 40,000. In this case you would win and you would pay CFA 40,000 for the stove. If you think carefully about this method, I think you will see that you should bid the maximum price you would be willing to pay. If you offer less than your maximum, you could lose the opportunity to get the stove. If you offer more than you can afford, you could win the auction and have to pay more than you would like.

Did you understand the rules of the auction?

[0] No [1] Yes [**If "no," explain and ask again, do not start until respondent answers "Yes."**]

How can a household win the auction?

[0] Incorrect answer [1] Correct answer (By having the highest bid for the stove on the auction sheet—or a similar response.)

D3. What price does the winning household have to pay for the item?

[0] Incorrect answer [1] Correct answer (The second-highest price in the box—or a similar response.)

[**If the answer is incorrect, explain and ask again, do not start until the answer is correct.**]

D4. What is your bid for the [**NAME OF STOVE**]? [**Write the proposed amount on the paper, fold it and put it in the auction box.**]

## F.2 Auction Arm "D"

I will now offer you the opportunity to participate in an auction for a *Jambar* stove and a *Jumbo Zama* stove. First, I'll give you a little demonstration of the stove, after that I'll explain the rules of the sale.

**[CONDUCT STOVE DEMONSTRATION]**

Now, for the sale of these stoves, I am going to ask you to write your bid—the maximum price you would pay to buy the *Jambar* and/or *Jumbo Zama* stoves—on this auction sheet, which will contain your name and your proposed bids. All households surveyed in your village will be invited to submit their proposal as well. All bids will be put in this box. Once all the bids are together, we will open the box and make a public reading of the proposals. The household that writes the highest bid for the *Jambar* stove, will win the auction of the *Jambar* stove. The household that writes the biggest proposal for the *Jumbo Zama* stove, will win the auction of the *Jumbo Zama* stove. The winning household(s) will be able to purchase their respective stoves at a price equal to the second-highest price for the respective stove. For example, suppose that you offer CFA 50,000 for this stove and that it is the highest proposal in the village. Assume the second best proposal is CFA 40,000. In this case you would win and you would pay CFA 40,000 for this stove. If you think carefully about this method, I think you should bid the maximum price you be willing to pay for each stove. If you offer less than your maximum, you could lose the opportunity to get the stove. If you offer more than you can afford, you could win the auction and have to pay more than you would like. (If the same household wins both auctions, the household will have the choice to buy either one or both stoves at a price equal to the second-highest bid for each stove.)

Did you understand the rules of the auction?

[0] No [1] Yes [**If "no," explain and ask again, do not start until respondent answers "Yes."**]

How can a household win the auction?

[0] Incorrect answer [1] Correct answer (By having the highest bid for the stove on the auction sheet—or a similar response.)

D3. What price does the winning household have to pay for the item?

[0] Incorrect answer [1] Correct answer (The second-highest price in the box—or a

similar response.)

**[If the answer is incorrect, explain and ask again, do not start until the answer is correct.]**

D4. What is your bid for the *Jambar* stove? **[Write the proposed amount on the paper, fold it and put it in the auction box.]**

D5. What is your bid for the *Jumbo Zama* stove? **[Write the proposed amount on the paper, fold it and put it in the auction box.]**

# Appendix G

## Additional tables

Table G.1: Testing for discontinuous changes at RGGVY Phase I threshold in 2001

| Outcome variable (2001) | (1) $\mathbb{1}$ (Village pop. (2001) > 300) Coef. | (2) Std. Err. | (3) $N$ | (4) Adj. $R^2$ | (5) Mean of outcome |
|---|---|---|---|---|---|
| Number of households | -0.08 | (61.96) | 7649 | 0.64 | 53.97 |
| Females (% of population) | -0.01 | (16.43) | 7649 | 0.28 | 48.73 |
| Ages 0–6 (% of population) | 0.04 | (35.86) | 7649 | 0.36 | 17.78 |
| Scheduled Caste/Tribe (% of population) | -0.57 | (338.47) | 7649 | 0.28 | 36.02 |
| Literate (% of population) | -0.01 | (6.49) | 7649 | 0.36 | 45.01 |
| All workers (% of population) | -1.00 | (1.93) | 7649 | 0.38 | 43.98 |
| Agricultural workers (% of population) | -0.38 | (228.06) | 7649 | 0.32 | 37.22 |
| Non-agricultural workers (% of population) | -0.62 | (2.87) | 7649 | 0.15 | 6.76 |
| Area (Hectares) | -14.02 | (101.07) | 7649 | 0.36 | 158.07 |
| Irrigated area (% of total area) | -0.65 | (387.19) | 7324 | 0.40 | 35.67 |
| Primary schools (per 1,000 people) | -0.10 | (0.40) | 7649 | 0.27 | 1.97 |
| Community health workers (per 1,000 people) | 0.05 | (0.22) | 7649 | 0.10 | 0.20 |
| $\mathbb{1}$ (Bus facilities) | 0.01 | (4.77) | 7649 | 0.22 | 0.17 |
| $\mathbb{1}$ (Postal facilities) | 0.02 | (0.13) | 7649 | 0.15 | 0.18 |
| $\mathbb{1}$ (Approach: Paved road) | 0.00 | (4.37) | 7649 | 0.10 | 0.37 |
| $\mathbb{1}$ (Power supply) | 0.03 | (0.08) | 7649 | 0.35 | 0.66 |

Column (1) reports the value of $\hat{\beta}_1$ obtained from estimating the following regression specification on our main analytical sample of single-habitation villages located in RGGVY Phase I districts: $y_{vds}^{2001} = \beta_0 + \beta_1 T_{vds} + \beta_2 \tilde{P}_{vds}^{2001} + \beta_3 T_{vds} \tilde{P}_{vds}^{2001} + \gamma_d + \epsilon_{vds}$, where $y_{vds}^{2001}$ represents an outcome variable for village $v$ in district $d$ in state $s$ in 2001, $T_{vds}$ is a binary variable that equals one if the population of village $v$ in 2001 is greater than 300, $\tilde{P}_{vds}^{2001}$ is the population running variable, and $\gamma_d$ represents a district fixed-effect. Standard errors—in column (2)—are clustered at the district level and inferred from $p$-values obtained using the free step-down resampling methodology of Westfall and Young (1993). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table G.2: RD estimates of impact of electrification on total population in 2011

| | | (1) | (2) | (3) |
|---|---|---|---|---|
| | | Total population (2011) | | |
| | | All | Male | Female |
| $\hat{\beta}_1$ | $\mathbb{1}$ (Village pop. (2001) > 300) | 3.39** | 2.22** | 1.16 |
| | | (1.70) | (0.94) | (0.91) |
| $\hat{\beta}_2$ | $\mathbb{1}$ (Village pop. (2001) > 300) × | 2.27 | 6.53* | -4.45 |
| | $\mathbb{1}$ (Village in guar-growing district) | (6.29) | (3.77) | (2.91) |
| District FEs | | Yes | Yes | Yes |
| Census (2001) controls | | Yes | Yes | Yes |
| $N$ | | 7649 | 7649 | 7649 |
| Adjusted $R^2$ | | 0.54 | 0.54 | 0.49 |
| Mean of outcome | | 349.24 | 178.92 | 170.32 |

This table shows results from estimating Equation (1.9). Each regression includes all single-habitation villages in RGGVY Phase I districts with a 2001 population within a fifty-person bandwidth of RGGVY's 300-person eligibility threshold. Estimates associated with the population running variable $\left(\tilde{P}_{vds}^{2001}\right)$ are omitted. Following Correia (2015), six singleton observations are excluded. Standard errors—in parentheses—are clustered at the district level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table G.3: Differences between villages in guar and non-guar-growing districts in 2001

| Outcome variable (2001) | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | All RGGVY Phase I villages | | | RD sample villages | | |
| | Non-guar | Guar | $p$-value of difference | Non-guar | Guar | $p$-value of difference |
| Total population | 1390.86 | 1502.61 | 0.056* | 299.99 | 306.22 | 0.284 |
| | (1654.30) | (1455.01) | | (29.18) | (28.44) | |
| Number of households | 247.03 | 231.91 | 0.142 | 54.07 | 48.87 | 0.385 |
| | (316.25) | (225.51) | | (11.68) | (8.78) | |
| Females (% of population) | 48.62 | 48.26 | 0.940 | 48.74 | 47.89 | 0.484 |
| | (2.91) | (2.40) | | (3.04) | (2.85) | |
| Age 0–6 (% of population) | 18.07 | 19.87 | 0.940 | 17.75 | 19.09 | 0.972 |
| | (4.17) | (3.47) | | (4.57) | (4.16) | |
| Scheduled Caste/Tribe (% of population) | 31.65 | 27.24 | 0.599 | 36.33 | 20.77 | 0.285 |
| | (27.65) | (22.97) | | (34.78) | (26.04) | |
| Literate (% of population) | 44.74 | 47.00 | 0.462 | 44.94 | 48.81 | 0.434 |
| | (14.48) | (13.35) | | (16.33) | (13.87) | |
| Total workers (% of population) | 41.61 | 46.14 | 0.219 | 43.91 | 47.12 | 0.434 |
| | (12.86) | (10.32) | | (14.14) | (12.05) | |
| Agricultural workers (% of population) | 33.79 | 37.82 | 0.847 | 37.17 | 39.94 | 0.742 |
| | (14.10) | (13.32) | | (15.27) | (13.96) | |
| Non-agricultural workers (% of population) | 7.81 | 8.31 | 0.729 | 6.75 | 7.18 | 0.972 |
| | (7.57) | (7.59) | | (7.89) | (9.32) | |
| Area (Hectares) | 358.69 | 1428.10 | 0.219 | 148.41 | 648.16 | 0.486 |
| | (756.26) | (2316.45) | | (224.87) | (1161.81) | |
| Irrigated area (% of total area) | 38.36 | 21.09 | 0.940 | 35.97 | 21.21 | 0.972 |
| | (33.84) | (25.04) | | (33.69) | (27.24) | |
| Primary schools (per 1,000 people) | 1.27 | 1.45 | 0.628 | 1.95 | 3.01 | 0.678 |
| | (2.24) | (6.64) | | (1.81) | (1.06) | |
| Community health workers (per 1,000 people) | 0.15 | 0.10 | 0.940 | 0.20 | 0.11 | 0.910 |
| | (1.10) | (0.60) | | (0.83) | (0.59) | |
| $\mathbb{1}$ (Bus facilities) | 0.27 | 0.60 | 0.092* | 0.17 | 0.32 | 0.393 |
| | (0.44) | (0.49) | | (0.37) | (0.47) | |
| $\mathbb{1}$ (Postal facilities) | 0.40 | 0.64 | 0.219 | 0.18 | 0.36 | 0.678 |
| | (0.49) | (0.48) | | (0.38) | (0.48) | |
| $\mathbb{1}$ (Approach: Paved road) | 0.53 | 0.60 | 0.219 | 0.37 | 0.34 | 0.434 |
| | (0.50) | (0.49) | | (0.48) | (0.48) | |
| $\mathbb{1}$ (Power supply) | 0.73 | 0.89 | 0.940 | 0.65 | 0.84 | 0.972 |
| | (0.45) | (0.31) | | (0.48) | (0.36) | |
| $N$ | 182051 | 6232 | | 7507 | 148 | |

This table reports mean and standard deviations (in parentheses) for villages located in guar- and non-growing districts of India. Columns (1) and (2) report these values for our full sample of habitation-matched villages in RGGVY Phase I districts; column (4) and (5) report these values for our main analytical sample of single-habitation villages. Columns (3) and (6) report the $p$-value for $\hat{\beta}_1$ obtained from estimating the following regression specification on the relevant sample: $y_{vds}^{2001} = \beta_0 + \beta_1 G_{ds} + \gamma_s + \epsilon_{vds}$, where $y_{vds}^{2001}$ represents an outcome variable for village $v$ in district $d$ in state $s$ in 2001, $G_{ds}$ is a binary variable that equals one if village $v$ is located in a guar-growing district, and $\gamma_s$ represent a state fixed-effect. Standard errors (not shown) are clustered at the district level; $p$-values are obtained using the free step-down resampling methodology of Westfall and Young (1993). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table G.4: Placebo RD estimates of impact of electrification on labor-market outcomes

| | | (1) | (2) | (3) | (4) |
|---|---|---|---|---|---|
| | | All workers | Ag. workers | Non-ag. workers | Non-workers |
| | | (% of 2011 population) | | | |
| $\hat{\beta}_1$ | $\mathbb{1}$ (Village pop. (2001) > 300) | 0.29 | -0.18 | 0.55 | -0.29 |
| | | (0.56) | (0.71) | (0.43) | (0.56) |
| $\hat{\beta}_2$ | $\mathbb{1}$ (Village pop. (2001) > 300) × | -0.63 | 1.74 | -2.39 | 0.63 |
| | $\mathbb{1}$ (Village in guar-growing district) | (1.51) | (2.21) | (1.80) | (1.51) |
| District FEs | | Yes | Yes | Yes | Yes |
| Census (2001) controls | | Yes | Yes | Yes | Yes |
| $N$ | | 6992 | 6992 | 6992 | 6992 |
| Adjusted $R^2$ | | 0.38 | 0.45 | 0.32 | 0.38 |
| Mean of outcome | | 48.23 | 39.94 | 8.28 | 51.77 |

This table shows results from estimating Equation (1.9) on a sample of single-habitation villages located in *non*-RGGVY Phase I districts with a Census 2001 population within a fifty-person bandwidth around RGGVY's 300-person eligibility threshold. Outcome variables for regressions reported in columns (1)–(4) are constructed using data from the Primary Census Abstract tables of the 2011 round of the Indian Census. Specifically, "agricultural workers" represents a village-level sum of main and marginal cultivators and agricultural laborers, while "non-agricultural workers" represents a village-level sum of main and marginal household-industry and "other" workers. Estimates associated with the population running variable $\left(\tilde{P}^{2001}_{vds}\right)$ are omitted. Following Correia (2015), 21 singleton observations are excluded. Standard errors—in parentheses—are clustered at the district level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table G.5: RD estimates with multiple hypothesis test adjustment

| Outcome variable | (1) $\hat{\beta}_2$ | (2) Adj. $p$-value |
|---|---|---|
| All workers (% of population) | 0.14 | 0.997 |
| Male | -0.13 | 0.996 |
| Female | 0.07 | 0.997 |
| Agricultural workers (% of population) | -6.39* | 0.095 |
| Male | -2.85 | 0.203 |
| Female | -3.25 | 0.265 |
| Non-agricultural workers (% of population) | 5.60** | 0.043 |
| Male | 2.30 | 0.296 |
| Female | 3.22 | 0.265 |
| Non-workers (% of population) | -0.14 | 0.997 |
| Male | 1.66 | 0.557 |
| Female | -1.65 | 0.557 |

Column (1) reports the estimated $\hat{\beta}_2$ coefficients from Tables 1.1 and 1.2. Column (2) reports corresponding $p$-values for this "family" of regressions, adjusted for multiple hypothesis testing using the free step-down resampling methodology of Westfall and Young (1993). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

## Table G.6: Propensity-score estimation using logistic regression

| Village-level characteristic | (1) 𝟙 (NGO village) | (2) 𝟙 (NGO village) | (3) 𝟙 (NGO village) |
|---|---|---|---|
| Area (km$^2$) | 0.00055 | 0.000096 | 0.000096 |
| | (0.00058) | (0.00011) | (0.00011) |
| Area$^2$ | −0.000000025 | | |
| | (0.00000010) | | |
| Total population | 0.00078 | −0.00058$^{*}$ | −0.00060$^{**}$ |
| | (0.00047) | (0.00030) | (0.00030) |
| Scheduled Caste population (proportion) | 0.26 | 1.06$^{***}$ | 1.06$^{***}$ |
| | (0.37) | (0.28) | (0.28) |
| Scheduled Tribe population (proportion) | 7.02 | −4.41 | −4.59 |
| | (4.84) | (5.07) | (5.14) |
| Population density | −0.057$^{*}$ | | |
| | (0.030) | | |
| Number of primary schools | 0.41$^{**}$ | 0.59$^{***}$ | 0.58$^{***}$ |
| | (0.18) | (0.12) | (0.12) |
| Number of middle schools | −0.16 | 0.080 | 0.11 |
| | (0.29) | (0.22) | (0.22) |
| Number of secondary schools | −0.058 | | |
| | (0.60) | | |
| 𝟙 (Medical facilities) | −0.66$^{**}$ | | |
| | (0.33) | | |
| Number of health centres | 0.63 | 0.60 | 0.57 |
| | (0.87) | (0.62) | (0.63) |
| Number of primary health centres | −0.44 | −0.23 | −0.20 |
| | (0.96) | (0.73) | (0.73) |
| Number of telephone connections | −0.099$^{**}$ | −0.093$^{*}$ | −0.094$^{*}$ |
| | (0.046) | (0.055) | (0.056) |
| 𝟙 (Bus services) | −0.50 | 0.65$^{***}$ | 0.62$^{***}$ |
| | (0.31) | (0.24) | (0.24) |
| 𝟙 (Credit societies) | 0.31 | 0.72$^{**}$ | 0.71$^{**}$ |
| | (0.47) | (0.34) | (0.34) |
| 𝟙 (Approach to village: paved road) | −0.29 | −0.43$^{*}$ | −0.41$^{*}$ |
| | (0.30) | (0.23) | (0.23) |
| Distance from nearest town (km) | −0.017$^{**}$ | −0.016$^{***}$ | −0.015$^{***}$ |
| | (0.0081) | (0.0045) | (0.0045) |
| Forest area (hectares) | 0.00019 | 0.00026 | 0.00036 |
| | (0.00035) | (0.00033) | (0.00034) |
| 𝟙 (Tap water) | | | −0.0054 |
| | | | (0.31) |
| 𝟙 (Electricity for all purposes) | | | −0.0058 |
| | | | (0.34) |
| Constant | −19.2 | −2.63$^{***}$ | −2.62$^{***}$ |
| | (594.4) | (0.20) | (0.33) |
| Observations | 1,960 | 1,965 | 1,903 |
| Pseudo $R^2$ | 0.51 | 0.079 | 0.077 |
| Sub-district fixed effects | Yes | No | No |

This table presents results from logistic regressions of an indicator for whether our partner NGO had operated in village $i$ in the past—represented by 𝟙 (NGO village)—on a set of village-level characteristics from the 2001 round of the Indian Census. Standard errors in parentheses. Our final model is shown in column (3). For this model, we initially restrict our sample to all Census-designated villages in the Bageshwar and Nainital districts of the state of Uttarakhand with non-zero or non-missing values for total population. We then exclude six villages where pretesting activities occurred with an alternative NGO partner (details available upon request). The final estimation sample for the model presented in column (3), thus, consists of all remaining villages with non-missing values for the village-level characteristics used for estimation. $^{*}$ $p < 0.10$, $^{**}$ $p < 0.05$, $^{***}$ $p < 0.01$.

Table G.7: Comparison of NGO and non-NGO villages using selected 2011 Census variables

| | (1) $\mathbb{1}$ (NGO village) | (2) $p$ value | (3) Adjusted $p$ value | (4) $R^2$ | (5) $N$ |
|---|---|---|---|---|---|
| *Number of dwelling rooms (%)* | | | | | |
| No exclusive room | 1.22 | 0.259 | 0.933 | 0.035 | 38 |
| One | −8.70 | 0.054* | 0.658 | 0.099 | 38 |
| Two | −1.80 | 0.715 | 0.999 | 0.0037 | 38 |
| Three | 4.63 | 0.178 | 0.891 | 0.050 | 38 |
| Four | −1.03 | 0.782 | 0.999 | 0.0022 | 38 |
| Five | 1.56 | 0.408 | 0.979 | 0.019 | 38 |
| Six or more | 4.13 | 0.306 | 0.941 | 0.029 | 38 |
| *Household size* | | | | | |
| One | −2.67 | 0.010** | 0.320 | 0.17 | 38 |
| Two | −2.57 | 0.025** | 0.487 | 0.13 | 38 |
| Three | −2.36 | 0.129 | 0.857 | 0.063 | 38 |
| Four | −2.69 | 0.146 | 0.873 | 0.058 | 38 |
| Five | 0.89 | 0.648 | 0.999 | 0.0059 | 38 |
| Six to eight | 7.87 | 0.011** | 0.334 | 0.16 | 38 |
| Nine or greater | 1.52 | 0.174 | 0.891 | 0.051 | 38 |
| Tap water from treated source (%) | −1.99 | 0.879 | 0.999 | 0.00065 | 38 |
| *Main source of lighting (%)* | | | | | |
| Electricity | 12.1 | 0.004*** | 0.197 | 0.21 | 38 |
| Kerosene | −11.1 | 0.004*** | 0.201 | 0.21 | 38 |
| *Type of fuel used for cooking (%)* | | | | | |
| Fuelwood | −0.33 | 0.958 | 0.999 | 0.000079 | 38 |
| LPG | 0.58 | 0.924 | 0.999 | 0.00025 | 38 |
| Electricity | −0.058 | 0.324 | 0.946 | 0.027 | 38 |
| Number of households availing of banking services | 2.04 | 0.711 | 0.999 | 0.0039 | 38 |
| *Asset ownership (%)* | | | | | |
| Radio | −1.37 | 0.803 | 0.999 | 0.0017 | 38 |
| Television | 9.72 | 0.145 | 0.873 | 0.058 | 38 |

Column (1) presents the estimated $\beta_1$ coefficients for the specified Census outcome variable from a regression model of the form: $Y_i = \beta_0 + \beta_1 \cdot \mathbb{1}\,(\text{NGO village}) + \nu_i$, where $Y_i$ represents a village-level characteristic for village $i$ in the 2011 Census round, $\mathbb{1}$ (NGO village) represents an indicator for whether our partner NGO had operated in village $i$ in the past, and $\nu_i$ represents a normally distributed error component. Column (2) shows the corresponding $p$ value—derived from heteroscedasticity robust standard errors—associated with each estimated coefficient. Column (3) shows $p$ values obtained using the free step-down resampling methodology of Westfall and Young (1993), as operationalized by Jones et al. (2018). The unit of analysis is the Census-designated village ($N$ = 38). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table G.8: Pre-trends for selected village-level variables from the 2001 and 2011 Census

| Village-level characteristic | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | $\mathbb{1}$ (Census 2011) × $\mathbb{1}$ (NGO village) | | $R^2$ | $N$ |
| Number of households | 5.05 | (13.0) | 0.037 | 76 |
| Total population | 34.3 | (66.4) | 0.032 | 76 |
| Total population (females) | 13.8 | (32.2) | 0.014 | 76 |
| Total population (males) | 20.5 | (35.6) | 0.051 | 76 |
| Total population (Scheduled Caste/Scheduled Tribe) | 11.8 | (58.6) | 0.013 | 76 |
| Number of primary schools | −0.00 | (0.21) | 0.017 | 76 |
| Number of other educational facilities | 0.16 | (0.44) | 0.026 | 76 |
| Number of primary health centres | 0.053 | (0.053) | 0.040 | 76 |
| Number of community health workers | −0.00 | (0.074) | 0.027 | 76 |
| $\mathbb{1}$ (Tap water) | 0.11 | (0.072) | 0.081 | 76 |
| $\mathbb{1}$ (Tubewell) | −0.053 | (0.053) | 0.040 | 76 |
| $\mathbb{1}$ (Bus services) | 0.11 | (0.17) | 0.097 | 76 |
| $\mathbb{1}$ (Electricity for agricultural use) | −0.26 | (0.17) | 0.092 | 76 |
| $\mathbb{1}$ (Electricity for domestic use) | −0.053 | (0.089) | 0.050 | 76 |
| $\mathbb{1}$ (Approach to village: paved road) | 0.11 | (0.21) | 0.16 | 76 |
| $\mathbb{1}$ (Post office) | 0.11 | (0.16) | 0.14 | 76 |
| Total irrigated land area (hectares) | 2.78 | (5.52) | 0.043 | 76 |
| Total unirrigated land area (hectares) | 1.44 | (12.7) | 0.011 | 76 |

Column (1) presents the estimated $\beta_3$ coefficients for the specified Census outcome variable from a regression model of the form: $Y_{it} = \beta_0 + \beta_1 \cdot \mathbb{1}\,(\text{Census 2011}) + \beta_2 \cdot \mathbb{1}\,(\text{NGO village}) + \beta_3\,[\mathbb{1}\,(\text{Census 2011}) \times \mathbb{1}\,(\text{NGO village})] + \nu_{it}$, where $Y_{it}$ represents a village-level characteristic for village $i$ in Census round $t$, $\mathbb{1}$ (Census 2011) represents an indicator for the 2011 Census round (the 2001 Census round is the omitted category), $\mathbb{1}$ (NGO village) represents an indicator for whether our partner NGO had operated in village $i$ in the past, and $\nu_{it}$ represents a normally distributed error term. Heteroscedasticity robust standard errors in parentheses shown in column (2). $^{*}$ $p < 0.10$, $^{**}$ $p < 0.05$, $^{***}$ $p < 0.01$

Table G.9: Baseline improved-stove ownership does not predict intervention-stove purchase

| | (1) | (2) |
|---|---|---|
| | $\mathbb{1}$ (Purchased intervention ICS) | |
| $TREATMENT_j$ | 0.51*** | 0.50*** |
| | (0.036) | (0.037) |
| $\mathbb{1}$ (Owns an improved stove at baseline) | 0.00 | 0.0015 |
| | (–) | (0.0073) |
| $TREATMENT_j \times \mathbb{1}$ (Owns an improved stove at baseline) | 0.051 | 0.042 |
| | (0.057) | (0.057) |
| Constant | −0.00*** | −0.11** |
| | (0.00) | (0.045) |
| Mean dep. (control) | 0.00 | 0.00 |
| Observations | 943 | 943 |
| Adjusted $R^2$ | 0.23 | 0.24 |
| Household-level controls | No | Yes |

The outcome variable is an indicator that equals one if household $i$ in hamlet $j$ purchased at least one of the two ICS promoted during the intervention. Baseline household-level controls for household size, number of children under five, awareness of existence of cleaner stoves and fuels, and total traditional-fuel collection time per day are included in column (3). "Traditional fuel" includes crop residue, dung, fuelwood, leaves, and household waste (trash); missing observations for total time spent collecting traditional fuel for 32 households are replaced with the sample mean value. Standard errors (in parentheses) are clustered at the hamlet level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table G.10: Comparing impacts on ownership of improved stoves and LPG stoves

| | (1)<br>$\mathbb{1}$ (Owns improved stove) | (2)<br>$\mathbb{1}$ (Owns LPG stove) |
|---|---|---|
| $POST_1$ | −0.048 | −0.034 |
| | (0.042) | (0.037) |
| $POST_2$ | 0.12*** | 0.12*** |
| | (0.044) | (0.042) |
| $TREATMENT_j \times POST_1$ | 0.38*** | 0.12*** |
| | (0.066) | (0.045) |
| $TREATMENT_j \times POST_2$ | 0.28*** | 0.046 |
| | (0.065) | (0.050) |
| $NGO_j \times POST_1$ | 0.15** | 0.15** |
| | (0.068) | (0.060) |
| $NGO_j \times POST_2$ | 0.040 | 0.065 |
| | (0.083) | (0.082) |
| $TREATMENT_j \times NGO_j \times POST_1$ | −0.10 | −0.12* |
| | (0.092) | (0.070) |
| $TREATMENT_j \times NGO_j \times POST_2$ | −0.047 | −0.056 |
| | (0.10) | (0.090) |
| Mean dep. (baseline non-NGO control) | 0.36 | 0.33 |
| Observations | 2,829 | 2,829 |
| Adjusted $R^2$ | 0.57 | 0.65 |
| Household fixed-effects | Yes | Yes |

The outcome variable in column (1) is an indicator that equals one if household $i$ in hamlet $j$ reports owning at least one improved stove in survey round $t$; the results reported in column (1) are identical to those reported in column (2) of Table 2.5. Similarly, as in Table 2.5, "improved stove" includes stoves fuelled by biogas, electricity, LPG, kerosene, and commercially available efficient biomass cookstoves; we also include the two ICS promoted as part of the promotion intervention in this definition. In column (2), the outcome variable is an indicator that equals one if household $i$ in hamlet $j$ reports owning an LPG stove. Standard errors (in parentheses) are clustered at the hamlet level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table G.11: Comparing impacts on fuel-collection time for all fuels and traditional fuels

| | (1) | (2) |
|---|---|---|
| | Fuel-collection time (minutes per day) | |
| | All fuels | Traditional fuels |
| $POST_1$ | −19.0 | −11.8 |
| | (18.5) | (19.0) |
| $POST_2$ | $-30.7^{**}$ | $-42.0^{***}$ |
| | (12.5) | (11.3) |
| $TREATMENT_j \times POST_1$ | 19.2 | 15.9 |
| | (25.4) | (24.8) |
| $TREATMENT_j \times POST_2$ | 12.8 | 14.0 |
| | (15.9) | (14.7) |
| $NGO_j \times POST_1$ | 53.3 | 50.1 |
| | (33.2) | (31.7) |
| $NGO_j \times POST_2$ | $48.5^{*}$ | $53.3^{**}$ |
| | (29.0) | (24.9) |
| $TREATMENT_j \times NGO_j \times POST_1$ | $-95.9^{**}$ | $-86.5^{**}$ |
| | (40.4) | (38.8) |
| $TREATMENT_j \times NGO_j \times POST_2$ | $-60.8^{*}$ | $-54.5^{*}$ |
| | (32.1) | (28.2) |
| Mean dep. (baseline non-NGO control) | 113.6 | 104.2 |
| Observations | 2,829 | 2,829 |
| Adjusted $R^2$ | 0.029 | 0.038 |
| Household fixed-effects | Yes | Yes |

Notes: the outcome variable for fuel-collection time in column (1) is derived from self-reported data on time spent (per day, week, or month) collecting fuelwood, crop residue, leaves, dung, biomass pellets, kerosene, LPG, biogas, and—if relevant—any other fuel used by the household; the results in column (1) are identical to those presented in column (4) of Table 2.6. In column (2), fuel-collection time is restricted to only traditional fuels (fuelwood, crop residue, leaves, and dung). Standard errors (in parentheses) are clustered at the hamlet level. $^{*}$ $p < 0.10$, $^{**}$ $p < 0.05$, $^{***}$ $p < 0.01$.

Table G.12: Impact of joint stove auction on willingness to pay with cluster adjustment

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | *Jambar* bid amount (CFA) | | | *Jumbo Zama* bid amount (CFA) | | |
| 𝟙 (Joint stove auction) | -228.6 | -67.07 | -339.8 | -2260.5*** | -3288.7*** | -3494.1*** |
| | (0.748) | (0.920) | (0.670) | (0.004) | (0.002) | (0.006) |
| 𝟙 (Long survey arm) | 1456.3** | 973.7 | 1292.5** | | | |
| | (0.014) | (0.138) | (0.016) | | | |
| Region FEs | Yes | Yes | Yes | Yes | Yes | Yes |
| "True valuation" FEs | No | Yes | Yes | No | Yes | Yes |
| Household- and village-level controls | No | No | Yes | No | No | Yes |
| $N$ | 982 | 982 | 982 | 317 | 317 | 317 |
| Adjusted $R^2$ | 0.047 | 0.271 | 0.313 | 0.088 | 0.219 | 0.300 |
| Mean of outcome | 3272.3 | 3272.3 | 3272.3 | 8873.8 | 8873.8 | 8873.8 |

This table shows results from estimating Equations (3.2) and (3.3) using ordinary least squares on the full sample of bid data from *Jambar* and *Jumbo Zama* sealed-bid, second-price auctions in columns (1) and (4), respectively. All models include region fixed-effects. Columns (2) and (4) also include fixed-effects for households that won an auction but declined to purchase the respective stove at the next-highest price. Columns (3) and (6) include controls for all household- and village-level characteristics shown in Table 3.1. Missing values in control variables for (i) reported village population are replaced with region-level means; (ii) relative wealth perception, household size, and age of household are replaced with village-level means; and (iii) households' bank account access, ownership of the *Jambar* stove, and household head's ability to read are replaced with zeros Additional binary variables that equal one for any household for which missing values are replaced in this way are also included in the estimation. $p$-values—in parentheses—are derived using the wild cluster bootstrap-$t$ procedure due to Cameron et al. (2008), with clustering at the village level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table G.13: Impact of joint stove auction on willingness to pay for randomly selected households

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | *Jambar* bid amount (CFA) | | | *Jumbo Zama* bid amount (CFA) | | |
| 𝟙 (Joint stove auction) | -232.4 | -65.17 | -345.3 | -2260.5*** | -3288.7*** | -3494.1*** |
| | (631.5) | (670.0) | (632.3) | (772.2) | (897.1) | (797.1) |
| 𝟙 (Long survey arm) | 1529.4*** | 1057.2* | 1319.0** | | | |
| | (525.9) | (530.8) | (517.0) | | | |
| Region FEs | Yes | Yes | Yes | Yes | Yes | Yes |
| "True valuation" FEs | No | Yes | Yes | No | Yes | Yes |
| Household- and village-level controls | No | No | Yes | No | No | Yes |
| $N$ | 937 | 937 | 937 | 317 | 317 | 317 |
| Adjusted $R^2$ | 0.043 | 0.268 | 0.298 | 0.067 | 0.199 | 0.240 |
| Mean of outcome | 3226.2 | 3226.2 | 3226.2 | 8873.8 | 8873.8 | 8873.8 |

This table shows results from estimating Equations (3.2) and (3.3) using ordinary least squares on bid data from *Jambar* and *Jumbo Zama* sealed-bid, second-price auctions conducted with only randomly selected households in columns (1) and (4), respectively. All models include region fixed-effects. Columns (2) and (4) also include fixed-effects for households that won an auction but declined to purchase the respective stove at the next-highest price. Columns (3) and (6) include controls for all household- and village-level characteristics shown in Table 3.1. Missing values in control variables for (i) reported village population are replaced with region-level means; (ii) relative wealth perception, household size, and age of household are replaced with village-level means; and (iii) households' bank account access, ownership of the *Jambar* stove, and household head's ability to read are replaced with zeros Additional binary variables that equal one for any household for which missing values are replaced in this way are also included in the estimation. Standard errors—in parentheses—are clustered at the village level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table G.14: Impact of joint stove auction on willingness to pay using Tobit regression

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | *Jambar* bid amount (CFA) | | | *Jumbo Zama* bid amount (CFA) | | |
| $\mathbb{1}$ (Joint stove auction) | -371.9 | -169.7 | -597.0 | -2346.3*** | -3460.0*** | -3602.3*** |
| | (741.5) | (786.5) | (776.7) | (881.4) | (1015.1) | (969.9) |
| $\mathbb{1}$ (Long survey arm) | 2582.9*** | 1893.5*** | 2307.8*** | | | |
| | (652.6) | (662.4) | (638.2) | | | |
| Region FEs | Yes | Yes | Yes | Yes | Yes | Yes |
| "True valuation" FEs | No | Yes | Yes | No | Yes | Yes |
| Household- and village-level controls | No | No | Yes | No | No | Yes |
| $N$ | 982 | 982 | 982 | 317 | 317 | 317 |
| Pseudo $R^2$ | 0.004 | 0.018 | 0.021 | 0.004 | 0.012 | 0.019 |
| Mean of outcome | 3272.3 | 3272.3 | 3272.3 | 8873.8 | 8873.8 | 8873.8 |

This table shows results from estimating Equations (3.2) and (3.3) using a Tobit regression model on the full sample of bid data from *Jambar* and *Jumbo Zama* sealed-bid, second-price auctions in columns (1) and (4), respectively. All models include region fixed-effects. Columns (2) and (4) also include fixed-effects for households that won an auction but declined to purchase the respective stove at the next-highest price. Columns (3) and (6) include controls for all household- and village-level characteristics shown in Table 3.1. Missing values in control variables for (i) reported village population are replaced with region-level means; (ii) relative wealth perception, household size, and age of household are replaced with village-level means; and (iii) households' bank account access, ownership of the *Jambar* stove, and household head's ability to read are replaced with zeros Additional binary variables that equal one for any household for which missing values are replaced in this way are also included in the estimation. Standard errors—in parentheses—are clustered at the village level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

# Appendix H

## Additional figures

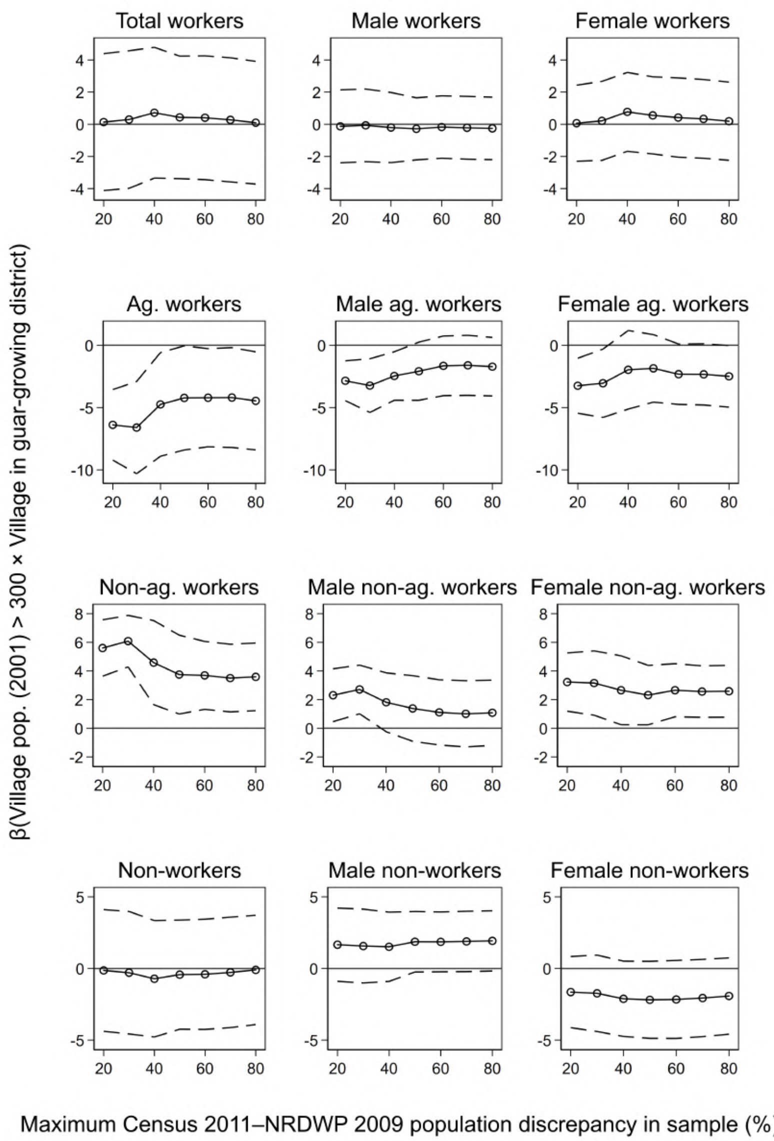


Figure H.1: Sensitivity of results to varying Census-NRDWP population discrepancy rates. This figure shows how the results reported in Tables 1.1 and 1.2 for the estimated value of $\hat{\beta}_2$ evolves as we relax the Census 2011-NRDWP 2009 population discrepancy threshold we impose during our fuzzy matching procedure to validate matches (see Appendix C). Markers represent point estimates; dashed lines indicate 90 percent confidence intervals.

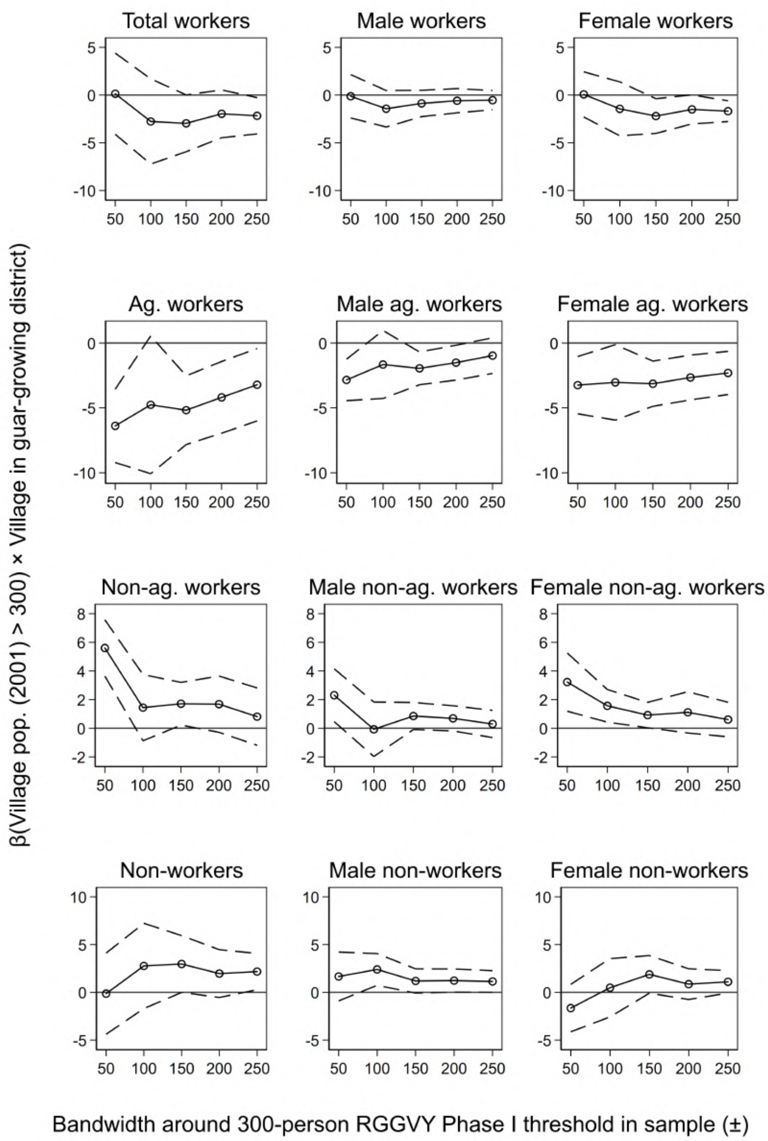


FIGURE H.2: Sensitivity of results to varying RD bandwidths. This figure shows how the results reported in Tables 1.1 and 1.2 for the estimated value of $\hat{\beta}_2$ evolves as we vary the population bandwidth around RGGVY's 300-person eligibility threshold to identify our analytical sample. Markers represent point estimates; dashed lines indicate 90 percent confidence intervals.

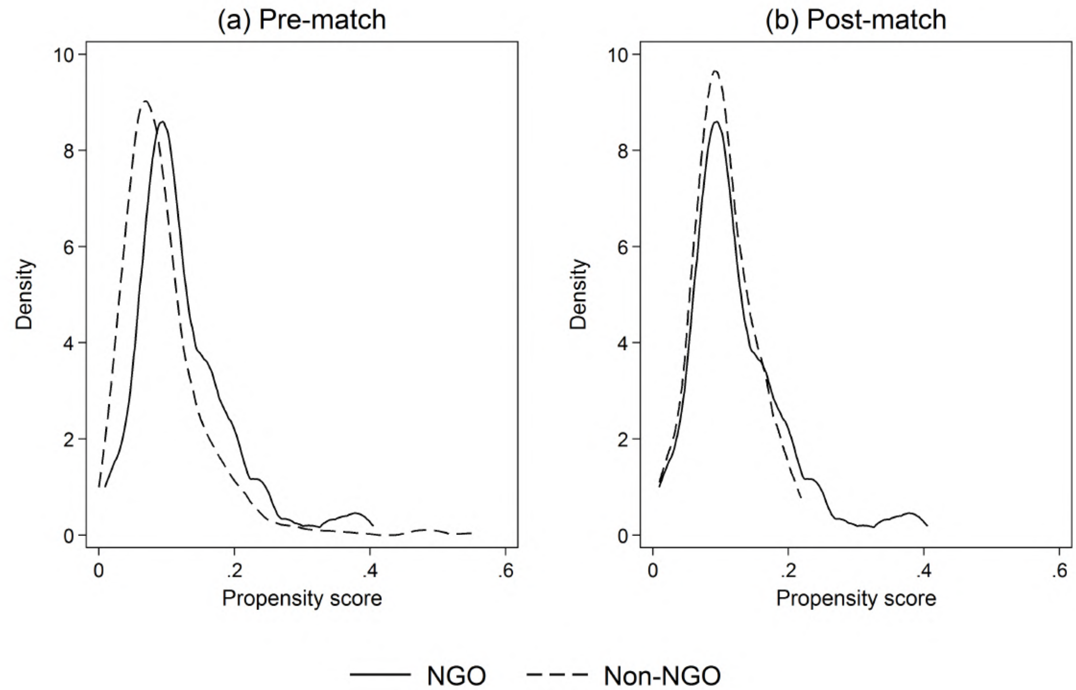


Figure H.3: Distribution of predicted propensity scores of NGO and non-NGO villages. This figure presents the distribution of predicted propensity scores using the model outlined in column (3) of Table G.6 before (panel *a*) and after (panel *b*) the propensity-score matching exercise. Prior to matching, we restrict our sample to villages in nine sub-districts of Bageshwar and Nainital districts of the state of Uttarakhand for implementation-related logistical reasons; the distribution of propensity scores for all villages in these sub-districts ($N_{\text{NGO}}$ = 97 and $N^{\text{Unmatched}}_{\text{Non-NGO}}$ = 536) is shown in panel (a). In panel (b), the distribution of propensity scores for only those non-NGO villages that are matched to at least one NGO village ($N^{\text{Matched}}_{\text{Non-NGO}}$ = 74) is shown.

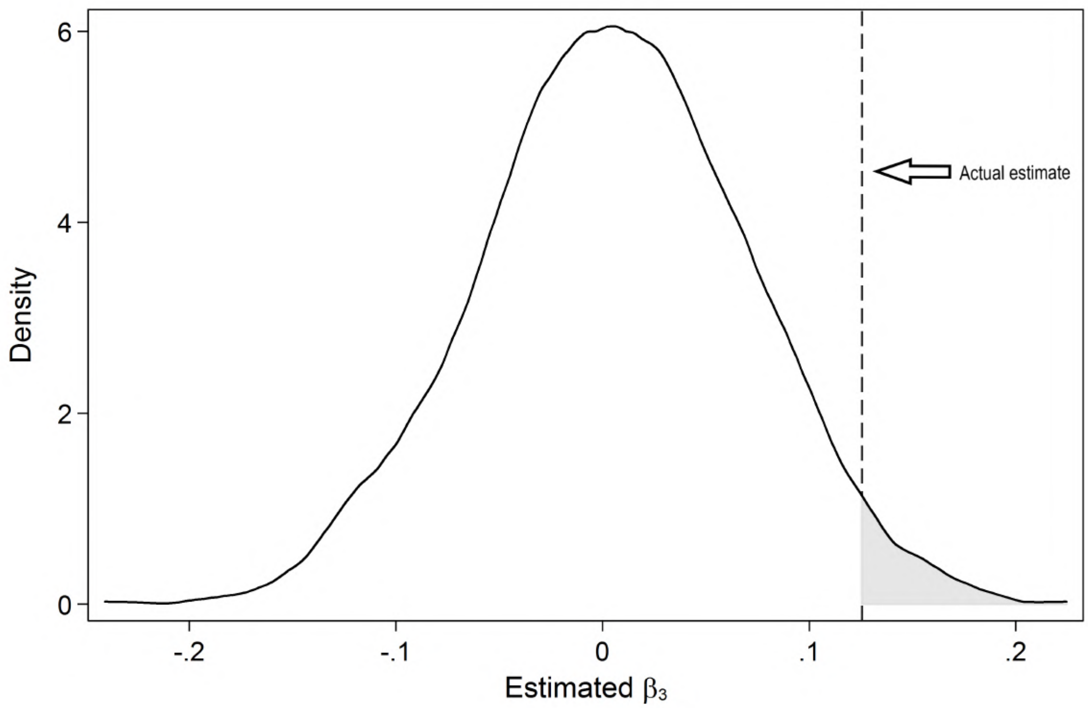


Figure H.4: Randomization-based inferential procedure applied to village-level stratum allocation. This figure plots the distribution of 1,000 estimated $\beta_3$ coefficients from a randomization inference procedure (Athey and Imbens, 2017) applied to village-level NGO stratum allocation to estimate Equation (2.19). We randomly assign each village in the sample to placebo NGO and non-NGO strata, and estimate the specification presented in column (3) of Table 2.3 to obtain a placebo "NGO effect" estimate for heterogeneity in purchase of intervention ICS. This procedure is repeated 1,000 times to obtain a distribution of placebo effects. The vertical line indicates the magnitude of our actual estimated 'NGO effect'. Approximately three percent of placebo estimates are larger than the actual estimated effect (shaded area).

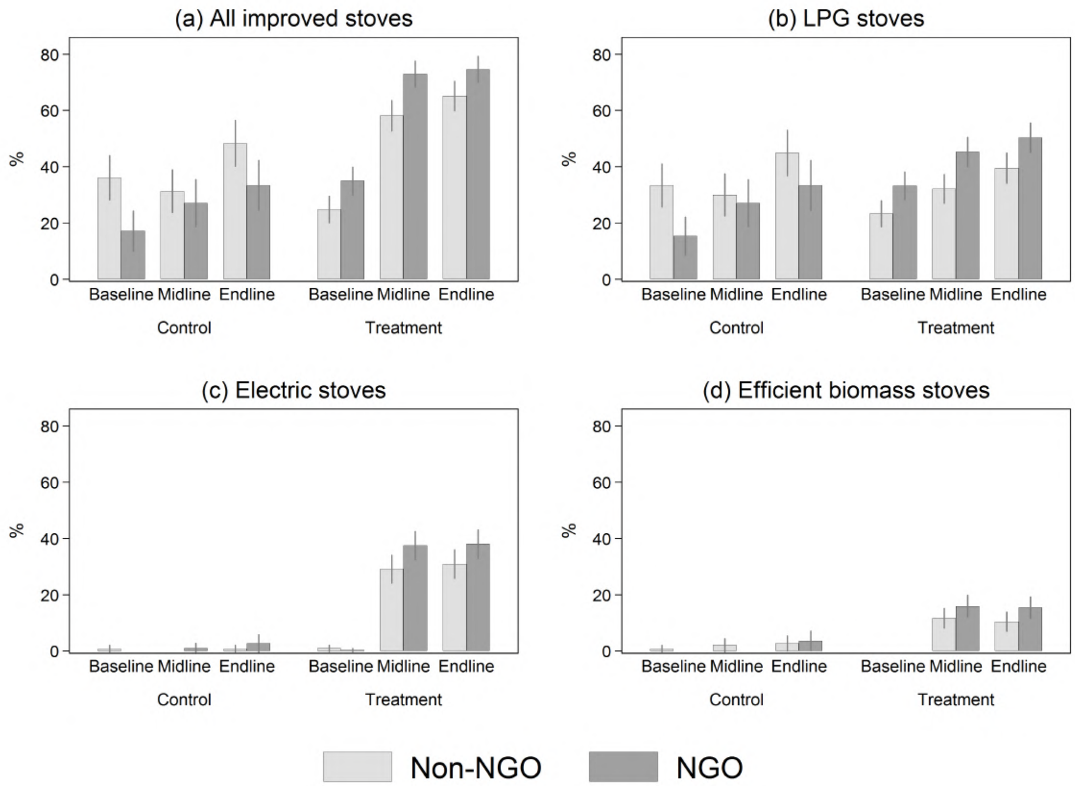


FIGURE H.5: Trends in ownership of selected improved stoves. This figure presents mean ownership rates of all improved stoves as well as LPG, electric and improved biomass variants separately for treatment/control and NGO/non-NGO communities during each survey round. "Improved stove" in panel (a) includes stoves fuelled by biogas, electricity, LPG, kerosene, and commercially available efficient biomass cookstoves; we also include the two ICS promoted as part of the promotion intervention in this definition. Electric stoves (panel *c*) and efficient biomass stoves (panel *d*) include the respective ICS promoted as part of the promotion intervention. Error bars represent 95 percent confidence intervals for the means. Baseline survey activities occurred approximately one year before the intervention; midline and endline surveys occurred approximately three and fifteen months, respectively, after the intervention (see Figure 2.1).

# Bibliography


Abadie, A., A. Diamond, and J. Hainmueller (2010): "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program," *Journal of the American Statistical Association*, 105, 493–505.

Abadie, A. and J. Gardeazabal (2003): "The Economic Costs of Conflict: A Case Study of the Basque Country," *American Economic Review*, 93, 113–132.

Acemoglu, D., S. Johnson, A. Kermani, J. Kwak, and T. Mitton (2016): "The value of connections in turbulent times: Evidence from the United States," *Journal of Financial Economics*, 121, 368–391.

Adair-Rohani, H., J. Lewis, J. Mingle, and S. Gumy (2016): "Burning Opportunity: Clean Household Energy for Health, Sustainable Development, and Wellbeing of Women and Children," Tech. rep., World Health Organization.

Adhvaryu, A., A. V. Chari, and S. Sharma (2013): "Firing Costs and Flexibility: Evidence from Firms' Employment Responses to Shocks in India," *Review of Economics and Statistics*, 95, 725–740.

Adukia, A., S. Asher, and P. Novosad (2018): "Educational Investment Responses to Economic Opportunity: Evidence from Indian Road Construction," Working paper.

Aggarwal, S. (2018): "Do rural roads create pathways out of poverty? Evidence from India," *Journal of Development Economics*, 133, 375–395.

Agricultural and Processed Food Products Export Development Authority (2011): "APEDA Agri Exchange Ready Reckoner Series: Guargum," Tech. rep., Ministry of Commerce & Industry, retrieved from `http://apeda.gov.in/apedawebsite/six_head_product/Guargum_final_Profile.pdf` on April 16, 2017.

Aldashev, G., M. Limardi, and T. Verdier (2015): "Watchdogs of the Invisible Hand: NGO monitoring and industry equilibrium," *Journal of Development Economics*, 116, 28–42.

Aldashev, G. and C. Navarra (2014): "Development NGOs: Basic Facts," Université de Namur Working Paper № 2014/09.

ALLCOTT, H. (2015): "Site Selection Bias in Program Evaluation," *Quarterly Journal of Economics*, 130, 1117–1165.

AMITI, M. AND B. S. JAVORCIK (2008): "Trade costs and location of foreign firms in China," *Journal of Development Economics*, 85, 129–149.

ANAND, U. (2015): "India has 31 lakh NGOs, more than double the number of schools," *The Indian Express*, retrieved from `http://bit.ly/2pqeKti` on July 11, 2017.

ASHER, S., T. GARG, AND P. NOVOSAD (2018): "The Ecological Impact of Transportation Infrastructure," World Bank Policy Research Working Paper № 8507.

ASHER, S. AND P. NOVOSAD (2018): "Rural Roads and Local Economic Development," Working paper.

ATHEY, S. AND G. IMBENS (2017): "The Econometrics of Randomized Experiments a," in *Handbook of Field Experiments*, Elsevier, 73–140.

AUNG, T. W., G. JAIN, K. SETHURAMAN, J. BAUMGARTNER, C. REYNOLDS, A. P. GRIESHOP, J. D. MARSHALL, AND M. BRAUER (2016): "Health and Climate-Relevant Pollutant Concentrations from a Carbon-Finance Approved Cookstove Intervention in Rural India," *Environmental Science & Technology*, 50, 7228–7238.

BAILIS, R., R. DRIGO, A. GHILARDI, AND O. MASERA (2015): "The carbon footprint of traditional woodfuels," *Nature Climate Change*, 5, 266–272.

BANERJEE, A., R. BANERJI, J. BERRY, E. DUFLO, H. KANNAN, S. MUKERJI, M. SHOTLAND, AND M. WALTON (2017): "From Proof of Concept to Scalable Policies: Challenges and Solutions, with an Application," *Journal of Economic Perspectives*, 31, 73–102.

BANERJEE, A., E. DUFLO, N. GOLDBERG, D. KARLAN, R. OSEI, W. PARIENTE, J. SHAPIRO, B. THUYSBAERT, AND C. UDRY (2015): "A multifaceted program causes lasting progress for the very poor: Evidence from six countries," *Science*, 348, 1260799.

BANERJEE, A. V., A. H. AMSDEN, R. H. BATES, J. N. BHAGWATI, A. DEATON, AND N. STERN (2007): *Making Aid Work*, The MIT Press.

BANERJEE, A. V. AND E. DUFLO (2007): "The Economic Lives of the Poor," *Journal of Economic Perspectives*, 21, 141–167.

BANERJEE, S. G., D. BARNES, K. SINGH, BIPULAND MAYER, AND H. SAMAD (2014): *Power for All: Electricity Access Challenge in India*, World Bank.

BARATI, R. AND J.-T. LIANG (2014): "A review of fracturing fluid systems used for hydraulic fracturing of oil and gas wells," *Journal of Applied Polymer Science*, 131.

BARNES, D. F. (2005): "Transformative Power: Meeting the Challenge of Rural Electrification," Tech. rep., World Bank, Knowledge Exchange Series № 2.

Barnwal, P. (2017): "Curbing Leakage in Public Programs: Evidence from India's Direct Benefit Transfer Policy," Working paper.

Barron, M. and M. Torero (2017): "Household electrification and indoor air pollution," *Journal of Environmental Economics and Management*, 86, 81–92.

Beckwith, R. (2012): "Depending On Guar For Shale Oil And Gas Development," *Journal of Petroleum Technology*, 64, 44–55.

Bellarby, J. (2009): *Well Completion Design, Volume 56 (Developments in Petroleum Science)*, Elsevier Science.

Bengtsson, N. (2013): "Catholics versus Protestants: On the Benefit Incidence of Faith-Based Foreign Aid," *Economic Development and Cultural Change*, 61, 479–502.

Bensch, G., M. Grimm, and J. Peters (2015): "Why do households forego high returns from technology adoption? Evidence from improved cooking stoves in Burkina Faso," *Journal of Economic Behavior & Organization*, 116, 187–205.

Bensch, G. and J. Peters (2015): "The intensive margin of technology adoption – Experimental evidence on improved cooking stoves in rural Senegal," *Journal of Health Economics*, 42, 44–63.

Bensch, G., J. Peters, and M. Sievert (2017): "The lighting transition in rural Africa — From kerosene to battery-powered LED and the emerging disposal problem," *Energy for Sustainable Development*, 39, 13–20.

Berge, L. I. O., K. Bjorvatn, K. S. Juniwaty, and B. Tungodden (2012): "Business Training in Tanzania: From Research-driven Experiment to Local Implementation," *Journal of African Economies*, 21, 808–827.

Bernard, T., A. de Janvry, S. Mbaye, and E. Sadoulet (2017): "Expected Product Market Reforms and Technology Adoption by Senegalese Onion Producers," *American Journal of Agricultural Economics*, 99, 1096–1115.

Bernard, T. and M. Torero (2015): "Social Interaction Effects and Connection to Electricity: Experimental Evidence from Rural Ethiopia," *Economic Development and Cultural Change*, 63, 459–484.

Beyene, A. D., R. A. Bluffstone, S. Dissanayake, Z. Gebreegziabher, P. Martinsson, A. Mekonnen, and M. Toman (2015): "Can improved biomass cookstoves contribute to REDD+ in low-income countries? Evidence from a controlled cooking test trial with randomized behavioral treatments," World Bank Policy Research Working Paper № 7394.

Bold, T., M. Kimenyi, G. Mwabu, A. Ng'ang'a, and J. Sandefur (2013): "Scaling up what works: experimental evidence on external validity in Kenyan education," Center for Global Development Working Paper No. 321.

Bonan, J., S. Pareglio, and M. Tavoni (2017): "Access to modern energy: a review of barriers, drivers and impacts," *Environment and Development Economics*, 22, 491–516.

Brandt, M., C. Romankiewicz, R. Spiekermann, and C. Samimi (2014): "Environmental change in time series – An interdisciplinary study in the Sahel of Mali and Senegal," *Journal of Arid Environments*, 105, 52–63.

Brass, J. N. (2012): "Why Do NGOs Go Where They Go? Evidence from Kenya," *World Development*, 40, 387–401.

Brooks, N., V. Bhojvaid, M. Jeuland, J. Lewis, O. Patange, and S. Pattanayak (2016): "How much do alternative cookstoves reduce biomass fuel use? Evidence from North India," *Resource and Energy Economics*, 43, 153–171.

Burlig, F. and L. Preonas (2016): "Out of the Darkness and Into the Light? Development Effects of Rural Electrification in India," Energy Institute at Haas Working Paper № WP-268R.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2008): "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 90, 414–427.

Cameron, L. and M. Shah (2017): "Scaling Up Sanitation: Evidence from an RCT in Indonesia," IZA Discussion Paper No. 10619.

Carlton, D. W. (1983): "The Location and Employment Choices of New Firms: An Econometric Model with Discrete and Continuous Endogenous Variables," *The Review of Economics and Statistics*, 65, 440.

Cavallo, E., S. Galiani, I. Noy, and J. Pantano (2013): "Catastrophic Natural Disasters and Economic Growth," *Review of Economics and Statistics*, 95, 1549–1561.

Chen, X. and W. D. Nordhaus (2011): "Using luminosity data as a proxy for economic statistics," *Proceedings of the National Academy of Sciences*, 108, 8589–8594.

Chudzikowski, R. J. (1971): "Guar gum and its applications," *Journal of the Society of Cosmetic Chemists*, 22, 43–60.

Coffman, M. and I. Noy (2011): "Hurricane Iniki: measuring the long-term economic impact of a natural disaster using synthetic control," *Environment and Development Economics*, 17, 187–205.

Correia, S. (2015): "Singletons, cluster-robust standard errors and fixed effects: A bad mix," Technical note.

Crump, R. K., V. J. Hotz, G. W. Imbens, and O. A. Mitnik (2009): "Dealing with limited overlap in estimation of average treatment effects," *Biometrika*, 96, 187–199.

Deaton, A. (2010): "Instruments, Randomization, and Learning about Development," *Journal of Economic Literature*, 48, 424–455.

Demont, M., P. Rutsaert, M. Ndour, and W. Verbeke (2013): "Reversing Urban Bias in African Rice Markets: Evidence from Senegal," *World Development*, 45, 63–74.

Devarajan, S., S. Khemani, and M. Walton (2013): "Can Civil Society Overcome Government Failure in Africa?" *The World Bank Research Observer*, 29, 20–47.

Dickinson, K. L., S. R. Patil, S. K. Pattanayak, C. Poulos, and J.-C. Yang (2015): "Nature's Call: Impacts of Sanitation Choices in Orissa, India," *Economic Development and Cultural Change*, 64, 1–29.

Dinkelman, T. (2011): "The Effects of Rural Electrification on Employment: New Evidence from South Africa," *American Economic Review*, 101, 3078–3108.

Doll, C. N., J.-P. Muller, and J. G. Morley (2006): "Mapping regional economic activity from night-time light satellite imagery," *Ecological Economics*, 57, 75–92.

Duque, V., M. Rosales-Rueda, and F. Sanchez (2018): "How do early-life shocks interact with subsequent human-capital investments? Evidence from administrative data," Working paper.

Edlin, A. S. and C. Shannon (1998): "Strict Monotonicity in Comparative Statics," *Journal of Economic Theory*, 81, 201–219.

Elsner, M. and K. Hoelzer (2016): "Quantitative Survey and Structural Classification of Hydraulic Fracturing Chemicals Reported in Unconventional Gas Production," *Environmental Science & Technology*, 50, 3290–3314.

Esarey, J. and A. Menger (2018): "Practical and Effective Approaches to Dealing With Clustered Data," *Political Science Research and Methods*, 1–19.

Fetter, T. R. (2018): "Fracking, Toxics, and Disclosure," `http://sites.duke.edu/trfetter/research`.

Fetter, T. R., A. Steck, C. Timmins, and D. Wrenn (2018): "Learning by Viewing? Social Learning, Regulatory Disclosure, and Firm Productivity in Shale Gas," National Bureau of Economic Research (NBER) Working Paper № 25401.

Foster, A. D. and M. R. Rosenzweig (2010): "Microeconomics of Technology Adoption," *Annual Review of Economics*, 2, 395–424.

Fruttero, A. and V. Gauri (2005): "The Strategic Choices of NGOs: Location Decisions in Rural Bangladesh," *Journal of Development Studies*, 41, 759–787.

Gaba, K. M., B. Min, A. Thakker, and C. Elvidge (2016): "`nightlights.io`: Twenty Years of India Lights," Accessed on September 9, 2016.

Glennerster, R. (2015): "A movement grows up. Running Randomized Evaluations: A Practical Guide." Retrieved from `http://runningres.com/blog/2015/7/13/a-movement-grows-up` on April 24, 2016.

Goppers, K. (2006): "The Auas High-voltage Transmission Line in Namibia Supported by a Swedish Concessionary Credit," Tech. rep., Swedish International Development Cooperation Agency, Sida, sida Evaluation 06/53.

Grant, L. E. and K. K. Grooms (2017): "Do Nonprofits Encourage Environmental Compliance?" *Journal of the Association of Environmental and Resource Economists*, 4, S261–S288.

Grimm, M., A. Munyehirwe, J. Peters, and M. Sievert (2017): "A First Step up the Energy Ladder? Low Cost Solar Kits and Household's Welfare in Rural Rwanda," *The World Bank Economic Review*, 31, 631âĂŞ649.

Gronau, R. (1977): "Leisure, Home Production, and Work–the Theory of the Allocation of Time Revisited," *Journal of Political Economy*, 85, 1099–1123.

Grossman, G., M. Humphreys, and G. Sacramone-Lutz (2016): "Information Technology and Political Engagement: Mixed Evidence from Uganda," Working paper.

Grossman, M. (1972): "On the Concept of Health Capital and the Demand for Health," *Journal of Political Economy*, 80, 223–255.

Hanna, R., E. Duflo, and M. Greenstone (2016): "Up in Smoke: The Influence of Household Behavior on the Long-Run Impact of Improved Cooking Stoves," *American Economic Journal: Economic Policy*, 8, 80–114.

Henderson, J. V. and Y. S. Lee (2015): "Organization of Disaster Aid Delivery: Spending Your Donations," *Economic Development and Cultural Change*, 63, 617–664.

Henderson, J. V., A. Storeygard, and D. N. Weil (2012): "Measuring Economic Growth from Outer Space," *American Economic Review*, 102, 994–1028.

Herman, M. L., G. L. Head, T. E. Fogarty, and P. M. Jackson (2003): *Managing Risk in Nonprofit Organizations: A Comprehensive Guide*, Wiley, chap. 7, 133–148.

Hoeffler, S. and D. Ariely (1999): "Constructing Stable Preferences: A Look Into Dimensions of Experience and Their Impact on Preference Stability," *Journal of Consumer Psychology*, 8, 113–139.

Holloway, G., C. Nicholson, C. Delgado, S. Staal, and S. Ehui (2000): "Agroindustrialization through institutional innovation Transaction costs, cooperatives and milk-market development in the east-African highlands," *Agricultural Economics*, 23, 279–288.

Hooper, L. G., Y. Dieye, A. Ndiaye, A. Diallo, C. S. Sack, V. S. Fan, K. M. Neuzil, and J. R. Ortiz (2018): "Traditional cooking practices and preferences for stove features among women in rural Senegal: Informing improved cookstove design and interventions," *PLOS ONE*, 13, e0206822.

Hymowitz, T. (1972): "The trans-domestication concept as applied to Guar," *Economic Botany*, 26, 49–60.

Independent Evaluation Group (2008): "The welfare impact of rural electrification: A reassessment of the costs and benefits," Tech. rep., World Bank.

International Energy Agency (2011): "World Energy Outlook 2011: Energy for All," Tech. rep., OECD/IEA.

——— (2015): "WEO-2015 Special Report: India Energy Outlook," Tech. rep., OECD/IEA.

Ito, T. (2009): "Caste discrimination and transaction costs in the labor market: Evidence from rural North India," *Journal of Development Economics*, 88, 292–300.

Jack, W. and T. Suri (2014): "Risk Sharing and Transactions Costs: Evidence from Kenya's Mobile Money Revolution," *American Economic Review*, 104, 183–223.

Jayachandran, S., J. de Laat, E. F. Lambin, C. Y. Stanton, R. Audy, and N. E. Thomas (2017): "Cash for carbon: A randomized trial of payments for ecosystem services to reduce deforestation," *Science*, 357, 267–273.

Jeuland, M., V. Bhojvaid, A. Kar, J. Lewis, O. Patange, S. K. Pattanayak, N. Ramanathan, I. H. Rehman, J. S. Tan Soo, and V. Ramanathan (2015a): "Preferences for improved cook stoves: Evidence from rural villages in north India," *Energy Economics*, 52, 287–298.

Jeuland, M. and S. K. Pattanayak (2012): "Benefits and Costs of Improved Cookstoves: Assessing the Implications of Variability in Health, Forest and Climate Impacts," *PLoS ONE*, 7, e30338.

Jeuland, M., S. K. Pattanayak, and R. Bluffstone (2015b): "The Economics of Household Air Pollution," *Annual Review of Resource Economics*, 7, 81–108.

Jeuland, M., S. K. Pattanayak, J. S. Tan Soo, and F. Usmani (2018a): "Preferences and the effectiveness of behavior-change interventions," Working paper.

Jeuland, M., J.-S. T. Soo, and D. Shindell (2018b): "The need for policies to reduce the costs of cleaner cooking in low income settings: Implications from systematic analysis of costs and benefits," *Energy Policy*, 121, 275–285.

Jones, D., D. Molitor, and J. Reif (2018): "What Do Workplace Wellness Programs Do? Evidence from the Illinois Workplace Wellness Study," National Bureau of Economic Research (NBER) Working Paper № 24229.

Joshi, S. (2003): "Cost/Benefits of Horizontal Wells," in *SPE Western Regional/AAPG Pacific Section Joint Meeting*, Society of Petroleum Engineers.

Khandker, S. R., D. F. Barnes, and H. A. Samad (2013): "Welfare Impacts of Rural Electrification: A Panel Data Analysis from Vietnam," *Economic Development and Cultural Change*, 61, 659–692.

King, G., E. Gakidou, N. Ravishankar, R. T. Moore, J. Lakin, M. Vargas, M. M. Téllez-Rojo, J. E. H. Ávila, M. H. Ávila, and H. H. Llamas (2007): "A "politically robust" experimental design for public policy evaluation, with application to the Mexican Universal Health Insurance program," *Journal of Policy Analysis and Management*, 26, 479–506.

Kranton, R. E. (1996): "Reciprocal Exchange: A Self-Sustaining System," *American Economic Review*, 86, 830–851.

Kumar, P., R. K. Rao, and N. H. Reddy (2016): "Sustained uptake of LPG as cleaner cooking fuel in rural India: Role of affordability, accessibility, and awareness," *World Development Perspectives*, 4, 33–37.

Kuravadi, N. A., S. Verma, S. Pareek, P. Gahlot, S. Kumari, U. K. Tanwar, P. Bhatele, M. Choudhary, K. S. Gill, V. Pruthi, S. K. Tripathi, K. S. Dhugga, and G. S. Randhawa (2013): "Guar," in *Agricultural Sustainability: Progress and Prospects in Crop Research*, Elsevier, 47–60.

Lambe, F. and A. Atteridge (2012): "Putting the Cook Before the Stove: a User-Centred Approach to Understanding Household Energy Decision-Making: A Case Study of Haryana State, Northern India," Tech. rep., Stockholm Environment Institute.

Lardoux, S. and E. V. de Walle (2003): "Polygamie et fécondité en milieu rural sénégalais," *Population*, 58, 807.

Lederman, D. and G. Porto (2015): "The Price Is Not Always Right: On the Impacts of Commodity Prices on Households (and Countries)," *The World Bank Research Observer*, lkv013.

Lee, K., E. Miguel, and C. Wolfram (2018): "Experimental Evidence on the Economics of Rural Electrification," Working paper.

Lenz, L., A. Munyehirwe, J. Peters, and M. Sievert (2017): "Does Large-Scale Infrastructure Investment Alleviate Poverty? Impacts of Rwanda's Electricity Access Roll-Out Program," *World Development*, 89, 88–110.

Levine, D. I., T. Beltramo, G. Blalock, C. Cotterman, and A. M. Simons (2018): "What Impedes Efficient Adoption of Products? Evidence from Randomized Sales Offers for Fuel-Efficient Cookstoves in Uganda," *Journal of the European Economic Association.*

Lewis, J. J., V. Bhojvaid, N. Brooks, I. Das, M. A. Jeuland, O. Patange, and S. K. Pattanayak (2015): "Piloting Improved Cookstoves in India," *Journal of Health Communication*, 20, 28–42.

Lewis, J. J., J. W. Hollingsworth, R. T. Chartier, E. M. Cooper, W. M. Foster, G. L. Gomes, P. S. Kussin, J. J. MacInnis, B. K. Padhi, P. Panigrahi, C. E. Rodes, I. T. Ryde, A. K. Singha, H. M. Stapleton, J. Thornburg, C. J. Young, J. N. Meyer, and S. K. Pattanayak (2016): "Biogas Stoves Reduce Firewood Use, Household Air Pollution, and Hospital Visits in Odisha, India," *Environmental Science & Technology*, 51, 560–569.

Lewis, J. J. and S. K. Pattanayak (2012): "Who Adopts Improved Fuels and Cookstoves? A Systematic Review," *Environmental Health Perspectives*, 120, 637–645.

Lin, L., S. K. Pattanayak, E. O. Sills, and W. D. Sunderlin (2012): "Site selection for forest carbon projects," in *Analysing REDD+: Challenges and choices*, ed. by A. Angelsen, M. Brockhaus, W. D. Sunderlin, and L. Verchot, Bogor, Indonesia: Center for International Forestry Research (CIFOR), chap. 12, 209–230.

Lipscomb, M., A. M. Mobarak, and T. Barham (2013): "Development Effects of Electrification: Evidence from the Topographic Placement of Hydropower Plants in Brazil," *American Economic Journal: Applied Economics*, 5, 200–231.

Litzow, E., S. K. Pattanayak, and T. Thinley (2017): "Evaluating Rural Electrification: Illustrating Research Gaps with the Case of Bhutan," Environment for Development (EfD) Discussion Paper № 17-19.

MacLean, L. M., J. N. Brass, S. Carley, A. El-Arini, and S. Breen (2015): "Democracy and the Distribution of NGOs Promoting Renewable Energy in Africa," *The Journal of Development Studies*, 51, 725–742.

Mahadevan, M. (2019): "The Price of Power: Costs of Political Corruption in Indian Electricity," Working paper.

Martin, P. and C. A. Rogers (1995): "Industrial location and public infrastructure," *Journal of International Economics*, 39, 335–351.

Maurya, N. K. (2014): "A Critical Evaluation of the State Finances of the Uttarakhand Governments: 2002-03 to 2011-12," Tech. rep., Fourteenth Finance Commission, Government of India.

McCrary, J. (2008): "Manipulation of the running variable in the regression discontinuity design: A density test," *Journal of Econometrics*, 142, 698–714.

Meeks, R., K. R. Sims, and H. Thompson (2016): "Waste Not: Can Biogas Deliver Sustainable Development?" *Environmental and Resource Economics*, in press.

Micklewright, J. and A. Wright (2003): "Private Donations for International Development," WIDER Discussion Paper № 2003/82.

Miguel, E. and M. Kremer (2004): "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities," *Econometrica*, 72, 159–217.

Miller, G. and A. M. Mobarak (2015): "Learning About New Technologies Through Social Networks: Experimental Evidence on Nontraditional Stoves in Bangladesh," *Marketing Science*, 34, 480–499.

Min, B., Z. O'Keeffe, and F. Zhang (2017): "Whose Power Gets Cut? Using High-Frequency Satellite Images to Measure Power Supply Irregularity," World Bank Policy Research Working Paper № 8131.

Ministry of Power (2006): "The Rural Electrification Policy," Extraordinary Gazette of India, 2006, No. 197.

——— (2012): "Rajiv Gandhi Grameen Vidyutikaran Yojana," `http://pib.nic.in/newsite/PrintRelease.aspx?relid=83765`.

Mobarak, A. M., P. Dwivedi, R. Bailis, L. Hildemann, and G. Miller (2012): "Low demand for nontraditional cookstove technologies," *Proceedings of the National Academy of Sciences*, 109, 10815–10820.

Mortimer, K., C. B. Ndamala, A. W. Naunje, J. Malava, C. Katundu, W. Weston, D. Havens, D. Pope, N. G. Bruce, M. Nyirenda, D. Wang, A. Crampin, J. Grigg, J. Balmes, and S. B. Gordon (2017): "A cleaner burning biomass-fuelled cookstove intervention to prevent pneumonia in children under 5 years old in rural Malawi (the Cooking and Pneumonia Study): a cluster randomised controlled trial," *The Lancet*, 389, 167–175.

Mudgil, D., S. Barak, and B. S. Khatkar (2011): "Guar gum: processing, properties and food applications—A Review," *Journal of Food Science and Technology*, 51, 409–418.

Nadel, S. and L. Pritchett (2016): "Searching for the Devil in the Details: Learning About Development Program Design," Center for Global Development (CGD) Working Paper № 434.

National Rainfed Area Authority (2014): "Potential of Rainfed Guar (Cluster beans) Cultivation, Processing and Export in India," Tech. rep., Ministry of Agriculture & Farmers Welfare, retrieved from `http://nraa.gov.in/pdf/Rainfed-guar-final-pdf.pdf` on April 16, 2017.

NIEHAUS, P. AND S. SUKHTANKAR (2013): "The marginal rate of corruption in public programs: Evidence from India," *Journal of Public Economics*, 104, 52–64.

O'DELL, K., S. PETERS, AND K. WHARTON (2014): "Women, Energy and Economic Empowerment: Applying a gender lens to amplify the impact of energy access," Tech. rep., Deloitte Consulting LLP.

ORR, I. (2016): "Hydraulic Fracturing: A Game-Changer for Energy and Economies," in *Alternative Energy and Shale Gas Encyclopedia*, John Wiley & Sons, Inc., 700–719.

PATTANAYAK, S. K., M. JEULAND, J. J. LEWIS, V. BHOJVAID, N. BROOKS, A. KAR, L. LIPINSKI, L. MORRISON, O. PATANGE, N. RAMANATHAN, I. H. REHMAN, R. THADANI, F. USMANI, M. VORA, AND V. RAMANATHAN (2016): "Cooking Up Change in the Himalayas: Experimental Evidence on Cookstove Promotion," Duke Environmental and Energy Economics Working Paper № EE 16-03.

PATTANAYAK, S. K., R. A. KRAMER, AND J. R. VINCENT (2017): "Ecosystem change and human health: implementation economics and policy," *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372, 20160130.

PATTANAYAK, S. K. AND A. PFAFF (2009): "Behavior, Environment, and Health in Developing Countries: Evaluation and Valuation," *Annual Review of Resource Economics*, 1, 183–217.

PATTANAYAK, S. K., J.-C. YANG, K. L. DICKINSON, C. POULOS, S. R. PATIL, R. K. MALLICK, J. L. BLITSTEIN, AND P. PRAHARAJ (2009): "Shame or subsidy revisited: social mobilization for sanitation in Orissa, India," *Bulletin of the World Health Organization*, 87, 580–587.

PAYNE, J. W., J. R. BETTMAN, AND E. J. JOHNSON (1992): "Behavioral Decision Research: A Constructive Processing Perspective," *Annual Review of Psychology*, 43, 87–131.

PEREIRA, M. G., J. A. SENA, M. A. V. FREITAS, AND N. F. DA SILVA (2011): "Evaluation of the impact of access to electricity: A comparative analysis of South Africa, China, India and Brazil," *Renewable and Sustainable Energy Reviews*, 15, 1427–1441.

PETERS, J., J. LANGBEIN, AND G. ROBERTS (2018): "Generalization in the Tropics – Development Policy, Randomized Controlled Trials, and External Validity," *The World Bank Research Observer*, 33, 34–64.

PINKOVSKIY, M. AND X. SALA-I-MARTIN (2016): "Lights, Camera … Income! Illuminating the National Accounts-Household Surveys Debate," *The Quarterly Journal of Economics*, 131, 579–631.

PINTO, A. (2016): "A Hard Bargain? A cost-benefit analysis of an improved cookstove program in India," Master's thesis, Duke University, `http://hdl.handle.net/10161/12495`.

Pritchett, L. and J. Sandefur (2014): "Context Matters for Size: Why External Validity Claims and Development Practice do not Mix," *Journal of Globalization and Development*, 4.

Rai, D. K. (2015): "Trends and Economic Dynamics of Guar in India," Indian Council for Research on International Economic Relations Working Paper No. 311.

Ravallion, M. (2009): "Evaluation in the Practice of Development," *The World Bank Research Observer*, 24, 29–53.

Rhodes, E., R. Dreibelbis, E. Klasen, N. Naithani, J. Baliddawa, D. Menya, S. Khatry, S. Levy, J. Tielsch, J. Miranda, C. Kennedy, and W. Checkley (2014): "Behavioral Attitudes and Preferences in Cooking Practices with Traditional Open-Fire Stoves in Peru, Nepal, and Kenya: Implications for Improved Cookstove Interventions," *International Journal of Environmental Research and Public Health*, 11, 10310–10326.

Rosen, S. (1974): "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *Journal of Political Economy*, 82, 34–55.

Rosenbaum, P. R. and D. B. Rubin (1984): "Reducing Bias in Observational Studies Using Subclassification on the Propensity Score," *Journal of the American Statistical Association*, 79, 516–524.

Rosenthal, J., K. Balakrishnan, N. Bruce, D. Chambers, J. Graham, D. Jack, L. Kline, O. Masera, S. Mehta, I. R. Mercado, G. Neta, S. Pattanayak, E. Puzzolo, H. Petach, A. Punturieri, A. Rubinstein, M. Sage, R. Sturke, A. Shankar, K. Sherr, K. Smith, and G. Yadama (2017): "Implementation Science to Accelerate Clean Cooking for Public Health," *Environmental Health Perspectives*, 125.

Schaner, S. (2016): "The Cost of Convenience?" *Journal of Human Resources*, 52, 919–945.

Shannon, A. K., F. Usmani, S. K. Pattanayak, and M. Jeuland (2018): "The Price of Purity: Willingness to Pay for Air and Water Purification Technologies in Rajasthan, India," *Environmental and Resource Economics*, in press.

Sharma, B. P., S. K. Pattanayak, M. Nepal, P. Shyamsundar, and B. S. Karky (2015): "REDD+ Impacts: Evidence from Nepal," South Asian Network for Development and Environmental Economics Working Paper No. 95-15.

Singh, S. K. (2014): "An Analysis of Guar Crop in India," Tech. rep., United States Department of Agriculture.

Singhal, S. and R. Nilakantan (2016): "The economic effects of a counterinsurgency policy in India: A synthetic control analysis," *European Journal of Political Economy*, 45, 1–17.

Smith, K. R. (2014): "In praise of power," *Science*, 345, 603–603.

——— (2017): "Why both gas and biomass are needed today to address the solid fuel cooking problem in India: A challenge to the biomass stove community," *Energy for Sustainable Development*, 38, 102–103.

Somanathan, E. and R. Bluffstone (2015): "Biogas: Clean Energy Access with Low-Cost Mitigation of Climate Change," *Environmental and Resource Economics*, 62, 265–277.

Srivastava, N. and R. Srivastava (2010): "Women, Work, and Employment Outcomes in Rural India," *Economic and Political Weekly*, 45, 49–63.

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

Tapscott, C. E. (2015): "An evaluation of flow and transport properties for hydraulic fracturing fluids in porous medium systems," Master's thesis, The University of North Carolina at Chapel Hill.

Thom, C. (2000): "Use of grid electricity by rural households in South Africa," *Energy for Sustainable Development*, 4, 36–43.

Thombare, N., U. Jha, S. Mishra, and M. Siddiqui (2016): "Guar gum as a promising starting material for diverse applications: A review," *International Journal of Biological Macromolecules*, 88, 361–372.

Tobin, J. (1958): "Estimation of Relationships for Limited Dependent Variables," *Econometrica*, 26, 24.

Tollefson, J. (2013): "Secrets of fracking fluids pave way for cleaner recipe," *Nature*, 501, 146–147.

Tully, S. (2006): "The Human Right to Access Electricity," *The Electricity Journal*, 19, 30–39.

Undersander, D., D. H. Putnam, A. R. Kaminski, K. A. Kelling, J. D. Doll, E. S. Oplinger, and J. L. Gunsolus (1991): "Guar," in *Alternative Field Crops Manual*, University of Wisconsin-Exension, Cooperative Extension; University of Minnesota Center for Alternative Plant & Animal Products and the Minnesota Extension Service.

United Nations (2018): "Accelerating SDG 7 achievement: Policy briefs in support of the first SDG 7 review at the UN High-Level Political Forum 2018," Tech. rep., United Nations Department of Economic and Social Affairs, retrieved from `https://sustainabledevelopment.un.org/content/documents/18041SDG7_Policy_Brief.pdf` on October 18, 2018.

Usmani, F., J. Steele, and M. Jeuland (2017): "Can economic incentives enhance adoption and use of a household energy technology? Evidence from a pilot study in Cambodia," *Environmental Research Letters*, 12, 035009.

Van Houtven, G. L., S. K. Pattanayak, F. Usmani, and J.-C. Yang (2017): "What are Households Willing to Pay for Improved Water Access? Results from a Meta-Analysis," *Ecological Economics*, 136, 126–135.

Venkataraman, C. (2005): "Residential Biofuels in South Asia: Carbonaceous Aerosol Emissions and Climate Impacts," *Science*, 307, 1454–1456.

Vickrey, W. (1961): "Counterspeculation, auctions, and competitive sealed tenders," *The Journal of Finance*, 16, 8–37.

Vivalt, E. (2015): "Heterogeneous Treatment Effects in Impact Evaluation," *American Economic Review*, 105, 467–470.

——— (2017): "How Much Can We Generalize from Impact Evaluations?" Working paper.

Werker, E. and F. Z. Ahmed (2008): "What Do Nongovernmental Organizations Do?" *Journal of Economic Perspectives*, 22, 73–92.

Westfall, P. H. and S. S. Young (1993): *Resampling-Based Multiple Testing: Examples and Methods for p-Value Adjustment*, Wiley-Interscience.

Wheeler, D. and A. Mody (1992): "International investment location decisions," *Journal of International Economics*, 33, 57–76.

World Bank (2000): "Reducing the Cost of Grid Extension for Rural Electrification," Tech. rep., Energy Sector Management Assistance Programme (ESMAP) № ESM 227.

——— (2018a): "Tracking SDG7: The Energy Progress Report," Tech. rep., `http://trackingsdg7.esmap.org/`.

——— (2018b): "World Development Indicators," `https://data.worldbank.org/`.

World Health Organization (2009): "Country profile of Environmental Burden of Disease: Senegal," Tech. rep., World Health Organization, Geneva.

Wysokinska, A. (2017): "Institutions or Culture? Lessons for development from two natural experiments of history," LSE Institute of Global Affairs Research Paper № 03/2017.

# Biography

Faraz Usmani is an environmental economist with research interests at the intersection of environmental, energy and development economics. He earned a Doctor of Philosophy from Duke University in 2019, a Master of Arts from Yale University in 2012, and a Bachelor of Arts from New York University in 2010.

He was a Visiting PhD Fellow at the United Nations University World Institute for Development Economics Research (UNU-WIDER) from February to April 2019. Between 2016 and 2019, he was a Doctoral Student Fellow at the Duke University Energy Initiative, and a Doctoral Scholar at the Duke Global Health Initiative (DGHI).

As of this writing, his peer-reviewed journal publications are as follows:

- Shannon, A. K., F. Usmani, S. K. Pattanayak, and M. Jeuland (2018): "The Price of Purity: Willingness to Pay for Air and Water Purification Technologies in Rajasthan, India," *Environmental and Resource Economics*, in press.
- Van Houtven, G. L., S. K. Pattanayak, F. Usmani, and J.-C. Yang (2017): "What are Households Willing to Pay for Improved Water Access? Results from a Meta-Analysis," *Ecological Economics*, 136, 126–135.
- Usmani, F., J. Steele, and M. Jeuland (2017): "Can economic incentives enhance adoption and use of a household energy technology? Evidence from a pilot study in Cambodia," *Environmental Research Letters,* 12, 035009.

After completing his doctorate, he will join the Charles H. Dyson School of Applied Economics and Management at Cornell University as an Atkinson Postdoctoral Fellow.