Paying for urban services: utility bills, prepaid metering and spending patterns of the poor
Pre-analysis plan
August 24, 2018

PIs: B. Kelsey Jack, Kathryn McDermott, Anja Sautmann

1. Introduction

This pre-analysis plan provides an overview of the study design, data collection, hypotheses and analytical framework for *Paying for urban services: utility bills, prepaid metering and spending patterns of the poor*. Details of the study are described below.

The rest of this document is structured as follows. In Section 2, we describe the study background and motivation, sample, and data sources for analysis. In Section 3, we describe the treatments and how they will be assigned, and present the hypotheses that will be tested through our design. Section 4 describes the empirical strategy and details of analysis.

2. Study background and design

2.1 Motivation

We will investigate how urban households in South Africa make decisions about electricity purchases on prepaid electricity meters. The purpose is to study why prepaid payment affects the electricity utilization patterns of poor households (relative to monthly billing).

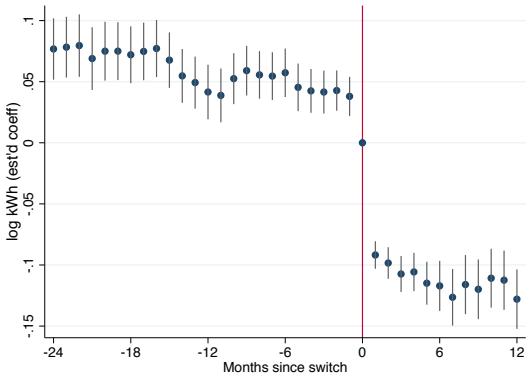

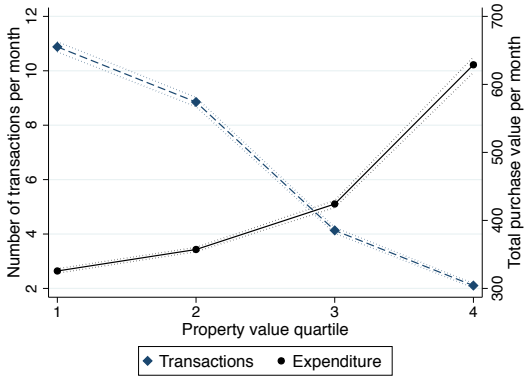

The research is motivated by two empirical facts observed in the study setting in Cape Town. First, as shown in the left panel of figure 1, electricity use falls by around 12 percent

when households are switched from monthly billing to prepaid metering.[1] Second, as shown in the right panel of figure 1, low-income customers on prepaid metering purchase electricity in small quantities and at very high frequencies (up to every 3 days), reminiscent of the purchasing patterns of poor consumers in other domains.[2]

These patterns may reflect liquidity and other constraints of poor households, or deliberate, unconstrained choices - with very different welfare implications: if poor households have to walk to a sales point and wait in line to buy electricity every 3 days because they cannot buy larger amounts at one time, incurring transaction costs and potentially a frequent risk of running out of electricity when they miscalculate their consumption, switching from billing to prepaid-metering imposes an additional burden on them. By contrast, if they choose to buy every 3 days in order to control their own electricity consumption in favor of savings or goods other than electricity, prepaid metering may improve their lives overall. Understanding the welfare consequences of different ways of paying for services such as electricity is an important step toward designing financial products that improve consumer welfare while also improving revenue recovery for service providers.

2.2 Data sources

The study will rely on two main data sources:

1) Administrative records from the City of Cape Town obtained through an existing non-disclosure agreement.
Two types of data from administrative records will be used:

a) Vending records of prepaid electricity purchases. Each electricity purchase is recorded in the vending data, and includes the monetary value spent, the kWh purchased, the date and location of the purchase, the customer tariff, and other details.
b) Data on property value, GIS location, and water consumption data of the households, which will be used for sampling and analysis.

2) Primary data collection through household surveys.

Households in the study will be surveyed twice:

a) Survey round 1 – A household representative involved in electricity purchase decisions will be recruited and surveyed. The survey will collect basic information about household demographics, income and spending flows, ownership of electricity-powered durables, and electricity purchasing patterns.

---

[1] Jack, B Kelsey and Grant Smith. 2016. "Charging Ahead: Prepaid Electricity Metering in South Africa" *NBER Working Paper* No. 22895.
[2] Jack, B Kelsey and Grant Smith. 2015. "Pay as You Go: Prepaid Metering and Electricity Expenditures in South Africa." *American Economic Review*, 105(5): 237-41.

We discuss how all variables will be used in the analysis in Section 4. A first round of treatments will be administered during survey round 1.

b) Survey round 2 – The same household (and, ideally, the same respondent) as in round 1 will be surveyed a second time, between one and two months after the first survey round. We discuss how we will handle attrition in Section 4. The survey will collect information about any shocks or changes to household composition or income since the round 1 survey. In addition, the second survey will administer a follow-up series of experimental treatments.

In both surveys, we will obtain information about electricity use from the household's electricity meter. Specifically, surveyors will enter codes into the meter to retrieve total consumption and/or total uploads and the meter balance. This information will be used to estimate round 1 treatment effects on electricity usage in the interval between survey rounds. We will also use the meter codes to retrieve information on recent zero balance events experienced by the household, where possible.

2.3 Outcomes of interest

Our main outcomes of interest are electricity purchases, obtained from the administrative data, and usage, obtained from customer meters during the survey, and preferences over transfers, obtained during the second survey round. We briefly discuss how we construct each.

Purchase data are recorded at the transaction level. For some analysis, these will be converted into (a) a daily panel and (b) and monthly panel. Households that do not make a purchase in the relevant window are assigned a zero for purchase quantities, resulting in a balanced panel over the window of time that a customer is supplied electricity by the City of Cape Town on a prepaid electricity meter.

Usage data will be obtained from the household's meter. During the first survey round, we will observe the total amount the household has ever loaded on to the meter and/or the total amount they have ever consumed through the meter, along with the meter balance. During the second survey round, we will observe the same. We will calculate usage in between survey rounds using the difference in the total lifetime uploads and/or the difference in the total lifetime consumption, along with the meter balance. The net of these is consumption between the first and second survey round.

2.4 Sample

The study population will be sampled from the City of Cape Town's records of all prepaid electricity customers. Using these administrative data, we will identify a sample of low and middle income urban households according to the following inclusion criteria:

- Purchased electricity four times or more in at least two different months between March 2018 and June 2018
- Purchased at least once in the first half of July 2018
- Are on either lifeline or domestic electricity tariff in July 2018

- We are able to match the household to 2015 GIS data (account is at least 3 years old)
- Electricity meter installed before June 2017
- No debt payments via the prepaid meter in July 2018 (household may have other debt with the city that is not paid through the electricity vending system)
- One of four different types of meter from which total consumption to date / total kWh loaded to date can be retrieved using short codes.

Eligible meters will be grouped by geographic neighborhood. Neighborhoods will be eliminated if safety concerns are serious, <10% of or < 100 meters in the neighborhood are eligible. Each neighborhood will be further divided into sub-areas for surveying purposes.

Surveyors will check the following screening conditions to determine eligibility:

You are 18 years or older.
You have access to the electricity meter on this property.
You regularly buy or are responsible for paying for electricity for this household.
You have a working cell phone on you that belongs to you and you know the phone number for and are willing to share the number with me.

The final sample will be determined by a combination of ex ante sampling using administrative data, availability for screening, and outcomes of the screening questions.

Our target sample is 800 households, excluding the pure control. The final sample will be determined by field logistics.

3. Treatments and hypotheses

3.1 Treatments

Round 1 treatment – along with the first survey round, households will be randomly assigned to one of the following four groups.

1) 100 Rand in cash handed over at the end of the survey.
2) One electricity token with the same value as the cash transfer. The token will be uploaded on to the household meter by the subject, with assistance from the survey enumerator.
3) Two electricity tokens with the same total value as the cash transfer. One of these tokens will be provided to the household at the time of the survey and uploaded by the subject onto the household meter with assistance by the survey enumerator. The second token will be sent by text message approximately three days later.
4) Survey-only control.

All tokens are sent to the subject directly by text message. All households in the first survey will receive R20 of electricity as compensation for participating in the survey. This amount will also be uploaded onto the household meter by the subject in the

presence of the survey enumerator. Thus, survey-only households also upload a token with the surveyor.

Round 2 choices – Households will be partially randomly assigned (see below) to receive two of the following *choice sets* as part of the second survey.

a) Cash versus one electricity token
b) Cash versus two electricity tokens (with one token sent 3 days later)
c) One versus two electricity tokens (with one token sent 3 days later)

Round 2 choice sets will be presented as multiple price lists. Respondents are asked to make a series of choices between two possible options in each choice set. The highest possible value that can be obtained remains constant, while the value of the other option increases or decreases. We randomized both the order within each choice (left or right choice on the screen) and the order from first to last within each choice set. One choice within one of the choice sets is drawn for implementation. The choice sets are depicted in the Appendix.

In addition, a pure control group will be followed in the administrative data throughout the study. The pure control consists of all eligible households according to the first stage sampling rules (excluding eligibility criteria for the survey respondent), who were never selected for surveying (anyone selected for surveying, regardless if the survey took place or not, is excluded from the pure control). These households will help identify time-fixed effects and effects of the survey on the survey-only control group.

3.2 Treatment assignment

Randomization will be stratified by survey team and geographic area. Surveys will be collected using handheld devices. Surveyors are not aware of a household's treatment at the start of the survey and will not be able to manipulate treatment assignment.

In the second round, treatments will be based on round 1 treatment assignment. Specifically, households will be assigned to two out of the four possible round 2 choice conditions. Following the numbering above, assignment will be as follows:

- Round 1 treatment 1 will receive round 2 choices c and randomly drawn a or b
- Round 1 treatment 2 will receive round 2 choices b and randomly drawn a or c
- Round 1 treatment 3 will receive round 2 choices a and randomly drawn b or c
- Round 1 control (4) will receive a random set of two choices

In other words, respondents who received a round 1 treatment will always see the choice set involving the two transfers that they did not experience in round 1, plus one of the two choice sets involving the transfer they did experience in round 1. The sample of households will be assigned evenly across the four round 1 groups.

3.3 Predictions

We will use the round 1 treatments to test for treatment effects on electricity purchase outcomes. We will use choices under the round 2 treatments as a revealed preference measure of preferences. Our study is designed to test the following models of behavior, each of which is associated with particular predictions. The predictions build on the idea that any household that is unconstrained will avoid transaction costs as much as possible. A household that is liquidity constrained will accept paying transaction costs in exchange for liquidity. A self-control constrained household will seek out transaction costs in order to reduce temptation/provide commitment.

Null: transaction costs are small and do not affect consumption, and households are not liquidity constrained and have no self-control issues. The correlation of transfer frequency and size with wealth as well as the reduction in consumption occur for other reasons, e.g. comfort with small vs. large numbers and frequency of shopping, and greater awareness of prices. Household consumption and preferences are unaffected by the difference in transfer delivery, i.e. smaller, high-frequency purchases or larger, lower frequency purchases.

- Round 1 treatment effects: electricity consumption behavior will be unaffected by the delivery difference in round 1 treatments. Cash and two types of electricity transfers have the same effect on overall consumption (as long as the household adjusts electricity purchasing at least once in response to the cash transfer).
- Round 2 choices: respondents will always choose the highest value option in their choice set. If households are not liquidity constrained, but experience transaction costs, they will prefer electricity over cash to reduce their transaction costs.

A1: constrained neoclassical model. Households are liquidity constrained; high-frequency purchases reflect a high cost of locking up liquidity on the electricity meter (balanced against the transaction cost of higher frequency purchases) and consumption is reduced due to resulting high transaction costs (note that otherwise we would just see "instant purchasing" whenever electricity is needed).

- Round 1 treatment effects:
    - transfers in electricity (treatments 2 and 3) will lead to greater electricity use than transfers in cash (treatment 1), both because of lower transaction costs and because of reduced flexibility in expenditures.
    - the size and timing of the electricity transfer does not affect electricity use (treatments 2 versus 3).
- Round 2 choices:
    - respondents will have a *positive* willingness to pay for cash over electricity (choices a and b)
    - respondents will always choose the highest value option when choosing among different electricity transfers (choice c), and weakly prefer one large transfer over multiple smaller transfers.

A2: Households use small purchases as a tool for constraining household-level consumption.

- Round 1 treatment effects:

    - transfers in electricity (treatments 2 and 3) will lead to greater electricity use than transfers in cash (treatment 1), both because of lower transaction costs and because of control issues
    - smaller and more spread out transfers of electricity will lead to a smaller increase in use than a single large transfer (treatments 2 versus 3)
- Round 2 choices:
    - respondents may be willing to pay for cash or electricity (choices a and b)
    - respondents will have a positive willingness to pay for smaller and more spread out transfers (choice c)

4. Empirical strategy

In this pre-analysis plan, we outline our reduced form tests for measuring treatment effects and testing our main predictions. A structural model will help quantify the welfare implications of these results.

4.1 Estimation

Round 1 treatment effects

We start by describing analysis using the administrative dataset to identify treatment effects in a panel data setting. We will test for treatment effects of the round 1 treatments on electricity purchase outcomes. We estimate impacts using a panel specification with household and time period fixed effects, and using cross sectional analysis. We cluster standard errors at the unit of randomization, the household.

Event-time analysis:
We will analyze two outcomes and two time intervals; purchase amounts and frequency at the month (30 days) and week (7 days) level, measured from the survey date. We will use household and time fixed effects. To investigate dynamics, we will estimate separate treatment effects by week or month since the survey.

Calendar-time analysis:
We will construct week of the year and month of the year measures of electricity purchases (frequency and total amount) and use month or week fixed effects, and household fixed effects in our analysis. For the survey month or week, treatment will be a continuous variable corresponding the fraction of the time period that follows the survey.

We will allow for differential trends and surveyor effects by interacting time fixed effects with geographic strata indicators (geographic strata are nested within survey team strata; we cannot distinguish individual surveyors as teams work down a list together). If these interactions do not explain a significant share of the variation in the data, we will continue without these indicators.

Cross sectional analysis:

We will analyze effects on consumption patterns in the cross section (recovered from meter data, as described in Section 2). All cross-sectional estimates will include controls from survey round 1 to improve precision. Specifically, we will control for household size, household structure (split into main household and backyard dwellers), past mean monthly energy purchases, past purchase frequency per month, income variables (levels and timing/predictability, coded categorically into independent cells by each), respondent education, as well as geographic controls and survey-team fixed effects.

Heterogeneous treatment effects:
We will examine heterogeneity in the response to treatment based on round 1 characteristics associated with either the neoclassical model or the alternative model described above.

We will proceed in three steps:

First, we will examine heterogeneity by pre-survey frequency of purchasing. We will split the sample in high-frequency and low-frequency purchasing households and study if consumption differences between cash and electricity transfers or between the two types of electricity transfers are larger for high-frequency households. Our model predicts that high-frequency households are the most constrained, either by credit constraints, or by self-control problems.

Second, we will construct measures of liquidity constraints, self-control issues, and intrahousehold coordination issues, as well as level of transaction costs. We will use proxies that exhibit sufficient variation in the round 1 survey data.

Proxies will be selected or combined into an index from the following data.

- Number of household members, their contribution to electricity purchases, the presence of backyard dwellers who consume electricity from the same meter, and presence of "temptation goods" (TV, laptop, tablet, etc.) indicate intra-household conflict and potential temptation levels.
- Monthly or weekly income and expenses serve as a measure of how liquidity constrained the household is likely to be. The type of income and expenses (and its variability) indicates how likely the household is going to be affected by shocks that require liquidity. Variability will additionally be obtained from the income and spending measures from both survey rounds.
- Information on households running out of electricity entirely in the past week.
- Information on where the household buys electricity and if the vendor charges for buying and other transaction-cost relevant information
- The reasons that households state for running out of electricity (assumption is that this incurs a high transaction cost). Answers about reasons such as

It was too late to go and buy
It was someone else's turn to buy
To prevent someone else from using more electricity
Didn't have money

  indicate transaction costs and types of constraints.

- Safety is another indicator of transaction costs:

Suppose you run out of electricity at 9:30PM tonight. Would you walk to your closest shop to buy some electricity?
Would you buy electricity in another way? If so, how?

- Questions such as

Do you ever let the meter run out of electricity so that you yourself will not be able to use any more electricity?
Do you ever let the meter run out so that others connected to the meter will not be able to use more electricity?
Do people in your household use more electricity than you think they should?
What exactly are they doing to use so much electricity?

  are used to determine the level of self-control or household-control problems.

- The following questions are a measure of liquidity constraints:

If there was an emergency today - something unforeseen or unexpected - for example, your car broke down or you needed to buy some medication for a family member-, what is the largest amount you would be able to spend on this emergency?
If this emergency cost R1,000, how many days would it take to get the money together?
A friend or family member asked for R100 to pay for transport today so that they could go for a job interview. If you wanted to give them the cash, do you have enough cash available?

Third, using these proxies, we will examine if the consumption and purchasing response to different types of transfers differs for individuals with high transaction costs, self-control issues, and liquidity constraints, by interacting the constructed proxies with round 1 treatment indicators.

Hawthorne effects and other confounds:
Finally, we will analyze the effect of being surveyed (Hawthorne effects) and receiving the small compensatory transfer for survey participation by randomly assigning a synthetic survey date to the pure control group, and examining differences with the sample control group.

Round 2 choices

We will test model predictions on preferences using observed round 2 choices. From our multiple price list experiments, we will estimate willingness to pay between any two choices A and B, $WTP_{AB,}$ for our sample (see below). We will observe two WTP realizations per subject. We will test whether willingness to pay for smaller vs. larger tokens is positive, and whether WTP for cash is positive, controlling for round 1 "habituation" effects (i.e. if subjects have a preference for the option they received previously). If round-1 consumption changes are correlated with proxies for transaction costs and for liquidity or self-control, we will further test if willingness to pay for cash and small tokens is affected by the same proxies (controlling for round 1 treatment assignment). Lastly, we will test

(a) if willingness to pay for cash vs. electricity is stronger for households with *higher* frequency of purchasing, as predicted by liquidity constraints; and
(b) if willingness to pay for small tokens vs. one large token is stronger for households who used *significantly more* electricity when receiving one rather than two tokens in round 1 of the survey.

We will take two different approaches to estimating willingness to pay. First, we will simply impute WTP as the mid-point of its interval, which is uniquely determined by those multiple price list choices that are internally consistent. Second, we will use a binary-choice approach where the multiple price list choice is predicted by the option that has higher value to the decision maker plus an error term (see Dean & Sautmann 2018). This is possible even if choices in the multiple price list are internally not consistent, but requires some assumptions on how the respondent makes those binary choices. We will control for default bias with a fixed utility value for choosing in round 2 the randomly assigned transfer from round 1.

4.2 Selection and attrition

Selection into our sample will occur at the first survey and attrition from our sample will occur at the second survey.

Selection (round 1)

Our study makes no claim of representativeness in sampling, as described in Section 2, so the primary concern with selection is internal validity. We will therefore estimate whether selection is correlated with treatment by regressing an indicator for participation in the first survey (conditional on being in the study sample) on treatment indicators. Because surveys are administered electronically, and treatments are chosen after a respondent has agreed to participate in the survey, this type of selection is unlikely. We will also include the interaction of the same set of covariates we use in our outcome analysis with the treatment indicators to check if there was differential selection.

Attrition (round 2)

A bigger concern is that attrition from the study is correlated with round 1 treatment. It is to be expected that control households might leave the survey at a slightly higher rate than any of the three treatment arms. The main concern is a correlation with assignment into one of the three treatment arms. In particular, it may mean that we do not obtain an internally valid estimate of the difference in electricity consumption between treatment arms (which requires round 2 meter observations). Moreover, if those who drop out are those who have the lowest preference for the specific treatment they received, this may bias willingness-to-pay estimates towards indifference between treatment arms. Given that the transfers we give are quite large, we do not expect this to be a major concern, as few households will be on the margin where the difference in delivery method matters for their decision to participate. However, we will test for this by

regressing an indicator for participation in the round 2 survey (conditional on being in the round 1 survey) on round 1 treatment. We will also test whether attrition is correlated with purchasing behavior as measured by the administrative data (from which we expect little to no attrition), as a check if the consumption response to the treatments could be plausibly correlated with attrition. If we cannot rule out biased attrition, we will estimate bounds on the impacts on usage and the willingness to pay for different treatments, using weights inverse to the rate of being included into the sample.
Lastly, in order to correct for any differential attrition effects on round 2 assignment, the randomization into the different willingness-to-pay measures will be stratified by the same geographical areas of round 1 and by round-1 treatment.
We will also include the interaction of the same set of covariates we use in our outcome analysis with the treatment indicators to check if there was differential attrition.

4.3 Missing values

Three types of missing values are relevant for our analysis.

First, households will be missing collapsed outcome data if they do not make a purchase in a specified window of time. We address this by replacing these observations with zeros.

Second, attrition will generate missing electricity usage measures, MPL responses and round 2 survey data. We discuss how we will handle attrition above.

Third, households may choose not to respond to some survey questions. These missing values may be associated with either control variables or heterogeneity variables. We will rely on a combination of approaches to address missing values: (1) we will analyze models with and without these covariates to determine sensitivity to controls and the potential bias from dropping households with missing values, (2) we will impute missing values where appropriate and (3) we will "dummy-out" missing observations with indicators for missing values.