

How Do Customers Utilize Real-Time Usage Feedback? Evidence from Singapore

Frank A. Wolak
Director, Program on Energy and Sustainable Development
Professor, Department of Economics
Stanford University
Stanford, CA 94305-6072
wolak@zia.stanford.edu

October 2014

Abstract

Real-time household-level feedback has the potential to improve the efficiency of electricity consumption. This feedback allows a household to obtain a better understanding of the relationship between different electricity consuming actions and changes in its monthly electricity bill, which should improve the household's ability to assess the costs and benefits of each additional electricity consuming action throughout the monthly billing cycle. This paper studies the impact of providing a household with real-time usage feedback on its monthly energy consumption. The Singapore Energy Market Authority (EMA) implemented an Intelligent Energy System Pilot in which over 1,700 households were provided with real-time meters and in-home display (IHD) units which provided information on each household's real-time electricity consumption. To assess the impact of the real-time meters and IHDs, the monthly consumption of these households is compared to a matched group of households that were not provided with these devices for a more than 20 month time period before and after this intervention. This data is used to estimate the average treatment effect associated with a randomly selected household having each device. We find that having a IHD unit leads to a reduction in electricity consumption of about 5% relative to the control group. This saving is equivalent to about 180 KWh annually for the average household in the sample which translates into roughly 50 Singapore dollars at the relevant retail electricity price. These results strongly support more widespread use of real-time customer feedback technologies as a cost effective strategy for meeting Singapore's residential electricity demand.

1 Introduction

Electricity demand is derived from the consumer's demand for services from electricity consuming capital equipment. A consumer demands hours of lighting, air conditioning, television viewing, and the use of other electricity-consuming appliances. However, few, if any consumers, understand how minutes of use of each of these electricity consuming appliances translates into kilowatt-hours (KWhs) of electricity use. With real-time feedback on their electricity consumption, a household can determine how use of each electricity consuming appliance translates into KWhs of electricity consumption. If the household knows the price of retail electricity, it can convert this magnitude into dollars on its monthly electricity bill.

If a household understands how much an hour of use of an electricity consuming appliance costs, it has the opportunity to become a more efficient electricity consumer in the sense that it can compare the benefit from an additional hour of use of this appliance to a more accurate estimate of the cost of this action. In order to receive this information a household must have a real-time or interval meter that reports its electricity consumption for very short time intervals throughout the day. In addition, this high-frequency electricity consumption information must be conveyed to households in an easily digestible manner.

The Singapore Energy Market Authority's (EMA) Intelligent Energy System (IES) Pilot is one such mechanism for providing this real-time information to households. Over 1,700 households were chosen to receive a real-time meter and an in-home display (IHD) unit that presents real-time information on the household's electricity consumption and the cost of this electricity consumption, as well as information on the household's cumulative electricity consumption and the total cost of this consumption at daily, weekly and monthly historical time intervals. By the above logic, customers armed with this information can become more efficient electricity consumers relative to similar households that do not have access to this information.

To investigate the validity of this hypothesis about the impact of providing real-time consumption feedback on a household's electricity consumption, a randomly selected set of more than 2,000 similar households were chosen that did not have a real-time meter or IHD installed on their premises to serve as the control group to measure the impact of this intervention on a household's electricity consumption. Data on the billing cycle-level electricity consumption was collected from both treatment and control customers before and after the intervention (the installation of real-time meters and IHDs for the customers in the treatment group) for a time interval of at least 20 months.

A difference-in-difference estimation methodology is applied to this data to compute the average treatment effect associated with the installation of an interval meter and an IHD. Across a variety of treatment and control group sample selection procedures the impact of the installation of an IHD is found to reduce a household's monthly electricity consumption by 5 percent, which amounts to an approximate annual saving for customers in our sample of 180 KWh annually. Valuing this annual KWh savings at the average retail price for 2012 of 0.279 Singapore dollars per KWh, implies an annual savings of approximately 50 Singapore dollars. Assuming a 10-year life for the real-time

meter and IHD and a discount rate of 5 percent, implies a discounted present value of savings for the life of the two devices of 385 Singapore dollars.

Extrapolating these results to all Singapore residential electricity consumers implies that rolling out real-time meters with IHDs to all Singapore households would more than pay for the cost this policy and therefore yield significant net benefits to the Singapore economy. In 2012, total household electricity sales in Singapore was 6,640 gigawatt-hours (GWh). Assuming a 5 percent reduction in this value is associated installing real-time meters and IHDs in all Singapore households, implies a reduction of 332 GWh. Valuing this at 0.279 Singapore dollars per KWh implies an annual savings of 92.6 million Singapore dollars. Assuming these annual savings last for the assumed 10-year life of the real-time meter and IHD, and applying a 5 percent discount rate to these savings yields a discounted present value of savings of more than 710 million Singapore dollars.

2 IES Pilot Treatment and Control Samples

This section describes the data collected for the treatment and control groups. The raw data is compiled from monthly billing cycle-level consumption data for each household. This data is converted to calendar-month values for each customer. The unit of analysis is a customer-month. For each billing cycle, the customer's consumption is converted to an average daily value. The average daily consumption for each calendar month is computed as day-weighted average of the average daily consumption values for all of the billing cycles that have days in the calendar month. For example, if one billing cycle has 10 days within the calendar month and a second billing cycle has the remaining 21 days of the 31-day month, the average daily consumption for the calendar month is $10/31$ times the average daily consumption for the first billing cycle plus $21/31$ times the average daily consumption for the second billing cycle in the calendar month. The raw billing cycle-level data start in June 2010 and runs through February 2013 for customers in both the treatment and control groups.

The treatment period is assumed to start with the first complete calendar month that the household's real-time meter starts functioning, which for our purposes is the date from which it started recording non-zero consumption values. Because there was a rollout period for the installation of the real-time meters for the customers in the control group, the treatment date differs by households and ranges between May and November of 2012. Although the date when the interval meter starts functioning is assumed to be observed by the above logic, the exact date when the IHD started functioning is unobserved. However my understanding is that generally both devices were installed during the same calendar month for the customers that have both devices. Therefore, the first complete calendar month that the customer's real-time meter starts functioning is assumed to be the first month that the IHD device starts working for customers with an IHD.

The dataset is unbalanced in the sense that the number of calendar months of data available differs across customers. This is the result of the fact that customers have different start and end dates for their billing cycles and that some households move out of a dwelling and other

customers move into the dwelling during the sample period. Nevertheless, all customers have monthly consumption observations of at least twelve calendar months and all customers in the treatment group are observed before and after the intervention for a minimum of 3 calendar months. Because this program was not a voluntary opt-in intervention, these differences in the length of time that data is available for each household is accounted for by household-level fixed effects in econometric model estimated to recover the average treatment effects.

3 Experimental Design

This section describes the experimental design underlying the measurement framework employed to compute the average treatment effects associated with installing a real-time meter and IHD in a household’s dwelling. Two approaches were employed to compute matched treatment and control groups such that the pre-intervention distribution of the average daily consumption for households in the control group is not statistically different from the distribution of average daily consumption for customers in the treatment group for each month in pre-sample period. The first sample first selects households from the treatment and control groups to match based on observable characteristics. The second sample uses all of the available data from households in the treatment and control group in the analysis. From both of these samples, two additional datasets are compiled to create a sample of treatment and control households where the distributions of monthly consumption for the treatment and control groups are statistically indistinguishable for all months during the pre-intervention period.

The key observable difference between households is the type of premise that the household occupies. There are three types of premises observed in the data, called HDB03, HDB04 and HDB05. The sample selection process draws households from the treatment and control groups to ensure that the distributions of household types across the treatment and control groups are the same. The Appendix describes the procedure employed to select a matched sample of households from the treatment and control groups to ensure that frequency of the three different household types are not statistically different between the treatment households and control households samples. Table 1 compares the distribution of customers in the control and treatment groups by type of premise using presence of an real-meter to define treatment status. Table 2 repeats this exercise using presence of an IHD unit to define treatment status. Note that all households with IHDs also have interval meters. Below each table is the value of the Chi-Squared statistic associated with the test of the null hypothesis that the percentage of households in each dwelling type is the same for the treatment and control samples. Consistent with the design of our sample selection procedure, for both treatments—a real-time meter and an IHD—the null hypothesis is not rejected at a 0.05 level of significance.

Using this balanced sample of treatment and control households to estimate the average treatment effects regression ensures that the estimation results are not impacted by differences in the distribution of premise types across the treatment and control groups. Table 3 shows distribution

of IHD ownership among customers with real-time meters, with 96.08% of the households provided with interval meters having IHD units.

The next step in the sample selection process is to select households for both groups so that the pre-intervention distribution of consumption for treated households is not statistically different from that of the control group households. This approach relies on the two-sample Kolmogorov-Smirnov test of equality of two distributions. Let Q_{im}^k be the average daily consumption of customer i of type k during month m , where k denotes one of control or treatment. As described above, all monthly consumption is expressed in terms of average daily consumption during that month to account for the fact that there are different numbers of days during different months of the year.

Suppose there are M months during the pre-intervention period, which ends during the month that first meter was installed for any consumer in the sample (May 2012). Suppose there are N_k customers in group k . Define the empirical distribution of average daily consumption in month m for group k as:

$$F_m^k(t) = \frac{1}{N_k} \sum_{i=1}^{N_k} I(Q_{im}^k \leq t)$$

where $I(X \leq t)$ is an indicator variable that takes on the value 1 if X is less than or equal to t and zero otherwise. Under the assumption that the Q_{im}^k , $i = 1, 2, \dots, N_k$ are independent and identically distributed within month m for each group k with population distribution equal to G_{mt}^k , we can perform the hypothesis test: $H : G_{mt}^k(t) = G_{mt}^h(t)$ versus $K : G_{mt}^k(t) \neq G_{mt}^h(t)$ using the two-sample Kolmogorov-Smirnov statistic

$$KS = \sup_t |F_{mt}^k(t) - F_{mt}^h(t)|$$

The Appendix describes the procedure employed to obtain sample of households that meets the criteria of failing to reject the null hypothesis of the equality of the average daily consumption distributions for the treatment and control groups during all months of the pre-intervention period. Table 4 reports the K-S statistic and the associated 1% critical value for month m for the test that the distribution of average daily consumption for month m for the control group is equal to that of the treatment group during the pre-intervention period. The table also gives number of observations in both the treatment and control groups. For the months July 2011 to April of 2012, the test statistic is never greater than the critical value, indicating that a size $\alpha = 0.01$ test of the null hypothesis would not be rejected for any of these months during the pre-sample period.

These results indicate that for this sample the customers from the treatment and control groups are matched on observables (premise type) and their pre-intervention consumption distributions are statistically indistinguishable. Hence we can consider these matched control and treatment samples to be produced by a quasi-experimental research design. For comparison, the regression analysis will also be performed with the larger sample of control and treatment group customers that are matched with treatment customers on the premise type distribution for which we reject the equality of the monthly consumption distributions for the treatment and control groups for at least one month of the pre-intervention period.

Tables 5 to 7 reports the same results as Tables 1,2 and 4 for the sampling scheme that uses all available treatment and control households, rather than a matched sample that equalizes the distributions of households across dwelling types for the treatment and control groups. Tables 5 and 6 demonstrate that the null hypothesis of the equality of the percentages of households across dwelling types for the treatment and control groups is rejected for both definitions of the treatment group for this larger sample. However, as Table 7 demonstrates it is still possible to find a subsample of these treatment and control groups where the null hypothesis equality of the pre-intervention monthly distributions of average daily consumption for the treatment and control groups is not rejected for any month during the pre-sample period. The resulting sample of treatment and control households and the entire set of treatment and control households are the final two datasets with use in our regression analysis.

4 Empirical Strategy

This section presents the difference-in-difference econometric modeling framework used to estimate the average treatment effect of each intervention on the customer’s electricity consumption. Let Q_{imt} be the average daily electricity consumption in month m for customer i in period t . Where period $t = 0$ denotes pre-intervention period and $t = 1$ is the post intervention period for the i^{th} customer.

We include fixed effects for each individual as well as for each month of the sample period. Hence any common patterns in consumption across the groups due to seasonal changes in climate (for example more demand for air conditioning in summer months leading to higher consumption) will not affect the estimate of the elements of $\beta = (\beta_{meter}, \beta_{IHD})'$, the parameter of interest. γ_m denotes the vector of fixed effects for each month of the sample of M months, δ_i denotes the fixed effect for each household. The full econometric specification follows:

$$\text{Log}(Q_{imt}) = \alpha + \beta_{meter}\mathbf{1}[meter = 1 \times t = 1] + \beta_{IHD}\mathbf{1}[IHD = 1 \times t = 1] + \delta_i + \gamma_m + \epsilon_i$$

β_{meter} is the average treatment effect of installing a meter and β_{IHD} is the average treatment effect of installing an IHD device.

Because our treatment effects model is simply estimating the difference in means of Q_{imt} for $t = 1$ and $t = 0$, the difference between the post-intervention and pre-intervention average consumption, an informative graphical presentation of the results is possible. Figure 1 plots the histogram of $Y_i = Q_{im1} - Q_{im0}$ for control and treatment groups for the matched treatment and control sample based on dwelling type used to construct Tables 1, 2, and 4. The distribution of Y_i for the treatment sample is almost a uniformly negative shift across all percentiles of the distribution of Y_i for the control sample. The regression result in Table 8 demonstrates that the average treatment effect associated with IHD unit is precisely estimated to be -7.76%, which corresponds to about -0.78 KWh/day or 285 KWh annually at a daily average consumption of approximately 10 KWh.

Table 8 also reports the results of estimating the difference-in-difference model with the larger

treatment and control sample where the null hypothesis of equality of the monthly pre-intervention distributions of average daily consumption is rejected for some months in the pre-intervention period. The estimation results are nearly identical to those estimated for the sample where this null hypothesis is not rejected. For both samples, the coefficient on indicator variable for the existence of a real-time meter in the customer's dwelling is not statistically different from zero.

Table 9 reports the results of estimating the difference-in-difference model for larger unmatched sample analyzed in Tables 5 to 7. The second column reports the case that all monthly observations for the treatment and control households are utilized. The third column report the estimates for the case that the treatment and control samples are selected to ensure that the null hypothesis of equality of distributions of daily average consumption across the treatment and control households is not rejected for any month during the pre-intervention period. The average treatment effect is -7.82% for complete sample results and -7.79% for the smaller sample results. Once again, the coefficient on the existence of a real-time meter is not statistically different from zero.

Because coefficient on the interval meter indicator variable is not statistically different from zero, we re-estimate the model dropping this variable for our four datasets. These results are reported in Table 10. The coefficient estimate on the IHD indicator is precisely estimated and equal to -0.05, implying a roughly 5 percent reduction in the household's daily average electricity consumption associated with the household having an IHD device installed in the dwelling.

5 Discussion

Assuming a 10 KWh daily average consumption for our sample period, the estimated annual electricity consumption reduction is approximately 186 KWh. The average retail price during 2012 is 0.279 Singapore dollars per KWh. This implies roughly 50 Singapore dollar per year savings for households with IHDs. Assuming a 10-year life for the real-time meter and IHD and a discount rate of 5 percent, implies a discounted present value of savings for the life of the two devices of 386 Singapore dollars.

Extrapolating these results to all Singapore residential electricity consumers implies that rolling out real-time meters with IHDs to all Singapore households would more than pay for the cost this policy and therefore yield significant net benefits to the Singapore economy. In 2012, total household electricity sales in Singapore was 6,640 gigawatt-hours (GWh). Five percent of that figure is 332 GWh. Valuing these annual savings at 0.279 Singapore dollars per KWh implies an annual savings of 92.6 million Singapore dollars. Assuming these annual savings last for the assumed 10-year life of the real-time meter and IHD, and applying a 5 percent discount rate to these savings yields a discounted present value of savings of more than 710 million Singapore dollars.

These results reinforce the importance of actionable information provided in a timely fashion to encouraging more efficient consumption of electricity by households. Kahn and Wolak (2012) found similar magnitudes of electricity savings from providing households with information about the nonlinear pricing schedules that households faced are used to compute their monthly electricity bill

and how electricity consuming actions both increase and decrease the households monthly electricity bill. The installation of real-time meters and IHD device leaves open the opportunity to capture further demand-side savings. By implementing dynamic pricing plans where the price a household pays for electricity varies with the hourly wholesale price can increase the potential savings that consumers can achieve from these technologies. As Wolak (2011) and (2007) demonstrates, dynamic pricing programs can produce 10 to 20 percent reductions in electricity demand during peak hours of the day and these programs are only technologically feasible if the household have a real-time meter.

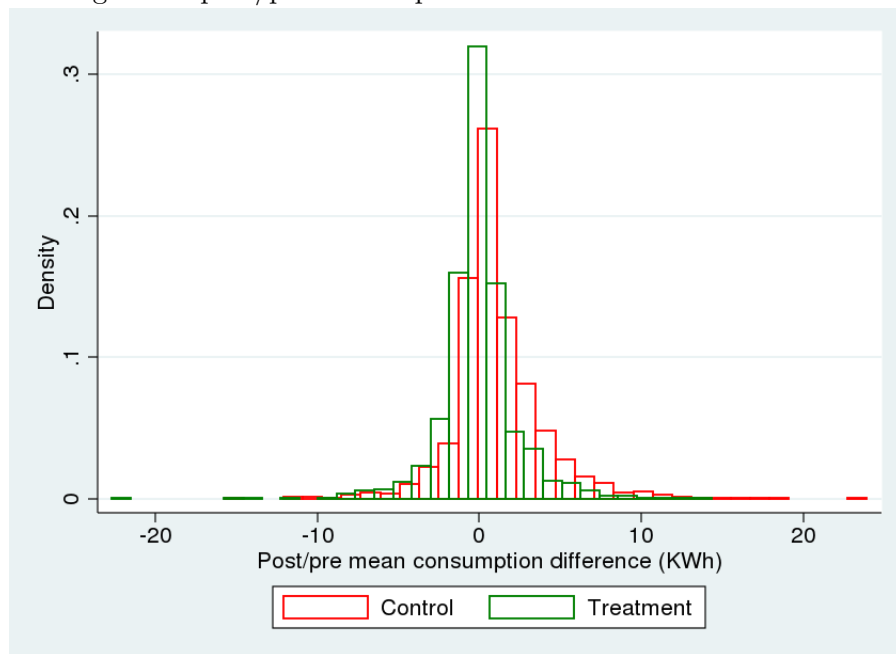
6 Conclusions

This paper evaluates the efficacy of real-time usage feedback to households. EMA's IES Pilot provides the ideal environment to study the impact of installing IHD units on the households electricity consumption. A simple difference-in-difference estimator applied to a variety of samples of treatment and control households shows that IHD units are associated with a 5.0% reduction in consumption. This translates into a 50 dollar per year saving in electricity consumption per household. Scaling these savings to all households in Singapore yields substantial aggregate benefits associated with the widespread adoption of these devices. The results of this analysis are consistent with those obtained from other information provision experiments. They also suggest that even greater savings are possible if dynamic pricing programs were adopted for households with interval meters.

References

- Kahn, M. & Wolak, F.A. 2013. “Using information to improve the effectiveness of non-linear pricing: Evidence from a field experiment”, <http://web.stanford.edu/group/fwolak/cgi-bin/>.
- Wolak, F.A. 2007. “Residential Customer Response to Real-Time Pricing: The Anaheim Critical-Peak Pricing Experiment” available at <http://web.stanford.edu/group/fwolak/cgi-bin/>.
- Wolak, F.A. 2011. “Do Residential Customers Respond to Hourly Prices? Evidence from a Dynamic Pricing Experiment”, *AER* 101:3, pp 83-87

Figure 1: Histogram of post/pre consumption difference for control & treatment groups



7 Matched On Premise Type Sample

Table 1: Distribution of premise types (based on interval meters)

Premise type	Meter = 0		Meter = 1		Total	
	freq	%	freq	%	freq	%
HDB03	112	5.44	91	4.98	203	5.23
HDB04	1,059	51.43	911	49.89	1,970	50.71
HDB05	888	43.13	824	45.13	1,712	44.07
Total	2,059	100.00	1,826	100.00	3,885	100.00

Pearson $\chi^2(2) = 1.7159$ Pr = 0.424

Table 2: Distribution of premise types (based on IHD units)

Premise type	Meter=0		IHD= 1		Total	
	freq	%	freq	%	freq	%
HDB03	112	5.44	77	4.42	189	4.97
HDB04	1,059	51.43	853	48.97	1,912	50.30
HDB05	888	43.13	812	46.61	1,700	44.73
Total	2,059	100.00	1,742	100.00	3,801	100.00

Pearson $\chi^2(2) = 5.6756$ Pr = 0.059

Table 3: Distribution of In-home display units within group that has meters

	IHD=0	IHD=1	Total
Count	84	1,742	1,826
%	3.92	96.08	100

Table 4: Two-Sample Kolmogorov-Smirnov test before intervention

Month	Control	Treatment	K-S stat	1% critical val
July 2011	1628	1655	.0533	.0569
Aug 2011	1664	1685	.0407	.0563
Sept 2011	1714	1701	.0525	.0558
Oct 2011	1741	1712	.0346	.0555
Nov 2011	1824	1729	.0321	.0547
Dec 2011	1891	1750	.0232	.0541
Jan 2012	2010	1755	.0493	.0533
Feb 2012	2056	1764	.0491	.0529
Mar 2012	2058	1775	.0268	.0528
Apr 2012	2059	1789	.0245	.0527

8 Sample Not Matched on Dwelling Type

Table 5: Distribution of premise types (based on interval meters)

Premise type	Meter = 0		Meter = 1		Total	
	freq	%	freq	%	freq	%
HDB03	65	3.11	91	4.98	156	3.99
HDB04	1,147	54.96	911	49.89	2,058	52.59
HDB05	875	41.93	824	45.13	1,699	43.42
Total	2,087	100.00	1,826	100.00	3,913	100.00

Pearson $\chi^2(2) = 15.5879$ Pr = 0.000

Table 6: Distribution of premise types (based on IHD units)

Premise type	Meter = 0		IHD= 1		Total	
	freq	%	freq	%	freq	%
HDB03	65	3.11	77	4.42	142	3.71
HDB04	1,147	54.96	853	48.97	2,000	52.23
HDB05	875	41.93	812	46.61	1,687	44.06
Total	2,087	100.00	1,742	100.00	3,829	100.00

Pearson $\chi^2(2) = 15.6265$ Pr = 0.000

Table 7: Two-Sample Kolmogorov-Smirnov test before intervention

Period	control N	Treatment N	K-S stat	1% cutoff
Jul2011	1628	1655	0.053	0.057
Aug2011	1669	1685	0.042	0.056
Sept2011	1728	1701	0.055	0.056
Oct2011	1756	1712	0.040	0.055
Nov2011	1839	1729	0.035	0.055
Dec2011	1897	1750	0.023	0.054
Jan2012	2033	1755	0.053	0.053
Feb2012	2082	1764	0.050	0.053
Mar2012	2084	1775	0.040	0.053
Apr2012	2085	1789	0.018	0.053

Table 8: Average Treatment Effects (Sample Matched on Dwelling Type)

	(K-S rejects)	(K-S fails to reject)
	Log(cons)	Log(cons)
$\mathbf{1}(Meter) * \mathbf{1}(t > T)$	0.0341 (1.48)	0.0440** (1.98)
$\mathbf{1}(IHD) * \mathbf{1}(t > T)$	-0.0780*** (-3.32)	-0.0776*** (-3.49)
Observations	164,841	74,037

T = Intervention, May 2012

t statistics in parentheses* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 9: Average Treatment Effect (Sample Not Matched on Dwelling Type)

	(K-S rejects)	(K-S fails to reject)
	Log(Cons)	Log(Cons)
$\mathbf{1}(Meter) * \mathbf{1}(t > T)$	0.0279 (1.19)	0.0275 (1.22)
$\mathbf{1}(IHD) * \mathbf{1}(t > T)$	-0.0782*** (-3.29)	-0.0779*** (-3.46)
Observations	263,775	74,353

T = Intervention, May 2012

t statistics in parentheses* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Average Treatment Effect of IHD and Real-Time Meter Combination

	Matched on premise		Not matched on premise	
	KS fails to reject	KS rejects	KS fails to reject	KS rejects
	Log(cons)	Log(cons)	Log(cons)	Log(cons)
$\mathbf{1}(IHD) * \mathbf{1}(t > T)$	-0.0533*** (-5.42)	-0.0537*** (-8.04)	-0.0502*** (-8.06)	-0.0503*** (-9.51)
Observations	72,456	163,260	72,772	262,194

T = Intervention month

t statistics in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix: Data Preparation

The following transformations on the original data set were performed to prepare it for analysis as a quasi-random assignment of consumers into treatment and control groups:

- The data was provided in three separate files.
 - The first file contained information on whether the customer was provided with an interval meter and/or IHD unit. It also contained hourly consumption data for each consumer. From this hourly data, it was possible to determine the first date from which the meter started recording non-zero consumption. This was assumed to be the treatment date since presumably the consumer started getting direct feedback about consumption only from this data and not simply when the meter was installed but not functional. Separately there is also a variable that indicates the date of meter installation, but we did not use this as the treatment date since in most cases there was a period of a week or two during which the meter recorded zero consumption. For control group customers there is not treatment date as such, however before May 2012 was used as a uniform month to indicate the pre-treatment period because no meters were installed before this month.
 - A second file contained information on energy consumption (in KWh) by consumer for both the treatment and control groups. Each observation represented consumption for a billing cycle. This consumption data was used to compute our dependent variable Q_{imt} .
 - The third file provided data on other characteristics for each consumer, like premise-type, zip code and street number.
 - All three files identified each customer uniquely by a contract account number, which was used to match data across the three files. Hence we created a single file by taking the treatment information ($meter = 1/0$, $IHD = 1/0$) from the first file, consumption information (Q_{imt}) from the second file and other characteristics (premise type, zipcode) from the third file. Any ids that could not be matched were discarded.
- Working with this merged file, we first did some routine clean up operations - making sure all the variables are in the appropriate format for analysis, deleting duplicates, renaming variables to more convenient and user-friendly names and so on.
- **Anomalies** - In some cases the first observation for a customer was of a short billing cycle (maybe a week or so) before a more regular monthly pattern. We dropped such observations because we cannot compute a daily average consumption value for this billing cycle because we do not have complete data for the billing cycle. Particularly for treatment group customers, in some cases the monthly consumption recorded was zero until after the treatment date. For such cases we dropped the customer because it would lead to a false increase in mean

consumption post-treatment. The panel data-set is not balanced, although most customers are present for the entire span of 20 months. Although we retain this character, we dropped customers for whom we did not have at least 3 observations both pre and post treatment.

- The next step was to compute the weighted daily average consumption value for each calendar month in the sample as discussed in the text. The first billing cycle data was from June 10, 2010 - July 10, 2011. However because we don't observe the data for June 2011 completely we start the analysis from July 2011. Similarly although the last observation in the data-set ends at March 10, 2013, the last month for analysis is February 2013.
- Next, the task was to match the treatment and control groups on the distribution of premise-type. This is important to ensure that a skew in premise-type is not creating some of the difference in means between the two groups. We randomly sampled from the control group to match the probability of a customer belonging to a certain premise type with that of the treatment group. We will explain this more concretely. There are three premise types for which we had observations from both groups. Suppose that P_i^T is the probability of a customer in the treatment group belonging to premise i , where $i = 1, 2, 3$. Similarly, P_i^C is the probability of a customer in the control group belonging to premise i . We would like $P_i^T = P_i^C$ for all i . But this is not the case in the original data-set. Hence we find sampling probabilities \tilde{P}_i for each premise type i such that

$$\frac{\tilde{P}_i P_i^C}{\sum_{i=1}^3 \tilde{P}_i P_i^C} = P_i^T$$

This gives us 3 equations in 3 unknowns that we can solve for each \tilde{P}_i . Using these sampling probabilities we randomly sample customers from the control group and form the matched control group sample 1.

- Although we now had control and treatment samples that were matched on observables, we created a further sub-sample that also matches the distribution of consumption across control and treatment groups in the pre-treatment period. For the purpose of this analysis we used May 2012 as the uniform treatment month across all observations since the first meters were installed in that month. As described in the text, we used the two-sample Kolmogorov-Smirnov statistic to test if the distribution of consumption across the groups was statistically different. We randomly sampled from the control group until we found a sub-sample for which we failed to reject the KS statistic test at the 1% level.