DISSERTATION

THREE ESSAYS ON WATER, POLLUTION, AND ENERGY ECONOMICS

Submitted by

Salvador Lurbé

Department of Agriculture and Resource Economics

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Spring 2021

Doctoral Committee:

Advisor: Dale Manning

Jesse Burkhardt
Jordan Suter
Jana Anderson

Copyright by Salvador Lurbé 2021

All Rights Reserved

ABSTRACT

THREE ESSAYS ON WATER, POLLUTION, AND ENERGY ECONOMICS

This dissertation contains three chapters related to the economics of Water, Pollution, and Energy. In Chapter one we investigate how the demand for water responds to conservation efforts based on social comparisons, specifically if the message a household receives affects the way it responds. Using ex post power tests, we demonstrate the need for a significant increase in sample size to apply causal identification strategies to identify heterogeneous impacts using Randomized Controlled Trials (RCTs) that are not specifically designed to identify such effects. Alternatively, RCTs could be designed specifically to identify heterogeneous treatment effects. In Chapter 2 we quantitatively test how household electricity use in rural Rwanda responds to electricity reliability. We examine technology adoption, technology disadoption, and the quantity of electricity purchased. For each model, we focus on the association between the decisions being made and the reliability of the electricity service, which is either experienced or observed depending on whether the household has adopted the electricity technology. We find that poor electricity reliability is a barrier to initial technology adoption and is associated with short-term disadoption decisions, but does not lead to permanent disadoption. The data suggest that households are short-sighted and that households can learn from peers' experiences with the service. Our research suggests that poor electricity reliability can limit willingness to pay for electrification in rural areas of the developing world, where electricity access is lagging behind development goals. In Chapter 3 we study the effects of pollution on crime by looking at the association between pollution, specifically at Particulate Matter (PM) and Ozone (O3), and counts of aggregated crime types and Anti-social Behaviour (ASB) in the UK. We primarily focus our analysis on ASB, as the literature has identified costs associated with them in the UK, but has overlooked its association with pollution. Using a fixed effects model, we find an association between pollution and some crime types, especially those that are economically motivated. We find weaker evidence of an association between pollution and offenses associated with aggressive and violent behavior (including ASB), and we discuss potential mechanisms in the context of the rational choice crime model. We conclude that one potential mechanism could be a decrease in the utility of the non-punishable alternative activity, or a decrease of the offender's value of the future costs associated with being caught.

ACKNOWLEDGEMENTS

Throughout the writing of this dissertation I have received a great deal of support and assistance. I would first like to thank my supervisor, Dr. Dale Manning, for offering me his time and his knowledge but fundamentally his patience, and my co-advisor Dr. Jesse Burkhardt, for his time, his knowledge and for the emotional support when I most needed it. I would also like to thank the rest of my committee, Dr. Jordan Suter and Dr. Jana Anderson, and Dr. Chris Goemans, our great Graduate Coordinator. It is also important to me to acknowledge my colleagues and friends in the program: together we built a community that made us better.

Last but not least, I would like to thank my family. To my wife Jimena, without whom this journey would have never happened. Jimena, you are always there for me. To my parents Mario and Gabriela and my siblings Leandro and Juliana, their wise counsel, their support, their example and their love were key to this and to any achievement I've had in my life.

A special mention to my dog Minga and her furry friends Copper, Charlie and Olive, who were there with me making the long days bearable.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1

Social Comparison and Residential Water Use: Evidence from a Randomized Controlled Trial

1.1 Introduction

As providers of natural resources and energy face increasing scarcity, interest in incentivizing consumer-level conservation has surged. Because of political and legal constraints, such as cost-recovery pricing and equity concerns, many firms and utilities find it difficult to use pricing policies as the sole mechanism for promoting conservation. Instead, utilities typically use a portfolio approach to manage demand, increasingly relying on non-pecuniary interventions to influence behavior. For example, many utilities across the U.S. are now distributing Home Water and Energy Reports (HWERs) to their customers, providing information on how to conserve electricity and water along with a comparison of resource use relative to others.

Previous studies in residential electricity show that receiving HWERs reduces electricity consumption amongst treated households by 2% to 6% [Allcott, 2011,0, Allcott and Kessler, 2019, Ferraro and Price, 2013, Henry et al., 2019, Jessoe et al., 2017]. In the context of water use, previous research has shown that HWERs can reduce consumption by as much as 7% [Brent et al., 2015, Torres and Carlsson, 2018]. Furthermore, some studies have shown that effects appear to persist over time, with no reduction in effect after two years of continual treatment [Allcott, 2011, Brent et al., 2015], while others show declining effects over time [Ferraro and Price, 2013]. Allcott and Rogers [2014] expand on this idea and find that effects become more persistent as the intervention continues over time and Jessoe et al. [2017] demonstrate that treatment effects can spillover beyond the targeted sector. However, the mechanisms behind these results remain unclear. One hypothesis is that households simply respond to being provided information, or to the knowledge that their consumption is being monitored. A competing hypothesis is that households respond to social comparison, either to the injunctive message (e.g., being categorized as "Great", "Good", or "Take Action" relative to one's peers, what we define as the "categorization effect") or to the knowledge that they are using more or less electricity or water than their neighbors. In other words, it remains unclear whether observed treatment effects in HWER studies are due to increased awareness of water use (e.g., simply receiving a letter in the mail or receiving it with the household consumption relative to the peers consumption) or due to the injunctive (normative) message associated with their consumption. Moreover, because water use monitoring is typically limited to monthly water bills, we do not know whether households are conserving indoor water, outdoor water, or both.

There have been attempts in the literature to test the mechanisms explaining the observed reductions. Using quasi-experimental data derived from a randomized controlled trial (RCT), Allcott [2011] finds no significant effect of the message received. Ferraro et al. [2011] show that receiving a social comparison along with conservation tips and a personally addressed letter has a larger effect than only receiving tips and a personal letter. This provides evidence that a social comparison incentivizes conservation, but it does not reveal the impact of the message received. Understanding the mechanism through which HWERs reduce consumption would inform policy makers and utility managers on how to better craft non-price interventions to maximize conservation in a cost effective way. If categories do play a role in how households react to the HWER, they can be designed to induce larger reductions. If categories do not play a role and the response is to information provision, utilities could save money by simply providing households with conservation tips. Likewise, if water conservation is occurring primarily in the growing season as a consequence of reduced outdoor water use, utilities will need to develop alternative strategies to reduce indoor water use during the non-growing season.

To study the heterogeneous effect of normative messages on conservation, it is key to separate the effect of the normative message from the degree of high usage, as high users are more likely to receive strong normative messages. To study the heterogeneous responses to HWERs, Brent et al. [2020] used an experiment in which the normative message was not a function of the usage in absolute terms but in relative terms. In other words, the message was a function of the household's percent change and its position in the distribution of the percent changes of similar households in the same cohort. This way, their experimental design decouples the type of message a household receives from their pre-treatment consumption. With this design, Brent et al. [2020] show that the normative message is a major driver of customer response.

In this paper, we investigate whether the observed treatment effect of HWERs is driven by a "categorization" effect and whether households are conserving indoor or outdoor water. First, we test whether households respond more the HWERs during the growing season or non-growing season. Second, we test whether households respond differently to HWERs based on receiving a message that they are in the categories of "great," "good," or should "take action" when compared to similar households. To do so, we evaluate two years of a RCT of a Home Water Report (HWR) administered by a relatively large city in the Western US. We begin by estimating the Average Treatment Effect (ATE) using a difference-in-differences method, assessing the overall success of the non-pecuniary intervention. We then use quantile regression to examine the heterogeneous relationship between the response to the treatment and the quantiles of consumption. Finally, we evaluate whether being classified in different normative categories has an effect on the response to the treatment. To do so, we employ a regression discontinuity approach as in Allcott [2011], a triple differences approach, and a new approach which we call "Time Variant Categorization", to determine if households respond differently to receiving different normative messages.

Consistent with previous literature [Allcott, 2011,0, Allcott and Kessler, 2019, Brent et al., 2015, Delmas et al., 2013, Ferraro and Price, 2013, Ferraro et al., 2011, Henry et al., 2019, Jessoe et al., 2017, Torres and Carlsson, 2018], we find an ATE of -2.4%. Based on the quantiles of consumption, we find that treatment effects are negative for all quantiles and significant across the entire consumption distribution, except for those consumers at the very top or at the very bottom of the distribution. Consistent with Allcott [2011], we find that the categorization effect is not significant but we perform ex post power analyses that reveal low power when applying common econometric techniques within the RCT designed to estimate average effects. We also provide evidence that most of the reduction in consumption happens during the growing season, consistent with findings from previous research, such as Ferraro and Price [2013] and Mini et al. [2015]. These latter results suggest that much of the treatment effect is occurring in outdoor water use, and HWERs might be incentivizing households to make larger capital intensive investments in xeriscaping.

Our results confirm the success of HWERs in decreasing resource consumption, both on average and across a wide spectrum of users. However, we fail to produce evidence that the message received impacts responses to the treatment. This suggests that future research designs are still required to understand the behavioral mechanisms through which social comparisons affect resource use.

This work contributes to two key literatures. First, it builds on a growing literature that measures the experimental effects of behavioral interventions in resource economics. Previous studies have shown that Home Energy Reports (HERs) decrease mean electricity use [Allcott, 2011,0, Allcott and Kessler, 2019, Ferraro et al., 2011, Henry et al., 2019, Jessoe et al., 2017] and HWRs decrease water consumption [Brent et al., 2015, Ferraro and Price, 2013, Torres and Carlsson, 2018]. However, some studies have found no significant treatment effects given certain conditions: Burkhardt et al. [2019c] could not find evidence of a reduction in electricity consumption during peak hours on hot summer days in Austin, Texas, while Myers and Souza [2020] did not find evidence of a reduction in electricity consumption in university dorms, where residents do not pay for the service. Furthermore, Delmas et al. [2013] review the literature on the effect of information strategies on energy conservation, including social comparisons, concluding that strategies based on social comparisons are, on average, more effective than the rest of the information provision strategies aiming to lower consumption. Our paper provides further evidence of the success of HWRs to decrease consumption.

Second, this paper builds on the literature on RCTs. RCTs have gained more traction in economics research, especially in the development field, where the work of Abhijit Banerjee, Esther Duflo and Michael Kremer received the 2019 Nobel Prize in Economics for using experiments to assess policies to reduce poverty. Even when some researchers argue that evidence from randomized experiments have no special priority [Deaton, 2010], most researchers agree that randomized experiments are more reliable for identifying causal effects [Banerjee et al., 2016, Freedman, 2006]. The theoretical foundation for using randomization as a means to identify causal effects began when Fisher [1925] proposed that, to ensure no observable characteristics of the units are reflected in the assignment, units should be allocated randomly to different treatment groups. This way, differences between treatment and control units should reflect the impact of the treatment. As economic research has increasingly relied on RCTs to study causal effects, this has been specially true to assess the effect of non-pecuniary interventions. Furthermore, Arimura et al. [2016] demonstrate the importance of controlling for simultaneity, common unobserved characteristics, and nonrandom group formation when identifying the effect of social norms on individual behavior. Otherwise one may overestimate the influence of social norms on individual environmental behavior. Our work shows the limits of using experimental data to answer a question when the experiment was not originally designed for it. Palm-Forster et al. [2019] identified underpowered designs and multiple hypothesis testing as key challenges when assessing agri-environmental programs, and they emphasised the need for a more detailed planning during the experimental design, including a power analysis. In this work, we empirically show that the power of the test for heterogeneous effects is extremely low given the relatively small sample collected as a result of the power needed to estimate average effects.

1.2 Experiment and Hypotheses

1.2.1 Experimental Design and Context

In this study, we evaluate an RCT that began in September 2014, and for which we observe two years of data, ending in September 2016.[1] For the experiment, 7,000 households were randomly chosen for the treatment group and 4,000 households for the control group. Households in the

[1] The Utility implemented the RCT and allowed us to evaluate the results. Due to a data sharing agreement, we are not allowed to disclose the study location.

treatment group received a Home Water Report (HWR) every two months, starting at the beginning of the experiment. During our sample period, households in the treatment group received 12 HWRs (two treatment years). On the other hand, households in the control group did not receive HWRs. HWRs include a water score, a comparison to other similar households, and conservation tips specially crafted for each household type, or cohort, described below. Figure 1.1 is an example of the Water Score received by a household while Figure 1.2 is an example of the conservation tips.

Figure 1.1: Example of a water score.

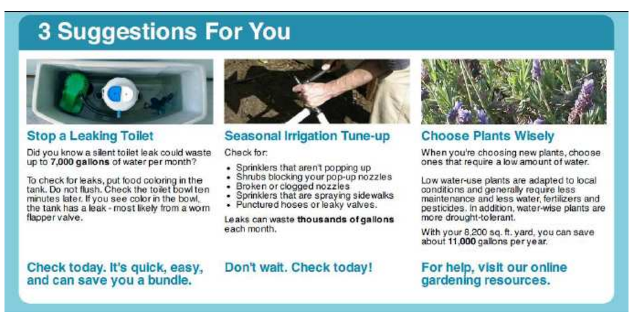

Figure 1.2: Example of conservation tips.

To provide households with water scores, each household was assigned to a cohort of similar houses. Each cohort is fixed across time and was defined by the household's number of occupants, the size of the irrigable area, and the type of residence (e.g., single family home or multi-unit). For each cohort-month, the 52nd and the 20th percentiles of consumption were calculated. Each household's consumption was then compared to the 52nd and 20th percentiles of the cohort's consumption and the household was given a water score based on their relative consumption. For example, if the household's consumption was less than the cohort-specific 20th percentile of consumption, the utility assigned a water score of 1 and provided the message "Great". If the household's consumption was greater than the cohort-specific 20th percentile of consumption but less than the 52nd percentile, the utility assigned a water score of 2 and provided the message "Good". Finally, if the household's consumption was greater than the cohort-specific 52nd percentile of consumption, the utility assigned a water score of 3 and provided the message "Take Action". Messages are calculated every month, but treated households observe water scores once every two months (one is calculated but not observed). Unfortunately, the utility did not keep data on the cohort composition, which means that even though we observe the 20th percentile and the median (50th percentile) consumption for each cohort, it is not possible for us to back out the 52nd percentile of consumption. In other words, we do not observe the cutoff between water scores 2 and 3. Based on this limitation, we primarily focus our analysis of a "categorization effect" on the distinction between "Great" and "Good", but we explore several methods for identifying a "categorization effect" between the "Good" and "Take Action" cohorts.

The experimental design allows us to test three main hypotheses. Our first hypothesis is that receiving a new HWR every two months reduces water consumption among the treated houses relative to the control houses on average. If true, then the results would suggest that households in our study respond to social comparison, information provision, or both. Furthermore, we test if the reduction in water consumption happens mostly during the summer, as has been found in other contexts Kenney et al. [2008]. Our second hypothesis is that more efficient households, or households that use less water than similar households in their cohort during the pre-treatment period, respond less to the treatment than households that use more water than similar households in their cohort during the pre-treatment period. Allcott [2011] showed this was true in the context of residential energy conservation. Finally, we hypothesize that the treatment message matters, meaning that households that receive a "Take Action" message will respond more than households that receive a "Good" message and those that receive a "Good" message will respond more than those that receive a "Great" message.

1.2.2 Data

Our primary data consists of monthly water consumption in gallons per day.[2] We also observe each household's cohort's 20th percentile and median consumption and their designated water score. Pre- and post-treatment summary statistics are as shown in Table 1.1.[3]

Table 1.1: Pre and Post Treatment Consumption

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Pre-treatment consumption (gpd)	79,208	290.585	236.630	2	135.4	372	4,179
Post-treatment consumption (gpd)	237,624	271.434	258.159	1.515	115.714	332	5,180

Notes: The pre-treatment consumption we observe is from January 2014 to August 2014, while the post-treatment consumption is from September 2014 to September 2016.

An RCT is built on the assumption that, in the absence of a treatment, both treatment and control groups would have followed the same trend [Angrist and Pischke, 2008, Deaton and Cartwright, 2017, Ferraro and Hanauer, 2014]. Therefore, our primary identification assumption is based on the treatment and the control groups being random samples from the same population (in this case, the universe of residential households in the municipality for which water is provided). While we assume this assumption was originally true, some households opted out of the experiment or were removed from it because they moved, resulting in 3,713 households in the treatment group and 6,441 in the control group. Our data do not include information about those households that left, which could be a source of selection bias if households that opted out have characteristics that influence how they would respond to the treatment. Furthermore, we flagged 90 households

[2] The utility reports total consumption over the billing cycle. As billing cycles are different across customers, we normalize consumption by dividing the billing cycle water use by the number of days in the cycle.

[3] The pre-treatment period is January 2014 to August 2014 and the post-treatment period is September 2014 to September 2016.

in the control group and 163 households in the treatment group as having non-valid consumption values (0s, negative values or extremely large values), which were not included in the analysis.

To test the validity of the parallel trends assumption after attrition and cleaning, we first compare sample pre-treatment means and standard deviations across the treatment and control groups (e.g., 291.45 gpd for the control group and 290.09 gpd for the treated group). The results of these tests are as shown in Table 1.2. For the two sided t-test for differences in pre-treatement means, the calculated p-value is 0.43. Thus, we fail to reject the null hypothesis that the mean pre-treatment consumption of the treatment and control groups are equal. For the two sided F-test for standard deviations, the p-value is 0.64, so we also failed to reject the null hypothesis that the ratio of pre-treatment standard deviations for treatment and control is equal to 1.

To further test the parallel trends assumption, we use local county assessor data (which we do not identify to protect the identity of the utility) to compare the characteristics of the households in the treatment and control groups. If we look at house types, 98.05% of the houses in the control group and 98.18% in the treatment group are Single Family Residential, while the rest of the samples are a mixture of mostly "Patio Home Single Family Residential", "Townhouses" and "Townhouse - 1/2 Duplex". Looking at a more detailed description, 39.37% of the houses in the control group and 40.22% in the treatment group are houses with two stories, 33.91% of the houses in the control group and 33.49% in the treatment group are ranches, 16.18% of the houses in the control group and 16.48% in the treatment group are "split level" houses, and 6.58% of the houses in the control group and 6.04% in the treatment group are "bi-Level two stories" houses. The mean square footage (std. errors in parenthesis) is 1,774.3 (542) for houses in the control group and 1,792.2 (540) in the treatment group. The mean lot size in square foot (std. errors in parenthesis) are 9,162 (5,064) and 9,348 (23,689) for control and treatment groups respectively. If we leave an outlier in the control group out of this test (1,756,115 square foot), in the control group, the mean square footage (and standard error) for the treatment group are 9,039 (4,701).

Having shown that means and standard deviations of consumption cannot be proven to be statistically different for treatment and control groups, and that the characteristics of the households are similar across both groups, we have no evidence to suggest that the groups were not the result of random sampling. These findings provide evidence in favor of the assumption of parallel trends, which is the same as assuming orthogonality of the treatment to other potentially confounding factors affecting the water consumption, such as household size. As the expected value of the differences between the means of all other variables affecting both groups is then zero, before presenting our more complete model for average treatment effect we start the analysis by performing a simple difference-in-differences means test to determine if there are statistically significant responses to the treatment. To do so, we compare the difference in treated and control consumption after treatment relative to before treatment. The results are reported in the bottom right quadrant of Table 1.2. We find a statistically significant reduction of 6.74 gallons per day.

Table 1.2: Difference-in-differences

GROUPS	(1) Pre-treatment (8 months)	(2) Post-treatment (24 months)	(3) Difference
Control	291.45	276.57	14.88
	(1.43)	(0.90)	(1.70)
Treatment	290.09	268.47	21.62
	(1.04)	(0.65)	(1.22)
Difference	1.36	8.10	-6.74
	(1.77)	(1.11)	(2.09)

Notes: Values are means by group. Standard deviations in parentheses. Units are gallons per day.

Across the sample, the distribution of the observed normative categories is as expected given the percentiles used for the cutoffs: we observe 45,632 "Great" messages (16.7%), 96,814 "Good" messages (35.4%) and 131,023 "Take Action" messages (47.9%). On the other hand, the mean consumption for those who receive a "Great" in a given month was 93 gallons per day, 191 gallons per day for those who received a "Good" and 428 gallons per day for those who received a "Take Action". Summary statistics are as shown in Table 1.3.

The mean number of occupants in the household is quite similar across messages, with averages of 3.05, 2.94 and 2.79 for the groups receiving a "Great", "Good" and "Take Action" messages

Table 1.3: Summary Statistics Conditional on the Assigned Waterscore

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Consumption if "Great" (gpd)	44,086	92.853	49.024	1.515	63.548	107.143	400.323
Consumption if "Good" (gpd)	94,969	191.175	112.377	30.938	115.667	237.097	910.323
Consumption if "Take Action" (gpd)	128,272	427.846	310.274	58.065	202.188	569.667	5,180.000

Notes: The pre-treatment consumption we observe is from January 2014 to August 2014, while the post-treatment consumption is from September 2014 to September 2016.

respectively. Furthermore, the size of the irrigable area is also very similar across households receiving different normative messages. About 50% of the houses that received any normative message have a "Large Outdoor Area", about 40% have a "Medium Outdoor Area", about 5% have a "Small Outdoor Area", and finally about 5% have a "Extra Large Outdoor Area". As the number of occupants and the size of the irrigable area define the cohort that a household belongs to, this balance across messages is expected.

1.3 Methodology and Results

The methodology and results are presented in three sections. First we estimate the Average Treatment Effect and the summer versus non-summer treatment effect. Second, we estimate heterogeneous treatment effects based on pre-treatment consumption patterns. Finally, we explore several specifications to test for the existence of a categorization effect.

1.3.1 Test of Hypothesis 1: Average Treatment Effect

To estimate the average treatment effect, we use a difference-in-differences model as follows:

$$log(gpd)_{it} = \beta_1 P_t + \beta_2 T_i + \tau T_i P_t + \mu_{my} + \upsilon_i + \epsilon_{it}, \tag{1.1}$$

where $log(gpd)_{it}$ is the logarithm of household i's water consumption in month t, T_i is a treatment indicator equal to 0 for the control group and 1 for the treatment group, P_t is a post-treatment indicator equal to 0 in the pre-treatment period and 1 in the post-treatment period, υ_i is a household fixed effect, and μ_{my} is a month-by-year fixed effect.[4]. Finally, ϵ_{it} is a normally distributed error term. The parameters are estimated using ordinary least squares and standard errors are clustered at the household level to allow for intra-household dependence across months. Finally, if we assume there was no contamination between the treatment and control groups [Angrist and Pischke, 2008, Cox, 1958], otherwise known as the "Stable unit treatment value assumption" (SUTVA), the coefficient associated with the interaction between T_i and P_t (τ), is an unbiased estimate of the Average Treatment Effect.

The use of water in western US cities varies considerably across summer and winter months, as irrigation plays an important role in total household water use during the summer months. Reducing water consumption is therefore different in summer months than during the rest of the year, and we can expect the results of the intervention to be different as well. We use a difference-in-difference-in-difference (DDD) estimator to compare the results of the intervention during summer months (from May to September) with the rest of the year, as shown in Equation 1.2.[5]

$$\begin{gathered} log(gpd)_{it} = \beta_0 + \beta_1 P_t + \beta_2 T_i + \beta_3 I(t \in [5,6,7,8,9]) \\ +\beta_4 T_i I(t \in [5,6,7,8,9]) + \beta_5 P_t I(t \in [5,6,7,8,9]) + \beta_6 T_i P_{it} \\ +\beta_7 T_i P_{it} I(t \in [5,6,7,8,9]) + \mu_{my} + \upsilon_i + \epsilon_{it} \end{gathered} \tag{1.2}$$

P_t and T_i and ϵ_{it} are as defined for the previous model, while $I(t \in [5,6,7,8,9])$ indicates if month t is between May and September. In this setting, β_7 can be interpreted as the expected proportional impact of HWRs between May and September compared to the rest of the year.

[4]The unit of observation is the household-month. In the post-treatment period treated households receive the treatment every two months. To account for this, we assign the treatment in period t to periods t and $t+1$.

[5]While we realize that May through September extend into the spring and fall, we use the term summer throughout for consistency.

Table 1.4 presents the results of estimating Equation 1.1 in columns 1, 2 and 3. In column 4 we present the results of estimating Equation 1.2. The coefficients on T and P are omitted due to multicollinearity with the fixed effects. Column 1 presents our preferred specification with both household and month-by-year fixed effects, household fixed effects are removed in column 2, and month-by-year fixed effects are removed in column 3. The estimated treatment effect (ATE) is stable across specifications, again indicating that the treatment is not correlated with omitted household characteristics.

Table 1.4: Primary Results

	(1) ATE	(2) ATE	(3) ATE	(4) Summer
$T_i P_{i,t}$	-0.0243***	-0.0235***	-0.0243***	-0.015471***
	(0.00464)	(0.00834)	(0.00466)	(0.006308)
$T_i I(t \in [5,6,7,8,9])$	-	-	-	0.012400
	-	-	-	(0.012135)
$P_{i,t} I(t \in [5,6,7,8,9])$	-	-	-	0.235238***
	-	-	-	(0.006083)
$T_i P_{i,t} I(t \in [5,6,7,8,9])$	-	-	-	-0.023615***
	-	-	-	(0.007517)
Household FE	Y	N	Y	Y
Month by year FE	Y	Y	N	Y
Std. errors in parentheses, clustered by hh's.	Y	Y	Y	Y
Observations	318,832	318,832	318,832	318,832
R-squared	0.708	0.449	0.264	0.669

Notes: Standard errors in parentheses, clustered by household.
Star values: * 0.10 ** 0.05 *** 0.010.

We find a statistically significant Average Treatment Effect (ATE) of -2.4%, confirming our first hypothesis presented in Section 1.2. Incidentally, the treatment effect identified in Table 1.2 is -6.74/290, which is approximately -2.3%. To provide context, we compare this value to the change in price needed to achieve a similar reduction. According to Arbués et al. [2003], the elasticity of residential water consumption ranges from -0.1 to -1.8. Thus, our estimate suggests that to achieve similar reductions in consumption, prices would have to increase by 0.2% to 4.3%. Notably, Allcott [2011] finds that Home Energy Reports result in a 2% reduction in residential energy consumption.

It is important to note that we cannot directly extrapolate our results to other regions, but because households in our sample were randomly selected, we can extrapolate the treatment effect to the population from which the sample was drawn. In doing so, we find that if 50,000 households in the region were provided with HWRs, water consumption would decrease by 121 million gallons per year.

Finally, our results show (column 4 of Table 1.4) that the expected reduction in summer 2.3% larger than the expected reduction in non-summer (1.5%), indicating that most of the reduction happens during the hotter months of the year. This result is not surprising, given that the marginal cost of reducing water consumption is less during the summer, when water is used more widely for indoor and outdoor use. Furthermore, it is consistent with the reduction in water consumption observed by previous studies, such as Ferraro and Price [2013], Kenney et al. [2008] and Mini et al. [2015]. This result has some important policy implications. Assuming the results from column 4 of Table 1.4 are homogeneous across the distribution of customers (the difference in reduction between hot and colder months is not a function of the degree of high usage), we can calculate the expected savings by degree of high usage based on the disaggregated consumption in Table 1.5. The expected savings for higher use customers (households whose calculated water score equals 3), consuming an average of 236 gpd from October through April (colder months) and an average of 583 gpd from May through September, would be of 3 gpd during the colder months and of 22 gpd during the warmer months. For medium use customers (water score equals 2), the expected savings are 2 and 10 gpd during the colder and warmer months respectively, and finally 1 and 5 gpd for the lower use customers (water score equals 1). As we can expect the marginal cost of the program to be constant across households (mailing), the difference between marginal cost and marginal benefit is substantially larger for higher use customers during the summer months. In other words, the net benefits of the treatment are largest for the highest water users during the summer months.

Table 1.5: Average consumption by assigned water score and by colder and warmer months of the year

Water Score	Months September through May	Months October through April
1	134	74
2	273	127
3	583	236

Notes: Units are gallons per day.

1.3.2 Test of Hypothesis 2: Heterogeneous Treatment Effects

In the previous section we estimated the ATE across the sample of households in our data. In this section we use quantile regression to study heterogeneity in the treatment effect at different points in the conditional distribution of the outcome variable. We implement the conditional quantile estimator for panel data, developed by Powell [2016] on the normalized sample, using robust standard errors, on the difference-in-difference specification described in Equation 1.1.

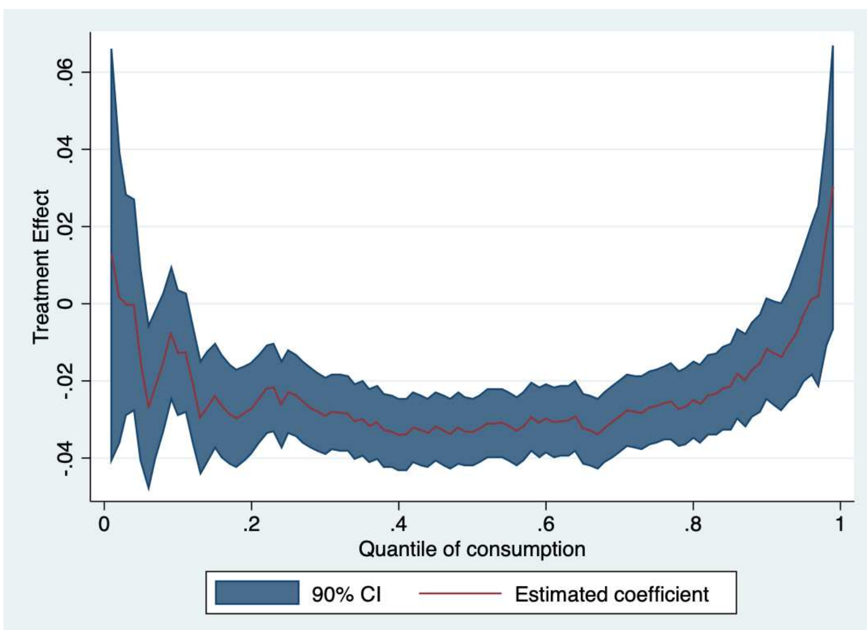

Figure 1.3: Quantile Regression: Treatment effects are percentages

Figure 1.3 presents the results of estimating the quantile regression. We find that quantile treatment effect point estimates are negative and statistically different from zero at a 90% confidence level across all quantiles, except for those at the very bottom or very top of the distribution. Point estimates range from 1% for those consumers in the lower quantiles of consumption to 3% for those in the larger quantiles, with a stable treatment effect of around -3% across quantiles between 0.1 and 0.9. However, we could not reject the hypothesis that coefficients at varying quantiles are statistically the same, indicating that even though we observe a U-shaped curve on the point estimates, the relationship between the household's consumption level and treatment effect appears constant. Moreover, consistent with Allcott [2011], we did not find evidence of a "Boomerang Effect", that would imply that the most efficient users increase their consumption in response to the HWRs.

1.3.3 Test of Hypothesis 3: Categorization Effect

In this section, we perform a series of tests to explore our final hypothesis: whether the treatment effect is a function of the social comparison message (categorization) that a household in the treatment group receives. As the RCT was not designed for this purpose, categorization was not randomized across participants and is endogenous. To address this identification challenge, we employ several different strategies, which we outline below.

Great to Good

We begin by looking for a categorization effect between the "Great" and "Good" categories. We employ three methods to test for this effect. First, we use a regression discontinuity (RD) approach, where we test whether there is a change in the parameters of the functional form from one side of the cutoff to the other. Second, we use a triple difference estimator to test if, within a particular bandwidth, the treatment effect depends on being categorized as "Great" or "Good". Third, we exploit the exogenous recategorization based on peers' responses to the treatment.

Regression Discontinuity: An RD identification strategy is quasi-experimental [Lee and Lemieux, 2010] as it exploits the arbitrary nature of the treatment when assigning water scores. For example, we observe a clear cutoff between the "Great" to "Good" categories, which allows us to assume that, within a certain bandwidth, households on one side of the cutoff are statistically identical to those on the other side of the cutoff, and that the receipt of a particular normative message is simply a product of random variation in month-to-month consumption among households in the cohort. Hence, the only meaningful difference across households within this particular bandwidth is the water score they have received [Angrist and Pischke, 2008]. We exploit this assumption to test for a categorization effect on the cutoff from "Great" to "Good".

The model used for the estimation is as follows:

$$log(gpd)_{it} = \beta_0 + \rho I(D20_{i-t} > 0) + \beta_1 D20_{i-t} I(D20_{i-t} < 0)$$

$$+\beta_2 D20_{i-t} I(D20_{i-t} > 0) + \mu_{my} + \upsilon_i + \epsilon_{it} \tag{1.3}$$

$$\forall : T_i = 1, P_{it} = 1, |D20_{i-t}| < h, |D20_{i-t}| \neq 0$$

where $D20_{i-t}$ is the distance to the 20th percentile cutoff in the previous month, $t-1$. For example, if the 20th percentile cutoff is 100 gpd, and household i consumed 110 gpd in month $t-1$, $D20_{i-t} = 10$. The term $I(D20_{i-t} > 0)$ is an indicator function that is equal to 1 when $D20_{i-t} > 0$, and 0 otherwise. The term $I(D20_{i-t} < 0)$ is an indicator function that is equal to 1 when $D20_{i-t} < 0$, and 0 otherwise. If $D20_{i-t} > 0$, then the assigned water score in t is 2, if $D20_{i-t} < 0$, then the assigned water score in t is 1. The variable h is half the bandwidth in which we are assuming that households are statistically identical. Both $D20_{i-t}$ and h are measured in gallons per day. Finally, ϵ_{it} is the normally distributed error term.

To explain the mechanism behind the RD setting, we use the expected value operator in Equation 1.3, conditional on observing a positive or a negative distance to the cutoff calculated from the previous month's consumption:

$$E[log(gpd)_{it}|(D20_{i-t} < 0)] = \beta_0 + \beta_1 D20_{i-t} + \mu_{my} + \upsilon_i \tag{1.4}$$

$$E[log(gpd)_{it}|(D20_{i-t} > 0)] = \beta_0 + \rho + \beta_2 D20_{i-t} + \mu_{my} + \upsilon_i \tag{1.5}$$

To the left of the cutoff (Equation 1.4), the intercept is β_0 and the slope is β_1 while, to the right of the cutoff (Equation 1.5), the intercept is $\beta_0 + \rho$ and the slope is β_2. The parameters of interest are then ρ and the difference between β_1 and β_2. The parameter ρ can be seen as the discontinuity in the intercept while, if there is a significant difference between β_1 and β_2, the regression discontinuity is manifested as a change in the slope.

We also estimate an RD specification but without allowing for different slopes, as follows:

$$log(gpd)_{it} = \beta_0 + \rho I(D20_{i-t} > 0) + \beta_1 D20_{i-t} + \mu_{my} + \upsilon_i + \epsilon_{it} \tag{1.6}$$

$$\forall : T_i = 1, P_{it} = 1, |D20_{i-t}| < h, |D20_{i-t}| \neq 0$$

If we use the expected value operator in Equation 1.6, conditional on observing a positive or a negative distance to the cutoff calculated from the previous month's consumption:

$$E[log(gpd)_{it}|(D20_{i-t} < 0)] = \beta_0 + \beta_1 D20_{i-t} + \mu_{my} + \upsilon_i \tag{1.7}$$

$$E[log(gpd)_{it}|(D20_{i-t} > 0)] = \beta_0 + \rho + \beta_1 D20_{i-t} + \mu_{my} + \upsilon_i \tag{1.8}$$

To the left of the cutoff the intercept is β_0 while, to the right of the cutoff, the intercept is $\beta_0 + \rho$. The slope is, on both sides, β_1. The parameter of interest is then ρ, representing a location shift in the underlying populations on either side of the cut-off.

Table 1.6 presents, in columns 1 and 2, the results of estimating Equation 1.3 (allowing for different slopes). In columns 3 and 4 we present the results of estimating Equation 1.6 (without allowing for different slopes). In both cases, we present the results for two different values of the bandwidth h.

Table 1.6: Regression Discontinuity

	(1) h=15	(2) h=30	(3) h=15	(4) h=30
$D20_{i-t}I(D20_{i-t} < 0)$	0.00288**	0.00349***	-	-
	(0.00146)	(0.000546)	-	-
$D20_{i-t}I(D20_{i-t} > 0)$	0.00463***	0.00387***	-	-
	(0.00140)	(0.000438)	-	-
$I(D20_{i-t} > 0)$	-0.0105	-0.00364	-0.0107	-0.00405
	(0.0170)	(0.0111)	(0.0170)	(0.0111)
$D20_{i-t}$	-	-	0.00379***	0.00371***
	-	-	(0.00101)	(0.000348)
Different Slopes	Y	Y	N	N
Same Slope	N	N	Y	Y
Household FE	Y	Y	Y	Y
Month by year FE	Y	Y	Y	Y
Observations	9,756	19,735	9,756	19,735
R-squared	0.711	0.692	0.711	0.691

Notes: The dependent variable in each specification is log(gpd). The unit of observation is a household-month. Standard errors in parentheses, clustered by household.
Star values: * 0.10 ** 0.05 *** 0.010.

The results indicate that ρ is, for both specifications and for every considered bandwidth h, not statistically different from zero, indicating that there is no regression discontinuity manifested as a change in intercept. Slopes are, in both cases, positive. This result is as expected since it is capturing serial autocorrelation. For example, we expect that the distance to the cutoff in month $t-1$ for house i is correlated with the distance to the cutoff in month t for house i. Finally, when we allow for different slopes, the difference between the slope on the negative side of the cutoff and on the positive side of the cutoff (β_1 and β_2), is not statistically significant [6]. In summary, we

[6] When h=15, the differences between the estimated slopes is -0.0017 and the standard error for that difference is 0.002. When h=30, the difference is -0.0005 and the standard error is 0.0007

do not find evidence of a regression discontinuity indicating a categorization effect either in the shape of a "bump" in the linear fit (ρ) nor in a change in slope.

Triple Difference: With the RD approach, we relied on the assumption that treated households within a specified bandwidth of the discontinuity were statistically similar. On the other hand, with the difference-in-differences approach (Equation 1.1), we assume orthogonality of the treatment (receiving HWRs) to other potential confounders, and we rely on the RD assumption that leaving only those observations within a small bandwidth ensures a random categorization. The three differences include before versus after, treated versus control, and negative versus positive distances to the 20th percentile cutoff (water score equal to 1 or 2). We use the indicator variable $I(WS_{i-t} = 2)$ for the third difference ($I(WS_{i-t} = 2) = 1$ if $WS_{i-t} = 2$ and $I(WS_{i-t} = 2) = 0$ if $WS_{i-t} = 1$). As before, ϵ_{it} is the normally distributed error term. This strategy can potentially strengthen the results from the RD by incorporating the control group. The triple difference specification is as follows:

$$\begin{gathered}
log(gpd)_{it} = \beta_0 + \beta_1 P_t + \beta_2 T_i + \beta_3 I(WS_{i-t} = 2) \\
+\beta_4 T_i I(WS_{i-t} = 2) + \beta_5 P_t I(WS_{i-t} = 2) + \beta_6 T_i P_{it} \\
+\beta_7 T_i P_{it} I(WS_{i-t} = 2) + \mu_{my} + \upsilon_i + \epsilon_{it} \\
\forall : |D20_{i-t}| < h, |D20_{i-t}| \neq 0
\end{gathered} \tag{1.9}$$

Again, we use the expectation operator to explain the intuition behind our triple differences specification. For those who received a water score equal to 2 and were treated:

$$E[log(gpd)_{it}|T_i = 1, P_t = 0, WS_{i-t} = 2] = \beta_0 + \beta_2 + \beta_3 + \beta_4 + \mu_{my} + \upsilon_i \tag{1.10}$$

$$E[log(gpd)_{it}|T_i = 1, P_t = 1, WS_{i-t} = 2] = \beta_0 + \beta_1 + \beta_2 + \beta_3 + \beta_4$$

$$+\beta_5 + \beta_6 + \beta_7 + \mu_{my} + \upsilon_i \tag{1.11}$$

The difference between the pre- and post-treatment periods expected consumption is $\beta_1 + \beta_5 + \beta_6 + \beta_7$. For those who received a water score equal to 2 and were untreated:

$$E[log(gpd)_{it}|T_i = 0, P_t = 0, WS_{i-t} = 2] = \beta_0 + \beta_3 + \mu_{my} + \upsilon_i \tag{1.12}$$

$$E[log(gpd)_{it}|T_i = 0, P_t = 1, WS_{i-t} = 2] = \beta_0 + \beta_1 + \beta_3 + \beta_5 + \mu_{my} + \upsilon_i \tag{1.13}$$

The difference between the pre- and post-treatment periods expected consumption is $\beta_1 + \beta_5$, and the overall difference-in-differences for those who received a water score equal to 2 is $\beta_6 + \beta_7$. Finally, for those who received a water score equal to 1, the overall difference-in-difference is β_6. The triple difference parameter is then β_7, which represents the additional treatment effect from receiving a "Great" instead of a "Good" message.

The results of estimating Equation 1.9 are presented in Table 1.7 for different bandwidths h. The triple differences parameter is not statistically different from zero for any h considered. Therefore, these results do not support the existence of a categorization effect.

Time Variant Category: The time variant category (TVC) identification strategy exploits plausibly exogenous variation in the normative message a household receives. Specifically, we define a household as treated if the household received the message that their consumption was “Great” in month t and then received the message that their consumption was “Good” in month $t + 2$, not because the household’s consumption changed, but rather because of a shift in the 20th per-

Table 1.7: Difference in Difference in Difference

	(1) h=15	(2) h=20	(3) h=25	(4) h=30	(5) h=35
$I(WS_{i-t} = 2)$	0.00527	0.0250	0.0567	0.0513	0.0560*
	(0.0489)	(0.0406)	(0.0360)	(0.0327)	(0.0302)
$T_i P_{it}$	-0.0468	-0.0498	-0.0369	-0.0179	-0.0129
	(0.0477)	(0.0404)	(0.0357)	(0.0325)	(0.0304)
$T_i I(WS_{i-t} = 2)$	0.0266	0.0234	-0.0148	0.00711	0.0121
	(0.0643)	(0.0526)	(0.0461)	(0.0417)	(0.0385)
$P_t I(WS_{i-t} = 2)$	0.0414	0.0421	0.0232	0.0368	0.0445
	(0.0497)	(0.0413)	(0.0366)	(0.0333)	(0.0308)
$T_i P_{it} I(WS_{i-t} = 2)$	-0.0258	-0.0297	0.0170	0.00258	-0.00378
	(0.0653)	(0.0536)	(0.0470)	(0.0425)	(0.0393)
Household FE	Y	Y	Y	Y	Y
Month by year FE	Y	Y	Y	Y	Y
Observations	21,579	29,100	36,233	42,913	49,122
R-squared	0.704	0.693	0.684	0.676	0.669

Notes: The dependent variable in each specification is log(gpd). The unit of observation is a household-month. Standard errors in parentheses, clustered by household.
Star values: * 0.10 ** 0.05 *** 0.010.

centile cutoff.[7] In other words, a household could receive a different normative message from one treatment period to the next simply because the consumption of their peer cohort increased or decreased.

To estimate the effect of this treatment, we use a triple-difference strategy, with the three differences defined as follows: the first difference is the exogenous categorization change, where a household is defined as treated, $T_{it} = 1$, if the household received a water score of 2 in month t, but they would have received a water score of 1 had the 20th percentile cutoff not changed between $t - 2$ and t. On the other hand, a household is defined as not treated, $T_{it} = 0$, if the household received a water score of 2 in month t and they would have received the same water score of 2 with the previously observed cut off in $t - 2$. The second difference is before and after this treatment (P_{it}). The third difference is if the household is in the broader treatment or control group ($I(T_i = 1)$). Finally, ϵ_{it} is the normally distributed error term.

[7] Note that households receive water scores every two months but we observe consumption every month. Thus, the next water score would be received in $t + 2$, or the previous water score would have been received in $t - 2$.

The triple difference specification is as follows:

$$log(gpd)_{it} = \beta_0 + \beta_1 P_{it} + \beta_2 T_{it} + \beta_3 I(T_i = 1)$$
$$+\beta_4 T_{it} I(T_i = 1) + \beta_5 P_t I(T_i = 1) + \beta_6 T_{it} P_{it} \tag{1.14}$$
$$+\beta_7 T_{it} P_{it} I(T_i = 1) + \mu_{my} + \upsilon_i + \epsilon_{it}$$

Again, we can use the expectation operator to explain the intuition behind this equation. For those whose categorization changed exogenously ($T_{it} = 1$) and were treated ($I(T_i = 1) = 1$):

$$E[log(gpd)_{it}|T_i = 1, P_{it} = 0, T_{it} = 1] = \beta_0 + \beta_2 + \beta_3 + \beta_4 + \mu_{my} + \upsilon_i \tag{1.15}$$
$$E[log(gpd)_{it}|T_i = 1, P_{it} = 1, T_{it} = 1] = \beta_0 + \beta_1 + \beta_2 + \beta_3 + \beta_4$$
$$+\beta_5 + \beta_6 + \beta_7 + \mu_{my} + \upsilon_i \tag{1.16}$$

The difference of the expected consumption between these two groups is $\beta_1 + \beta_5 + \beta_6 + \beta_7$. For those whose categorization changed exogenously and were untreated:

$$E[log(gpd)_{it}|T_i = 0, P_{it} = 0, T_{it} = 1] = \beta_0 + \beta_3 + \mu_{my} + \upsilon_i \tag{1.17}$$
$$E[log(gpd)_{it}|T_i = 0, P_{it} = 1, T_{it} = 1] = \beta_0 + \beta_1 + \beta_3 + \beta_5 + \mu_{my} + \upsilon_i \tag{1.18}$$

The difference between these two groups is $\beta_1 + \beta_5$, and the overall difference-in-difference for those whose categorization changed exogenously is $\beta_6 + \beta_7$. Finally, for those whose categorization did not change exogenously, the overall difference-in-difference is β_6. The triple difference parameter is then β_7, which represents the extra treatment effect given by receiving the exogenous treatment based on changes in the consumption of the cohort.

We allow here for two different specifications: we compare the immediate response to the water scores (1 period) or the next two responses (2 periods). Results are displayed in Table 1.8. Again, we find no evidence of a categorization effect as the estimated coefficient for the triple difference variable is not significant.

Table 1.8: Time Variant Category

	(1) 1 Period	(2) 2 Periods
$P_t I(T_i = 1)$	0.0350	0.144***
	(0.0252)	(0.0412)
$T_{it} P_{it}$	0.0201	0.0147
	(0.0214)	(0.0351)
$T_{it} P_{it} I(T_i = 1)$	-0.0244	-0.0163
	(0.0318)	(0.0520)
Household FE	Y	Y
Month by year FE	Y	Y
Observations	4,383	2,188
R-squared	0.729	0.765

Notes: The dependent variable in each specification is log(gpd). The unit of observation is a household-month. Standard errors in parentheses, clustered by household.

Star values: * 0.10 ** 0.05 *** 0.010.

Good to Take Action

In the previous subsection, we looked for the presence of a "categorization effect" between the "Great" and "Good" categories. We did not find statistical evidence of such an effect. In the following section, we look for the presence of a "categorization effect" between the "Good" and "Take Action" categories. We do so using two estimation strategies. We begin with a "fuzzy" regression discontinuity approach, and then turn our attention to a triple differences approach.

Fuzzy RD: The consumption cutoff defining houses in the "Good" and "Take Action" categories is the 52nd percentile of consumption for each cohort, which is based on the per capita consumption of each household in the cohort. Unfortunately, we observe only the median (50th percentile). This limitation imposes challenges in a regression discontinuity setting as, from our perspective, the water score is not deterministically defined based on the distance to the median. Thus, we must assume it is probabilistically defined. In other words, being on one side or the other of the median changes the probability of being treated (getting a water score of 3). This setting lends itself to a Fuzzy Regression Discontinuity (FRD) model to determine if there is a statistically significant categorization effect.

The primary difference between a sharp RD approach (Equation 1.3) and a fuzzy regression discontinuity approach (FRD) is that we use use a 2-stage least squares model, where we instrument the treatment dummy variable $I(WS_{i-t} = 3)$ with the dummy variable indicating if the consumption is larger than the median $I(D50_{i-t} > 0)$ [Angrist and Pischke, 2008]. The model is then defined in equations 1.19 and 1.20, where both ϵ_{it} and ξ_{it} are normally distributed error terms:

Structural equation:

$$log(gpd)_{it} = \beta_0 + \rho I(WS_{i-t} = 3) + \beta_1 D50_{i-t} + \mu_{my} + \upsilon_i + \epsilon_{it} \tag{1.19}$$

$$\forall : T_i = 1, P_{it} = 1, |D50_{i-t}| < h, |D50_{i-t}| \neq 0$$

First stage:

$$I(WS_{i-t} = 3) = \theta_0 + \theta_1 D50_{i-t} + \theta_2 I(D50_{i-t} > 0) + \xi_{it} \tag{1.20}$$

Furthermore, as in Equation 1.3, we define a bandwidth h such that we can assume that households with consumption within the specified bandwidth, $|D50_{i-t}| < h$, are statistically identical. In practice, we weight the observations using $|D50_{i-t}|$ to give more weight to those observations closer to the fuzzy cut off. The weighting function we use is as follows:

$$w_{it} = \frac{max(0, h - |D50_{i-t}|)}{h} \tag{1.21}$$

The results of estimating Equation 1.19 are presented in Table 1.9. The coefficient of interest is the coefficient on $I(WS_{i-t} = 3)$. The coefficients are all negative but not statistically significant. This provides suggestive evidence that households that received the "Take Action" message reduced their water consumption more than households that received the "Good" message. However, it is likely that the effect is too small, relative to the overall treatment effect, to be identified with our data. To better understand the limitations of our RCT in identifying this effect, we estimate the Power of the Test in Section 1.4.

Table 1.9: Fuzzy Regression Discontinuity

	(1)	(2)	(3)	(4)	(5)
	h=15	h=20	h=25	h=30	h=35
$I(WS_{i-t} = 3)$	-1.007	-0.352	-0.187	-0.140	-0.145
	(1.950)	(0.627)	(0.301)	(0.185)	(0.127)
$D50_{i-t}$	0.00253*	0.00227**	0.00239***	0.00259***	0.00282***
	(0.00140)	(0.000983)	(0.000685)	(0.000521)	(0.000412)
Household FE	Y	Y	Y	Y	Y
Month by year FE	Y	Y	Y	Y	Y
Observations	9,124	12,713	16,265	19,499	22,785
R-squared	-0.030	-0.008	-0.002	0.001	0.003

Notes: The dependent variable in each specification is log(gpd). The unit of observation is a household-month. Standard errors in parentheses, clustered by household.
Star values: * 0.10 ** 0.05 *** 0.010.

Triple Differences: Finally, we estimate a triple differences specification to identify a categorization effect between the "Good" and "Take Action" categories. The three differences are pre- versus post-treatment, treated versus control, and water score equal to 2 ("Good") versus 3 ("Take Action"). This strategy can potentially improve upon the FRD strategy by incorporating data from the control group. The triple difference specification is the same as Equation 1.9, except $I(WS_{i-t} = 2)$ is replaced with $I(WS_{i-t} = 3)$.

The results of the triple differences estimation are presented in Table 1.10 for 5 different bandwidths, h. Again, the parameter of interest, β_7, is not statistically different from zero for any h considered. Therefore, these results do not support the existence of a categorization effect.

Table 1.10: Difference in Difference in Difference

	(1) h=15	(2) h=20	(3) h=25	(4) h=30	(5) h=35
$I(WS_{i-t} = 3)$	0.0773	-0.198**	-0.00178	-0.0439	0.0103
	(0.0703)	(0.0901)	(0.0667)	(0.0478)	(0.0366)
T_iP_{it}	-0.0926**	-0.0634*	-0.0708**	-0.0429	-0.0219
	(0.0426)	(0.0351)	(0.0303)	(0.0264)	(0.0243)
$T_iI(WS_{i-t} = 3)$	-0.246	0.00616	-0.0990	0.0200	0.0238
	(0.226)	(0.146)	(0.0897)	(0.0611)	(0.0489)
$P_tI(WS_{i-t} = 3)$	-0.0178	0.270***	0.0827	0.131***	0.0959**
	(0.0720)	(0.0906)	(0.0675)	(0.0484)	(0.0375)
$T_iP_{it}I(WS_{i-t} = 3)$	0.245	-0.0240	0.0874	-0.0291	-0.0367
	(0.227)	(0.146)	(0.0907)	(0.0622)	(0.0500)
Household FE	Y	Y	Y	Y	Y
Month by year FE	Y	Y	Y	Y	Y
Observations	12,579	19,143	26,041	32,819	39,648
R-squared	0.765	0.741	0.723	0.713	0.705

Notes: The dependent variable in each specification is log(gpd). The unit of observation is a household-month. Standard errors in parentheses, clustered by household.

Star values: * 0.10 ** 0.05 *** 0.010.

1.4 Ex Post Power of the Test Analysis

In this section, we use an ex post power test to determine if the non-statistically significant categorization effects can be interpreted as true null effects [Brown et al., 2018].

Deriving a statistical estimate requires understanding the probability of a false negative, defined as a Type II error. To do so, we need to simulate the distribution of an alternative hypothesis for a given true treatment effect. If the test had a high probability of detecting an increase in the treatment effect due to the categorization effect, a null finding could be interpreted as "no effect", but if the probability of detecting the categorization effect is low for relatively large effects, our experiment is not powerful enough to identify a reasonable categorization effect.

In general, posterior power is estimated using a closed interval test based on the marginal distribution of the regression coefficient β_{TE}, associated with the treatment effect we want to analyze:

$$H_0 : |\beta_{TE}| \in [\beta_{TE(Crit)}, \infty). \tag{1.22}$$

The cutoff, or critical bound of the test statistic, is $\beta_{TE(Crit)} = t_{(1-(\alpha/2),n-m)} \times \sigma_{\beta_{TE}}$, where n is the number of observations, m is the number of independent variables and $\sigma_{\beta_{TE}}$ is the estimated standard error of the coefficient from the linear regression. The probability that the null hypothesis is true is then the proportion of times the simulated β_{TE} estimates are within the interval defined in Equation 1.22.

We use both methods explained in Brown et al. [2018] to do Ex Post Power Simulation: Monte Carlo Resampling (similar to bootstrapped percentile t-test) and a Bayesian approach to characterize the entire posterior distribution of an estimated effect size. We perform the Ex Post Power Simulation Test on the categorization effects between "Great" and "Good", testing the power for both the RD and DDD settings, and on the categorization effect between "Good" and "Take Action" for the DDD setting.

1.4.1 Ex Post Power Simulation with Monte Carlo Resampling

For a sample size n and a given effect size $\bar{\beta}_{TE}$ (the one for which we want to estimate the power of the test), we resample with replacement from the original design matrix (the matrix of observed explanatory variables) and from the OLS residual vector (the estimated errors) $(X_{N(n)}, \hat{\epsilon}_{N(n)})$ to generate a bootstrap data set $(X^*_{N(n)}, \hat{\epsilon}^*_{N(n)})$. To keep the original structure in the variance covariance matrix, where we allowed for within household auto-correlation (clustered standard errors), we sample n households and, for each household we include all observations for a total of $N(n)$, which will be slightly different for each sample as the panel is unbalanced.

The next step is to calculate a new $y^*_{N(n)} = X^*_{N(n)}\bar{\beta}_{TE} + \hat{\epsilon}^*_{N(n)}$, to finally re estimate $\hat{\beta}^*_{TE}$ for each of the M bootstrap data sets. The power of the test given the true effect size is then the number of times that $|\hat{\beta}^*_{TE}| > \beta_{TE(Crit)}$ divided by the number of bootstrap samples M. Following Brown et al. [2018], we use $M = 10,000$.

1.4.2 Ex Post Power Simulation from Posterior Marginal Distributions

This Bayesian approach characterizes the entire posterior distribution of an estimated effect size. For the coefficient associated with the treatment effect, we assume a prior distribution with the mean equal to the effect size we want to test, and we look at the proportion of times the simulated posterior distribution met the interval test defined in Equation 1.22. The distributional assumptions for the Bayesian approach are as follows:

$$\mathbf{y_i} \sim \mathbf{N}(\mathbf{X_i}\beta, \sigma^2) \tag{1.23}$$
$$\pi(\beta_k) \sim \mathbf{N}(0, \sigma_k^2) \tag{1.24}$$
$$\forall k \text{ (k does not include } \beta_{TE})$$
$$\pi(\beta_{TE}) \sim \mathbf{N}(\bar{\beta}_{TE}, \sigma^2_{\beta_{TE}}) \tag{1.25}$$
$$\pi(\sigma^2) \sim \mathbf{N}(0, \sigma_0^2) \tag{1.26}$$

where the π's are prior distributions assigned to the model parameters. We use diffuse priors for the variance of the β_ks and σ^2s, setting σ_k^2 and σ_0^2 to 10,000. Strong priors are used on the treatment effect, anchoring β_{TE} to the posited effect size and using the square of the OLS standard error as prior for the variance.

Following Brown et al. [2018], we generate 5 Markov Chain Monte Carlo (MCMC) chains of 6,000 samples each, with a burn-in period of 5,000 and a thinning interval of 10 to reduce autocorrelation in the chain.

1.4.3 Results of the Power Analysis

Table 1.11 shows the results for the Ex Post Power Simulation with Monte Carlo Resampling on the RD between "Great" and "Good", given different true effects and the number of households resampled. Consistent with the estimated ATE of 2.4%, we tested the probability of identifying the categorization effect given true effect sizes of 1%, 2% or 3% on a sample of size n, where n is the original sample size.

Table 1.11: Ex Post Power Simulation with Monte Carlo Resampling for RD between "Great" and "Good"

Effect Size	Power
1%	17%
2%	51%
3%	83%

Notes: This table reports the probability of identifying a given true categorization effect of size 1%, 2%, and 3%.

Table 1.12 shows the results for the Ex Post Power Simulation from Posterior Marginal Distributions. We tested the probability of identifying the categorization effect given true effect sizes of 1%, 2% or 3%. In the RD setting to identify the categorization effect between "Great" and "Good", a true effect of 1% would mean that treated household's that got a "Good" reacted by consuming 1% less than those households that got a "Great". In other words, the estimated parameter ρ would be added to the Treatment Effect due to the categorization effect. Our results indicate that even with a categorization effect of 2%, almost as large as the ATE, it would have been a coin toss to detect it with the data.

The results for the DDD between "Great" and "Good" and between "Good" and "Take Action" suggest that we had a 0% probability of identifying a categorization effect given a true effect of 3%. The DDD parameter that we estimate between two categories represents the additional treatment effect from receiving one or the other message.

To understand why our estimation strategies perform so poorly in identifying the categorization effect, we start by defining the power of the test as a function of the effect size, the variance of the population and the sample size. First, as we "borrow" data from an RCT designed to identify the

Table 1.12: Ex Post Power Simulation from Posterior Marginal Distributions

	(1)	(2)	(3)
Effect Size	1%	2%	3%
RD "Great" to "Good"	6.1%	21%	51.6%
DDD "Great" to "Good"	0%	0%	0%
DDD "Good" to "Take Action"	0%	0%	0%

Notes: For a given true effect size (columns), we report the probability of identifying it with our sample size using different models (rows).

ATE of receiving HWRs, we do not have the exogenous random variation in categorization that would allow for a powerful analysis. Instead, as we rely on the bandwidth h to assume randomized categorization, we are left with considerably fewer observations than the original design. In other words, as we rely on the bandwidth h for randomization both for the DDD and the RD designs, we are using experimental data (with a sample size designed to answer a specific question) as the source of quasi-experimental data to answer a different question. As a result, if we assume that an estimated coefficient β_{TE} is significant at an $1-\alpha$ level if $|\beta_{TE}| > Cutoff$, a large $\sigma_{\beta_{TE}}$ pushes the cutoff for significance further from 0. Table 1.13 shows the size of these standard errors for each model and the corresponding cutoff for a 95% confidence level.

Table 1.13: Determinants of Power

	(1)	(2)
Model	Standard Error	Cutoff
RD "Great" to "Good"	0.0104	-0.020
DDD "Great" to "Good"	0.039	-0.077
DDD "Good" to "Take Action"	0.050	-0.098

Notes: We present the standard error calculated for the categorization effect parameter and the corresponding cut-off associated with 95% confidence level, for different identification models.

As we have such a large cutoff compared to the expected categorization effects, we can conclude that the power is low based on the estimated standard errors.

1.5 Conclusions and Discussions

On average, HWRs work as a means to incentivize conservation, as shown by the ATE estimate here, and most of the reduction happens during the warmer months of the year when households irrigate. This finding implies that households are conserving water by reducing outdoor water use during the growing season. Despite ample opportunity for reductions in indoor water use, such as shorter or fewer showers or loads of laundry, it appears that households only modestly reduce indoor water use in response to social comparison. Hence, households may need more information on how to conserve indoor water use.

Although our results demonstrate a U-shaped association between the point estimates of the treatment effect and the quantiles of pre-treatment consumption, our results could not prove that the pre-treatment quantile estimates are different from one another. We also found no evidence of a "boomerang effect", consistent with Allcott [2011]'s findings on HERs.

Finally, we found no evidence of the existence of a "categorization effect" between categories "Great" and "Good" or between "Good" and "Take Action". This result holds with RD and DDD estimation strategies.[8] We use a power test to show that identifying the categorization effect was improbable even for large effect sizes because the RCT we used was designed to identify different treatment effects. To account for possible confounders, randomization was used to make sure both treatment and control groups were random samples from the same population. To study the effect of a normative category on the household's response to social comparison, our strategies needed to address the endogeneity that came from the lack of randomization in the allocation of messages. We used, depending on the model, two different strategies that imposed strong limitations on the data availability. First, for the RD and for the DDD models, we relied on a bandwidth h to assume that, if a household was very close to the cutoff between categories, the message a household received was random. Second, for the TVC model we relied on households that received a message that was the result of exogenous changes in their cohort's consumption. This finding highlights

[8] We also tried these estimation strategies on the summer months.

the challenge of using experimental data (as the source of quasi-experimental data) to answer a different question than the experimental data were originally designed for.

Identifying a categorization effect remains an important unanswered question. If it were identified, it would imply that extending the range of the “Good” or "Take action” categories would result in a greater treatment effects. On the other hand, if the categorization effect is irrelevant, the mechanism driving the effect of HWRs might not be social comparison, but a response to information provision or simply to the knowledge that consumption is being monitored closely. Our work validates previous findings [Allcott, 2011,0,0, Allcott and Kessler, 2019, Brent et al., 2015, Ferraro and Price, 2013, Henry et al., 2019, Jessoe et al., 2017, Torres and Carlsson, 2018], with a similar methodology to Allcott [2011] but on a different dataset, and on HWRs instead of HERs. Finally our analysis includes a methodological extension based on the use of the triple difference estimation.

With a deeper understanding behind the mechanisms through which HWRs incentivize conservation, we expect to inform policy makers and utilities on how to better craft non-pecuniary interventions. Future work should use a RCT designed exclusively to detect the categorization effect, with randomized cutoffs, and with a RCT designed to test for the existence of alternative mechanisms that could explain the observed reduction in consumption that occurs as a result of information provision. Finally, even when we failed to identify a categorization effect, it is important to mention that we did not correct for multiple hypotheses testing. Future work should correct for the increased probability of incorrectly identifying a significant effect.

Chapter 2

Learning Through Experience: The Impact of Technology Reliability on Adoption and Use

2.1 Introduction

Worldwide, 1.1 billion people live without electricity, with nearly 70% in Sub-Saharan Africa [Outlook, 2017]. Only 14% of rural households in Sub-Saharan Africa, excluding South Africa, have access to electricity [UN, 2018]. The UN has made access to affordable, reliable, sustainable and modern energy one of its Sustainable Development Goals because energy services are critical ingredients for socioeconomic development [UN, 2018]. Anecdotal evidence from different domains suggests that low quality supply, or supply with uncertain quality, can lead to low technology adoption rates through increased cynicism and skepticism related to new technologies, especially when provided by unknown firms [Eder et al., 2015, Müggenburg et al., 2012, Shyu, 2013]. Despite qualitative descriptions of this phenomenon, there is limited empirical evidence of how experience with technology quality and reliability affects the adoption and continued use of the technology over time. In this study, we identify how observing or experiencing variation in electricity reliability (technology quality) affects technology adoption (electrification) and usage decisions. Furthermore, we test different learning models, including Bayesian learning, to explore how consumers use information for their decision making process.

To answer these empirical questions, we partner with a private-sector, solar based, mini-grid provider in rural Rwanda that sells pre-paid electricity days of direct current (DC) service. We observe electricity adoption, purchase, and use behavior at a high temporal resolution, particularly for a rural region of a developing country. Our rich data-set, observing about 5 years of individual daily usage, from May 2014 to February 2019, from more than 2,400 households from 75 rural villages, allows us to empirically examine the impact of learning about technology quality on initial adoption decisions and continued use over time.

We first model consumption decisions on the extensive margin (technology adoption), conceptualizing adoption and disadoption as an optimal stopping problem, approximating the underlying dynamic decision process with a reduced-form logit regression model as in [Sampson and Perry, 2018]. We test how experienced reliability affects these decisions using different models of learning, including testable hypotheses derived from Bayesian learning models as in [Lin et al., 2019] and [Maniloff, 2019]. We also model time to re-connect and the number of electricity days purchased after running out of pre-paid days, as a function of observed and experienced electricity reliability. We model electricity reliability as the blackout rate for a particular village.

We find evidence that the observed blackout rate (through neighbors' experiences) affects the timing of the decision to adopt, but households tend to place more weight on their most recent experiences with blackouts relative to blackouts from several periods prior.[9] We also find evidence that higher blackout rates increase the length of time households wait to re-purchase electricity after running out of pre-paid electricity (topping-up). These findings support the hypothesis that electricity reliability plays an important role in electrification decisions. However, we do not find evidence that consumers, once accustomed to electrification, choose to permanently disconnect as a consequence of poor reliability. Finally, we found no evidence of Bayesian learning in the sense that household beliefs do not respond to new information in the way Bayesian learning would predict.

This work contributes to several literatures. Most directly, it adds new evidence to the importance of reliability in electricity provision and demand in the developing world, where development initiatives have not always successfully improved societal welfare [Grimm et al., 2019, Lee et al., 2016]. The magnitude of the rural electrification problem (often referred to as a 'wicked prob-

[9]Throughout, we define experienced blackout rates as blackouts experienced by the focal household during a period in which they had electricity. In contrast, we define observed blackout rates as blackouts that occurred during a period in which the focal household did not have electricity. In other words, the household observed their peers or neighbors experiencing blackouts.

lem' [Gibson, 2017]) suggests that behavioral insights into the demand for electricity can have far-reaching impacts for rural areas of the developing world. Sedai et al. [2020] and Meles [2020] successfully showed that households in India and in Ethiopia, respectively, suffered welfare losses for each day of electricity lost due to blackouts, either through the loss of income or through expenditure on alternative energy sources. Our contribution goes beyond these findings, as we show that blackouts also affect the decision to adopt and use electricity technology.

Furthermore, this work contributes to a broader literature on technology adoption. It has been shown that household adoption of new technologies not only contributes to economic growth [Hall and Khan, 2003] and productivity improvements [Jorgenson, 2011], but it has the potential to create meaningful improvements in quality of life, for example through improvements in indoor air quality [Mobarak et al., 2012] or by facilitating labor market participation, particularly for women [Dinkelman, 2011]. Despite numerous benefits, the literature has revealed multiple mechanisms that deter technology adoption and its continued use [Carter et al., 2014, Hanna et al., 2016, Jack et al., 2015], and we add to this literature by showing that reliability is a key determinant of adoption.

Additionally, our work contributes to the vast literature on learning and peer effects. In recent work, Streletskaya et al. [2020] provides an overview of the linkages and complementary topics between technology adoption and behavioral economics, highlighting learning and social preferences as an area for fruitful future research. Rosenberg and Nathan [1982] detail three general types of learning in the context of agricultural technology adoption: (1) learning by doing, when technology suppliers learn how to improve their technology over time based on feedback and experiences of their customers, (2) learning by using, when farmers improve their mastery of the adopted technology over time through personal use experience, and (3) traditional learning based on obtaining information about a particular technology from various sources, including other farmers who have personal experience with the technology. In some cases, learning from peers can provide potential adopters with useful information [Oster and Thornton, 2012, Sampson and Perry, 2018], though heterogeneity in returns can limit peer learning [Munshi, 2004]. Further, peer experiences may not always lead to increased adoption. For example, Adhvaryu [2014] finds that misdiagnosis of malaria slows the adoption of malarial treatment by peer households in Tanzania. In our work, we show that not only do a household's own experiences with blackouts affect consumption decisions, but they are also affected by the reliability they observe through peers' use of the technology.

In the learning literature, actors are often assumed to be Bayesian updaters, particularly in research focusing on learning by using and on traditional learning [BenYishay and Mobarak, 2018, Fudenberg and Tirole, 1991, Lin et al., 2019, Maniloff, 2019]. Behavioral economics research provides evidence that people do not apply Bayesian updating rules when faced with new information in an experimental settings [Charness and Levin, 2005, Zizzo et al., 2000], while Lindner and Gibbs [1990] show that farmer behavior may not be explained through a Bayesian learning process. Similarly, Fischer et al. [1996] find that observed farmer learning is slower than predicted by traditional Bayesian approaches. We tested for Bayesian learning and we also could not find conclusive evidence of its presence.

Finally, Hall and Khan [2003] describes how the idea of adopting a new technology is similar to any other kind of investment under uncertainty, characterized by 1) uncertainty over future profit streams, 2) irreversibility that creates at least some sunk costs, and 3) the opportunity to delay. As such, inefficiencies associated with decisions about technology adoption can be studied in the real options framework suggested by Dixit et al. [1994]. For instance, Ansar and Sparks [2009] used this approach to explain the inclination of households and firms to require very high internal rates of return in order to make energy-saving investments. The existence of a peer effect through which potential customers can learn about technology reliability from their connected peers, potentially decreasing the uncertainty of future utility streams, created a further incentive to wait to adopt.

2.2 Empirical Context and Data

2.2.1 Context

We explore the impact of electricity reliability on adoption and use in the context of rural Rwanda. Rwanda's national electrification rate is estimated at 30% (12% in rural areas, 72% in

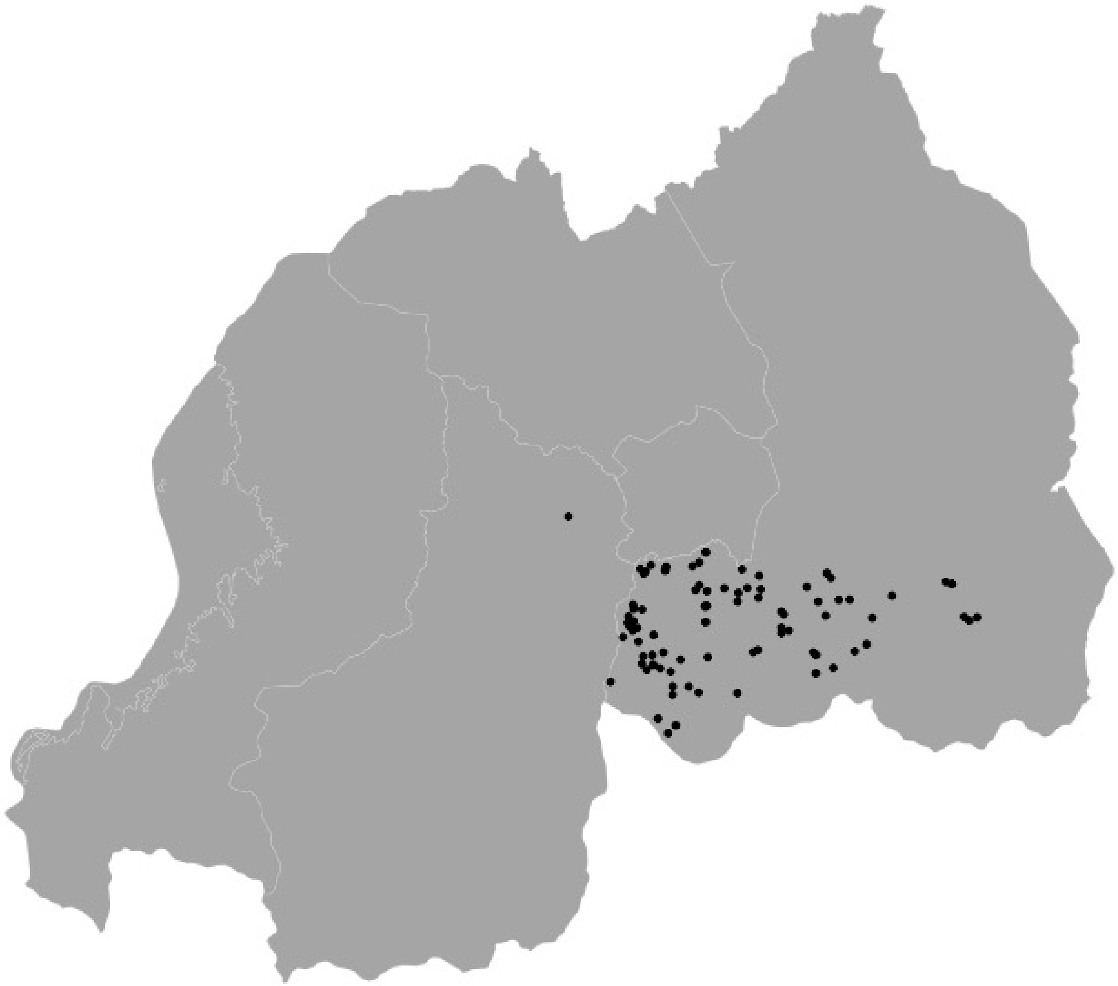

Figure 2.1: MeshPower Sites

urban areas) and, according to US-AID [2019], the government is targeting 100% electricity access by 2024. US-AID [2019] identified the biggest issues and bottlenecks the Rwandan government needs to address to achieve their goal: misalignment of power supply and demand, limited financing for off-grid companies and limited affordability of electricity solutions for rural households and businesses.

In this context, MeshPower provides electricity to 75 communities that are not covered by the country's national grid, mainly rural villages, as shown in Figure 2.1.

Figure 2.2: MeshPower Solar Panels in Base Station. Credit: meshpower.co.rw

MeshPower offers direct current (DC) electricity from their “base station”, typically in the center of a village. At the base station, MeshPower sets up solar panels linked to a secure battery storage unit, as shown in Figure 2.2.

DC electricity travels from the base station to each business or home via aerial cabling up to a distance of 200 metres, as shown in Figure 2.3. These cables do not only distribute electricity from the base stations to the customers, but are also used by MeshPower to gather the individual consumption information traveling from households to the base station.

MeshPower sells energy only, rather than hardware and asset contracts, providing a range of tariffs that cover both business and home use, plus different levels of intensity. For each household or business, MeshPower installs a station (as in Figure 2.4) from which the customer can pull electricity by plugging any device in the USB ports, or by using the LED lighting powered by the station.

Customers buy pre-paid "electricity days" as they need it and as they can afford it. Once their balance reaches zero, households can purchase additional electricity days (a "top-up"), or they can decide to wait and purchase electricity days later. MeshPower provides a range of services

Figure 2.3: MeshPower Cables Distributing DC Electricity. Credit: meshpower.co.rw

Figure 2.4: MeshPower DC Home Station. Credit: meshpower.co.rw

depending on the number of LED lights and USB outlets, with tariffs that range between 50 to 140 Rawandan Francs (RWFs) per day (between 0.1 and 0.2 USD per day, approximately).

Data

We observe approximately 5 years of individual daily usage, from May 2014 to February 2019, from more than 2,400 households from 75 rural villages. Specifically, for every connected household we observe the daily tariff, daily cash purchase of electricity and each day's final balance.

We observe around 775,000 electricity days consumed, with tariffs distributed as follows: 30% were for 100 RWF/day, 43% were for 50 RWF/day, 12% were for 80 RWF/day and 11% were for 140 RWF/day. Our data set includes more than 81,000 electricity purchases. The mean purchase was of around 700 RWF (less than $1), the median was 500 RWF, while the maximum purchase was 9,500 RWF. Most of the purchases were either 500 RWF (75%), 1,000 RWF (15%) or 1,500 RWF (1.7%). Most of the purchases (62%) occurred before the household had a zero balance (they did not let their electricity service run out) and the vast majority (88%) stayed without electricity for 5 days or less. If service was allowed to lapse, the distribution for the number of days it took to reconnect (top-up), including only those observations in which a household stayed at least one day without electricity before purchasing more electricity days (limited to 20 days or less), is displayed in Figure 2.5.

We also know when MeshPower began providing service to each village and from this, we can infer how long each household waited to adopt the electricity technology. The average number of days it took customers to connect to the MeshPower service is 281 days and the median is 161 days.

Even though we do not directly observe disconnections, we assume that a household disconnected from a site in month t if between t and $t+6$ months, we do not observe any electricity usage or purchase. Under this assumption, we observe 961 customers who later chose to disconnect and did so after an average of 409 days, with a minimum of 89 days and a maximum of 1,369 days.

We do not directly observe blackouts, our measure of reliability. Instead, we observe for each site and for each day, the total number of active customers and the number of customers that were

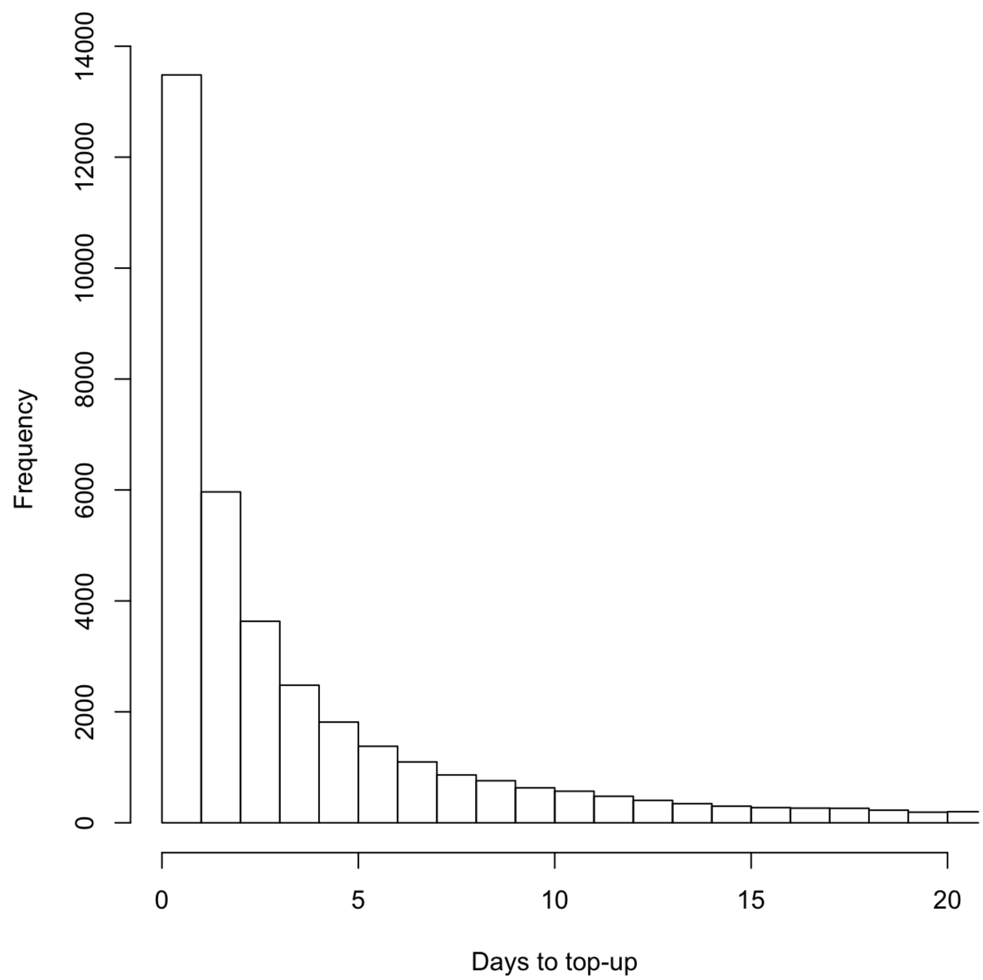

Figure 2.5: Distribution of the Number of Days to Repurchase Electricity Days

not charged for their electricity. MeshPower's policy is to not charge customers for electricity on a given day if during the day there is a miscommunication between the equipment installed in the household and the base station. It is worth noting that households do not bear the direct financial risk of an outage. Instead, dis-utility from outages is from not being able to reliably use electricity devices. As we do not observe the nature of the problem a household is facing on a given day and for which MeshPower decided not to charge the electricity day, we opted for a binary approach to define site-wise blackouts: if on any given day at a particular site, the fraction of households that were not charged for electricity was larger than 80%, it is safe to assume there was a system-wide blackout that day.[10] If we instead considered a continuous variable representing the percentage of households that were not charged on a given day, lower values could be indicative of the presence of endogenous, individually experienced blackouts, often the result of problems with the wires or with the customer's station. Under this set of assumptions, in the almost 5 years of daily data, we observe, per site, an average of 33 blackout days. The average percentage of days with blackouts across sites is 3.73%, The site that had the largest number of blackouts was Kavure, with 152 days with blackouts (45.5% of the observed days), and 3 other sites had more than 100 days with blackouts (10%, 11% and 13% of the days had a blackout) over the sample period.

Finally, as both electricity demand and solar electricity generation are affected by weather outcomes, we use modelled data from from the Climatic Research Unit (CRU) of the University of East Anglia to control for its effect. Specifically, we use mean temperature and accumulated rain at the month level.

Summary statistics for the explanatory variables and for response variables are as shown in Table 2.1.

2.3 Model of Experience, Learning, and Technology Adoption

In this section, we develop a simple utility based model to explain a household's decision to purchase an electricity day as a function of the household's beliefs about the system's *unreliability*,

[10]After discussing with Meshpower, we determined that 80% was a reasonable cutoff to assume a blackout occurred.

Table 2.1: Summary Statistics

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Blackout rate (% of days in a month with blackouts)	23,196	0.033	0.101	0	0	0.03	1
Variance of the blackout rate per site	23,196	0.011	0.026	0.0001	0.001	0.009	0.198
Mean temperature (C) per month	23,196	21.522	0.775	18.000	21.000	22.000	23.000
Mean rain (mm) per month	23,196	80.304	46.157	3.300	40.600	120.200	192.600
Connection (binary)	23,196	0.114	0.318	0	0	0	1
Dis-connection (binary)	15,492	0.068	0.252	0	0	0	1
Time to re-connect	39,763	10.214	29.485	1	1	8	180
Amount spent after having zero balance (Francs)	39,763	614.035	324.443	3	500	500	9,100

Notes: All variables are measured at the household-month level.

measured as the monthly blackout rate. Furthermore, we use the model to explain how the household decision maker can use the unreliability he or she observes or experiences to update those beliefs. In our model, households buy one electricity day at a time.

The decision maker in household h in village c buys an electricity day to use during period t if the expected net benefit from purchasing electricity is positive. In other words, households purchase electricity if the money valued expected utility $E(V)$ is larger than the cost C (Equation 1). In a pre-paid setting as the one we study, this decision is reversible: the decision maker can decide to stop purchasing electricity days and keep the balance at zero for as long as the household deems optimal.

Utility is a function of electricity services, which we model as a function of how unreliable (q_c) the service in village c is, and of the value of electricity γ_h, which is specific to each household h.

$$E[V(q_c, \gamma_h)] > C \tag{2.1}$$

We assume utility decreases with unreliability ($V_h'(q_c) < 0$) and it is convex ($V_h''(q_c) > 0$). Unreliability in village c in period t is then assumed to be random draw from a distribution with mean q_c and variance σ_c^2.

$$q_{c,t} \sim \mathcal{F}(q_c,\, \sigma_c^2) \tag{2.2}$$

As the true expected value q_c is unknown to consumers, households will act according to what they believe it will be for period t. Assuming consumers in village c share the same information, we define the belief about q_c for period t as $m_{c,t}$.

2.3.1 Learning models

We test three different learning models to explain how decision makers update their beliefs about the expected unreliability. First, we test whether households respond more to recent information than information from several periods prior. We call this model short term memory ($m_{c,t} = q_{c,t-1}$), in which a household's beliefs about the expected unreliability are exclusively based on what was observed or experienced during the previous period. Second, we test whether households consider information from the entire sample period. We call this model long term unweighted memory ($m_{c,t} = 1/(t-1)\sum_{t=0}^{t-1} q_{c,t}$), where a household's beliefs about expected unreliability are the simple average of what was observed or experienced since the beginning of the service in the village, $t = 0$. Finally, we test for Bayesian learning using the following two testable hypothesis (see appendix for theoretical foundations of these tests): (1) The larger the variance of the unreliability distribution (σ_c^2), the less the household learns from the observed or experienced realizations, as in Lin et al. [2019], and (2) The longer the site has been active, the less the household learns from the observed or experienced realizations, as in Maniloff [2019]. Finding evidence of these latter hypotheses would suggest that households behave in a way that is consistent with Bayesian learning.

2.4 Econometric Specifications

2.4.1 Extensive Margin: Technology Adoption and Disadoption

We first model the extensive margin of adoption for households that eventually connect, conceptualizing it as an optimal stopping problem in which household h in village c decides whether to connect ($a_{h,t} = 1$), or not ($a_{h,t} = 0$) in every month t, after the service becomes available, based on the household's beliefs about the expected unreliability (expected blackout rate) at the beginning of month t, $m_{c,t}$. As every household h living in village c observes the same blackout rate, we use the subscript c on $m_{c,t}$. We test different learning models, as described in Section 2.3.1, captured by the variable $m_{c,t}$: households could respond to the blackout rate observed in $t-1$, to the blackout rate observed in $t-2$, or to the simple average of the observed blackout rate from $t=0$, the beginning of the service period in a particular village, to $t-1$. We approximate the underlying dynamic decision process with a reduced-form logit model (as in Sampson and Perry [2018]). Thus, we model the odds of adopting as a function of neighboring households' experiences, expressed as the rate of blackouts in the previous time periods, as follows:

$$log(\frac{p(a_{h,t}=1)}{1-p(a_{h,t}=1)}) = \delta\ m_{c,t} + f(T_{c,t})\ + \alpha_h + \beta^T\ X_{c,t} + \eta_t. \tag{2.3}$$

Note that the probability of adoption depends on households not currently having the service. Therefore, $m_{c,t}$ captures what a household believes about the expected blackout rate from observing their connected neighbors experiences with the service in village c. The term $f(T_{c,t})$ is a flexible village specific time control to capture the trend in technology diffusion,[11] $X_{c,t}$ is a vector of weather variables and, to control for unobservables, we include household (α_h) and time (η_t) fixed effects, clustering standard errors at the village level. The time fixed effect is a month by year dummy, which captures time-varying unobservables that are constant across sites. On the other hand, the time trend is specific to each village to capture the time since electrification for each village.[12]

For robustness and for ease of calculating marginal effects, we also estimate the adoption and disadoption relationship using a simple linear model as follows:

$$p(a_{h,t}=1) = \delta\ m_{c,t} + f(T_{c,t})\ + \alpha_h + \beta^T\ X_{c,t} + \eta_t +\ \varepsilon_{h,t}. \tag{2.4}$$

[11] In all specifications, $f(T_{c,t})$ is a polynomial of degree three.

[12] In the Appendix, we also test a specification where instead of using a trend variable we use the number of connected households.

Using the linear model, we also test for evidence of Bayesian learning [13]. To test the first Bayesian hypothesis we study if a potential customer learns less from the unreliability realizations if σ_c^2, the variance of the blackout rate, is larger. In other words, we interact σ_c^2 with the unreliability information $m_{c,t}$. The model is as follows:

$$p(a_{h,t} = 1) = \delta\ m_{c,t} + \beta_1\ \sigma_c^2 + \beta_2\ \sigma_c^2 m_{c,t} + f\,(T_{c,t}) + \beta^T\ X_{c,t} + \alpha_h + \eta_t +\ \varepsilon_{h,t}. \tag{2.5}$$

To test the second Bayesian hypothesis we study if a potential customer learns less from the unreliability realizations if the site has been active for longer. In other words, we test whether the effect of new information is diminished as the household receives more information. The model is as follows:[14]

$$p(a_{h,t} = 1) = \delta\ m_{c,t} + f\,(T_{c,t}) + \beta\ m_{c,t}T_{c,t} + \beta^T\ X_{c,t} + \alpha_h + \eta_t +\ \varepsilon_{h,t} \tag{2.6}$$

Also on the extensive margin, we model the probability of disadoption ($d_{h,t} = 1$) for connected households. For households that dis-adopt, the experience with unreliability is first hand realizations of the blackout rate, as opposed to information from observing peer experiences. The logit and linear disadoption models are as follows:[15]

$$log(\frac{p(d_{h,t} = 1)}{1 - p(d_{h,t} = 1)}) = \delta\ m_{c,t} + f\,(T_{c,t}) + \alpha_h + \beta^T\ X_{c,t} + \eta_t \tag{2.7}$$

and

$$p(d_{h,t} = 1) = \delta\ m_{c,t} + f\,(T_{c,t}) + \alpha_h + \beta^T\ X_{c,t} + \eta_t +\ \varepsilon_{h,t}. \tag{2.8}$$

[13] We do not use the logit specification for the Bayesian models because of convergence problems.

[14] In all specifications, $f(T_{c,t})$ is a polynomial of degree three, while T is simply a linear time trend.

[15] In the Appendix, we also tested a specification where instead of using a trend variable we use the number of connected households. The results are largely unaffected.

To find evidence of Bayesian learning in the decision to dis-adopt, we use the following models:

$$p(d_{h,t} = 1) = \delta\ m_{c,t} + \beta_1\ \sigma_c^2 + \beta_2\ \sigma_c^2 m_{c,t} + f\left(T_{h,t}\right) + \beta^T\ X_{c,t} + \alpha_h + \eta_t +\ \varepsilon_{h,t} \tag{2.9}$$

and

$$p(d_{h,t} = 1) = \delta\ m_{c,t} + f\left(T_{h,t}\right) + \beta\ m_{c,t} T_{h,t} + \beta^T\ X_{c,t} + \alpha_h + \eta_t +\ \varepsilon_{h,t}. \tag{2.10}$$

To be clear, the main conceptual difference between the decisions to adopt and to dis-adopt is that, in Equations 2.7, 2.9 and 2.10, $m_{c,t}$ is not only observed by household h but also personally experienced. The underlying assumptions are then that every customer experiences the same rate of blackouts, and that every future customer observes the same rate of blackouts on their neighbors.

2.4.2 Extensive Margin: Reconnection Decisions

Due to the pre-paid nature of the service provided by the utility, we can study the decision of when to re-purchase electricity days once the balance reaches zero. In the context of pre-paid electricity, allowing the balance to reach zero can be seen as a temporary and reversible disadoption, which we model as a function of the experienced electricity unreliability during the previous consumption period, which we denote $i - 1$, and the observed electricity unreliability during the previous period with zero balance, which we denote i. Furthermore, we can also study if the experienced and observed electricity unreliability affects how many electricity days a household purchases once they decide to reconnect.

For prepaid electricity credit, we model the expected number of days between a zero balance and the subsequent purchase (top-up) for household h that took place in month t, denoted $\lambda_{h,t}$. In other words, $\lambda_{h,t}$ is a count variable modeling the number of days with a zero balance for households that let their service lapse for any amount of time. The time between a zero balance and subsequent purchase depends on the personally experienced blackout rate during the previous positive consumption period, $i-1$, denoted as $m_{c,i-1}$, and the observed blackout rate during the most recent time period with zero consumption, denoted $m_{c,i}$. To be clear, the subscript t denotes months and years of the sample, while the subscript i denotes periods that are relative to each household's consumption pattern. Assuming a Poisson distribution, we model the log of the expected number of days to purchase electricity days, $\lambda_{h,t}$, as:

$$\log(\lambda_{h,t}) = \delta\ m_{c,i-1} + \varphi\ m_{c,i} + \alpha_h + \beta^T\ X_{c,t} + \eta_t \tag{2.11}$$

To control for unobservables, we include household (α_h) and time (η_t) fixed effects. As we do not have access to daily weather realizations, the variables included in $X_{c,t}$ include weather realizations for the month t in which household h purchases electricity days after having zero balance. To allow for over dispersion and to explore for robustness, we test different distributional assumptions, including Gaussian (linear link function) and Negative Binomial (log link function).

We use the same linear specification for a Gaussian model to explain the amount spent on each electricity purchase in time period t for household h in constant local currency ($e_{h,t}$) as:

$$e_{h,t} = \delta\ m_{c,i-1} + \varphi\ m_{c,i} + \alpha_h + \beta^T\ X_{c,t} + \eta_t + \ \varepsilon_{h,i}. \tag{2.12}$$

As we did with adoption and disadoption decisions, we use both models to test for Bayesian learning. To test the first Bayesian hypothesis, we use Equations 2.19 and 2.20 to model how the variance of the blackout rates (σ_c^2) affects learning.

$$\log(\lambda_{h,t}) = \beta_1\ \sigma_c^2 + \beta_2\ m_{c,i-1} + \beta_3\ \sigma_c^2\ m_{c,i-1} + \varphi\ m_{c,i} + \varphi\ \sigma_c^2\ m_{c,i} + \alpha_h + \eta_t \tag{2.13}$$

$$e_{h,t} = \beta_1\ \sigma_c^2 + \beta_2\ m_{c,i-1} + \beta_3\ \sigma_c^2\ m_{c,i-1} + \varphi\ m_{c,i} + \varphi\ \sigma_c^2\ m_{c,i} + \alpha_h + \eta_t + \ \varepsilon_{h,i} \tag{2.14}$$

Finally, to test the second Bayesian hypothesis, we use Equations 2.21 and 2.22 to model the relationship between the time spent as a costumer ($T_{h,t}$) and learning ($\lambda_{h,i}$).

$$\log(\lambda_{h,t}) = \delta\ m_{c,i-1} + f\left(T_{h,t}\right) + \beta_0\ m_{c,i-1} T_{h,t} + \varphi\ m_{c,i} + \beta_1\ m_{c,i} T_{h,t} + \alpha_h + \eta_t \qquad (2.15)$$

$$e_{h,t} = \delta\ m_{c,i-1} + f\left(T_{h,t}\right) + \beta_0\ m_{c,i-1} T_{h,t} + \varphi\ m_{c,i} + \beta_1\ m_{c,i} T_{h,t} + \alpha_h + \eta_t +\ \varepsilon_{h,i} \qquad (2.16)$$

When estimating Equations 2.17, 2.18, 2.19, 2.20, 2.21 and 2.22, we cluster standard errors at the village level. For each model, our goal is to test the significance of the estimators of the coefficients capturing reliability (observed and experienced).

Due to the pre-paid nature of the service provided by the utility, we can study the decision of when to re-purchase electricity days once the balance reaches zero. In the context of pre-paid electricity, allowing the balance to reach zero can be seen as a temporary and reversible disadoption, which we model as a function of the experienced electricity unreliability during the previous consumption period ($i-1$), and observed electricity unreliability during the period in which consumption was zero due to having zero balance while waiting for the i^{th} purchase. Furthermore, we can also study if the experienced and observed electricity unreliability affects how many electricity days a household purchases once they decide to reconnect.

For prepaid electricity credit, we model the expected number of days ($\lambda_{h,i}$) between a zero balance and the subsequent i^{th} purchase for household h that took place in month t. The time between a zero balance and subsequent purchase depends on the personally experienced blackout rate during the previous positive consumption period, $i-1$, denoted as $m_{c,i-1}$, and the observed blackout rate during the most recent time period with zero consumption before the i^{th} purchase, denoted $m_{c,i}$. To be clear, the subscript t denotes months and years of the sample, while the subscript i denotes purchases that are relative to each household's consumption pattern. For instance, the i^{th}

purchase for household h happened in month t, after a time period in which household h had zero consumption (they choose to not purchase more credit after consuming the $(i-1)^{th}$ purchased balance). Assuming a Poisson distribution, we model the log of the expected number of days to purchase electricity days, $\lambda_{h,i}$, as:

$$\log(\lambda_{h,i}) = \delta\ m_{c,i-1} + \varphi\ m_{c,i} + \alpha_h + \beta^T\ X_{c,t} + \eta_t \tag{2.17}$$

To control for unobservables, we include household (α_h) and time (η_t) fixed effects. As we do not have access to daily weather realizations, the variables included in $X_{c,t}$ include weather realizations for the month t in which household h purchase electricity days for the i^{th} time after having zero balance. To allow for over dispersion and to explore for robustness, we test different distributional assumptions, including Gaussian (linear link function) and Negative Binomial (log link function).

We use the same linear specification for a Gaussian model to explain the amount spent on the i^{th} electricity purchase in time period t for household h in constant local currency ($e_{h,i}$) as:

$$e_{h,i} = \delta\ m_{c,i-1} + \varphi\ m_{c,i} + \alpha_h + \beta^T\ X_{c,t} + \eta_t +\ \varepsilon_{h,i}. \tag{2.18}$$

As we did with adoption and disadoption decisions, we use both models to test for Bayesian learning based on the two testable hypotheses we described. Based on the first testable Bayesian hypothesis, we use equations 2.19 and 2.20 to model how the variance of the blackout rates (σ_c^2) affects learning.

$$\log(\lambda_{h,i}) = \beta_1\ \sigma_c^2 + \beta_2\ m_{c,i-1} + \beta_3\ \sigma_c^2\ m_{c,i-1} + \varphi\ m_{c,i} + \varphi\ \sigma_c^2\ m_{c,i} + \alpha_h + \eta_t \tag{2.19}$$

$$e_{h,i} = \beta_1\ \sigma_c^2 + \beta_2\ m_{c,i-1} + \beta_3\ \sigma_c^2\ m_{c,i-1} + \varphi\ m_{c,i} + \varphi\ \sigma_c^2\ m_{c,i} + \alpha_h + \eta_t +\ \varepsilon_{h,i} \tag{2.20}$$

Finally, based on the second testable Bayesian hypothesis, we use equations 2.21 and 2.22 to model the relationship between the time spent as a costumer ($T_{h,t}$) and learning ($\lambda_{h,i}$).

$$\log(\lambda_{h,i}) = \delta\ m_{c,i-1} + f\left(T_{h,t}\right) + \beta_0\ m_{c,i-1}T_{h,t} + \varphi\ m_{c,i} + \beta_1\ m_{c,i}T_{h,t} + \alpha_h + \eta_t \tag{2.21}$$

$$e_{h,i} = \delta\ m_{c,i-1} + f\left(T_{h,t}\right) + \beta_0\ m_{c,i-1}T_{h,t} + \varphi\ m_{c,i} + \beta_1\ m_{c,i}T_{h,t} + \alpha_h + \eta_t + \ \varepsilon_{h,i} \tag{2.22}$$

When estimating equations 2.17, 2.18, 2.19, 2.20, 2.21 and 2.22, we cluster standard errors at the village level. For each model, our goal is to test the significance of the estimators of the coefficients capturing reliability (observed and experienced).

2.4.3 Identification Assumptions

An implicit assumption in all our models is that blackouts are exogenous to electricity purchase decisions, conditional on our control variables and fixed effects. If usage decisions were endogenous to blackouts, the estimated parameters would be biased.

One possible threat to identification is that the blackout rate might be a function of the number of households connected to the local grid at any given moment.[16] Our first approach is to look at the consumption path per site, counting the average daily number of households consuming electricity days per month, and at the rate of blackouts per month to visually inspect their relationship. In Figure 2.6 we show both curves for the 10 biggest sites in our sample, where it is visually clear that larger numbers of households using electricity are not temporally aligned with larger blackout rates.

[16]As a reminder, we do not directly observe blackouts. Instead, we observe, per day and site, the total number of active users and the number of users who were charged an electricity day based on the communication between the equipment in the household and the base station.

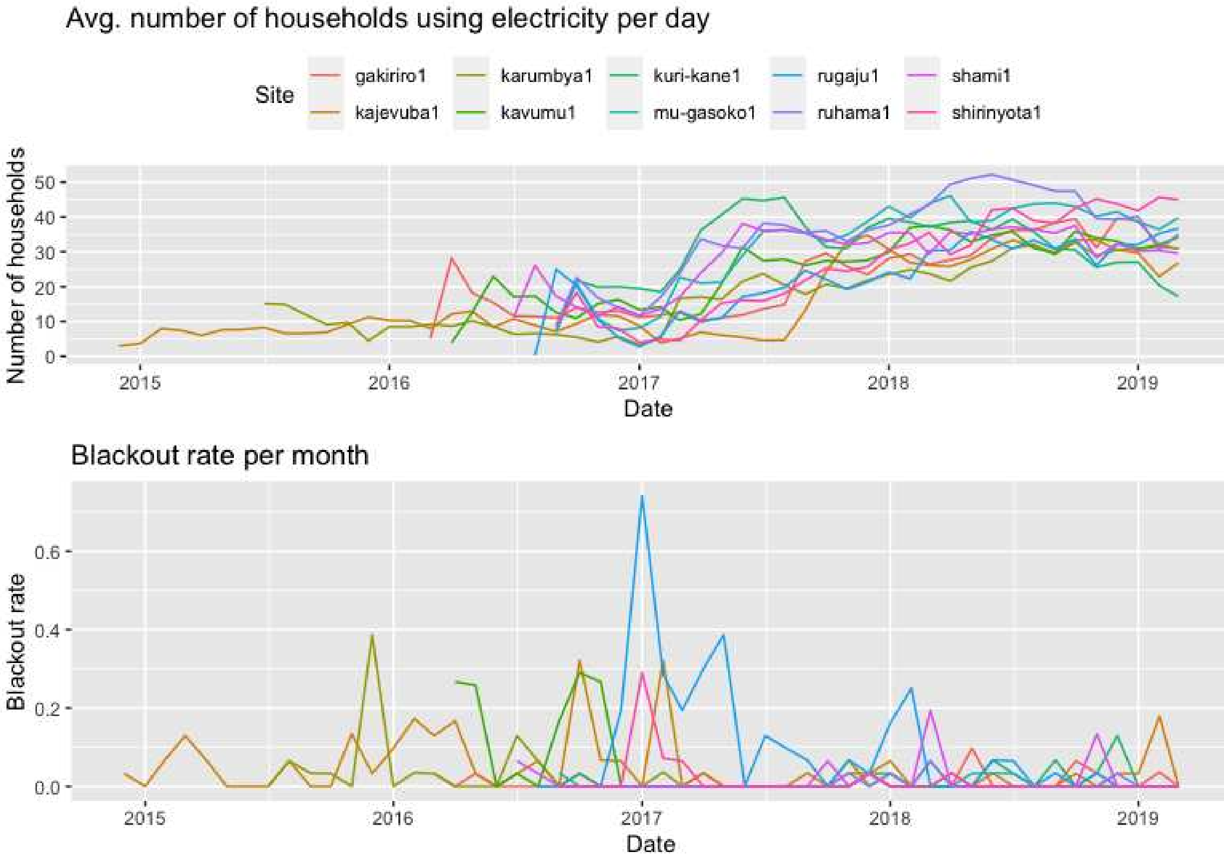

Figure 2.6: Number of connected households and blackout rates for the five biggest sites across time

Second, to further understand why blackouts happen and their association with the usage decisions, we use a linear model to explain the rate of blackouts (r) per month t in village c ($r_{c,t}$)) (number of days with blackouts divided by the number of days in the month) as a function of a polynomial of the number of connected households $f(H_{c,t})$, a trend variable $T_{c,t}$ to capture learning over time, and a vector of weather realizations $X_{c,t}$. We cluster standard errors by village.

$$r_{c,t} = \beta_0 + \ f(H_{c,t}) \ + T_{c,t} + \beta^T \ X_{c,t} + \varepsilon_{c,t} \tag{2.23}$$

The results in Table 2.2 show that blackouts are explained mostly by the length of the service (trend) and precipitation, but there is no clear evidence that the blackout rate is explained by the number of connected households. In other words, it seems that sites are becoming more reliable over time and that rain, likely correlated with storm activity, is associated with a higher blackout rate. There is also no evidence of reverse causality from adoption to the blackout rate. These results are important as they show that the dependent and independent variables in our main models are not simultaneously defined. If this were not the case, the estimated coefficients would be biased if they were not lagged.

Despite the previous test, one might still be concerned that contemporaneous congestion causes blackouts. To further test for this issue, we model whether the *daily* count of the number of houses that were not charged ($p_{c,t}$) in period t in village c is a function of the number of active households in the village on that day ($u_{h,t}$). If this were true, then the decision to adopt would be simultaneously determined with the experienced blackout rate and would lead to biased estimates. We should note, however, that since all of our models use the lagged blackout rate, much of the simultaneity concern is alleviated simply through the exploitation of different time periods. Nonetheless, we use a Poisson regression to estimate the association between the count of households that were not charged ($p_{c,t}$) in period t in village c and the number of active users ($u_{h,t}$) in period t in village c, as follows:[17]

[17]To properly identify the effect of congestion, we include time fixed effects (η_t), village fixed effects (α_h) and a flexible trend function $f\ (T_{c,t})$ to capture learning over time. As we do not have access to daily weather realization,

Table 2.2: Modelling Blackout Rate

	(1) OLS
Connected households (t)	-0.143*
	(0.085)
Connected households^2 (t)	-0.076
	(0.081)
Mean temperature (t) (C)	-0.002
	(0.002)
Mean rain (t) (mm)	8.126e(-05)**
	(3.718e(-05))
Trend	-7.792e(-04)***
	(1.808e(-04))
Clustered SE	Y
Observations	2212

Notes: Standard errors in parentheses. The number of observations in this table is lower than the primary regressions because this analysis is performed at the village level, not the household level.
Star values: * 0.10 ** 0.05 *** 0.010.

$$log(p_{c,t}) = \delta\ u_{c,t} + f\,(T_{c,t}) + \gamma_c + \eta_t + \ \varepsilon_{c,t} \tag{2.24}$$

With this model, our goal is to test the null hypothesis that the marginal effect of $u_{h,t}$ on the expected count of uncharged households, $p_{c,t}$, is less than one. If the marginal effect is less than one, then an increase in $u_{h,t}$ will not result in an increase in the overall blackout rate $p_{c,t}/u_{h,t}$, which is what we use to define our dependent variable in our main regressions. We show the results in Table 2.3.

One more household using electricity services is associated with an increase in 1.7% points in the number of connected households that were not charged (exp(0.01769906)=1.017857). The mean number of non charged households is 8.93 per month, therefore an increase of 1.7% would mean an average increase of 0.16 uncharged households.

we rely on the fixed effects under the assumption that villages are so close by that they are subject to the same weather. To account for correlated error terms, we use a sandwich estimator to calculate standard errors clustering by village.

Table 2.3: Modelling Uncharged households

	(1) Poisson
Number of active users	0.0177***
	(0.0001)
Household FE	Y
Month by year FE	Y
Clustered SE	Y
Observations	75,756

Notes: Standard errors in parentheses, clustered by Village. The number of observations in this analysis is larger than the primary regressions because this analysis is performed at the daily level rather than the monthly level.
Star values: * 0.10 ** 0.05 *** 0.010.

The 95% CI for the coefficient associated with total number of households is [0.0174, 0.0180], meaning one more connected household is associated with an increase in a range between [1.76%, 1.88%] (95% CI), or of between 0.157 and 0.162 (95% CI) non charged households. Therefore, an increase in $u_{h,t}$ will not result in an increase in the rate $p_{c,t}/u_{h,t}$ and we can conclude that daily or monthly congestion is not a significant factor in the determining the blackout rate.

2.5 Results

2.5.1 Technology Adoption

We estimate Equations 2.3 and 2.4 using a logit and a linear model respectively, testing the different learning models discussed in Sections 2.3.1 and 2.4.1. Table 2.4 displays the results of estimating a logit model while Table 2.5 displays the results of estimating a linear model. Columns 1 and 2 of each table include only the blackout rates in the previous 1 to 2 periods. Column 3, however tests whether longer term information is important by including the average blackout rate since servicing began in the village and $t-1$. Columns 4 and 5 include both short term and long term information, blackout rates in the previous two periods and average blackout rates since the beginning of the service provision, to test for the robustness of short versus longer sighted learning models.

Table 2.4: Probability of Adoption

	(1) Logit	(2) Logit	(3) Logit	(4) Logit	(5) Logit
Blackout rate $(t-1)$	-1.4118	-1.643*	-	-1.323*	-1.373*
	(1.2639)	(0.851)	-	(0.789)	(0.800)
Blackout rate $(t-2)$	-	1.293	-	-	0.140
	-	(0.692)	-	-	(0.821)
Avg. Blackout rate $(t-1)$	-	-	1.308*	-	-
	-	-	(0.781)	-	-
Avg. Blackout rate $(t-2)$	-	-	-	1.030	-
	-	-	-	(0.630)	-
Avg. Blackout rate $(t-3)$	-	-	-	-	0.995
	-	-	-	-	(0.681)
Household FE	Y	Y	Y	Y	Y
Month by year FE	Y	Y	Y	Y	Y
Village Time Trend	Y	Y	Y	Y	Y
Clustered SE	Y	Y	Y	Y	Y
Observations	23,196	22,113	23,196	21,030	21,030

Notes: Standard errors in parentheses, clustered by Village. The unit of analysis in these regressions is the household-month.

Star values: * 0.10 ** 0.05 *** 0.010.

Table 2.5: Probability of Adoption - Linear model

	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) OLS
Blackout rate $(t-1)$	-0.12599**	-0.094385**	-	-0.09605**	-0.09885**
	(0.04816)	(0.042962)	-	(0.04493)	(0.04504)
Blackout rate $(t-2)$	-	-0.008288	-	-	-0.02268
	-	(0.053984)	-	-	(0.05415)
Avg. Blackout rate $(t-1)$	-	-	-0.1355	-	-
	-	-	(0.1766)	-	-
Avg. Blackout rate $(t-2)$	-	-	-	0.11172	-
	-	-	-	(0.14122)	-
Avg. Blackout rate $(t-3)$	-	-	-	-	0.05280
	-	-	-	-	(0.14598)
Household FE	Y	Y	Y	Y	Y
Month by year FE	Y	Y	Y	Y	Y
Clustered SE	Y	Y	Y	Y	Y
Observations	23,196	22,113	23,196	21,030	21,030

Notes: Standard errors in parentheses, clustered by Village. The unit of analysis in these regressions is the household-month.

Star values: * 0.10 ** 0.05 *** 0.010.

Our results, in Table 2.4 (logit model) and in Table 2.5 (linear model), provide evidence that an increase in the rate of blackouts during the previous month decreases the probability of adoption. Furthermore, households seem to be making short sighted decisions as we do not find evidence that the rate of blackouts during $t-2$ is association with the probability to connect in month t, nor does the average blackout rate since the beginning of service. The results from the linear model allow for easy interpretation of the marginal effects. In an extreme case where, in a given month, every day there was a blackout, the probability to connect would decrease by approximately 10%, depending on the preferred specification.

Table 2.6 show the result for the model derived from the first hypothesis from the Bayesian Learning model, expressed in Equation 2.5. We found no evidence of the variance (σ_c^2) affecting the way new information is processed when making adoption decisions.

Table 2.6: Probability of Adoption - Bayesian Learning first testable hypothesis

	(1) OLS	(2) OLS	(3) OLS
Blackout rate $(t1)$	-0.13227*	-0.090790	-
	(0.07074)	(0.062733)	-
Blackout rate $(t2)$	-	-0.004999	-
	-	(0.060272)	-
Avg. Blackout rate $(t1)$	-	-	-0.1951
	-	-	(0.3146)
Variance per site:Blackout rate $(t1)$	0.45970	-0.253181	-
	(2.46786)	(2.101530)	-
Variance per site:Blackout rate $(t2)$	-	-0.228588	-
	-	(1.947711)	-
Variance per site:Avg. Blackout rate $(t1)$	-	-	4.4287
	-	-	(14.1587)
Household FE	Y	Y	Y
Month by year FE	Y	Y	Y
Clustered SE	Y	Y	Y
Observations	23,196	22,113	23,196

Notes: Standard errors in parentheses, clustered by Village. The unit of analysis in these regressions is the household-month.

Star values: * 0.10 ** 0.05 *** 0.010.

Table 2.7 shows the result of the regression model derived from the second hypothesis from the Bayesian Learning model, expressed in Equation 2.6. As with the first Bayesian hypothesis, we find no statistically significant evidence of Bayesian learning.

Table 2.7: Probability of Adoption - Bayesian Learning second testable hypothesis

	(1) OLS	(2) OLS	(3) OLS
Blackout rate $(t1)$	-0.178695	-0.081160	-
	(0.113390)	(0.095642)	-
Blackout rate $(t2)$	-	0.017542	-
	-	(0.088946)	-
Avg. Blackout rate $(t1)$	-	-	-0.23257
	-	-	(0.24386)
Time since opening:Blackout rate $(t1)$	0.004258	-0.001014	-
	(0.007719)	(0.006757)	-
Time since opening:Blackout rate $(t2)$	-	-0.001929	-
	-	(0.007123)	-
Time since opening:Avg. Blackout rate $(t1)$	-	-	0.02561
	-	-	(0.02470)
Household FE	Y	Y	Y
Month by year FE	Y	Y	Y
Clustered SE	Y	Y	Y
Observations	23,196	22,113	23,196

Notes: Standard errors in parentheses, clustered by Village. The unit of analysis in these regressions is the household-month.
Star values: * 0.10 ** 0.05 *** 0.010.

2.5.2 Probability of Disadoption

Next, we estimate equations 2.7 and 2.8 using a logit and a linear model respectively, testing the different learning models discussed in Sections 2.3.1 and 2.4.1. Our results, in Tables 2.8 and 2.9, provide no evidence that an increase in the blackout rate during the previous month increases the probability of permanent disadoption by connected customers, regardless of the learning model. The linear model does, however, provide some weak evidence that longer term blackout rates affect the decision to dis-adopt.

Table 2.8: Probability of Disadoption - Logit Models

	(1) Logit	(2) Logit	(3) Logit	(4) Logit
Blackout rate $(t-1)$	0.507	0.297	-	-
	(0.901)	(1.153)	-	-
Blackout rate $(t-2)$	-	0.695	-	-
	-	(1.307)	-	-
Avg. Blackout rate $(t-1)$	-	-	-0.053	-
	-	-	(0.449)	-
Avg. Blackout rate since connected $(t-1)$	-	-	-	0.266
	-	-	-	(3.332)
Household FE	Y	Y	Y	Y
Month by year FE	Y	Y	Y	Y
Clustered SE	Y	Y	Y	Y
Observations	15,492	15,009	15,492	15,492

Notes: Standard errors in parentheses, clustered by Village. The unit of analysis in these regressions is the household-month.

Star values: * 0.10 ** 0.05 *** 0.010.

Table 2.9: Probability of Disadoption - Linear Models

	(1) OLS	(2) OLS	(3) OLS	(4) OLS
Blackout rate $(t-1)$	0.06188	0.05439	-	-
	(0.04831)	(0.05148)	-	-
Blackout rate $(t-2)$	-	0.05413	-	-
	-	(0.05505)	-	-
Avg. Blackout rate $(t-1)$	-	-	0.3418*	-
	-	-	(0.1813)	-
Avg. Blackout rate since connected $(t-1)$	-	-	-	0.2942**
	-	-	-	(0.1323)
Household FE	Y	Y	Y	Y
Month by year FE	Y	Y	Y	Y
Clustered SE	Y	Y	Y	Y
Observations	15,492	15,009	15,492	15,492

Notes: Standard errors in parentheses, clustered by Village. The unit of analysis in these regressions is the household-month.

Star values: * 0.10 ** 0.05 *** 0.010.

Table 2.10 displays the result for the model derived from the first hypothesis from the Bayesian Learning model, expressed in Equation 2.9. Again, we find no conclusive evidence of the variance (σ_c^2) affecting adoption decisions.

Table 2.10: Probability of Disadoption - Bayesian Learning first testable hypothesis

	(1) OLS	(2) OLS	(3) OLS	(4) OLS
Blackout rate $(t1)$	0.02634	0.05062	-	-
	(0.06193)	(0.05824)	-	-
Blackout rate $(t2)$	-	-0.11431**	-	-
	-	(0.05120)	-	-
Avg. Blackout rate $(t1)$	-	-	0.05297	-
	-	-	(0.23607)	-
Avg. Blackout rate since connected $(t1)$	-	-	-	0.05608
	-	-	-	(0.17601)
Variance per site:Blackout rate $(t1)$	2.17324	-0.44493	-	-
	(3.42262)	(2.45472)	-	-
Variance per site:Blackout rate $(t2)$	-	11.01559***	-	-
	-	(2.01096)	-	-
Variance per site:Avg. Blackout rate $(t1)$	-	-	20.55544*	-
	-	-	(11.86071)	-
Variance per site:Avg. Blackout rate since connected $(t1)$	-	-	-	17.84118*
	-	-	-	(9.70546)
Household FE	Y	Y	Y	Y
Month by year FE	Y	Y	Y	Y
Clustered SE	Y	Y	Y	Y
Observations	15,492	15,009	15,492	15,492

Notes: Standard errors in parentheses, clustered by Village. The unit of analysis in these regressions is the household-month.

Star values: * 0.10 ** 0.05 *** 0.010.

Table 2.11 displays the result of the model derived from the second hypothesis from the Bayesian Learning model, expressed in Equation 2.10. We found no statistically significant evidence that the time as a costumer affects the way new information is processed when making disadoption decisions.

2.5.3 Time to Re-purchase Electricity (Reconnection)

In this section, we present the results of estimating the models developed in Section 2.4.2. We estimate Equation 2.17 using a series of distributional assumptions and levels at which we cluster standard errors. Furthermore, we also include a model in which we control for the number of times a household reached zero balance.

Table 2.11: Probability of Disadoption - Bayesian Learning second testable hypothesis

	(1) OLS	(2) OLS	(3) OLS	(4) OLS
Blackout rate $(t1)$	-0.009175	0.035522	-	-
	(0.086625)	(0.089462)	-	-
Blackout rate $(t2)$	-	-0.070432	-	-
	-	(0.057888)	-	-
Avg. Blackout rate $(t1)$	-	-	0.17466	-
	-	-	(0.16304)	-
Avg. Blackout rate since connected $(t1)$	-	-	-	0.10588
	-	-	-	(0.12404)
Time since opening:Blackout rate $(t1)$	0.006105	0.001247	-	-
	(0.005221)	(0.005076)	-	-
Time since opening:Blackout rate $(t2)$	-	0.010679***	-	-
	-	(0.003869)	-	-
Time since opening:Avg. Blackout rate $(t1)$	-	-	0.04204**	-
	-	-	(0.01899)	-
Time since opening:Avg. Blackout rate since connected $(t1)$	-	-	-	0.04282*
	-	-	-	(0.02162)
Household FE	Y	Y	Y	Y
Month by year FE	Y	Y	Y	Y
Clustered SE	Y	Y	Y	Y
Observations	15,492	15,009	15,492	15,492

Notes: Standard errors in parentheses, clustered by Village. The unit of analysis in these regressions is the household-month.

Star values: * 0.10 ** 0.05 *** 0.010.

Our results, in Table 2.12, indicate that an increase in the rate of blackouts during the waiting period (observed through peers' experiences, $m_{c,}$) and during the previous consumption period (experienced personally, $m_{c,i-1}$) increases the time it takes to reconnect.

Table 2.12: Time to Reconnection

	(1) OLS	(2) Poisson	(3) NB	(4) OLS	(5) Poisson	(6) NB
Blackouts rate previous consumption $(m_{c,i-1})$	6.988***	0.67348***	0.60299***	6.988***	0.67348***	0.60299***
	(1.882)	(0.02725)	(0.09654)	(2.000)	(0.02725)	(0.09654)
Blackouts rate while balance is zero $(m_{c,i})$	8.448***	0.87816***	1.41623***	8.448***	0.87816***	1.41623***
	(1.571)	(0.01730)	(0.06202)	(1.025)	(0.01730)	(0.06202)
Household FE	Y	Y	Y	Y	Y	Y
Month by year FE	Y	Y	Y	Y	Y	Y
Clustered SE Site	Y	Y	Y	N	N	N
Clustered SE Household	N	N	N	Y	Y	Y
Observations	39,763	39,763	39,763	39,763	39,763	39,763

Notes: Standard errors in parentheses, clustered by Village. The unit of analysis in these regressions is the household-waiting period.

Star values: * 0.10 ** 0.05 *** 0.010.

Given any distributional assumption and levels at which we cluster standard errors, an increase in the rate of blackouts during the waiting period (closer in time but not personally experienced) has a larger effect than an increase in the rate of blackouts during the previous consumption period. Furthermore, in Table 2.13 we show the results for a model in which we interact the two rates of blackouts with a variable counting the number of waiting periods, showing that the effect of the blackout rate during the waiting period slowly decreases the longer the household has been without service.

Table 2.13: Time to Reconnection - Controlling for the number of waiting periods

	(1) OLS	(2) Poisson	(3) NB
Blackouts rate previous consumption ($m_{c,i-1}$)	8.8912687***	0.6763657***	0.7262***
	(2.5994818)	(0.0386821)	(0.1403)
Blackouts rate while balance is zero ($m_{c,i}$)	11.7906212***	1.0654686***	1.720***
	(2.2382660)	(0.0263299)	(0.09262)
Number of waiting period	0.0003486	-0.0002325**	0.000435
	(0.0027223)	(0.0001029)	(0.0026)
Number of waiting period: Blackouts rate previous consumption	-0.0592812	-0.0003141	-0.0037
	(0.0362565)	(0.0010031)	(0.00297)
Number of waiting period: Blackouts rate while balance is zero	-0.0853785***	-0.0058621***	-0.00790***
	(0.0253058)	(0.0006486)	(0.001790)
Household FE	Y	Y	Y
Month by year FE	Y	Y	Y
Clustered SE Site	Y	Y	Y
Observations	39,763	39,763	39,763

Notes: Standard errors in parentheses, clustered by Village. The unit of analysis in these regressions is the household-waiting period.

Star values: * 0.10 ** 0.05 *** 0.010.

Tables 2.14 and 2.15 show the results of testing the first and second Bayesian hypotheses with respect to the time to re-purchase model. These results are not conclusive and not especially clear. On the one hand, it seems that given certain distributional assumptions, there is an association between the variance (σ_c^2) and the way a costumer uses new information about reliability. We expect this association to be negative, as the Bayesian Learning model assumes the larger the variance, the less valuable are new draws from a distribution. If we look only at significant coefficients, the relationship is negative for the older and personally experienced information and positive for the more recent and observed information. These results can be explained as randomness or there could be an underlying behavioral explanation yet to unfold.

Table 2.14: Time to Reconnection - Bayesian Learning 1st hypothesis

	(1) OLS	(2) Poisson	(3) NB
Blackouts rate previous consumption ($m_{c,i-1}$)	6.000**	0.6539***	5.825e-01***
	(2.457)	(3.276e-02)	(1.152e-01)
Blackouts rate while balance is zero ($m_{c,i}$)	6.265***	7.550e-01***	1.131e+00***
	(1.571)	(2.271e-02)	(7.908e-02)
Variance per Site : Blackouts rate previous consumption	9.170	-2.463**	1.854
	(129.992)	(1.117)	(4.639)
Variance per Site : Blackouts rate while balance is zero	157.122	7.752***	20.49***
	(122.270)	(0.9690)	(3.965)
Household FE	Y	Y	Y
Month by year FE	Y	Y	Y
Clustered SE Site	Y	Y	Y
Observations	39,763	39,763	39,763

Notes: Standard errors in parentheses, clustered by Village. The unit of analysis in these regressions is the household-waiting period.

Star values: * 0.10 ** 0.05 *** 0.010.

Table 2.15: Time to Reconnection - Bayesian Learning 2nd hypothesis

	(1) OLS	(2) Poisson	(3) NB
Blackouts rate previous consumption ($m_{c,i-1}$)	-0.965360	-2.571e-01***	-0.1514423
	(3.360258)	(5.938e-02)	(0.1943158)
Blackouts rate while balance is zero ($m_{c,i}$)	11.253962***	1.228***	1.7401439***
	(3.344357)	(3.598e-02)	(0.1253054)
Days since first connection	0.020157**	2.020e-03***	0.0018602***
	(0.008979)	(2.014e-04)	(0.0006379)
Days since first connection : Blackouts rate previous consumption	0.016061**	1.811e-03***	0.0014686***
	(0.006280)	(9.895e-05)	(0.0003370)
Days since first connection : Blackouts rate while balance is zero	-0.005438	-6.607e-04***	-0.0006778***
	(0.005203)	(6.282e-05)	(0.0002155)
Household FE	Y	Y	Y
Month by year FE	Y	Y	Y
Clustered SE Site	Y	Y	Y
Observations	39,763	39,763	39,763
39,763	39,763	39,763	

Notes: Standard errors in parentheses, clustered by Village. The unit of analysis in these regressions is the household-waiting period.

Star values: * 0.10 ** 0.05 *** 0.010.

2.5.4 Amount Purchased

When we look at the amount purchased (Equation 2.18) rather than at the time it takes to re purchase pre-paid electricity, we find no evidence (Table 2.16) of an association between the rate of blackouts during the waiting period or during the consumption period and the amount purchased (Results displayed in Table 2.16).

Table 2.16: Amount Purchased

	(1) OLS
Blackouts rate previous consumption ($m_{c,i-1}$)	3.963
	(32.065)
Blackouts rate while balance is zero ($m_{c,i}$)	-9.812
	(15.723)
Household FE	Y
Month by year FE	Y
Clustered SE	Y
Observations	39,763
R-squared	0.3212

Notes: Standard errors in parentheses, clustered by Village. The unit of analysis in these regressions is the household-waiting period.

Star values: * 0.10 ** 0.05 *** 0.010.

Finally, Table 2.17 displays the results of the Bayesian learning tests on the amount purchased.

2.5.5 Results Summary

Our results provide evidence that electricity service unreliability can negatively impact electricity service adoption rates and electricity usage. First of all, an increase in the blackout rate during the previous month decreases the probability of adoption. Households respond significantly more to more recent information. Specifically, households respond to the blackout rate in $t-1$ but not the blackout rate two months prior ($t-2$). We also find no evidence that longer term information, such as the simple average of the blackout rate since the beginning of service, affects the probability of adoption. We also did not find evidence that an increase in the blackout rate during the

Table 2.17: Amount Purchased - Bayesian Learning 1st and 2nd hypotheses

	(1) OLS	(2) OLS
Blackouts rate previous consumption ($m_{c,i-1}$)	-3.967	12.51024
	(30.469)	(79.78164)
Blackouts rate while balance is zero ($m_{c,i}$)	-29.814	-15.84701
	(21.220)	(37.83225)
Variance per Site : Blackouts rate previous consumption	1338.655	-
	(923.559)	-
variance per Site : Blackouts rate while balance is zero	1621.398	-
	(1232.512)	-
Days since first connection : Blackouts rate previous consumption	-	-0.01753
	-	(0.13912)
Days since first connection : Blackouts rate while balance is zero	-	0.01226
	-	(0.06118)
Household FE	Y	Y
Month by year FE	Y	Y
Clustered SE Site	Y	Y
Observations	39,763	39,763

Notes: Standard errors in parentheses, clustered by Village. The unit of analysis in these regressions is the household-waiting period.
Star values: * 0.10 ** 0.05 *** 0.010.

previous month increases the probability of disadoption of connected customers, regardless of the learning model.

To look for evidence of Bayesian learning in both adoption and disadoption models, we tested the interaction between the household specific experience (time since site availability or time as a customer) and the different ways we assume a household can learn, finding no clear evidence of a decline in the value given to observed information based on experience. We also tested the interaction between the variance observed in the site and the different ways we assume a household can learn, finding no clear evidence of a decline in the value given to observed information based on the site specific variance of reliability.

An increase in the blackout rate during the waiting period (observed through peers' experiences) and during the previous consumption period (experienced personally) increases the time it takes to reconnect, with a larger effect for an increase in the blackout rate during the waiting period (closer in time but not personally experienced). Finally, we find no evidence of a negative association between the blackout rate during the waiting period or the previous consumption period and the amount purchased.

2.6 Conclusions and Discussion

Our findings on adoption decisions are aligned with what was predicted in our theoretical model: expected utility is negatively affected by unreliability, which causes households to purchase fewer electricity days. Furthermore, we confirm that households learn about technology quality through their own unreliability experiences and through observation of peer experiences with unreliability.

There are several possible explanations for the lack of statistical evidence on the association between long term disadoption decisions and experienced reliability that are consistent with the theoretical model. First of all, the value given to having access to electricity (γ_h) could increase as the household has access to electricity for a certain period of time, even though this was not a feature of our model. If this is the case, we might not have the power to detect the smaller effect reliability would have on the probability to disadopt. Furthermore, because of the fixed costs, the cost of electricity (C) is declining in the time since the household became a customer. These findings are consistent with what we observed in a surveyed sample of households from MeshPower's service area (see Appendix), where the majority of customers who dis-adopted stated that they did so because they had access to other sources of electricity or because of their economic situation. Overall, the customers' perception of the reliability of MeshPower is as good or better than expected. However, the model studying the time it takes a household to re-purchase electricity yielded the most robust results. As we discussed before, it can be seen as a temporal and reversible disconnection, and we found evidence of the rate of blackouts, not only experienced but also observed, affecting these decisions.

We found substantial evidence that households place more weight on recent information than information from several periods prior. This suggests that improving the service reliability can quickly change peoples' perceived benefits, so demonstrations, free trials, etc. could be an effective way to improve the adoption of electricity services. We did not find evidence of Bayesian learning, consistent with previous literature using experimental settings [Charness and Levin, 2005, Zizzo et al., 2000].

We provide empirical evidence to support the claim in Lee et al. [2018] that reliability can be a factor in explaining the difficulty of assuring universal electricity adoption in rural areas of the developing world [Eder et al., 2015, Müggenburg et al., 2012, Shyu, 2013]. Furthermore, we provide a new context in which peer experience can undermine adoption efforts of apparently beneficial technology, as in Adhvaryu [2014]. Our results also suggest that the low willingness to pay (WTP) for electricity services found by Lee et al. [2016] and by Grimm et al. [2019] might be driven in part by low reliability.

Our results are complementary to Sedai et al. [2020] and Meles [2020], as they quantified the cost of blackouts for connected household but they did not look at the effect of blackouts on the decisions at the extensive or intensive margins of electricity consumption.

Next, evidence that households learn from neighbors suggests the existence of an option value to waiting to adopt [Ansar and Sparks, 2009]. Because information on the quality of electricity services is a public good, there exists an economic argument for subsidizing initial connections that produce information. Furthermore, if experiences with technology from one firm inform beliefs about other firms' products, this collective reputation (Winfree and McCluskey [2005]) can lead to underinvestment in reliability and slower adoption over time. This suggests a benefit to establishing clear technology standards that lead to socially optimal levels of investment in quality, particularly when the private sector plays a large role in rural electrification efforts. Nonetheless, our findings indicating that adoption decisions are affected mostly by recent reliability outcomes, increasing the noise associated with the public information, suggest that learning from experience can exacerbate the inefficiency of adoption paths.

Finally, our study is not without limitations. For instance, we did not have access to data related to the intensity of the electricity usage or to actual outages. Future studies should rely on higher frequency data to understand the short term effects of outages on the intensive margin. On the other hand, to properly study how households learn from peers, having access to geo-localized data at the household level would allow researchers to properly capture the effects of the neighboring customers on non-customers.

Chapter 3

The Effect of Pollution on Crime and Anti-Social Behavior in the United Kingdom

3.1 Introduction

Scientific literature has identified the impact of short-term pollution exposure on crime in the United States (US) [Burkhardt et al., 2019b, Herrnstadt et al., 2020, Lu et al., 2018] and in the United Kingdom (UK) [Roth et al., 2020], shedding light on a previously overlooked cost of pollution. Causal mechanisms explaining the association between crime and pollution are not fully understood. Research in epidemiology indicates a biological mechanism such that pollution exposure can have short-term effects on cognitive skills, anxiety, and certain behaviors associated with criminal or violent activities [Kioumourtzoglou et al., 2017, Lu et al., 2018, Power et al., 2015]. While the economic literature has quantitatively demonstrated the phenomena, there is little evidence of the mechanisms driving the relationship.

To better understand the relationship between criminal activity and pollution, this study uses dis-aggregated counts of offenses from the UK to determine which types of crimes are associated with pollution, including offenses that are not formally crimes. Our primary hypotheses are a natural evolution of previous work on this topic. We hypothesize that pollution affects criminal behavior through two channels. First, Burkhardt et al. [2019a] finds that changes in $PM_{2.5}$ only affect the propensity for violent crimes with an emphasis on assaults, a form of impulsive and extremely aggressive behavior. This work combined with previous work in epidemiology suggests that fine particulate matter air pollution ($PM_{2.5}$) can induce biological processes, like systemic inflammation, which could potentially exacerbate aggressive behavior [Brook et al., 2004, Cunningham et al., 2009, Donaldson et al., 2001]. We hypothesize that if increased aggression is the key mechanism behind the crime-pollution relationship, then we should find a similar relationship between pollution and a lesser form of aggression, namely anti-social behaviour (ASB).[18] Second, in contrast to Burkhardt et al. [2019a], Roth et al. [2020] find that pollution in the UK affects not only violent crime, but also crimes that are economically motivated with larger effects on those that are spontaneous. Based on the rational choice theory of crime model proposed by Becker [1968], Roth et al. [2020] derives a simple utility based decision model and concludes that the mechanism explaining the association between crime and pollution is that pollution exposure increases the offender's discount rate, which reduces the present expected costs of future punishment. Our second hypothesis follows this alternative strand of research and as such, we test whether we can identify the effect of pollution on economically motivated crimes at a more aggregate level using similar data from the UK. Combined, the results of these tests shed further light on the mechanisms driving the association between pollution and crime.

Using a fixed effects model to explain the relationship between pollution and offenses, by month and by local authority district (LAD), we found evidence of an association between pollution and the rates of "Burglary" and "Robbery", and a relatively weak association with ASB and "Violence and Sexual Offenses". We could not find evidence of an association between pollution and "Public Order". Thus we find little evidence to support the hypothesis that pollution increases aggression in the UK. Instead, our results support previous findings that indicate pollution exposure in the UK increases the propensity to commit economically motivated crimes [Roth et al., 2020]. Our results indicate that pollution in the UK may affect cognitive functions causing a change in the parameters of an offender's utility function, but it does not appear to substantially influence the probability of aggression.

This paper contributes to the crime and the economics literature in several ways. First, we provide more empirical evidence on the association between pollution and crime. In the US, this

[18] In the UK, local police forces report not only crimes but, as defined by the Antisocial Behaviour Act from 2003, "behavior by a person which causes, or is likely to cause, harassment, alarm or distress to persons not of the same household as the person". As a punishment for ASB, one can receive a civil injunction, Community Protection Notice (CPN) or Criminal Behaviour Order (CBO), or an Anti-social behaviour order (ASBO) [Government, 2020]. ASB covers a wide range of unacceptable activity that causes harm to an individual, to their community or to their environment, including nuisance, rowdy or inconsiderate neighbors, vandalism, graffiti and flyer posting, environmental damage including littering, dumping of rubbish and abandonment of cars, etc.

association has been consistently identified using different data sets and aggregation levels. While Lu et al. [2018] document correlation between air pollutants and crime rates measured annually in almost 10,000 US cities, Herrnstadt et al. [2020] and Burkhardt et al. [2019a] identified this association using data for several US cities and for 99% of the US counties respectively. Finally, Burkhardt et al. [2019b] uses daily data to identify the association in 391 counties in the United States. Outside the US, Roth et al. [2020] used daily data to study the relationship for a two-year period in London. In our work, we replicate the findings of Roth et al. [2020] using a different data set, where offenses were aggregated to the month level and locations were "jittered" to favor anonymity [Tompson et al., 2015], serving future researchers efforts to understand crime events using the publicly available data instead of the data of limited availability that Roth et al. [2020] had access to.

Second, by studying the association between different offenses (crimes and ASB) and pollution, we shed light on potential mechanisms. Based on our theoretical model, we argue that pollution can decrease a criminal's will to perform alternatives activities, or it can change the difference between the reward and the discounted punishment. Both options are consistent with epidemiological literature indicating that pollution exposure can affect cognitive skills, anxiety, and certain behaviors associated with criminal or violent activities [Kioumourtzoglou et al., 2017, Lu et al., 2018, Power et al., 2015]. Interestingly, we find little evidence that ASB is associated with pollution levels in the UK. This finding is surprising in light of research on US crimes and pollution, and indicates further research in the US is warranted.

A third strand of literature has studied if certain demographic characteristics can influence the relationship between pollution and offenses. Burkhardt et al. [2019a] tested for heterogeneity in the association across socio-demographic and regional dimensions, concluding the most important explanatory factor in the relationship between pollution and crime is age, and the results are not ameliorated by higher incomes. Roth et al. [2020], on the other hand, used housing prices as a proxy for wealth and found a U-shape relationship where the effects are largest at the tails of the housing price distribution. Building on these findings, we tested if the percentage of young males (those more likely to commit aggressive acts) and the percentage of non-domestic electricity consumption (a proxy for how residential a LAD is) modify the association between offenses and crime. We found only weak evidence that these two demographics characteristics affect the rate of "Violence and Sexual Offenses".

3.2 Theoretical Model

Several in the economics literature have used utility theory to propose a rational choice model to explain crime, the most well-known being Becker [1968]. In this paper, we build on a model developed by Roth et al. [2020]. The primary assumption is that agent i would commit a crime (C) if the expected utility of committing the crime is larger than the utility of not committing the crime (NC), as in Equation 3.1.

$$E[U^i(C)] > U^i(NC) \tag{3.1}$$

If agent i decides to commit a crime, there is a given probability p of getting caught. Therefore:

$$E[U^i(C)] = p\ U^i(C)|Caught + (1-p)\ U^i(C)|Not\ Caught > U^i(NC) \tag{3.2}$$

Assuming that there is a reward for the crime (R^C) and a probable discounted future punishment ($\beta\ P$), the final mathematical model is as shown in Equation 3.3:

$$p\ U^i(R^C - \beta\ P) + (1-p)\ U^i(R^C) > U^i(NC) \tag{3.3}$$

We believe this model is valid for both crimes and ASB decisions. Though ASB are not formally crimes, a court may give anyone a civil injunction or a Community Protection Notice (CPN) if the court receives reports of persistent ASB from the police, a council or a landlord. For example, a court may order an offender to stay away from a particular place, stop spending time with certain people, work on improving behaviour, or fix damage caused to a victim's property. Failing to comply can result in more severe punishments, including imprisonment [Government, 2020].

Moreover, utility theory assumes that if an agent engages in certain behavior, it is because that agent receives utility from the behavior. Even when ASB rarely has an economic motivation that can be easily translated into personal well-being, the same can be said about certain crime types, such as assaults or public order crimes.

Based on the proposed model, we can propose mechanisms through which pollution might affect crime (and ASB). First, pollution might affect the left-hand side of the inequality in Equation 3.3 if it alters the difference between the reward and the discounted punishment. Second, pollution might affect the right-hand side of the inequality by decreasing the utility of the alternative activity to committing a crime (NC). Third, it might decrease the concavity of the utility function (risk seeking behaviour) and fourth, it might change the probability of getting caught.

To tease out the potential mechanisms, we consider three types of offenses: economically motivated crimes, crimes associated with aggressive and violent behavior, and offenses associated with aggressive and violent behavior that are not formally crimes. If the effect is homogeneous across the types of offenses, we believe we cannot disentangle the importance of the competing mechanisms. If the effect is not homogeneous across these three types, it is less likely that the mechanism is through the level of risk (p) or through the level of risk aversion (concavity of the utility function), as the changes would affect all crime types. If the effect is mostly associated with aggressive and violent behavior, it favors the idea that the mechanism is through increasing the utility of the criminal activity. If the effect is mostly on economically motivated crimes, it favors the idea that the mechanism is through the decrease of the utility of the alternative activity or through decreasing the discount rate of future punishment. Possible results and favored mechanisms are summarized in Table 3.1.

Table 3.1: Mechanism favored by each possible result

Possible result	Favored Mechanism
Homogeneous effect across offenses.	We cannot disentangle the competing mechanisms.
Heterogeneous effect across offenses	Less likely through level of risk (p) or risk aversion
Mostly on activities associated with aggressive behavior	More likely through increasing the utility of the criminal activity
Mostly on economically motivated crimes	More likely through the decrease of the utility of the alternative activity or through decreasing the discount factor of future punishment

In the Conclusions section, we discuss the mechanisms based on our quantitative results.

3.3 Data

3.3.1 Crime and ASB

Data on crime and ASB counts is publicly available from Police.uk, the site for open data about crime and policing in England, Wales and Northern Ireland. Police.uk centralizes the access to data reported from 45 police agencies, as shown in Figure 3.1.

Figure 3.1: Map of Police Agencies.

The data set is a collection of events. Each agency is responsible for reporting in their area, including geo-location of the events. Police.uk, seeking to preserve anonymity by rounding up the time stamps to the monthly level and adding noise to the incident coordinates by allocating each event to the closest "snap point" out of a list of 750,000 locations in the UK [Tompson et al., 2015]. Thus, we do not know the precise location or time each crime occurred.

From May 2013 onward, Police agencies report 16 Types of offenses: "Anti-social Behaviour", "Bicycle Theft", "Burglary", "Criminal Damage and Arson", "Drugs", "Other Crime", "Other Theft", "Possession of Weapons", "Public Order", "Robbery", "Shoplifting", "Theft from the person", "Vehicle Crime" and "Violence and Sexual Offenses".

To disentangle mechanism and test our hypotheses, we compare crime types that we believe are different in nature. As such, we focus on ASB (low-level aggressive behavior that is not technically a crime), "Public Order" (POC), "Violence and Sexual Offenses" (VSO) (higher-level aggressive behavior that is a crime), and "Burglary" (BUR) and "Robbery" (ROB) (economically motivated crimes).

As categories changed over time, the initial date for the different counts varies. Table 3.2 shows how and when the classification changed.

Table 3.2: Crime Types reporting across time.

Crime Type	December 2010	September 2011	April 2013	May 2013
Anti-Social Behaviour	Y	Y	Y	Y
Burglary	Y	Y	Y	Y
Robbery	Y	Y	Y	Y
Public Order	Y	Y	Y	Y
Violent Crime	Y	Y	Y	N
Violence and Sexual Offenses	N	N	N	Y

Police agencies began reporting ASB, BUR and ROB in December 2010. POC was first recorded in May 2013. "Violent Crime" was renamed VSO in May 2013, so we renamed the incidents classified as "Violent Crime" as VSO so that the panel starts in December 2010.[19]

UK Police agencies upload, in every month t, a list of the crime events that happened between $t - 2$ and $t - 38$. In other words, they upload 36 months of data with a lag of 2 months. This

[19]Violent crimes included sexual offenses prior to 2013.

means that every time they upload new data, Police.uk is uploading one new month and updating what happened in the previous 35 months. Crime events for one specific month are then reviewed 35 times after they were published for the first time.

To understand how reliable each version is, in Figure 3.2 we show how much a count changed, on average, every time it was updated.

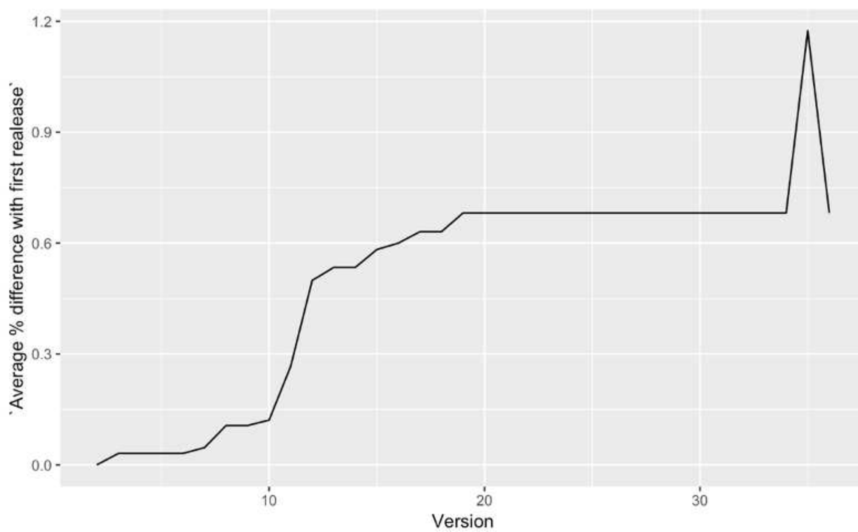

Figure 3.2: Average percentage difference with the first release as a count is updated

Figure 3.2 shows that counts do not change much on average. After 36 versions, the crime counts changed less than 0.7% relative to the original uploaded value. The spike we observe in the 35th version is driven by a large correction in one specific agency that was "re-corrected" in the next version (the 36th), ending with the same value as they had in the 34th version. Nevertheless, the average hides the fact that some counts did change considerably (up to 74%) while 92% of counts did not change. As our sample ends in December 2018 and we used the last available version, all our reported counts were validated at least 15 times.

We use LADs as the observational unit for two main reasons: previous research has shown that due to the spatial "jittering" performed by police.uk to preserve anonymity, the spatial accuracy of the data is very good at the LAD level, but not so much at a more disaggregated level such as the postal code [Tompson et al., 2015].[20] Second, at the LAD level, the UK Office for National Statistics (ONS) has demographic and socioeconomic data available that we make use of.

In the UK, there are 382 LADs (326 in England, 22 in Wales, 11 in Northern Ireland and 32 in Scotland), and their size and population differ across regions. Figure 3.3 shows the LADs across the UK, and the agency in which it is located.

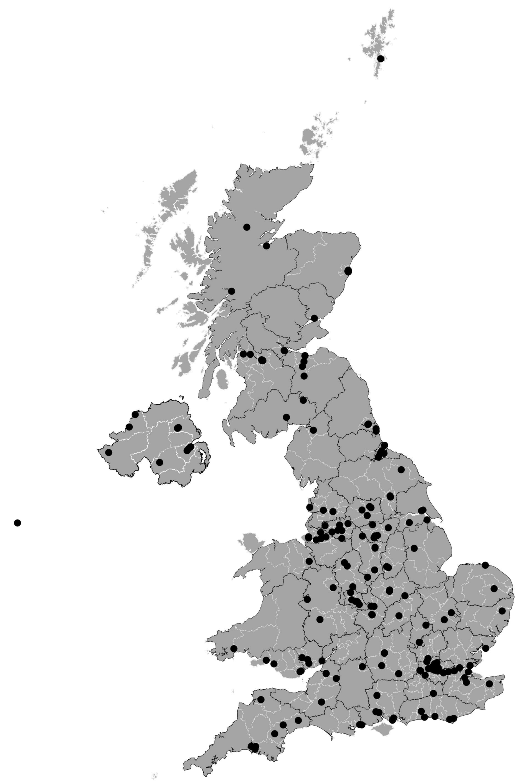

Figure 3.3: Police Agencies, Local Authority Districts (LADs) and Pollution Stations Map

3.3.2 Pollution

Pollution data are publicly available from the Department of Environment Food and Rural Affairs, measured hourly by a network of monitoring sites across the UK. To create LAD-specific observations, we use an average of the measurements from the 5 closest monitoring sites, weighted

[20]Postal codes are smaller than LADs.

based on the distance to the LAD's centroid. For our analysis we use the monthly average and the average of the daily maximums for $PM_{2.5}$ and Ozone.

Figure 3.3 shows the location of the pollution stations across the UK, shedding light on the relative distance of each LAD to the closest pollution station. Site specific PM pollution is at best very weakly correlated with pollution measurements that are conducted more than 20 kilometers away from a particular site [EPA, 1997]. Therefore, we remove every LAD for which the closest pollution station is further than 20 kilometers away.

Table 3.3: Summary Statistics - Distance to Pollution Stations

Statistic		Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Distance (km)	Every LAD (Population)	14.1	12.8	0.3	4.8	19.2	83.0
Distance (km)	Pollution station within 20 km	8.8	5.6	0.3	3.7	13.1	20.0

The first row of Table 3.3 displays summary statistics of the distance to the closest pollution station for all LADs. As the 75^{th} percentile is very close to 20 kilometers, we are including approximately 75% of the LADs in our analysis. The second row of Table 3.3 displays the summary statistics of the distance to the closest pollution station once we exclude those LADs that where further than 20 kms away from the closest pollution station. On average, the nearest pollution monitoring station is 8km away from a LAD centroid.

Following [Roth et al., 2020], we transform the pollution variables (measured in μg/m^3) to an Air Quality Index (AQI), according to the US EPA formula. AQI is an index that ranges from 0 to 500, measuring the level of Health Concern. An AQI value less than 50 represents good air quality with little to no potential to affect public health. An AQI value between 51 and 100 indicates a moderate health concern. An AQI between 101 and 151 represents unhealthy air quality for sensitive groups. An AQI between 151 and 200 indicates everyone may experience some adverse health effects. Finally, an AQI between 201 and 300 is very unhealthy for all group and an AQI over

300 indicates the air quality is hazardous for all groups.[21]. As in [Roth et al., 2020], we calculate the monthly AQI for the monthly daily mean concentration of each pollutant ($PM_{2.5}$ and O3), and use the monthly AQI for the largest of the 2, under the assumption that the "binding" pollutant is the one with the largest AQI. To test for robustness and to further understand the association, we also we look at the percentage of days in a given month the mean daily AQI and the max daily AQI has surpassed each threshold across levels of health concern.

3.3.3 Weather

Weather data are publicly available from the national meteorological service for the UK (The Met Office), measured by a network of weather stations across the UK. As we did with pollution, to create LAD-specific observations we use an average of the 5 closest weather monitoring stations, weighted by the distance to the LAD's centroid. For our analysis, we use the monthly average of the daily max temperature, the monthly average of the daily min temperature and the monthly average daily precipitation. Figure 3.4 shows that the density of weather stations across the UK is larger than for pollution station.

Weather data are available up to the end of 2018, defining the end of the time series used in the analysis.

3.3.4 Demographics

We use demographic data from the UK Office for National Statistics. Specifically, we use LAD level population data (gender and ages), and on electricity consumption, disaggregated by LAD, domestic and non-domestic. These two variables are important to account for time variant characteristics at the LAD level that can affect crime rates.

[21] https://airnow.gov/index.cfm?action=aqibasics.aqi

Figure 3.4: Location of Weather Stations Across the UK

3.3.5 Summary statistics

Pooling all the LADs together, Table 3.4 shows the summary statistics for the count for each crime type, including ASB, per 10,000 inhabitants and for pollution and control variables.

Table 3.4: Summary Statistics

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Anti Social Behaviour (# per 10.000 inhabitants)	23,108	19.17	12.59	0.06	10.65	24.00	120.75
Public Order (# per 10.000 inhabitants)	16,457	2.34	2.37	0.08	1.02	2.80	34.46
Violence and Sexual Offences (# per 10.000 inhabitants)	23,112	8.60	5.74	0.58	4.93	10.43	83.70
Burglary (# per 10.000 inhabitants)	23,094	5.53	2.75	0.06	3.58	7.02	72.60
Robbery (# per 10.000 inhabitants)	20,931	0.76	0.91	0.03	0.22	0.94	11.49
$PM_{2.5}$ AQI	23,112	69.59	15.50	30.00	59.00	77.00	153.00
O3 AQI	23,112	30.42	6.18	11.00	26.00	35.00	50.00
AQI	23,112	69.59	15.50	30.00	59.00	77.00	153.00
Mean $PM_{2.5}$ (μg/m^3)	23,112	12.00	4.47	4.09	8.97	13.88	35.26
Max $PM_{2.5}$ (μg/m^3)	23,112	20.87	7.05	7.03	16.05	24.34	58.58
Mean O3 (μg/m^3)	23,112	42.87	11.70	10.32	34.85	51.08	79.00
Max O3 (μg/m^3)	23,112	64.35	13.28	23.11	55.50	73.13	108.14
Mean Max temperature (C)	23,112	13.13	5.08	2.03	8.57	17.54	26.15
Mean Rain (mm)	23,112	4.04	9.15	0.00	1.29	3.33	185.99
Young Male Population (%)	23,112	0.08	0.02	0.03	0.07	0.08	0.15
Non-domestic Electricity Consumption (%)	23,112	0.94	0.02	0.87	0.93	0.96	0.99

Table 3.5 displays summary statistics for the percentage of days in a month AQI (from $PM_{2.5}$, O3 and both) surpasses each cutoff (50, 100, 150, 200 and 300), looking at both the average AQI per day and the max AQI per day.

Table 3.5 shows that the vast majority of days had a max AQI of less than 100, which is the cutoff for what it is considered "unhealthy", but the vast majority of days had a mean AQI of less than 50, which is the limit for what it is considered good air quality. As we previously discussed, the binding pollutant is almost always $PM_{2.5}$.

3.4 Identification Strategy and Empirical Specification

Identifying the effect of Pollution on Crime and aggressive behaviour using observational data has been a challenging endeavour. In the present setting, there are four major threats to identification: 1) omitted variable bias, 2) measurement error, 3) sample selection, and 4) misspecification

Table 3.5: Summary Statistics - AQI counts

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
% of days max $PM_{2.5}$ AQI > 50	23,112	78.19	17.55	3.00	67.70	93.30	100.00
% of days max $PM_{2.5}$ AQI > 100	23,112	10.80	13.50	0.00	0.00	16.67	80.65
% of days max $PM_{2.5}$ AQI > 150	23,112	2.34	5.04	0.00	0.00	3.20	52.00
% of days max $PM_{2.5}$ AQI > 200	23,112	0.03	0.33	0.00	0.00	0.00	7.00
% of days max $PM_{2.5}$ AQI > 300	23,112	0.003	0.10	0.00	0.00	0.00	3.00
% of days max O3 AQI > 50	23,112	1.54	4.19	0.00	0.00	0.00	39.00
% of days max O3 AQI > 100	23,112	0.08	0.51	0.00	0.00	0.00	13.00
% of days max O3 AQI > 150	23,112	0.00	0.04	0.00	0.00	0.00	3.00
% of days max O3 AQI > 200	23,112	0.00	0.00	0.00	0.00	0.00	0.00
% of days max O3 AQI > 300	23,112	0.00	0.00	0.00	0.00	0.00	0.00
% of days max AQI > 50	23,112	78.19	17.54	3.00	67.70	93.30	100.00
% of days max AQI > 100	23,112	10.84	13.47	0.00	0.00	16.67	80.65
% of days max AQI > 150	23,112	2.34	5.04	0.00	0.00	3.20	52.00
% of days max AQI > 200	23,112	0.03	0.33	0.00	0.00	0.00	7.00
% of days max AQI > 300	23,112	0.003	0.10	0.00	0.00	0.00	3.00
% of days mean $PM_{2.5}$ AQI > 50	23,112	31.82	21.05	0.00	16.13	43.33	100.00
% of days mean $PM_{2.5}$ AQI > 100	23,112	2.45	5.37	0.00	0.00	3.20	45.00
% of days mean $PM_{2.5}$ AQI > 150	23,112	0.18	1.00	0.00	0.00	0.00	11.00
% of days mean $PM_{2.5}$ AQI > 200	23,112	0.00	0.00	0.00	0.00	0.00	0.00
% of days mean $PM_{2.5}$ AQI > 300	23,112	0.00	0.00	0.00	0.00	0.00	0.00
% of days mean O3 AQI > 50	23,112	0.002	0.09	0.00	0.00	0.00	3.00
% of days mean O3 AQI > 100	23,112	0.00	0.00	0.00	0.00	0.00	0.00
% of days mean O3 AQI > 150	23,112	0.00	0.00	0.00	0.00	0.00	0.00
% of days mean O3 AQI > 200	23,112	0.00	0.00	0.00	0.00	0.00	0.00
% of days mean O3 AQI > 300	23,112	0.00	0.00	0.00	0.00	0.00	0.00
% of days mean AQI > 50	23,112	31.82	21.05	0.00	16.13	43.33	100.00
% of days mean AQI > 100	23,112	2.45	5.37	0.00	0.00	3.20	45.00
% of days mean AQI > 150	23,112	0.18	1.00	0.00	0.00	0.00	11.00
% of days mean AQI > 200	23,112	0.00	0.00	0.00	0.00	0.00	0.00
% of days mean AQI > 300	23,112	0.00	0.00	0.00	0.00	0.00	0.00

of the functional form of the regression, but they can all be summarized as endogeneity (correlation between an independent variable and the error term).

Omitted variable bias means that the coefficients of the variables included in the model that are correlated with the omitted variable are biased. In this specific setting, crime is likely correlated with many explanatory variables that we do not observe, and those are likely correlated with some of our predictors. A panel data model allows us to control for time-varying police force specific unobservables.

The impact of measurement error depends on which variable is measured with error and the nature of the mismeasurement. A mismeasured dependent variable with purely random measurement error results in unbiased parameter estimates, but the standard errors of the parameter estimates will be larger than without the measurement error. Random measurement error of an independent variable results attenuation bias. As crime data depends on reporting, the dependent variable (count) is noisy. If there is consistent under reporting, it will bias coefficients and increase standard errors. However, we have no reason to believe that crime reporting measurement error is systematic and we find that the corrected crime reports are not significantly different from the first crime reports (Figure 3.2). Likewise, pollution is certainly measured with error, with the mostly likely cause being daily changes in wind direction, which is clearly random. Hence, our estimates, to the extent we can control for other endogeneity challenges, will likely suffer from a small amount of attenuation bias and additional variance.

Sampling bias arises when the sample is not representative of the population. As we use only those LADs in U.K. that have a pollution monitoring station within 20 km of the LAD centroid, we are left with a non-random sample of LADs. In particular, LADs in our sample our more populated than LADs that do not fit our omission criteria. Hence, we will account for this selection challenge when interpreting our coefficient estimates.

A functional form misspecification generally means that the model does not account for important nonlinearities, causing bias in the remaining parameter estimates. As the "right" specification is unknown, we can only test for different functional form assumptions and use statistical methods to test their fit. To address this challenge we estimate a series of models to determine the robustness of the results to various functional forms.

Most of the empirical work on this topic has relied on the panel structure of the data to estimate models with geographic and time fixed effects, with some of them testing the robustness of the results with an instrumental variable approach using wind direction [Roth et al., 2020] or forest fire smoke [Burkhardt et al., 2019a].

We rely on a panel data structure with LADs as the observational unit (as in Equation 3.4) to model the log of the monthly (t) crime rate (events per 10,000 inhabitants) ($log(y_{i,t})$) at the LAD level (i) in a given Police Force (j), as a function of AQI ($AQI_{i,t}$). We use a set of covariates to control for weather and for socioeconomic and demographic characteristics (vector $X_{i,t}$). As LADs are nested within police agencies, we include police force-by-year-by-month fixed effects ($\theta_{j,t}$) to control for police force specific unobservables that can change every month such as police enforcement, reporting issues, and changes in the overall agency policies. These fixed effects also capture police force specific seasonality. Finally, $\epsilon_{i,t}$ is the normally distributed error term.

$$log(y_{i,t}) = \beta AQI_{i,t} + \gamma^T X_{i,t} + \theta_{j,t} + \epsilon_{i,t} \tag{3.4}$$

Based on the unreliability of pollution measurement, we used the inverse of the distance to the closest pollution monitoring station as the weight in a weighted regression. Finally, we cluster standard errors at the police force level to account for correlated residuals not only at the LAD level but also at the Agency level.

To study heterogeneous effects, we use a second specification where we interact AQI with the % of young males in the population and with the % of non-domestic electricity consumption.

Lastly, including Police Force-by-year-by-month fixed effects imposes constraints on our data set, as the City of London Police report crimes on a unique LAD. The one to one matching between LAD and Police Force forces us to remove the City of London from the sample. While we recognize the limitation of this decision, we feel it is more important to control for confounding variables via our fixed effects than to include London in the sample at the expense of weaker identification.

3.5 Results

In this section, we first examine the association between AQI and crime and ASB counts, as in Equation 3.4. Our primary model estimates a linear association with AQI measured as a continuous variable. We test alternative functional forms to capture non-linear relationships including a categorical model where we use a binned version of the AQI measurements. To further understand the mechanism and to test robustness, we also estimate models using the percentage of days per month the maximum or mean AQI surpassed the AQI health concern cut-offs outlined in Section 3.3.2. Finally, we look at the heterogeneity in the association between AQI on offenses.

3.5.1 Average effects

Table 3.6: Regression Results

	Dependent Variable:				
	Anti-social behaviour rate	Public order rate	Violence and sexual offences rate	Burglary rate	Robbery rate
	(1)	(2)	(3)	(4)	(5)
AQI	0.0065	0.0016	0.0031	0.0117***	0.0184**
	(0.0039)	(0.0043)	(0.0023)	(0.0025)	(0.0075)
Young Male Population (%)	9.2821***	9.0164***	5.3925***	6.4694***	19.1214***
	(1.2552)	(1.5206)	(1.2356)	(1.3850)	(2.5785)
Mean Max temperature (C)	0.0360	0.0242	−0.0137	0.0622	−0.0041
	(0.0660)	(0.0726)	(0.0585)	(0.0497)	(0.0966)
Mean Rain (mm)	−0.0004	0.0016	0.0013	0.0003	−0.0036*
	(0.0017)	(0.0024)	(0.0013)	(0.0015)	(0.0019)
Non-domestic Electricity (%)	2.4505**	3.0016*	3.3297***	−0.9243	2.7960
Consumption	(1.1787)	(1.5982)	(0.9355)	(1.4174)	(1.9072)
Police District by Year by Month FE	Y	Y	Y	Y	Y
Adjusted R^2	0.5753	0.7008	0.4261	0.4422	0.6771
Observations	21,572	15,225	21,576	21,558	19,475

Notes: All dependent variables are in logs. Standard errors in parentheses, clustered by police district
Star values: * 0.10 ** 0.05 *** 0.010.

Table 3.6 displays the results of our primary linear specification. We find statistical evidence of a linear and positive association between pollution (AQI) and crime for BUR and ROB. We find no evidence of a linear association between AQI and PUC or VSO. However, Figure 3.5 shows the point estimate and the confidence interval (CI) for the betas associated with AQI in each regression. The thicker line represents approximately a 50% CI (plus/minus one standard deviation) while the thinner line represents approximately a 95% CI (plus/minus two standard deviations). This figure indicates that there is potentially a relationship between ASB and AQI, but it is not statistically significant at conventional levels.

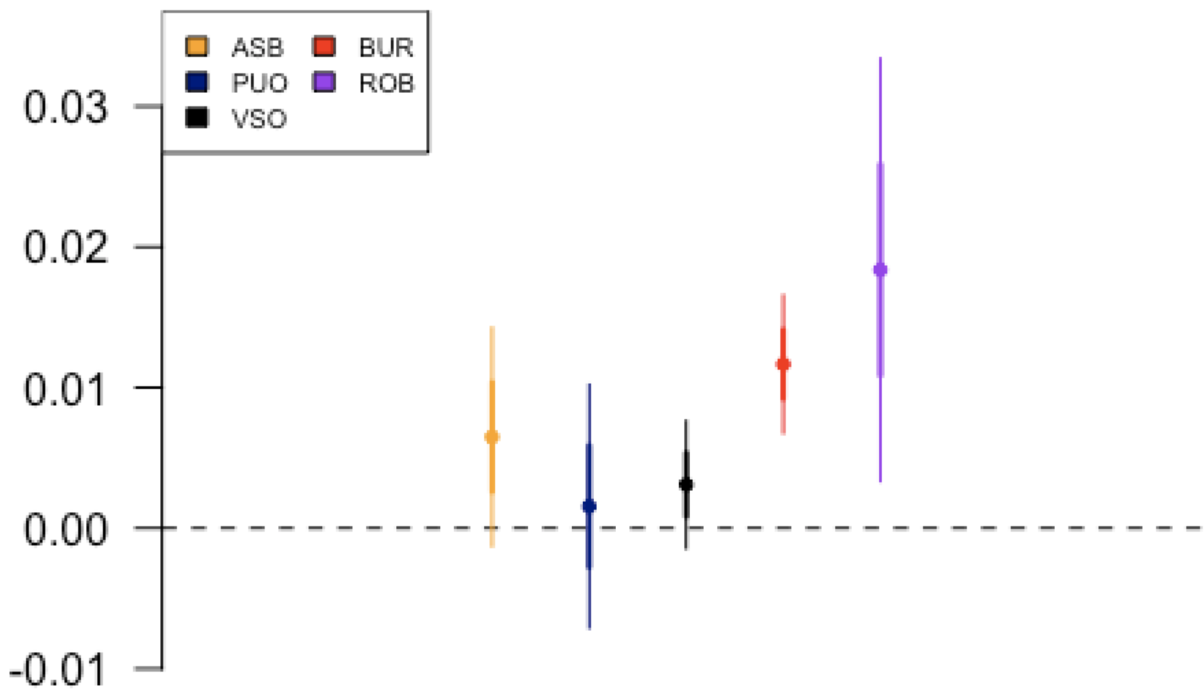

Figure 3.5: Point estimates and 50% (thicker line) and 95% (thinner line) CIs for the coefficients associated AQI as a linear variable

Assuming that the underlying association might not be linear, we choose a very flexible approach where we create quantile bins for the AQI measurement, regressing the crime rate on these categorical variable. Based on Table 3.4, we make the cuts at the 25^{th} percentile (AQI=59), median

(AQI=66) and 75^{th} percentile (AQI=77). The results of these regressions are displayed in Table 3.7 and in Figure 3.6.

Table 3.7: Regression results - Non-linear

	Dependent Variable:				
	Anti-social behaviour rate	Public order rate	Violence and sexual offences rate	Burglary rate	Robbery rate
	(1)	(2)	(3)	(4)	(5)
AQI bins(59,66]	0.0626	0.0453	0.0359	0.0765***	0.2281*
	(0.0537)	(0.0676)	(0.0302)	(0.0267)	(0.1225)
AQI bins(66,77]	0.1273	0.0781	0.0761	0.1122***	0.3459*
	(0.0771)	(0.0831)	(0.0466)	(0.0396)	(0.1926)
AQI bins(77,153]	0.1777*	0.0391	0.0826	0.1976***	0.4848**
	(0.1002)	(0.0821)	(0.0594)	(0.0537)	(0.2291)
Young Male Population (%)	9.2544***	9.0103***	5.3866***	6.4032***	19.0183***
	(1.2664)	(1.5305)	(1.2341)	(1.4079)	(2.6257)
Mean Max temperature (C)	0.0357	0.0247	−0.0140	0.0608	−0.0045
	(0.0666)	(0.0727)	(0.0581)	(0.0527)	(0.0946)
Mean Rain (mm)	−0.0004	0.0016	0.0014	0.0005	−0.0035*
	(0.0017)	(0.0025)	(0.0013)	(0.0015)	(0.0019)
Non-domestic Electricity Consumption (%)	2.4436**	2.9927*	3.3402***	−0.9264	2.7430
	(1.1552)	(1.5873)	(0.9360)	(1.4341)	(1.8670)
Police District by Year by Month FE	Y	Y	Y	Y	Y
Adjusted R^2	0.5751	0.7010	0.4261	0.4383	0.6765
Observations	21,572	15,225	21,576	21,558	19,475

Notes: All dependent variables are in logs. Standard errors in parentheses, clustered by police district
Star values: * 0.10 ** 0.05 *** 0.010.

The quantile binned results show that ASB is significantly affected by AQI levels in the uppermost bin, relative to AQI levels in the lowest bin. Figure 3.6 shows that the impact of AQI on crimes is increasing in the AQI bin for all crimes except public order. If we focus on ASB, we can observe that an AQI in the third quantile is close to being significant at the 95% level (p-value is 0.054) compared to an AQI in the first quantile, providing some stronger evidence of the association between AQI and the rate of ASB.

Our final analysis is based on the levels of health concern associated with AQI values. In this analysis we replace AQI with the percentage of days per month in which the mean or the max AQI

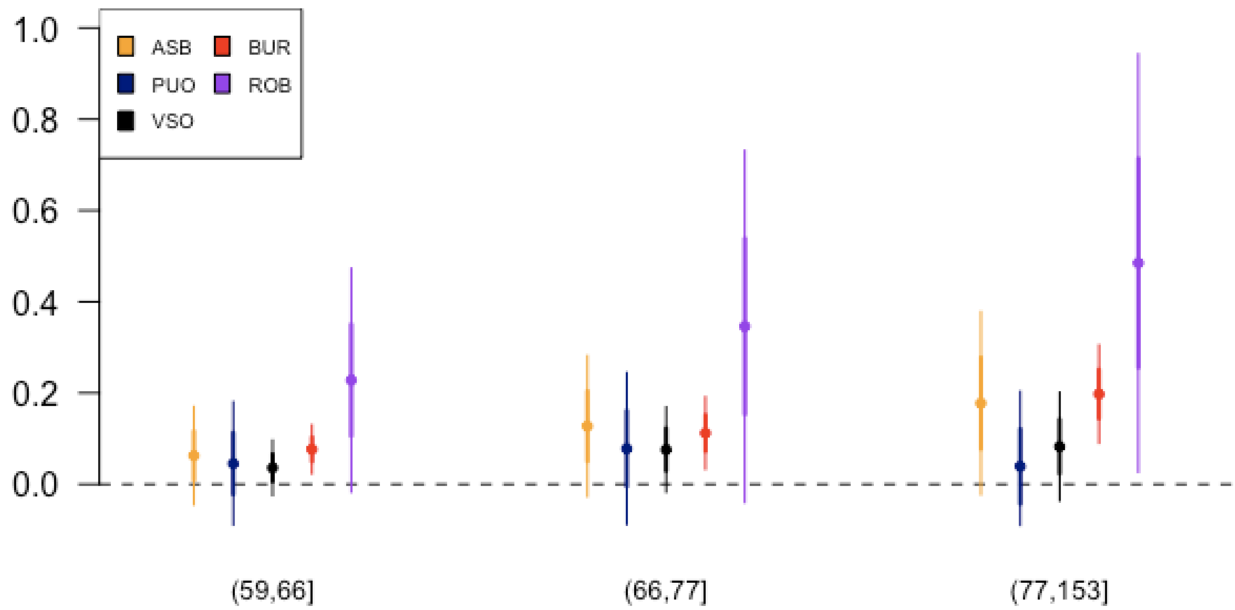

Figure 3.6: Point estimates and 50% (thicker line) and 95% (thinner line) CIs for the coefficients associated with AQI quantiles. Numbers show the range of AQIs in the 2nd, 3rd and 4th quantile

(independently) was larger than the cut-offs (50, 100, and 150). We do not test the percentage of days with AQI larger than 200 and 300 because we observe very few months where this value was more than zero. We test for $PM_{2.5}$ AQI, O3 AQI and total AQI. Note that the binding pollutant was almost exclusively $PM_{2.5}$. Results are shown in table 3.8. Each element of this table represents the coefficient and standard error associated with the pollution variable in an independent regression using all of the controls and fixed effects specified in Equation 3.4.

Results are consistent with the above estimates. The evidence of an association between ASB rates and pollution is weak, as we only find a statistically significant effect when considering the percentage of days where mean AQI (or mean $PM_{2.5}$ AQI) is larger than 50. We did not find this association with larger cutoffs, implying we begin observing effects at relatively low AQI levels. Additionally, the fraction of days above higher AQI cutoffs is relatively small, reducing the statistical power of each regression. For PUO rates, we begin observing an effect at larger mean levels of AQI and $PM_{2.5}$ AQI (100 and 150), with increasing significance but not larger magnitudes for larger cutoffs.

Table 3.8: Regression results (AQI Health Concern Cut-offs)

	Dependent Variable:				
	Anti-social behaviour rate	Public order rate	Violence and sexual offences rate	Burglary rate	Robbery rate
	(1)	(2)	(3)	(4)	(5)
% of days max AQI > 50	0.3716	1.69575	0.5209*	0.3458**	0.4238
	(0.3422)	(1.8083)	(0.2718)	(0.1614)	(0.7912)
% of days max AQI > 100	0.4843	2.2537	0.6134	0.5169	1.2937
	(0.4510)	(2.0078)	(0.5033)	(0.3750)	(0.8591)
% of days max AQI > 150	0.8994	4.8424	0.8582	0.6393	2.0328
	(0.8966)	(3.7162)	(0.8564)	(0.7176)	(1.2764)
% of days mean AQI > 50	1.3730**	5.7562	1.3457**	0.5537**	2.5694*
	(0.6842)	(4.0536)	(0.6378)	(0.2195)	(1.3440)
% of days mean AQI > 100	1.7248	10.3333*	0.9574	0.3555	4.5254**
	(1.5188)	(6.0330)	(1.4169)	(0.5754)	(1.7077)
% of days mean AQI > 150	1.6090	9.2468**	1.8012	0.2027	6.5395***
	(1.0479)	(4.0804)	(1.7238)	(0.9080)	(1.6224)
% of days max $PM_{2.5}$ AQI > 50	0.3677	1.6723	0.5189*	0.3450**	0.4164
	(0.3400)	(1.7925)	(0.2698)	(0.1598)	(0.7860)
% of days max $PM_{2.5}$ AQI > 100	0.5073	2.4418	0.6645	0.4975	1.3550
	(0.4467)	(2.1197)	(0.4998)	(0.3809)	(0.9127)
% of days max $PM_{2.5}$ AQI > 150	0.9287	4.9203	0.8650	0.6267	2.0615
	(0.8953)	(3.7062)	(0.8564)	(0.7087)	(1.2925)
% of days mean $PM_{2.5}$ AQI > 50	1.3730**	0.0453	1.3457**	0.5537**	2.5694*
	(0.6842)	(4.0536)	(0.6378)	(0.2195)	(1.3440)
% of days mean $PM_{2.5}$ AQI > 100	1.7248	10.3333*	0.9574	0.3555	4.5254**
	(1.5188)	(6.0330)	(1.4169)	(0.5754)	(1.7077)
% of days mean $PM_{2.5}$ AQI > 150	1.6090	9.2468**	1.8012	0.2027	6.5395***
	(1.0479)	(4.0804)	(1.7238)	(0.9080)	(1.6224)
% of days max O3 AQI > 50	-1.0730	-4.4117	-1.5446**	-0.2683	-2.4360
	(0.7971)	(3.1267)	(0.6489)	(0.4187)	(1.9451)
% of days max O3 AQI > 100	-1.2962	-6.6763	-2.1177**	1.0004	-1.8062
	(1.8925)	(4.3120)	(0.9918)	(1.2211)	(3.7622)
% of days max O3 AQI > 150	-12.3600	-12.3720	-2.2075	5.3364	-9.8818
	(7.1740)	(10.4970)	(5.4745)	(8.1774)	(15.7435)
% of days mean O3 AQI > 50	-7.1241	6.0155	7.6289	5.6438	-3.9559
	(5.9750)	(6.1869)	(5.1977)	(6.3013)	(10.8736)
Police District by Year by Month FE	Y	Y	Y	Y	Y

Notes: Each element of this table represents the coefficient and standard error (in parentheses) associated to the pollution variable.
Each element is an independent regression, with specification as presented in Equation 3.4.
All dependent variables are in logs. Standard errors (in parentheses) clustered by police district.
Star values: * 0.10 ** 0.05 *** 0.010.

For VSO and BUR rates, the association is similar to what we observed for ASB, with the difference that max AQI, instead of only mean AQI, can also play a role. Finally, the evidence of an association between the ROB rates and AQI is more robust. The larger the percentage of days with mean AQI larger than each cut-off, the larger the robbery rate, with the effect and significance increasing with the cut-off. An interesting takeaway from these models is that mean daily AQI seems to play a more important role than max daily AQI, indicating that the effect is more likely a result of sustained exposure to pollution than to short-term daily peaks.

3.5.2 Heterogeneous Effects

Building on previous literature and our theoretical model, we now turn to evaluating heterogeneity in the relationship between crimes and pollution [Burkhardt et al., 2019a, Roth et al., 2020]. In particular, we interact the the % of young males in the total population and the % of non-domestic electricity consumption with the pollution variable. Are hypotheses are twofold. First, we suspect that the fraction of young males in the population will increase the association as most crimes are committed by males. Second, we hypothesize that the fraction of non-domestic electricity consumption, a proxy for industrial processes, which produce emissions and are often associated with lower income areas will also increase the association. This latter affect could work through two channels in our theoretical model. Higher rates of industrialization could lead to higher pollution levels, which could increase the discount rate, β, in Equation 3.3. Likewise, if industrial areas are relatively low income, then the reward for economically motivated crimes would be relatively higher in these areas (R^C in Equation 3.3). Table 3.9 displays the results for these regression models.

We find virtually no evidence of either variable modifying the relationship between pollution and crime. This could be due to a lack of statistical power, or a true null effect. We do find an increase in the treatment effect for VOS associated with larger % of young male population and

Table 3.9: Regression Results - Heterogeneous Impacts

	Dependent Variable:				
	Anti-social behaviour rate	Public order rate	Violence and sexual offences rate	Burglary rate	Robbery rate
	(1)	(2)	(3)	(4)	(5)
AQI	−0.0163	−0.0488	−0.0317*	0.0222	−0.0236
	(0.0219)	(0.0337)	(0.0170)	(0.0265)	(0.0365)
Young Male Population (%)	12.0521***	9.5613***	2.7030	7.4199***	18.9990***
	(1.4177)	(2.8490)	(1.8068)	(1.8576)	(5.1666)
Non-domestic Electricity Consumption (%)	0.5152	−0.6548	0.9558	−0.2124	−0.3494
	(2.4289)	(2.9557)	(1.7460)	(2.1218)	(3.6155)
Mean Max temperature (C)	0.0349	0.0250	−0.0128	0.0614	−0.0032
	(0.0656)	(0.0722)	(0.0584)	(0.0487)	(0.0965)
Mean Rain (mm)	−0.0004	0.0015	0.0013	0.0003	−0.0036*
	(0.0017)	(0.0025)	(0.0013)	(0.0015)	(0.0019)
AQI*Young Male Population (%)	−0.0404	−0.0080	0.0394*	−0.0140	0.0022
	(0.0260)	(0.0428)	(0.0205)	(0.0313)	(0.0652)
AQI*Non-domestic Electricity Consumption (%)	0.0277	0.0542	0.0337*	−0.0101	0.0445
	(0.0216)	(0.0388)	(0.0175)	(0.0267)	(0.0321)
Police District by Year by Month FE	Y	Y	Y	Y	Y
Adjusted R^2	0.5756	0.7009	0.4266	0.4422	0.6771
Observations	21,572	15,225	21,576	21,558	19,475

Notes: All dependent variables are in logs. Standard errors in parentheses, clustered by police district
Star values: * 0.10 ** 0.05 *** 0.010.

with larger % of non-domestic electricity consumption, however, the effects are nearly canceled out by the coefficient on AQI.[22]

3.6 Conclusions and Discussion

We find evidence of an association between pollution exposure and BUR and ROB rates for nearly all AQI levels. On the other hand, we find that higher levels of $PM_{2.5}$ are associated with an increase in ASB rates. Our results are consistent with Roth et al. [2020], who find that pollution affects not only violent crime, but also other crimes including those that are economically motivated.

Our theoretical model allows us to examine four different mechanisms through which pollution could affect crime, each of which could be caused by a biological process identified in the epidemiological literature. First, pollution can affect the difference between the reward and the discounted punishment ($R^C - \beta P$). Second, pollution can decrease the utility of an alternative activity (NC), or increase the perceived utility of a successful crime (R^C). Third, pollution can reduce the concavity of the utility function (risk seeking behaviour), and fourth, it can change the probability of getting caught.

Because we observe an association for some crimes and not others, it seems unlikely that the mechanism is a change in risk perceptions, ruling out the third possible mechanism. With respect to mechanism four, the probability of getting caught could increase with pollution if pollution affects the cognitive function of the offender, making it easier for police to catch the offender. Given that we and Roth et al. [2020] find a stronger association for economically motivated crimes, we believe the more likely explanation is that pollution either decreases the utility of the alternative, or increases the offender's discount rate through cognitive impairment, as Roth et al. [2020] discussed.

[22] We also interact these variables with the quantile AQI bins but find no statistical relationships. The results are reported in the Appendix.

An alternative explanation to the heterogeneous effects could be related to the clearance rate, which is effectively expressed in our model as the probability of getting caught. If the heterogeneity in the effects was correlated with the clearance rate associated with each offense type, the mechanism could rely on both the discount rate associated with the future punishment and the probability of getting caught. According to studies in the US, the clearance rate is larger for violent crimes than it is for property crimes [Baughman, 2020, FBI, 2010]. As we can then expect similar results in the UK, we believe the clearance rate for the considered offenses should be as follows, in decreasing order: VSO, BUR (property crimes where threats or violence is used), ROB, PUO (violent crime where the victim is not an individual) and ASB. Looking at our results, neither the magnitude or significance of the estimated effects follow this order. Therefore, we do not believe clearance rates associated with each crime type explain the heterogeneous effects we observe, dismissing this possible mechanism.

In contrast, we find virtually no evidence that the effect is moderated by several key characteristics including the fraction of young males in the population and the fraction of the non-domestic energy consumption. Furthermore, when we modelled the percentage of days a month mean or maximum AQI exceeded certain thresholds, we concluded that daily mean AQI seems to play a more important role than the daily maximum AQI, indicating that the effect is likely due to sustained exposure to pollution more than to daily peaks.

In summary, our work combined with the work of Roth et al. [2020] and compared to previous work in the US, indicates the relationship between pollution and crime differs by region. Whereas we find that pollution tends to affect economically motivated crimes in the UK, Burkhardt et al. [2019b] and Burkhardt et al. [2019a] find that pollution largely impacts spontaneous and aggressive crimes in the US. We hypothesized that if increased aggression is the key mechanism behind the crime and pollution relationship, then we should find a similar relationship between pollution and a lesser form of aggression, namely ASB. It turns out that ASB is only weakly associated with high levels of pollution in the UK, providing evidence that the UK based relationship is not driven by cognitive impairment leading to increased aggression. On the other hand, we replicate the results of Roth et al. [2020] showing that pollution does impact economically motivated crime rates, and this is likely due to a change in the offender's discount rate or a increase in the utility of committing a crime and/or a decrease in the utility of alternative activity.

Finally, our results can inform policy makers tasked with managing ASB levels in the UK, as we have identified its association with pollution which was previously overlooked. By replicating Roth et al. [2020], we also provide evidence that crime.uk data, even with its temporal aggregation and spatial jittering, can serve to answer questions aimed at understanding crime patterns.

References

A. Adhvaryu. Learning, misallocation, and technology adoption: evidence from new malaria therapy in tanzania. *The Review of economic studies*, 81(4):1331–1365, 2014.

H. Allcott. Social norms and energy conservation. *Journal of Public Economics*, 95(9-10):1082–1095, 2011. URL `https://EconPapers.repec.org/RePEc:eee:pubeco:v:95:y:2011:i:9-10:p:1082-1095`.

H. Allcott. Site selection bias in program evaluation. *The Quarterly Journal of Economics*, 130 (3):1117–1165, 2015.

H. Allcott and J. B. Kessler. The welfare effects of nudges: A case study of energy use social comparisons. *American Economic Journal: Applied Economics*, 11(1):236–76, 2019.

H. Allcott and T. Rogers. The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review*, 104(10):3003–37, 2014.

J. D. Angrist and J.-S. Pischke. *Mostly harmless econometrics: An empiricist's companion*. Princeton university press, 2008.

J. Ansar and R. Sparks. The experience curve, option value, and the energy paradox. *Energy Policy*, 37(3):1012–1020, 2009.

F. Arbués, M. Á. Garcıa-Valiñas, and R. Martınez-Espiñeira. Estimation of residential water demand: a state-of-the-art review. *The Journal of Socio-Economics*, 32(1):81–102, 2003.

T. H. Arimura, H. Katayama, and M. Sakudo. Do social norms matter to energy-saving behavior? endogenous social and correlated effects. *Journal of the Association of Environmental and Resource Economists*, 3(3):525–553, 2016.

A. V. Banerjee, E. Duflo, and M. Kremer. The influence of randomized controlled trials on development economics research and on development policy. *The State of Economics, The State of the World*, 2016.

S. B. Baughman. How effective are police? the problem of clearance rates and criminal accountability. *Alabama Law Review, Forthcoming*, 2020.

G. S. Becker. Crime and punishment: An economic approach. In *The economic dimensions of crime*, pages 13–68. Springer, 1968.

A. BenYishay and A. M. Mobarak. Social Learning and Incentives for Experimentation and Communication. *The Review of Economic Studies*, 86(3):976–1009, 07 2018. ISSN 0034-6527. 10.1093/restud/rdy039. URL `https://doi.org/10.1093/restud/rdy039`.

D. A. Brent, J. H. Cook, and S. Olsen. Social comparisons, household water use, and participation in utility conservation programs: Evidence from three randomized trials. *Journal of the Association of Environmental and Resource Economists*, 2(4):597–627, 2015.

D. A. Brent, C. Lott, M. Taylor, J. Cook, K. Rollins, and S. Stoddard. What causes heterogeneous responses to social comparison messages for water conservation? *Environmental and Resource Economics*, 77(3):503–537, 2020.

R. D. Brook, B. Franklin, W. Cascio, Y. Hong, G. Howard, M. Lipsett, R. Luepker, M. Mittleman, J. Samet, S. C. Smith Jr, et al. Air pollution and cardiovascular disease: a statement for healthcare professionals from the expert panel on population and prevention science of the american heart association. *Circulation*, 109(21):2655–2671, 2004.

J. P. Brown, D. M. Lambert, and T. R. Wojan. The effect of the conservation reserve program on rural economies: deriving a statistical verdict from a null finding. *American Journal of Agricultural Economics*, 101(2):528–540, 2018.

J. Burkhardt, J. Bayham, A. Wilson, J. D. Berman, K. O'Dell, B. Ford, E. V. Fischer, and J. R. Pierce. The relationship between monthly air pollution and violent crime across the united states. *Journal of Environmental Economics and Policy*, pages 1–18, 2019a.

J. Burkhardt, J. Bayham, A. Wilson, E. Carter, J. D. Berman, K. O'Dell, B. Ford, E. V. Fischer, and J. R. Pierce. The effect of pollution on crime: Evidence from data on particulate matter and ozone. *Journal of Environmental Economics and Management*, 98:102267, 2019b.

J. Burkhardt, K. Gillingham, and P. K. Kopalle. Experimental evidence on the effect of information and pricing on residential electricity consumption. Technical report, National Bureau of Economic Research, 2019c.

M. R. Carter, R. Laajaj, and D. Yang. Subsidies and the persistence of technology adoption: Field experimental evidence from mozambique. Technical report, National Bureau of Economic Research, 2014.

G. Charness and D. Levin. When optimal choices feel wrong: A laboratory study of bayesian updating, complexity, and affect. *American Economic Review*, 95(4):1300–1309, 2005.

D. R. Cox. Planning of experiments. 1958.

C. Cunningham, S. Campion, K. Lunnon, C. L. Murray, J. F. Woods, R. M. Deacon, J. N. P. Rawlins, and V. H. Perry. Systemic inflammation induces acute behavioral and cognitive changes and accelerates neurodegenerative disease. *Biological psychiatry*, 65(4):304–312, 2009.

A. Deaton. Instruments, randomization, and learning about development. *Journal of economic literature*, 48(2):424–55, 2010.

A. Deaton and N. Cartwright. Understanding and misunderstanding randomized controlled trials. *Social Science & Medicine*, 2017.

M. A. Delmas, M. Fischlein, and O. I. Asensio. Information strategies and energy conservation behavior: A meta-analysis of experimental studies from 1975 to 2012. *Energy Policy*, 61:729–739, 2013.

T. Dinkelman. The effects of rural electrification on employment: New evidence from south africa. *American Economic Review*, 101(7):3078–3108, 2011.

R. K. Dixit, A. K. Dixit, and R. S. Pindyck. *Investment under uncertainty*. Princeton university press, 1994.

K. Donaldson, V. Stone, A. Seaton, and W. MacNee. Ambient particle inhalation and the cardiovascular system: potential mechanisms. *Environmental health perspectives*, 109(suppl 4): 523–527, 2001.

J. M. Eder, C. F. Mutsaerts, and P. Sriwannawit. Mini-grids and renewable energy in rural africa: How diffusion theory explains adoption of electricity in uganda. *Energy Research & Social Science*, 5:45–54, 2015.

EPA. Guidance for network design and optimum site exposure for pm2.5 and pm10. 1997.

FBI. Percent of crimes cleared by arrest, or exceptional means, in 2010. 2010.

P. J. Ferraro and M. M. Hanauer. Advances in measuring the environmental and social impacts of environmental programs. *Annual review of environment and resources*, 39:495–517, 2014.

P. J. Ferraro and M. K. Price. Using nonpecuniary strategies to influence behavior: evidence from a large-scale field experiment. *Review of Economics and Statistics*, 95(1):64–73, 2013.

P. J. Ferraro, J. J. Miranda, and M. K. Price. The persistence of treatment effects with norm-based policy instruments: evidence from a randomized environmental policy experiment. *American Economic Review*, 101(3):318–22, 2011.

A. Fischer, A. Arnold, and M. Gibbs. Information and the speed of innovation adoption. *American Journal of Agricultural Economics*, 78(4):1073–1081, 1996.

R. A. Fisher. Theory of statistical estimation. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 22, pages 700–725. Cambridge University Press, 1925.

D. A. Freedman. Statistical models for causation: what inferential leverage do they provide? *Evaluation review*, 30(6):691–713, 2006.

D. Fudenberg and J. Tirole. Perfect bayesian equilibrium and sequential equilibrium. *journal of Economic Theory*, 53(2):236–260, 1991.

H. D. Gibson. *Taming a Wicked Problem: Energy Access Planning From an Energy-Poor Perspective*. PhD thesis, 2017.

U. Government. Punishments for antisocial behaviour. `https://www.gov.uk/civil-injunctions-criminal-behaviour-orders`, 2020.

M. Grimm, L. Lenz, J. Peters, and M. Sievert. Demand for off-grid solar electricity–experimental evidence from rwanda. 2019.

B. H. Hall and B. Khan. Adoption of new technology. Technical report, National bureau of economic research, 2003.

R. Hanna, E. Duflo, and M. Greenstone. Up in smoke: the influence of household behavior on the long-run impact of improved cooking stoves. *American Economic Journal: Economic Policy*, 8 (1):80–114, 2016.

M. L. Henry, P. J. Ferraro, and A. Kontoleon. The behavioural effect of electronic home energy reports: Evidence from a randomised field trial in the united states. *Energy Policy*, 132:1256–1261, 2019.

E. Herrnstadt, A. Heyes, E. Muehlegger, and S. Saberian. Air pollution and criminal activity: Microgeographic evidence from chicago. *American Economic Journal: Applied Economics*, 2020.

B. K. Jack, P. Oliva, C. Severen, E. Walker, and S. Bell. Technology adoption under uncertainty: Take-up and subsequent investment in zambia. Technical report, National Bureau of Economic Research, 2015.

K. Jessoe, G. E. Lade, F. Loge, and E. Spang. Spillovers from behavioral interventions: Experimental evidence from water and energy use. *Working Paper*, 2017.

D. W. Jorgenson. Innovation and productivity growth. *American Journal of Agricultural Economics*, 93(2):276–296, 2011.

D. S. Kenney, C. Goemans, R. Klein, J. Lowrey, and K. Reidy. Residential water demand management: lessons from aurora, colorado 1. *JAWRA Journal of the American Water Resources Association*, 44(1):192–207, 2008.

M.-A. Kioumourtzoglou, M. C. Power, J. E. Hart, O. I. Okereke, B. A. Coull, F. Laden, and M. G. Weisskopf. The association between air pollution and onset of depression among middle-aged and older women. *American journal of epidemiology*, 185(9):801–809, 2017.

Lee, Miguel, and Wolfram. Energy institute blog, Mar 2018. URL `https://energyathaas.wordpress.com/2018/03/12/does-solving-energy-poverty-help-solve-poverty-not-quite/`.

D. S. Lee and T. Lemieux. Regression discontinuity designs in economics. *Journal of economic literature*, 48(2):281–355, 2010.

K. Lee, E. Miguel, and C. Wolfram. Experimental evidence on the demand for and costs of rural electrification. Technical report, National Bureau of Economic Research, 2016.

W. Lin, D. L. Ortega, V. Caputo, and J. L. Lusk. Personality traits and consumer acceptance of controversial food technology: A cross-country investigation of genetically modified animal products. *Food Quality and Preference*, 76:10–19, 2019.

R. K. Lindner and M. Gibbs. A test of bayesian learning from farmer trials of new wheat varieties. *Australian Journal of Agricultural Economics*, 34(1):21–38, 1990.

J. G. Lu, J. J. Lee, F. Gino, and A. D. Galinsky. Polluted morality: Air pollution predicts criminal activity and unethical behavior. *Psychological science*, 29(3):340–355, 2018.

P. Maniloff. Can learning explain deterrence? evidence from oil and gas production. *Journal of the Association of Environmental and Resource Economists*, 6(5):853–881, 2019.

T. H. Meles. Impact of power outages on households in developing countries: Evidence from ethiopia. *Energy Economics*, 91:104882, 2020.

C. Mini, T. Hogue, and S. Pincetl. The effectiveness of water conservation measures on summer residential water use in los angeles, california. *Resources, Conservation and Recycling*, 94: 136–145, 2015.

A. M. Mobarak, P. Dwivedi, R. Bailis, L. Hildemann, and G. Miller. Low demand for nontraditional cookstove technologies. *Proceedings of the National Academy of Sciences*, 109(27): 10815–10820, 2012.

H. Müggenburg, A. Tillmans, P. Schweizer-Ries, T. Raabe, and P. Adelmann. Social acceptance of picopv systems as a means of rural electrification—a socio-technical case study in ethiopia. *Energy for Sustainable Development*, 16(1):90–97, 2012.

K. Munshi. Social learning in a heterogeneous population: technology diffusion in the indian green revolution. *Journal of development Economics*, 73(1):185–213, 2004.

E. Myers and M. Souza. Social comparison nudges without monetary incentives: Evidence from home energy reports. *Journal of Environmental Economics and Management*, page 102315, 2020.

E. Oster and R. Thornton. Determinants of technology adoption: Peer effects in menstrual cup take-up. *Journal of the European Economic Association*, 10(6):1263–1293, 2012.

E. A. Outlook. From poverty to prosperity. *Paris: International Energy Agency*, 2017.

L. H. Palm-Forster, P. J. Ferraro, N. Janusch, C. A. Vossler, and K. D. Messer. Behavioral and experimental agri-environmental research: methodological challenges, literature gaps, and recommendations. *Environmental and resource economics*, 73(3):719–742, 2019.

D. Powell. Quantile regression with nonadditive fixed effects. *RAND Corporation*, 2016.

M. C. Power, M.-A. Kioumourtzoglou, J. E. Hart, O. I. Okereke, F. Laden, and M. G. Weisskopf. The relation between past exposure to fine particulate air pollution and prevalent anxiety: observational cohort study. *bmj*, 350:h1111, 2015.

N. Rosenberg and R. Nathan. *Inside the black box: technology and economics*. cambridge university press, 1982.

S. Roth, M. Bondy, and L. Sager. Crime is in the air: The contemporaneous relationship between air pollution and crime. 2020.

G. S. Sampson and E. D. Perry. The role of peer effects in natural resource appropriation–the case of groundwater. *American Journal of Agricultural Economics*, 101(1):154–171, 2018.

A. K. Sedai, R. Nepal, and T. Jamasb. Flickering lifelines: Electrification and household welfare in india. 2020.

C.-W. Shyu. End-users' experiences with electricity supply from stand-alone mini-grid solar pv power stations in rural areas of western china. *Energy for Sustainable Development*, 17(4): 391–400, 2013.

N. A. Streletskaya, S. D. Bell, M. Kecinski, T. Li, S. Banerjee, L. H. Palm-Forster, and D. Pannell. Agricultural adoption and behavioral economics: Bridging the gap. *Applied Economic Perspectives and Policy*, 42(1):54–66, 2020.

L. Tompson, S. Johnson, M. Ashby, C. Perkins, and P. Edwards. Uk open source crime data: accuracy and possibilities for research. *Cartography and geographic information science*, 42 (2):97–111, 2015.

M. M. J. Torres and F. Carlsson. Direct and spillover effects of a social information campaign on residential water-savings. *Journal of Environmental Economics and Management*, 92:222–243, 2018.

UN. Sustainable development goals. `https://www.un.org/sustainabledevelopment/energy/`, 2018. Accessed: 2019-04-30.

US-AID. Rwanda, power africa fact sheet. `https://www.usaid.gov/powerafrica/rwanda`, 2019. Accessed: 2019-04-30.

J. A. Winfree and J. J. McCluskey. Collective reputation and quality. *american Journal of agricultural Economics*, 87(1):206–213, 2005.

D. J. Zizzo, S. Stolarz-Fantino, J. Wen, and E. Fantino. A violation of the monotonicity axiom: Experimental evidence on the conjunction fallacy. *Journal of Economic Behavior & Organization*, 41(3):263–276, 2000.

Appendix

A Bayesian Learning

Among the ways an agent can learn from new information, literature has looked at the Bayesian update, which is a specific way of updating prior beliefs given new information. In our theoretical model, we assume technology quality in village c in period t is a random draw from a distribution with mean q_c (unknown to the consumers) and variance σ_c^2. For the Bayesian update model, we assume the random draw comes from a normal distribution for which the consumer has prior beliefs on the mean reliability for village c (q_c) and knows the exact value for the variance σ_c^2. The agent then updates his beliefs on q_c after every draw he observes or experiences from the random process. Therefore, in month t, the consumer h assumes the mean q_c follows a normal distribution with mean $\mu_{h,t}$ and variance $\sigma_{h,t}^2$.

$$q_c \sim \mathcal{N}(\mu_{h,t},\, \sigma_{h,t}^2) \tag{1}$$

Assuming that household h in month t observes or experiences $q_{c,t}$ (a random draw from q_c), the updated (posterior) beliefs follows a normal distribution with mean $\mu_{h,t+1}$ and variance $\sigma_{h,t+1}^2$. Under these assumptions, the updated mean and variance are as defined by the following updating equations:

$$\sigma_{h,t+1}^2 = \frac{\sigma_{h,t}^2 \times \sigma_c^2}{\sigma_{h,t}^2 + \sigma_c^2} \tag{2}$$

and

$$\mu_{h,t+1} = \frac{\sigma_c^2}{\sigma_c^2 + \sigma_{h,t}^2} \times \mu_{h,t} + \frac{\sigma_{h,t}^2}{\sigma_c^2 + \sigma_{h,t}^2} \times q_{c,t} \tag{3}$$

A.1 Testable hypothesis

To test for Bayesian Learning, we focus on modelling the probabilities to adopt and to dis-adopt and the time to reconnect, specifically testing two hypotheses derived from the theory:

1) Based on the updating rule, the larger the variance of the technology quality (σ_c^2), the less weight is given to new information ($q_{c,t}$), as in Lin et al. [2019].

2) Based on the updating rule, and on how $\sigma_{h,t+1}^2$ tends towards zero, the weight given to new information is less as t increases, as in Maniloff [2019].

B Extensive Margin: Technology Adoption and Disadoption Robustness

In this section of the appendix, we try a different specification where we replace the trend variable with the number of connected households as we could not include both variables in the same model due to multicollinearity issues. As shown in tables A1 and A2, qualitative results hold using the new specification.

Table A1: Probability of Adoption - Robustness

	(1)	(2)	(3)	(4)	(5)
	Logit	Logit	Logit	Logit	Logit
Blackout rate $(t-1)$	-1.455	-2.721*	-	-1.282*	-1.341*
	(1.087)	(1.464)	-	(0.781)	(0.794)
Blackout rate $(t-2)$	-	0.153	-	-	0.167
	-	(1.108)	-	-	(0.814)
Avg. Blackout rate $(t-1)$	-	-	0.087	-	-
	-	-	(5.576)	-	-
Avg. Blackout rate $(t-2)$	-	-	-	0.882	-
	-	-	-	(0.623)	-
Avg. Blackout rate $(t-3)$	-	-	-	-	0.841
	-	-	-	-	(0.674)
Household FE	Y	Y	Y	Y	Y
Month by year FE	Y	Y	Y	Y	Y
Clustered SE	Y	Y	Y	Y	Y
Observations	23,196	22,113	23,196	21,030	21,030

Notes:Standard errors in parentheses, clustered by Village. The unit of analysis in these regressions is the household-month.

Star values: * 0.10 ** 0.05 *** 0.010.

Table A2: Probability of Disadoption - Robustness

	(1)	(2)	(3)	(4)
	Logit	Logit	Logit	Logit
Blackout rate $(t-1)$	0.361	0.377	-	-
	(0.551)	(0.561)	-	-
Blackout rate $(t-2)$	-	0.537	-	-
	-	(0.793)	-	-
Avg. Blackout rate $(t-1)$	-	-	-2.467	-
	-	-	(2.272)	-
Avg. Blackout rate since connected $(t-1)$	-	-	-	-0.715
	-	-	-	(1.069)
Household FE	Y	Y	Y	Y
Month by year FE	Y	Y	Y	Y
Clustered SE	Y	Y	Y	Y
Observations	15,492	15,009	15,492	15,492

Notes:Standard errors in parentheses, clustered by Village. The unit of analysis in these regressions is the household-month.

Star values: * 0.10 ** 0.05 *** 0.010.

C Survey

We conducted a survey in 9 villages in Rwanda, starting on 10/01/2019 and ending on 12/17/2019, for a total of 2.5 months. We surveyed the villages of Gitaraga, Kajevuba, Bizenga, Koperative, Nyamabuye, Gihandagazi, Kavumu, Inunga and Cyabayagara. In each village, we selected 40 households: 20 of them were randomly chosen among the active MeshPower costumers in 2017, and 20 households were randomly selected among non-MeshPower customers in 2017. At the moment of surveying, we were able to survey 343 households, 123 were MeshPower customer and 220 were not. In total, given that some households where originally among the non-customer group but they were customers by the time we survey, 69 respondents were customers and chose not to be anymore. We asked respondents who used to be costumers why they left, and we asked all costumers what they think about the reliability of the service.

Out of the 69 respondents who were not customers any more, 30 stated that they left because they had access to other sources of electricity, 14 stated that they could not afford MeshPower service anymore, 5 moved to places where the service was not available while none of them stated that reliability was the reason why they left. 23 households stated "other reasons", without disclosing them. When we asked MP customers about their perception of the reliability of the MeshPower service, 43.26% answered that it was better than expected, 56.03% answered that it was as expected and only 0.71% answered it was worse than expected.

D Heterogeneous Impacts of Pollution on Crime

Table A3: Regression Results Non-linear - Heterogeneous Impacts

	Dependent Variable:				
	Anti-social behaviour rate	Public order rate	Violence and sexual offences rate	Burglary rate	Robbery rate
	(1)	(2)	(3)	(4)	(5)
AQI bins(59,66]	−0.164	−0.431	−0.693	−0.506	−0.955
	(0.540)	(0.550)	(0.522)	(0.755)	(0.938)
AQI bins(66,77]	−1.233	−1.091	−1.511*	−0.155	−2.298
	(0.808)	(0.985)	(0.868)	(1.064)	(1.479)
AQI bins(77,153]	−0.509	−1.295	−1.253*	0.641	−1.029
	(0.862)	(1.151)	(0.676)	(1.200)	(1.567)
Young Male Population (%)	9.976***	9.620***	5.644***	7.232***	19.399***
	(0.798)	(1.515)	(0.977)	(1.374)	(2.909)
Non-domestic Electricity Consumption (%)	1.759	2.215	0.770	−1.108	1.223
	(1.461)	(1.693)	(1.549)	(1.675)	(2.371)
Mean Max temperature (C)	0.035	0.024	0.009	0.060	−0.004
	(0.066)	(0.072)	(0.061)	(0.052)	(0.095)
Mean Rain (mm)	−0.0004	0.001	0.001	0.0004	−0.004*
	(0.002)	(0.003)	(0.002)	(0.001)	(0.002)
AQI bins(59,66]:Young Male Population (%)	−1.020	−1.710	−1.261*	−1.277	−0.236
	(1.142)	(1.088)	(0.735)	(0.800)	(1.934)
AQI bins(66,77]:Young Male Population (%)	−0.822	−0.036	−0.627	−1.447	−1.474
	(1.340)	(1.410)	(0.725)	(1.052)	(1.951)
AQI bins(77,153]:Young Male Population (%)	−1.180	−0.784	0.014	−0.698	0.366
	(1.084)	(1.427)	(0.825)	(1.261)	(2.479)
AQI bins(59,66]:Non-domestic Electricity Consumption (%)	0.330	0.651	0.933*	0.727	1.279
	(0.607)	(0.606)	(0.517)	(0.834)	(0.893)
AQI bins(66,77]:Non-domestic Electricity Consumption (%)	1.515*	1.244	1.808**	0.408	2.934*
	(0.876)	(1.132)	(0.880)	(1.108)	(1.508)
AQI bins(77,153]:Non-domestic Electricity Consumption (%)	0.832	1.484	1.512**	−0.413	1.577
	(0.895)	(1.312)	(0.669)	(1.253)	(1.404)
Police District by Year by Month FE	Y	Y	Y	Y	Y
Adjusted R^2	0.576	0.701	0.505	0.443	0.678
Observations	21,572	15,225	21,576	21,558	19,475

Notes: All dependent variables are in logs. Standard errors in parentheses, clustered by police district
Star values: * 0.10 ** 0.05 *** 0.010.