Information Frictions, Inertia, and Selection on Elasticity: A Field Experiment on Electricity Tariff Choice*

Koichiro Ito
University of Chicago and NBER

Takanori Ida
Kyoto University

Makoto Tanaka
GRIPS

This version: July 17, 2016

Abstract

We develop a discrete-continuous choice model to characterize the link between plan choice, switching frictions, and subsequent continuous choice of service utilization. We then test the model predictions by using a randomized controlled trial in electricity tariff choice. We find that both information frictions and inertia prevent consumers from switching to a tariff that is privately and socially beneficial. While interventions to mitigate these frictions increased overall switching rates, they also incentivized relatively inelastic consumers to switch. We characterize this phenomenon by selection on elasticity and show how it affects the optimal rate design in the presence of switching frictions.

*Ito: Harris School of Public Policy, University of Chicago, 1155 East 60th St., Chicago, IL 60637, and NBER (e-mail: ito@uchicago.edu). Ida: Graduate School of Economics, Kyoto University, Yoshida, Sakyo, Kyoto 606-8501, Japan (e-mail: ida@econ.kyoto-u.ac.jp). Tanaka: National Graduate Institute for Policy Studies, 7-22-1 Roppongi, Minato-ku, Tokyo 106-8677, Japan (e-mail: mtanaka@grips.ac.jp).

We thank Severin Borenstein, Meredith Fowlie, Catherine Wolfram, Frank Wolak, and seminar participants at UC Berkeley Energy Camp and University of Colorado, Boulder for helpful comments. We thank the Japanese Ministry of Economy, Trade and Industry, the city of Yokohama, Accenture, Tokyo Electric Power Company, Toshiba Corporation, and Panasonic Corporation for their collaboration on this study. We thank the New Energy Promotion Council for financial support.

1 Introduction

Recent economic studies find that consumers in many markets have strong default bias in plan choice. In health insurance markets, Einav et al. (2013); Handel (2013); Handel and Kolstad (2013) find that most consumers overlook benefits of alternative health insurance plans. Likewise, Benartzi and Thaler (2007) present that individuals usually keep their default retirement plans, leaving substantial gains that can be obtained by switching to alternative plans. Similar empirical evidence has been documented in many other markets, including cell phone (Miravete, 2003) and electricity (Hortaçsu, Madanizadeh and Puller, 2015; Cappers et al., 2015).

Why do consumers stick to their default plans when there is an obvious economic gain from switching to an alternative plan? Standard models of plan choice cannot fully explain this phenomenon, and the recent literature provides two potential reasons for such switching frictions. The first possibility is that consumers have information frictions when making plan choice. Handel and Kolstad (2013) present that health insurance consumers are not fully informed about the characteristics of available plans. If there are substantial information frictions, it could be an enough obstacle to prevent customers from switching to otherwise economically more attractive plans. The second possibility is a variety of switching costs, which include explicit switching costs such as a penalty fee for early termination and implicit switching costs—which we call by "inertia"—such as status-quo bias for a default plan. Hortaçsu, Madanizadeh and Puller (2015) show that most residential electricity customers in Texas did not switch to a new electricity provider even though information on alternative tariffs were readily available on a website, and explicit switching costs were low. This evidence suggests that the observed low switching rates may not be fully explained by information frictions by itself.

A key element in plan choice in health insurance, cell phone service, or energy service is that consumers' discrete decisions on plan choice and their subsequent continuous choice on utilization are interrelated (Einav et al., 2013). In this paper, therefore, we first develop a discrete-continuous choice model to characterize the link between plan choice, switching frictions, and subsequent continuous choice of service utilization. We then test the model predictions by using a randomized controlled trial in electricity tariff choice. Our field was a regulated residential electricity market in the city of Yokohama in Japan. Customers in the experiment were served by a single electric

utility company with a default time-invariant electricity tariff. We introduced a new tariff, which was time-variant to better reflect the marginal cost of electricity. Compared to the default tariff, the new tariff was designed to be socially more efficient as well as privately beneficial for about a half of consumers.[1] Therefore, under a standard plan choice model with no information frictions and switching costs, at least half of customers were expected to switch to the new tariff for their private benefits.

We randomly assigned 2153 households to one of 4 groups: control, baseline treatment, information treatment, and information + upfront incentive treatment groups. Consumers in the baseline treatment group received the option to switch to the new tariff. If they made no action, they continued to have the default tariff. Consumers in the information treatment group received the option to switch as well as information about their "expected gains from switching." This gain was calculated based on their 30-minute interval usage data in the past, assuming zero price elasticity. This information should be redundant for customers who do not have information frictions. In addition to the information treatment, customers in the third treatment group received an upfront incentive ($60) upon switching to the new tariff. We used these interventions to investigate how information frictions and inertia influence plan choice.

The second stage of the experiment was to examine subsequent consumer behavior followed by their switching decisions. Compared to the default tariff, customers under the new tariff had higher prices for peak hours—1 pm to 4 pm in weekdays—and a lower price for off-peak hours. By design, our experiment created partial compliance for each treatment group. In our context, both of the treatment effect on the treated (TOT) and the intention to treat (ITT) are of interest. The TOT estimate tells us the average price responsiveness of the compliers—those who switched to the new tariff—under each of the three interventions. The ITT estimate tells us the overall policy outcome as a product of the TOT estimate and the switching rate.

We begin by showing several findings from the first stage of the experiment—tariff switching decisions. First, the switching rate was 16% for the baseline treatment group, 31% for the information treatment group, and 48% for the information + upfront incentive treatment group. These switching rates are substantially lower than those predicted by a standard model of utility max-

[1] The tariff schedule is variable critical peak pricing (VCPP) similar to Ito, Ida and Tanaka (2015). Similar dynamic price schedules were tested or discussed in previous experiments including Wolak (2011); Joskow (2012); Joskow and Wolfram (2012); Jessoe and Rapson (2014).

imization. By the construction of the new tariff, over 50% of customers would have gained from switching, even with no change in their electricity usage. With the upfront incentive, almost all customers would have benefited from switching. The low switching rates imply that many customers left money on the table particularly when they were not fully informed about their expected gains from switching. Second, we examine whether consumers who have positive and larger expected gains from switching—so called "structural winners"—were more likely to switch. We find that the switching rate was homogeneous across the distribution of the expected gains, when customers were not informed about the expected gain. In contrast, when they were explicitly informed about the expected gains, the structural winners were more likely to switch than those who would lose from switching. Third, we test how risk preferences affected tariff switching decisions. We find that risk aversion and certainty premium negatively affected switching. The intervention of information provision mitigated this issue to some extent—the information intervention increased the switching rate of consumers who have higher risk aversion and certainty premium.

We then analyze the continuous choice on electricity usage by using household-level 30-minute interval consumption data. We show several findings from the continuous choice. First, the treatment effect (the change in electricity usage in response to the price change) is heterogeneous among consumers. This is because we find that the TOT estimates differ across our three treatment groups, which should not be the case if the treatment effect is fully homogeneous across consumers. Second, providing information and upfront incentives increased switching rates but also induced the selection of inelastic consumers. Third, this selection on elasticity can be partly explained by the indirect selection through the relationship between consumer characteristics and price elasticity. Fourth, because of the different selection on elasticity driven by the three treatments, the intention to treat estimates are statistically indifferent between the baseline treatment group and the information treatment group. That is, the overall policy outcome is the same between these two groups despite the fact that information provision increased the switching rate.

An advantage of the reduced-form analysis for the discrete choice and continuous choice is that it provides transparent analysis without imposing a specific structure on consumersÕ utility functions and their optimization. A limitation of this approach is that we cannot estimate the magnitudes of information frictions and inertia without having a structure on the utility maximization. Another limitation is that we cannot use reduced-form analysis to jointly estimate the discrete choice

and the continuous choice, while our model suggests that consumer preferences enter both of the decisions, and consumers jointly make the discrete and continuous decisions. For this reason, the final section of this paper structurally estimates consumer preferences using the discrete-continuous choice model based on the model developed in our model section. Given a set of assumptions behind the estimation including a functional form assumption on the utility function, the model allows us to jointly estimate the discrete and continuous decisions and identify the information frictions and inertia by exogenous variation created by our RCT.

We find economically and statistically significant amount of information frictions and inertia. We also find that these two frictions are strongly related to consumer characteristics such as their expected gains from switching and household income. We then test selection on elasticity using the estimates from the discrete-continuous choice model. Consumers in the baseline treatment group exhibit strong selection on elasticity—elastic consumers are more likely to switch to the new tariff. However, consumers who received information provision show quite different patterns of the selection on elasticity—inelastic consumers are more likely to switch. Consistent with our finding in the reduced-form analysis, part of the reason for this relationship is that consumers who have larger expected gains from switching, who are more likely to be induced to switch by the information provision, are less price-elastic. Finally, those who received the information provision and upfront incentive exhibit similar selection on elasticity to the one for the baseline group, but the selection on elasticity is weaker. One interpretation of this result is that the information provision generates selection of inelastic consumers because it attracts structural winners, but providing upfront cash incentive also attracts structural lowers as well. These two effects generate the selection on elasticity similar to the baseline treatment group but the resulting selection is weaker for the information + incentive treatment group.

This draft leaves a few more analyses for our further work. This includes welfare analysis based on the estimates of the discrete-continuous choice model. Such analysis could highlight the policy implications of information provision and upfront incentives in terms of consumer welfare and social welfare, which are key measures to consider the optimal rate design in the presence of switching frictions.

2 Tariff Choice and Electricity Demand with Switching Frictions

In this section, we develop a stylized model of electricity demand and tariff choice in the presence of information frictions and inertia. In the absence of information frictions and inertia, consumers choose an electricity tariff that simply maximizes their utility. Our model allows the possibility that consumers may have a cost to acquire information at switching (information frictions) as well as remaining implicit switching frictions due to state-dependence (inertia). The model guides us to construct key predictions for our empirical analysis. It shows how a consumer's discrete choice on electricity tariffs is related to the consumer's subsequent continuous choice on electricity use. It also shows how information frictions and inertia affect the degree to which consumers switch as well as the types of consumers who decide to switch.

2.1 A Discrete-Continuous Choice Model of Tariff Choice and Utilization

We model electricity tariff choice as a two-stage discrete-continuous choice problem.[2] Suppose that there are two electricity tariff schedules—a default tariff ($j = 0$) and a new tariff ($j = 1$). Electricity consumer i can remain with the default tariff or switch to the new tariff. In the first stage, the consumer decides whether to switch to the new tariff. In the second stage, the consumer optimizes electricity usage given the tariff chosen at the first stage. The consumer solves the two-stage problem by backward induction.

Tariff schedule j consists of two prices (p_{jh} for $h = 0, 1$)—the price of electricity per kilowatt hour for peak hours (p_{j1}) and that for off-peak hours (p_{j0}). For example, p_{j1} is the price for hours between 1 pm and 4 pm on weekdays, and p_{j0} is the price for other hours. Our model is a consumer-level optimization problem, so we omit i subscripts to simply notation. We denote electricity usage for consumer i for hour h by x_h.

Second stage (usage choice).—In the second stage, consumers optimize electricity usage given tariff choice j. Suppose that the second-stage utility function is quasi-linear and separable between peak and off-peak hours:

$$u_j(x_{j0}, x_{j1}, p_{j0}, p_{j1}; \theta) = \sum_{h=0}^{1} \Big(v(x_{jh}; \theta) - p_{jh} x_{jh} \Big). \tag{1}$$

[2] Our model is closely related to previous work on discrete-continuous choice models (Dubin and McFadden, 1984; Hanemann, 1984; Bento et al., 2009; Einav et al., 2013).

$v(x_{jh};\theta)$ is utility from electricity usage and θ is a vector of preference parameters. The optimal usage for hour h given tariff j—often called the *conditional* demand—is:

$$x^*_{jh}(p_{jh};\theta) = \arg\max_{x_{jh}} \; u_j. \tag{2}$$

Then the optimal utility given tariff choice j—the conditional indirect utility—is given by:

$$u^*_j(p_{j1}, p_{j2};\theta) = u_j(x^*_{j1}, x^*_{j2}, p_{j1}, p_{j2};\theta). \tag{3}$$

First stage (tariff choice).—In the first stage, consumers switch to the new tariff ($j = 1$) if switching leads to better conditional indirect utility compare to the utility obtained by remaining with the existing tariff ($j = 0$). We allow the possibility that consumers have information frictions (ι) when switching. Fully-informed consumers can calculate their gains from switching with no costs ($\iota = 0$). However, consumers may have information frictions ($\iota > 0$) that prevents them from switching. In addition, even if consumers are free from information frictions, there can be other switching frictions, for example, due to state-dependence. We define this friction—switching frictions aside from information frictions—by inertia (δ). Consumers switch to the new tariff ($j = 1$) if:

$$\Delta u \equiv u^*_1(p_{11}, p_{12};\theta) - \iota - \delta - u^*_0(p_{01}, p_{02};\theta) > 0. \tag{4}$$

An important distinction between discrete choice models and discrete-continuos choice models is that the preference parameters (θ) enter both of the discrete and continuous optimization problems in the later model. Intuitively, consumers consider their preferences for electricity usage (e.g. the price elasticity of demand) when they consume electricity (a continuous choice) as well as when they choose a tariff (a discrete choice). Therefore, these two decisions should be modeled as a joint optimization problem (Dubin and McFadden, 1984; Hanemann, 1984).

We are primarily interested in two questions regarding equation (4). First, the magnitude and distributions of information frictions (ι) and inertia (δ) provide key policy implications to improve switching rates. Second, in addition to the overall switching rates, an important question is "what types of customers are more likely to switch?" To shed light on these two factors, we provide results with a parametrized utility function that is often used in the literature on electricity demand.

2.2 Parametrization and Theoretical Predictions for Empirical Analysis

Suppose that the second-stage utility function is parameterized by:

$$u_j(x_{j1}, x_{j2}, p_{j1}, p_{j2}; \theta) = \sum_{h=0}^{1} \left(\alpha_h^{-1/\epsilon_h} \frac{x_{jh}^{1+1/\epsilon_h}}{1+1/\epsilon_h} - p_{jh} x_{jh} \right), \tag{5}$$

where $\theta = (\alpha, \epsilon)$ and $-1 < \epsilon \leq 0$. This parametrization assumes a quasi-linear utility function with a constant price elasticity ϵ_h. Note that the constant elasticity means that it is constant over the level of consumption for consumer i, but it is heterogenous across consumers. The first order conditions with respect to x_{jh} lead to the optimal usage given tariff j—the conditional demand—by:

$$x_{jh}^*(p_{jh}; \theta) = \alpha_h p_{jh}^{\epsilon_h}. \tag{6}$$

This demand function provides a log-log specification familiar in the literature of electricity demand, $\ln x_{jh}^* = \alpha_h' + \epsilon_h \ln p_{jh}$, where $\alpha' = \ln \alpha$. The parameters are the price elasticity of electricity demand (ϵ) and consumer-level demand shifter (α). One way to interpret the demand function is that the demand is consist of consumer-level factors that are unrelated to price (e.g. a household with larger houses or larger household sizes need to use more electricity) and factors that are determined by price and elasticity. This specification is similar to the model of health insurance coverage and utilization used by Einav et al. (2013). The conditional indirect utility is:

$$u_j^*(p_{j1}, p_{j2}; \theta) = -\sum_{h=0}^{1} \frac{\alpha_h p^{1+\epsilon_h}}{1+\epsilon_h}. \tag{7}$$

The ultimate policy outcome is determined by two factors—the overall switching rates and the price elasticity of demand for consumers who switch. Therefore, we examine two primary questions: 1) what are the obstacles/frictions that prevent consumers from switching? and 2) how consumer types affect their switching decisions? To answer these questions, we need to examine empirical evidence of the following parameters and comparative statics:

1. $\iota = \Delta u(\iota \neq 0) - \Delta u(\iota = 0)$. This parameter tells us the monetized value of information frictions. We are also interested in empirically estimating the distribution of ι across consumers, and how the heterogeneity is related to observable consumer characteristics.

2. $\delta = \Delta u(\delta \neq 0) - \Delta u(\delta = 0)$. This parameter tells us the monetized value of inertia aside from information frictions. We are also interested in empirically estimating the distribution of δ across consumers, and how the heterogeneity is related to observable consumer characteristics.

3. $\partial \Delta u / \partial \alpha$. This derivative tells us how consumer type that is unrelated to price-responsiveness affect switching decisions. For example, with the particular parametric utility function described above, it is possible to show that $\partial \Delta u / \partial \alpha_1 < 0$ if $p_{11} > p_{01}$. That is, given that the peak hour price is higher for tariff $j = 1$ than tariff $j = 0$, customers who have relatively higher electricity needs for peak hours are less likely to switch, holding other factors fixed. We want to empirically test this prediction.

4. $\partial \Delta u / \partial \epsilon$. This derivative tells us how consumer type unrelated to price-responsiveness affect switching decisions. For example, with the particular parametric utility function described above, it is possible to show that $\partial \Delta u / \partial \epsilon_1 < 0$ if $p_{11} > p_{01}$ and ϵ and α are in a certain range. That is, given that the peak hour price is higher for tariff $j = 1$ than tariff $j = 0$, inelastic customers are less likely to switch, holding other factors fixed. We want to empirically test this prediction.

The goal of our empirical analysis is to provide empirical evidence on these parameters and comparative statics by using a randomized controlled trial in the field. In the next section, we describe the experimental design.

3 Experimental Design

We used a field experiment and high-frequency data on household-level electricity usage to test several hypotheses. In this section, we describe our experimental design, data, treatments, and hypotheses.

3.1 Experimental Design and Data

Our field experiment was conducted for residents in the city of Yokohama in Japan in the summer of 2013. The experiment was implemented in collaboration with the Ministry of Economy, Trade

and Industry (METI), the city of Yokohama, Tokyo Electric Power Company (TEPCO), Toshiba Corporation, and Panasonic Corporation.

To invite as broad a set of households as possible, we provided generous participation rewards, which included free installations of an advanced meter and in-home display and a participation reward of 24,000 yen (approximately $240 in 2013). We contacted all residential electricity customers in the area by mail. Of these, 3,659 customers confirmed their participation. We excluded students, customers who had self-generation devices, and those without access to the internet. This process left us with 2153 households. Similar to previous field experiments in electricity demand (Wolak, 2006, 2011; Faruqui and Sergici, 2011; Jessoe and Rapson, 2014), our experiment was a randomized controlled trial (RCT) for self-selected participants, as opposed to a RCT for a purely randomly selected sample of the population. Therefore, it is important to carefully consider the external validity of the experiment, although the internal validity of the experiment is guaranteed by random assignment of treatments. To explore the external validity of our sample, we collected data from a random sample of the population in the corresponding geographical area. We analyze the observables between our sample and the random sample below.

3.2 Electricity Tariff Schedules

Most residential electricity customers in many countries still pay time-invariant prices. Figure 1 shows the default time-invariant tariff for customers in our experimental region. Time-invariant tariffs tend to be highly inefficient in electricity markets because the marginal cost of electricity usually depend on hour of the day.

To overcome the inefficiency, we introduced a new tariff, which intended to better reflect the marginal cost of electricity. Our tariff took a form of variable critical peak pricing (V-CPP), similar to the tariff schedule used in Ito, Ida and Tanaka (2015). Customers under the new tariff paid higher marginal prices during peak hours—1 pm to 4 pm on weekdays. They received day-ahead and same-day notices about whether their peak hour price is either 100 cents or 45 cents. Given a day, all customers under the new tariff paid the same price.

[Figure 1 about here]

The electricity provider in our experimental region was a regulated utility company. Therefore,

we made a revenue neutrality constraint when designing the new tariff. Specifically, we lowed off-peak prices for the new tariff. Compared to the default tariff, the off-peak price for the new tariff was lower by 3 cents.

By design, the new tariff created "structural winners" (Borenstein, 2013). Keeping electricity usage patterns constant, some customers are likely to gain from the new tariff because they use more electricity in off-peak hours. Such customers, called structural winners, can lower their bill by simply switching from the default flat tariff to the new tariff. Structural losers are the opposite types of customers. Some customers use more electricity in peak hours so that they are likely to lose money by switching to the new tariff unless they are very price elastic.

Figure 2 shows the distribution of structural winners and losers for customers in our experiment. We use each customer's pre-experimental data on usage in 30-minute intervals to calculate their bills under the default tariff and under the new tariff. We then calculate each customer's "expected gains from the new tariff." The distribution is bell-shaped with mean zero by construction.

The distribution implies that even with zero price elasticity, about a half of customers can lower their bill simply by switching to the new tariff. However, evidence from many electricity markets shows that residential electricity customers rarely switched from their default tariff to a new tariff. In our experiment, we consider a few treatments to address this problem.

[Figure 2 about here]

3.3 Treatment Groups

[Table 1 about here]

Table 1 shows our experimental design. We randomly assigned the 2,153 households to one of 4 groups: control (C), opt-in treatment (Z_1), opt-in & information treatment (Z_2), and opt-in & information & upfront incentive treatment groups (Z_3).

Control Group (C): The 697 customers in this group received an advanced electricity meter, an in-home display, and the participation reward. Other than that, this group did not receive any treatment.[3]

[3]These advanced electricity meters are sometimes called "smart meters," which record usage at 15-, 30-, or 60-minute intervals. Conventional electricity meters do not record high-frequency usage. These meters typically record only cumulative usage since the installation. Electric utility companies, therefore, need to know usage at the beginning and at the end of a billing cycle to know monthly or bi-monthly usage.

Baseline Treatment Group (Z_1): The 486 customers in this group received an advanced electricity meter, an in-home display, and the participation reward. In addition, we invite this group to opt-in the new tariff. We showed customers the new price schedule, and customers received an option to opt-in. If costumers made no action, they continued to have the default tariff.

Information Treatment Group (Z_2): The 468 customers in this group received an advanced electricity meter, an in-home display, and the participation reward. We invite this group to opt-in the new tariff. We showed customers the new price schedule. In addition, each customer received information about the customer's expected gain from switching. This information is exactly the same as Figure 2. Then, customers received an option to opt-in. If costumers made no action, they continued to have the default tariff.

Information + Incentive Treatment Group (Z_3): The 502 customers in this group received an advanced electricity meter, an in-home display, and the participation reward. We invite this group to opt-in the new tariff. We showed customers the new price schedule. Each customer received information about the customer's expected gain from switching. In addition, customers received an upfront incentive ($60) if they switched to the new tariff. The upfront incentive by $60 was designed to make almost everyone become structural winners. Then, customers received an option to opt-in. If costumers made no action, they continued to have the default tariff.

3.4 Data

The primary data for this study are high-frequency data on household electricity usage. Advanced electricity meters, often called "smart meters," were installed for all participating households, enabling us to collect household-level electricity usage at 30-minute intervals. In addition to the usage data, we conducted three surveys. We conducted the first survey prior to treatment assignment to collect risk preference and demographic information. We conducted the second survey upon completion of the experiment to explore the mechanism behind our findings. Finally, we conducted the third survey for a random sample of households in the area to investigate the external validity of our sample.

[Table 2 about here]

Table 2 presents the summary statistics of demographic variables and pre-experiment consumption data by treatment group. A comparison across control and treatment groups indicates statistical balance in observables because of random assignment of the groups. During the experimental period, very little attrition occurred in each group. Because this small attrition occurred at approximately the same rate in each group, it is unlikely to significantly bias our estimates.

3.5 Timeline of the Experiment

Figure 3 shows the timeline of the experiment. The first stage of the experiment was the pre-experimental period. We collected data on usage, demographic variables, and risk preferences. We used the pre-experimental data to implement blocking randomization to minimize standard errors. As explained in the previous section, customers were randomly assigned to one of the four groups.

[Figure 3 about here]

The second stage of the experiment was opt-in decisions by customers. In June 2013, we notified customers in the treatment groups that they could switch to the new tariff. At the same time, customers in the second and third treatment groups received the additional treatments. Customers had about a month to opt-in the new tariff.

The third stage of the experiment was the implementation of dynamic pricing for customers who switched to the new tariff. We informed that we would provide day-ahead and same-day messages for critical peak days. For a given treatment day, all customers on the new tariff paid the same critical peak price, which was either 45 or 100 cents/kWh. Across the treatment days, customers in this group experienced different critical peak prices.

4 Reduced-Form Estimation of Tariff Switching Decisions

In this section, we present experimental results on tariff switching decisions. Because we constructed the new tariff based on a revenue neutrality condition, at least 50% of customers could lower their expected payment by switching to the new tariff. If customers are risk-neutral and have neither information frictions nor switching cost, standard economic theory predicts that at least 50% of customers switch to the new tariff.

Figure 6 shows how many percent of customers in each treatment group switched to the new tariff. In the baseline treatment group—who received neither information provision nor an upfront incentive for switching, only 16% of customers switched. It implies that at least 34% of customers left money on the table in terms of their expected payment. In the information treatment group—who received information about their shadow bills, we find that the switch rate was increased to %30. It suggests that many customers did not switch even if they were explicitly informed that their expected payment would be lower if they switched to the new tariff. Finally, the information + incentive treatment group received a upfront incentive upon switching to the new tariff. With the incentive, 99% of customers could lower their expected payment by purely switching to the new tariff. However, we find that only 49% of customers actually switched.

[Figure 6 about here]

The findings in figure 6 provides descriptive evidence for two key factors in switching decisions. First, the results for the baseline treatment group and the information treatment group imply that customers are likely to have information frictions on their expected payoff from the new tariff. Second, the results for the information treatment group and the information + incentive treatment group suggest that there are likely to be switching costs in addition to information frictions. The upfront incentive made many customers decide to switch, while it still remained a half of customers staying in their default tariff.

4.1 Expected Gains from Switching

In addition to the aggregate switching patterns, we examine if switching rates systematically differed by consumer type. For a moment, consider a simple model in which all consumers have homogeneous price elasticity. This assumption implies that elasticity is uncorrelated with expected gains from switching. In this case, we expect that structural winners—those with a positive expected gain from switching—should be more likely to switch as long as they were fully informed about their expected gains.

Figure 5 presents the distributions of expected gains from switching for each treatment group. The shaded areas shows consumers who switched to the new tariff. For the baseline treatment group, we find that the switching rate is homogeneous across the expected gains from switching.

This implies either that most customers did not know their expected gains or that some customers who had negative expected gains believed that they were very price-elastic. The figure also suggests that many consumers left money on the table—the switching rate is very low even for those who had fairly large positive expected gains.

In contrast, the histogram for the information treatment group is quite different from the baseline treatment group. When customers were informed about their expected gains, the switching rate for structural winners increased. This implies that information provision was indeed effective to convince structural winners to switch. However, one can see that many structural winners did not switch even when they were informed about their positive expected gains. It suggests that consumers were likely to have switching costs beyond information frictions.

Finally, when consumers received both information and an upfront incentive to switch, the switching rates homogeneously increased over the distribution. This suggests that the information + incentive treatment was able to attract structural losers as well as structural winners for the new tariff.

Note that the ultimate policy outcome of dynamic pricing depends on both switching rates and price elasticity. Therefore, a high switching rate itself does not necessarily imply a success. For example, although we found higher switching rates for the information group and the information + incentive group than the baseline group, it is possible that switchers in the baseline group had higher price elasticity than those in the other groups. We test these possibilities in our econometric estimation.

[Figure 5 about here]

4.2 Risk Preference

Another potentially important factor for switching is risk preference. When consumers make tariff switching decisions, they have uncertainty about electricity usage in the future. This is true even when they have information on their *expected* usage because actual usage depends on uncertain factors such as weather. Therefore, consumers with different risk preference may show different switching pattern. In addition, our information and upfront incentive treatments can have different effects on risk-averse and risk-loving consumers.

[Table 3 about here]

To examine the relationship between risk preference and switching decisions, we elicited each customer's risk preference before they made a switching decision. We used a method developed by Callen et al. (2014). The elicitation was based on two series of questionnaires in Table 3. In the first task, subjects make a series of choices between a relatively safe option A and a relatively risky option B. Where the subject switches from preferring option A to option B determines her risk aversion parameter q. In the second task, we put uncertainty in option A, but kept option A being less risky than option B. This method let us to obtain another risk aversion parameter q'. The advantage of this method is that we can obtain two risk preference parameters, from which we can calculate a certainty premium parameter (cp) based on q and q'.

For each of the three risk preference parameters (q, q', and cp), we estimate an OLS regression:

$$Switch_i = \gamma_1 Risk_i + \gamma_2 Z_2 \cdot Risk_2 + \gamma_3 Z_2 \cdot Risk_3 + \epsilon_i, \tag{8}$$

where $Switch_{it}$ equals to one if consumer i switched to the new tariff, $Risk_i$ is a risk preference parameter, and Z_2 and Z_3 are indicator variables for treatment groups 2 and 3. Parameter γ_1 tells us how risk preference is associated with switching decisions. Parameters γ_2 and γ_3 show how the information and incentive treatments affected the relationship between risk preference and switching decision.

Table 4 shows the estimation results of equation (8) for risk parameters q and q' in columns 1 and 2, and certainty premium cp in column 3. The estimates of γ_1 are negative and statistically significant, suggesting that risk aversion and certainty premium are negatively associated with switching to the newt tariff. The estimates of γ_2 and γ_3 are positive and statistically significant that γ_3 is statistically insignificant for certainty premium. The positive coefficients imply that each of the two treatments—providing shadow bill information and an upfront incentive to switch—weakens both the correlation between risk aversion and switching and the correlation between certainty premium and switching. It suggests that the two treatments changed switching decisions particularly for those who were risk averse and certainty loving.

[Table 4 about here]

5 Reduced-Form Estimation of Electricity Demand

We let customers make tariff switching decisions by July 15, 2013. We then introduced the new tariff on July 20. Compared to the default tariff, consumers under the new tariff faced higher prices for peak-hours (1 pm to 4 pm) and a lower price for off-peak-hours. For now, we focus on the effects of the new tariff on peak-hour usage.

By design, our experiment created partial compliance for each treatment group. In our context, both of the treatment on treated (TOT) and the intention to treat (ITT) are of interest. The TOT estimate tells us the average price responsiveness of those who switched to the new tariff. The ITT estimate tells us the overall policy outcome of each treatment, which is the product of the TOT estiamte and the switching rate.

5.1 Treatment on Treated

We use Z_1, Z_2 and Z_3 as dummy variables for the three treatment groups and D_1, D_2 and D_3 as dummy variables for the compliers of each group—those who switched to the new tariff. To estimate the treatment on treated, we run two-stage least squares:

$$\ln x_{it} = \alpha_1 D_{1it} + \alpha_2 D_{2it} + \alpha_3 D_{3it} + \phi_i + \lambda_t + \eta_{it}, \tag{9}$$

where $\ln x_{it}$ is the natural log of electricity usage for household i in a 30-minute interval t, ϕ_i is household fixed effects, and λ_t is time fixed effects. The instruments are Z_1, Z_2 and Z_3, which are the initial random assignment of the treatment group.[4] The random assignment assures the orthogonality between the instruments and error term, allowing us to obtain the TOT estimates, α_1, α_2, and α_3. Including ϕ_i and λ_t does not affect the point estimates but slightly reduces standard errors because large part of variation in electricity usage can be explained by household fixed effects and time fixed effects. We cluster the standard errors at the household level to adjust for serial correlation. Recall that the treatment groups had increases in price during the treatment hours—1 pm to 4 pm. In this regression, we include only these hours to estimate the treatment effects on the treatment hours. We examine potential spillover effects for nontreatment hours in the following

[4] We use the natural log of usage for the dependent variable so that we can interpret the treatment effects approximately in percentage terms. The treatment effects in the exact percentage terms can be obtained by $\exp(\alpha) - 1$ and $\exp(\beta) - 1$. We report both of the log points and exact percentage terms.

analysis.

[Table 5 about here]

Column 1 of Table 5 shows the TOT estimates for all treatment days. The estimate is the largest for the baseline opt-in group. The point estimate implies a reduction in usage by 0.220 log points (20.0 percent). For the information treatment group, the TOT estimate is a reduction in usage by 0.09 log points (8.6 percent). That is, customers who were induced by the information provision were less price-responsive than customers who switched to the new tariff in the baseline treatment group. Finally, the TOT estimate for customers in the information + incentive group is a usage reduction by 0.132 log points (12.3 percent). The estimate for the baseline treatment group is different from that of each of the other treatment groups at the 1 percent statistical significance level. The estimate for the information group is different from that for the information + incentive group at the 5 percent statistical significance level.

The TOT estimates in Table 5 suggest that the three treatments induced different selection on price elasticity. The baseline treatment attracted relatively elastic customers. The information treatment induced relatively inelastic customers. Finally, the average elasticity of customers who were induced by the information + incentive treatment were between the elasticity for the first and second treatment groups.

Customers who switched to the new tariff had either 45 cents/kWh or 100 cents/kWh for their peak hour price, depending on the treatment day. In columns 2 and 3, we estimate the TOT separately for the days with the high peak hour price and the days with the low peak hour price. Results for columns 2 and 3 are consistent with those in column 1. In terms of the responses to prices, the high peak price produced larger reductions in usage.

5.2 Intention to Treat

For policymakers, one of the key outcomes of interest is the intention to treat, which shows the overall impact of each policy. In our experiment, the ITT is simply the product of the TOT and switching rate. We estimate ordinary least squares:

$$\ln x_{it} = \beta_1 Z_{1it} + \beta_2 Z_{2it} + \beta_3 Z_{3it} + \phi_i + \lambda_t + \eta_{it}, \tag{10}$$

where Z_1, Z_2 and Z_3 are the initial random assignment of the treatment groups, and β_1, β_2, and β_3 provide the ITT estimates for the three treatment groups.

Table 6 shows the results. The ITT estimate is the largest for the information + incentive group because this group had the highest switching rate and a relatively elastic TOT estimate. The estimate (-0.066) implies that the intervention induced a usage reduction by 0.066 log points (6.4 percent). Not surprisingly, the overall usage reduction is smaller than the ones obtained by mandatory critical peak pricing.[5] Yet, a 6% reduction in peak hour usage is still economically significant amount.

[Table 6 about here]

The ITT estimates for the baseline group and the information treatment group are -0.037 and -0.029, and the difference between the two estimates is statistically insignificant. Although the information provision increased the switching rate, the final policy outcome is statistically equivalent to the baseline treatment because information provision attracted inelastic customers. Columns 2 and 3 show that we find consistent results for the treatment days with the higher peak hour price and those with the lower peak hour price.

5.3 Mechanism Behind Selection on Elasticity

Why did our three treatments induce different selection on price elasticity? There are two possible mechanisms that can explain the selection on elasticity. The first mechanism is what we call *indirect selection on elasticity* through the association between the price elasticity and other consumer characteristic. For example, in the analysis of the tariff choice in Section 4, we find that the information and information + incentive treatments incentivized those who have positive expected gains from switching—structural winners—to switch. Suppose that structural winners are less price elastic than others. Then, the selection on elasticity found in Table 5 can be partly explained by the fact that structural winners, who are inelastic, were more likely to be induced to switch by one treatment than another. In this case, the selection on elasticity is indirectly induced by the association between the price elasticity and the expected gains from switching.

[5] For example, Ito (2014) finds that critical peak pricing similar to our experiment induced about a 15% reduction in usage when it is implemented as a mandatory tariff.

The second mechanism is what we call *direct selection on elasticity*. Even conditional on the indirect selection mechanism, different treatments may attract different types of consumers (i.e. consumers with different price elasticity) to switch. For example, consider two consumers who have exactly the same observable characteristics including the price elasticity, the expected gains from switching, and risk preferences. Suppose that one of them received information provision, while the other did not. If the two consumers made different switching decisions, the information provision must have had a direct impact on the selection on elasticity.

We examine the direct and indirect selection on elasticity by estimating the price elasticity as a function of observable consumer characteristics. This analysis can be done by investigating the heterogeneity in the TOT—how the treatment on treated (the price responsiveness of compliers) is related to consumer characteristics. We define Z_{it}, which equals to one if consumer i is assigned to either of the three treatment groups. Similarly, we define D_{it}, which equals to one if consumer i switched to the new tariff (i.e. complier) and the new price was in effect in time t. We then make the interaction variables of Z_{it} and V_{ki} and those of D_{it} and V_{ki}, where V_{ki} is a consumer's characteristics $k \in K$. We run two-stage least squares:

$$\ln x_{it} = \alpha_1 D_{1it} + \alpha_2 D_{2it} + \alpha_3 D_{3it} + \sum_{k \in K} \alpha_k D_{it} V_{ki} + \phi_i + \lambda_t + \eta_{it}. \tag{11}$$

The instruments are $Z_{1it}, Z_{2it}, Z_{3it}$ and $\sum_{k \in K} \alpha_k Z_{it} X_{ki}$. Note that this estimation provides the relationship between the LATE—the treatment on treated in our case—and covariates, which is not necessarily equivalent to the relationship between the ATE and covariates. Another way to interpret this estimation is that we examine if the differences between α_1, α_2, and α_3 remain/disappear when we control for the relationship between the treatment effects and covariates. If they indeed disappear, we can conclude that the selection on elasticity is induced fully by the indirect mechanism through the relationship between the elasticity and covariates.

Table 7 shows the results. For covariates V_{ki}, we use four variables—the expected gains from switching dollars, the quartile of the expected gains from switching, and the risk aversion and certainty premium obtained in Section 3. Panel A shows the results with the expected gains from switching in dollars. For a reference, column 1 shows the TOT estimates equivalent to the result in Table 5. In other columns, we estimate the linear relationship between the TOT and the expected

gains from switching. The positive and statistically significant interaction effect in column 2 implies that consumers with higher expected gains from switching are less price-elastic (note that the price-elasticity is estimated to be negative so that the positive coefficient of the interaction term makes it less elastic).

[Table 7 about here]

In Panel B, we estimate this relationship in a less parametric way by estimating the treatment variable interacted with the quartile of the expected gains from switching, with the first quartile being the excluded category. The estimates imply that the partial effect of the expected gains on the treatment effects seem to be monotonically positive. The results also indicate that the largest interaction effect comes from the fourth quantile. That is, consumers with positive and very large expected gains from switching are substantially less price-elastic than other consumers.

Likewise, columns 2 through 4 in Panels A and B indicate that risk aversion and certainty premium can be another indirect selection mechanisms that partly explain the selection on elasticity in Table 5. The estimates indicate that consumers with higher risk aversion and certainty premium are less price-elastic. In addition, the last column suggests that the effect of certainty premium dominates the effect of risk aversion. Recall that in Table , we find that providing information and upfront incentives encourage consumers that have risk aversion and certainty premium. That is, although these interventions help such consumers to switch to the new tariff, the newly nudged consumers are likely to be less price-elastic partly because of the association between the risk preferences and price elasticity.

Do the interactions between the treatment variable and covariates fully explain the selection on elasticity? First, compare the coefficients on the information treatment group and on the information + incentive treatment group. The two coefficients are statistically different at the 10% level in column 1 of Panel A. As we include the interaction terms from column 2 through the last column, the difference between the two coefficients become smaller. In column 4 of Panel 4, for example, the point estimates of the two coefficients become very close. These results imply that the difference in the selection on elasticity between these two treatment groups can be well explained by the mechanism through the associations between the price elasticity and the expected gains from switching as well as risk preferences. That is, compared to the intervention of information provision

only, the information + incentive treatment attracted marginal compliers, who are different from the compliers of the information provision treatment group in terms of these consumer characteristics. For example, in Figure 6, we find that the information + incentive treatment incentivized structural losers to switch. Our findings suggest that the primary reason why we see different selection on elasticity between these two groups is that structural losers are more likely to be relatively price-elastic.

In contrast, the differences in the price elasticity between the baseline treatment group and the other two groups do not completely disappear when we include the interaction variables. That is, we cannot fully explain the different selection on elasticity between these groups solely by the indirect selection mechanism given the covariates we have. It implies that the information provision treatment was likely to have a direct effect on the selection, which induced relatively less elastic consumers to switch.

In sum, the reduced-form analysis in this section and the previous section provides key findings and insights on the discrete choice of tariff and the subsequent continuous choice on electricity demand. We find that i) the treatment effect (the change in electricity usage in response to the price change) is heterogeneous among consumers (otherwise the TOT estimates should not differ among the three treatments), ii) providing information and upfront incentives increased switching rates but also induced the selection of inelastic consumers, iii) the selection on elasticity can be partly explained by the indirect selection through the relationship between consumer characteristics and price elasticity, and iv) as a result the intention to treat estimates are statistically indifferent between the baseline treatment group and the information treatment group.

An advantage of the reduced-form analysis is that it provides transparent analysis without imposing a specific structure on consumers' utility functions and their optimization. A limitation is that we cannot estimate the magnitudes of information frictions and inertia without having a structure on the utility maximization. Another limitation is that we cannot use reduced-form analysis to jointly estimate the discrete choice and the continuous choice, while our model in Section 2 suggests that consumer preferences enter both of the decisions, and consumers jointly make the discrete and continuous decisions. In the next section, we structurally estimate the discrete-continuous choice model developed in section 2.

6 Structural Estimation of a Discrete-Continuous Choice Model

The results from the previous sections suggest that customers self-select into the new tariff based on observed heterogeneity and unobserved heterogeneity. In this section, we examine how these two types of selection—selection on unobservables such as the price elasticity and selection on observables such as the expected gains from switching—interact with information frictions and inertia in plan choice. Our estimation framework is similar to Bento et al. (2009); Einav et al. (2013) in the sense that we aim to identify potentially heterogeneous consumer preferences by jointly estimating their discrete and continuous choices. The primary difference between our estimation and discrete-continuous choice estimation in previous studies is that we allow consumers to have information frictions and inertia and that we estimate these switching frictions by exogenous variation created by our RCT.

6.1 Parameterization

Second stage (usage choice).—In Section 2, we show that the demand function conditional on tariff choice j is characterized by $\ln x^*_{jh} = \alpha_h + \epsilon_h \ln p_{jh}$ for tariff j and hour h, with the the second-stage utility function in equation (5). In this section, we add subscript i to distinguish individual-specific variables and parameters from others. We also use subscript t to indicate each of half an hour interval data points. Our empirical specification for the demand estimation is,

$$\ln x^*_{ijt} = \alpha_{ih} + \epsilon_{ih} \ln p_{jh} + \lambda_t + \eta_{it}. \tag{12}$$

λ_t is time-fixed effects, which absorb demand shocks common to all consumers, such as changes in weather. We include an additive error term $\eta_{it} \sim \mathcal{N}(\mu_\eta, \sigma_\eta)$ to capture consumer-time specific unobservable factor that cannot be explained by the demand function.

We allow a random coefficient for the price elasticity, $\epsilon_{ih} \sim -\ln \mathcal{N}(\mu_\epsilon + z_i\beta_\epsilon, \sigma_\epsilon)$. That is, $-\epsilon_{ih}$ (note that we expect that the elasticity is negative) is log-normally distributed with mean $\mu_\epsilon + z_i\beta_\epsilon$ and standard deviation σ_ϵ. $z_i\beta_\epsilon$ is the mean-shifter, which explains how the price elasticity differs by observable variables z_i. For z_i, we use the expected gains from switching (g_i) and household income (y_i). That is, $z_i = (g_i, y_i)$, and $\beta_\epsilon = (\beta_{\epsilon g}, \beta_{\epsilon y})'$. These stochastic terms gives the likelihood

of observing that $x_{it} = x^*_{it}$ given tariff choice j, which we denote $\Pr(x_{it} = x^*_{it} | j \text{ chosen})$

First stage (tariff choice).—As we showed in Section 2, the conditional utility function—$u^*_{ij}(p_{j1}, p_{j2}; \theta_i)$—can be obtained by inserting x^*_{ijt} into the utility function. Then, the consumer's discrete choice of tariff switching can be described by equation (4). Note that $u^*_{ij}(p_{j1}, p_{j2}; \theta_i)$ is the utility function described by our specification of the utility function and x^*_{ijt}. One way to proceed is to interpret $u^*_j(p_{j1}, p_{j2}; \theta_i)$ to be each consumer's actual utility and not to include additive error terms to the utility function (Einav et al., 2013). Another way—this is more common in typical discrete choice models—is to include an additive error term to $u^*_j(p_{j1}, p_{j2}; \theta_i)$. That is, we could model that the actual indirect utility is $u^*_j(p_{j1}, p_{j2}; \theta) + \nu_{ij}$. In the context of electricity tariff choice, either approach seems to be reasonable. Electricity is homogeneous good, and consumers do not have different service or attributes between different tariffs expect for prices. Therefore, ν_{ij} can be less important than other contexts such as discrete choices of differentiated products, in which there is more likely to be unobservable factors between choices. Nonetheless, we include ν_{ij} to estimate our model, by assuming $\nu_{ij} \sim Logistic$ so that $\nu_{i0} - \nu_{i1} \sim EV(1)$.

We leverage our RCT to identify the parameters in the structural model. Our empirical specification of equation (4) is that consumer i switches to the new tariff if:

$$\Delta u_i \equiv u^*_{i1}(p_{11}, p_{12}; \theta_i) - \iota_i \cdot Z_1 - (\delta_i - 60 \cdot Z_3) - u^*_{i0}(p_{01}, p_{02}; \theta_i) > \nu_{i0} - \nu_{i1}. \tag{13}$$

Recall that $Z_1 = 1(i \in \text{baseline treatment group})$. The other two treatment groups ($Z_2 = 1$ or $Z_3 = 1$) received information intervention. The term $\iota_i \cdot Z_1$ indicates that we identify ι by using this exogenous variation induced by the experimental intervention. The two treatment groups ($Z_2 = 1$ or $Z_3 = 1$) . The information + incentive treatment group ($Z_3 = 1$) received a cash incentive ($60) upon switching to the new tariff. We interpret that this monetary incentive causes a level shift of monetized value of inertia (δ) by $60. This is why we include $(\delta_i - 60 \cdot Z_3)$ as a term for inertia. This exogenous shift—induced by the experimental variation—helps identify parameter δ.

We allow a random coefficient for the information frictions and inertia by $\iota_i \sim \ln \mathcal{N}(\mu_\iota + z_i \beta_\iota, \sigma_\iota)$ and $\delta_i \sim \ln \mathcal{N}(\mu_\delta + z_i \beta_\delta, \sigma_\delta)$. The same structures of the mean shifters are the same as that for

the random-coefficient term for the price elasticity. That is, $z_i = (g_i, y_i)$, which means that we estimate how the expected gains from switching (g_i) and household income (y_i) can explain the heterogeneity of ι and δ.

6.2 Estimation

Our model can be estimated by maximum simulated likelihood (MSM) estimation or Bayesian approaches such as Markov Chain Monte Carlo (MCMC) with Gibbs sampling. It is known that the estimators from the two approaches have the same properties asymptotically. In our context, either approach seems to be an equally reasonable procedure. Bayesian approaches can be appealing in terms of estimation speed, but our situation does not involve many choice sets so that the integrals of the maximum likelihood can be simulated with fairly reasonable numbers of simulation. For this reason, we use maximum simulated likelihood estimation with scrambled randomized Halton draws.

The likelihood function for our discrete-continuous choice model is,

$$\mathcal{L} = \prod_{i=1}^{N} \prod_{j=0}^{1} \prod_{t=1}^{T} \Big[\Pr_i(\text{choose } j) \cdot \Pr(x_{it} = x_{it}^* | j \text{ chosen}) \Big]^{d_j}, \tag{14}$$

where $d_j = 1(i \text{ is on tariff } j)$, $\Pr_i(\text{choose } 1) = \Pr(\Delta u_i > 0)$, and $\Pr_i(\text{choose } 0) = \Pr(\Delta u_i \leq 0)$. That is, the likelihood funcqtion is consist of two elements—i) the likelihood of the discrete choice described by $\Pr_i(\text{choose } j)$, and ii) the likelihood of the continuous choice conditional on the discrete choice described by $\Pr(x_{it} = x_{it}^* | j \text{ chosen})$.

6.3 Estimation Results

Table 8 shows the estimates of the discrete-continuous choice model. The mean of price elasticity for peak hours (e_{i,h_1}) is -0.13. This estimate is fairly close to the price elasticity of residential electricity demand estimated in the recent literature (Wolak, 2011; Ito, 2014; Jessoe and Rapson, 2014; Ito, Ida and Tanaka, 2015). The interaction terms imply that consumers that have larger expected gains from switching and those who have higher income are less price-elastic.[6] The standard deviation indicates that there is substantial unobserved heterogeneity in elasticity. For off-peak hours, we do not find statistically significant price elasticity estimates. One potential reason for the result is

[6] Note that the interaction variables are demeaned in the estimation.

that the price change for off-peak hours was not large in our experiment, and consumers did not respond to that small price change. It implies that we do not find "consumption shifting" between peak and off-peak hours, which is consistent with the findings in Ito, Ida and Tanaka (2015).

[Table 8 about here]

Column 3 shows that the mean of the information friction (ι) is \$74.9. We reject the hypothesis that consumers are fully informed ($\iota = 0$) at the 1% statistical significance level. The information friction is positively related to the expected gains from switching. That is, consumers with larger expected gains from switching—structural winners—have larger information frictions. The information friction is negatively associated with household income, which implies that higher-income households have lower information frictions.

Column 4 presents that the inertia, an implicit switching friction aside from information frictions—-(δ) is \$67.9. The magnitude of inertia is negatively related to the expected gains from switching. It is also negatively associated with household income, which implies that higher-income households have lower inertia at their switching decisions.

[Figure 7 about here]

In Figure 7, we test selection on price elasticity using the estimates from the structural model. We obtain an estimate of the consumer-level price elasticity from the model and compute the average switching rate for each of the quantile of the price elasticity. First, the result for the baseline treatment group suggests that price-elastic consumers are more likely to self-select into the new tariff than price-inelastic consumers. Note that the price elasticity is negative so that the lower quantiles imply consumers with more elastic price elasticities.

Second, consumers who received the information provision exhibit different selection on elasticity in the sense that inelastic consumers are more likely to switch. Consistent with our finding in the reduced-form analysis, part of the reason for this relationship is that consumers who have larger expected gains from switching, who are more likely to be induced to switch by the information provision, are less price-elastic.

Third, those who received the information provision and upfront incentive exhibit similar selection on elasticity to the one for the baseline group, but the slope of the line—the relationship

between the switching rates and the price elasticity—is less steep compared to the baseline treatment group. One interpretation of this result is that the information provision generates selection of inelastic consumers because it attracts structural winners, but providing upfront cash incentive also attracts structural lowers as well. These two effects generate the downward-sloping but less steep slope for the information + incentive treatment group.

7 Conclusion and Further Work

We develop a discrete-continuous choice model to characterize the link between plan choice, switching frictions, and subsequent continuous choice of service utilization. We then test the model predictions by using a randomized controlled trial in electricity tariff choice. We find that both information frictions and inertia prevent consumers from switching to a tariff that is privately and socially beneficial. While interventions to mitigate these frictions increased overall switching rates, they also incentivized relatively inelastic consumers to switch. This draft leaves a few more analyses for our further work. One of them is welfare analysis based on the estimates of the discrete-continuous choice model. Such analysis could highlight the policy implications of information provision and upfront incentives in terms of consumer welfare and social welfare, which are key measures to consider the optimal rate design in the presence of switching frictions.

References

Benartzi, Shlomo, and Richard H. Thaler. 2007. "Heuristics and biases in retirement savings behavior." *The journal of economic perspectives*, 81–104.

Bento, Antonio M., Lawrence H. Goulder, Mark R. Jacobsen, and Roger H. Von Haefen. 2009. "Distributional and efficiency impacts of increased US gasoline taxes." *The American Economic Review*, 667–699.

Borenstein, Severin. 2013. "Effective and Equitable Adoption of Opt-In Residential Dynamic Electricity Pricing." *Review of Industrial Organization*, 42(2): 127–160.

Callen, Michael, Mohammad Isaqzadeh, James D. Long, and Charles Sprenger. 2014. "Violence and Risk Preference: Experimental Evidence from Afghanistan." *American Economic Review*, 104(1): 123–48.

Cappers, Peter, Meredith Fowlie, Steve George, Anna Spurlock, Annika Todd, Michael Sullivan, and Catherine D.Catherine Wolfram. 2015. "Default Bias, FollowâĂŘOn Behavior and Welfare in Residential Electricity Pricing Programs." *Working Paper*.

Dubin, Jeffrey A., and Daniel L. McFadden. 1984. "An Econometric Analysis of Residential Electric Appliance Holdings and Consumption." *Econometrica*, 52(2): 345–62.

Einav, Liran, Amy Finkelstein, Stephen Ryan, Paul Schrimpf, and Mark R. Cullen. 2013. "Selection on Moral Hazard in Health Insurance." *The American economic review*, 103(1): 178.

Faruqui, Ahmad, and Sanem Sergici. 2011. "Dynamic pricing of electricity in the mid-Atlantic region: econometric results from the Baltimore gas and electric company experiment." *Journal of Regulatory Economics*, 40(1): 82–109.

Handel, Benjamin R. 2013. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review*, 103(7): 2643–82.

Handel, Benjamin R., and Jonathan T. Kolstad. 2013. "Health Insurance for "Humans": Information Frictions, Plan Choice, and Consumer Welfare." National Bureau of Economic Research Working Paper 19373.

Hanemann, W. Michael. 1984. "Discrete/Continuous Models of Consumer Demand." *Econometrica*, 52(3): 541–561.

Hortaçsu, Ali, Seyed Ali Madanizadeh, and Steven L. Puller. 2015. "Power to Choose? An Analysis of Consumer Inertia in the Residential Electricity Market." *NBER Working Paper*, , (20988).

Ito, Koichiro. 2014. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." *American Economic Review*, 104(2): 537–63.

Ito, Koichiro, Takanori Ida, and Makoto Tanaka. 2015. "The Persistence of Intrinsic and Extrinsic Motivation: Experimental Evidence from Energy Demand." *NBER Working Paper-Working Paper*, 20910.

Jessoe, Katrina, and David Rapson. 2014. "Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use." *American Economic Review*, 104(4): 1417–38.

Joskow, Paul L. 2012. "Creating a Smarter U.S. Electricity Grid." *Journal of Economic Perspectives*, 26(1): 29–48.

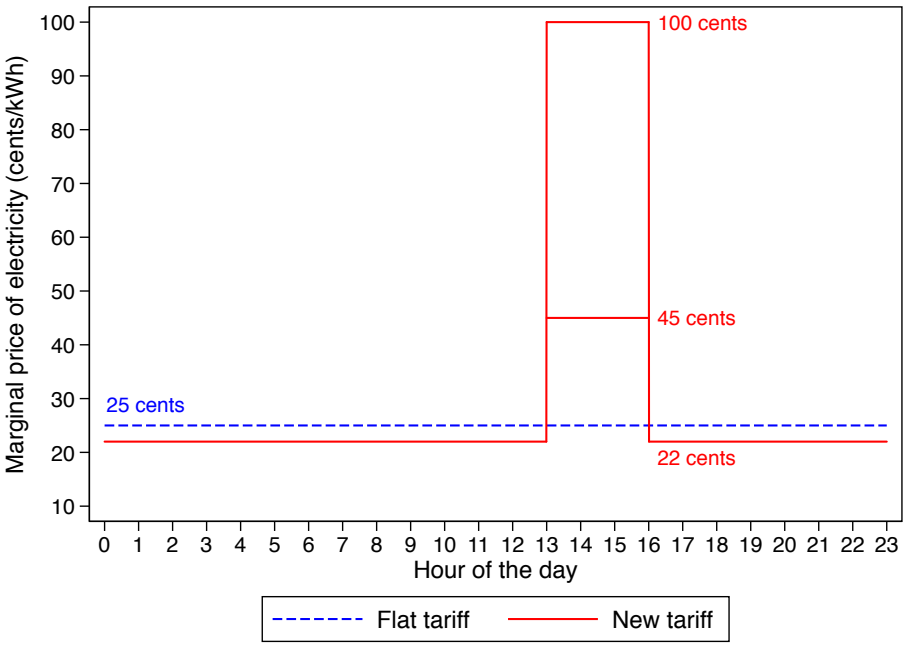
Joskow, Paul L., and Catherine D. Wolfram. 2012. "Dynamic Pricing of Electricity." *American Economic Review*, 102(3): 381–85.

Miravete, Eugenio J. 2003. "Choosing the Wrong Calling Plan? Ignorance and Learning." *American Economic Review*, 93(1): 297–310.

Wolak, Frank A. 2006. "Residential customer response to real-time pricing: the Anaheim critical-peak pricing experiment." *Working Paper*.

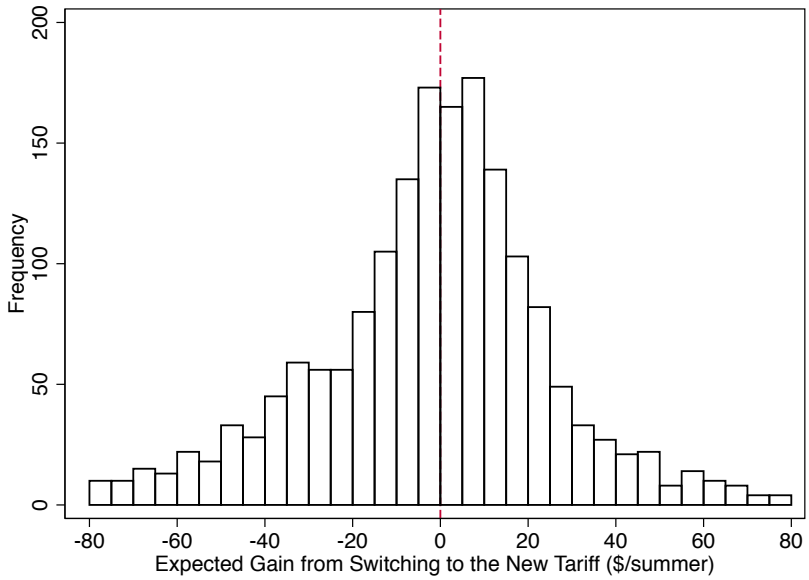
Wolak, Frank A. 2011. "Do Residential Customers Respond to Hourly Prices? Evidence from a Dynamic Pricing Experiment." *The American Economic Review*, 101(3): 83–87.

Figure 1: Electricity Rate Plans



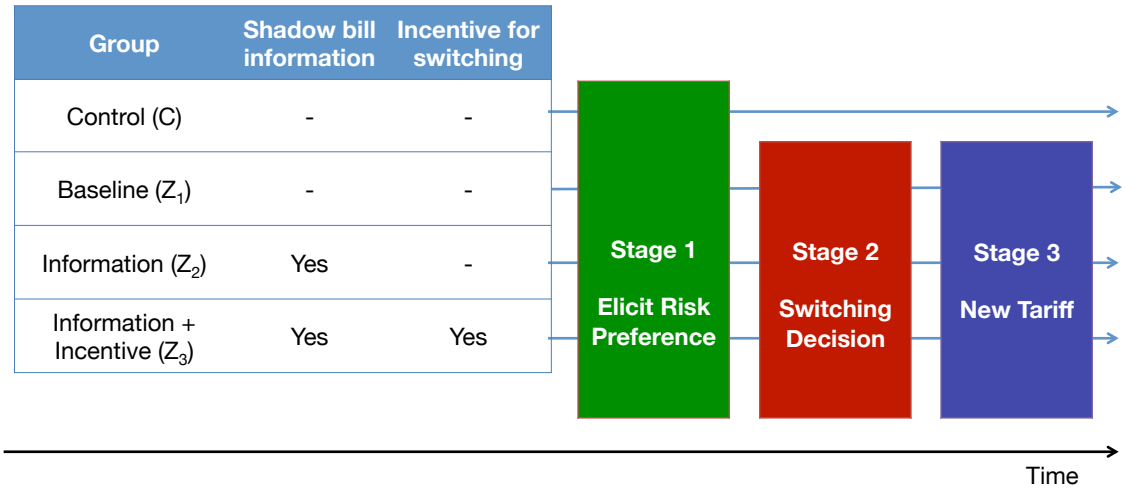
Note: This figure shows electricity tariff schedules for the default flat tariff and the new tariff. Customers who switched to the new tariff had 22 cents/kWh in off-peak hours and either 45 or 100 cents/kWh for peak hours, depending on the day. Peak hours were between 1 pm and 4 pm on weekdays. Customers on the new tariff received day-ahead and same-day notices about their peak hour prices.

Figure 2: The Distribution of Expected Gains from Switching



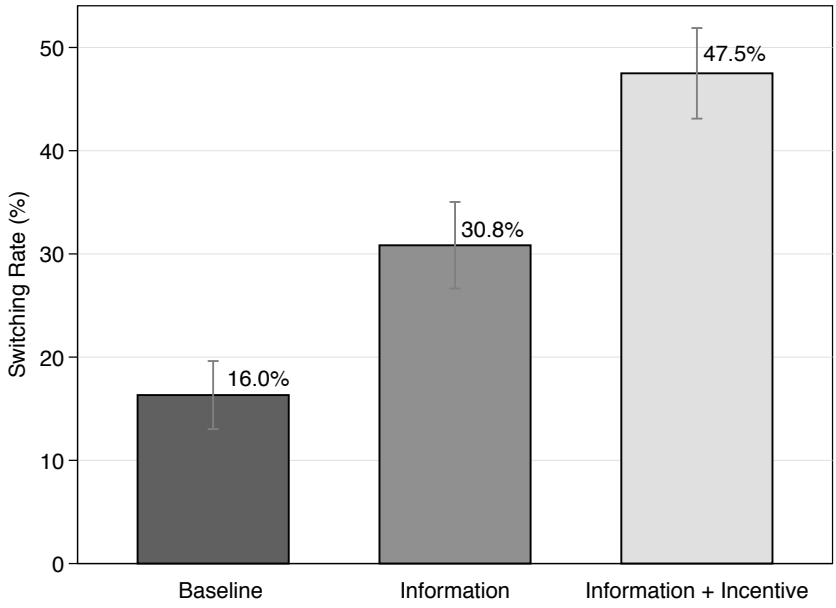
Note: This figure shows the histogram of expected gains from switching to the new tariff. The expected gains were calculated based on each customer's past usage, assuming zero price elasticity.

Figure 3: Timeline of the Experiment



Note: After we randomly assigned customers into four groups, we elicited risk preference parameters from all groups, allowed customers in the three treatment groups to make switching decisions, and implemented dynamic pricing for customers who switched to the new tariff.

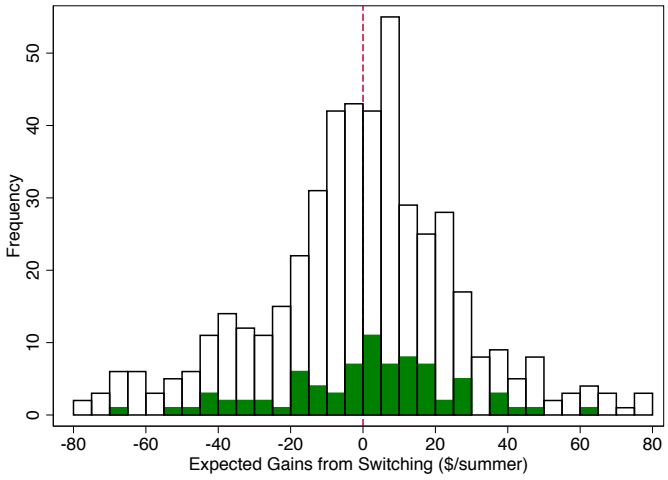
Figure 4: Switching Rate by Treatment Group



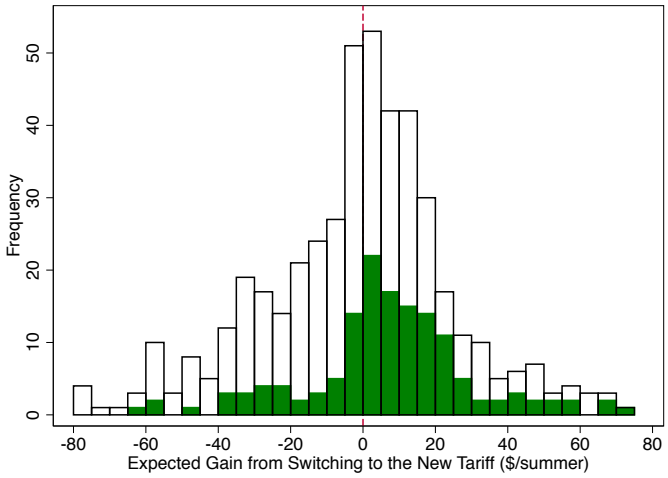
Note: The figure shows the percentage of customers who switched to the new price schedule for each treatment group. The bars are 95% confidence intervals.

Figure 5: Switching Decision by Expected Gains from Switching

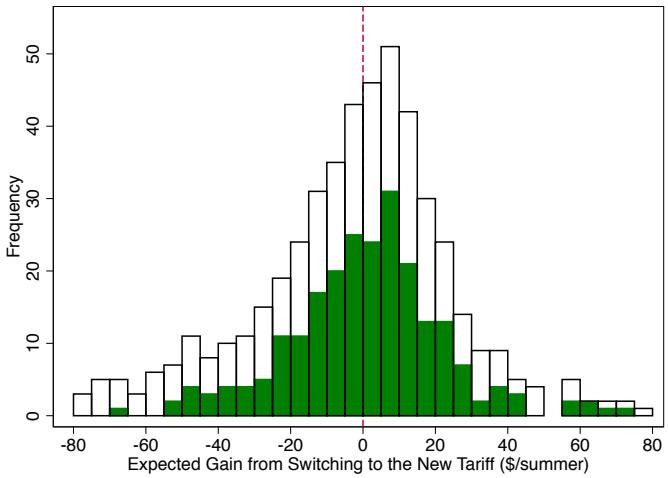
Panel A: Baseline treatment group (Z_1)



Panel B: Information treatment group (Z_2)

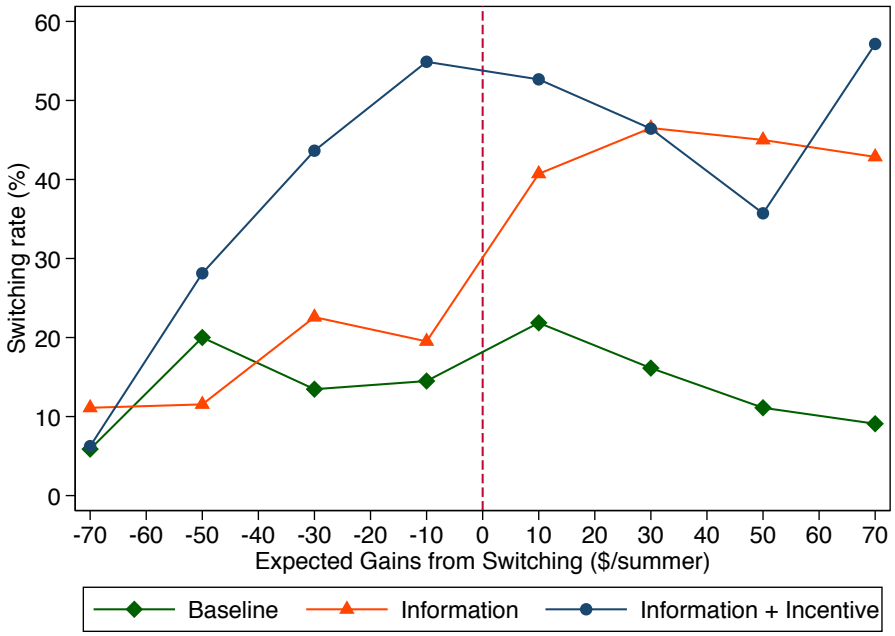


Panel C: Information + incentive treatment group (Z_3)



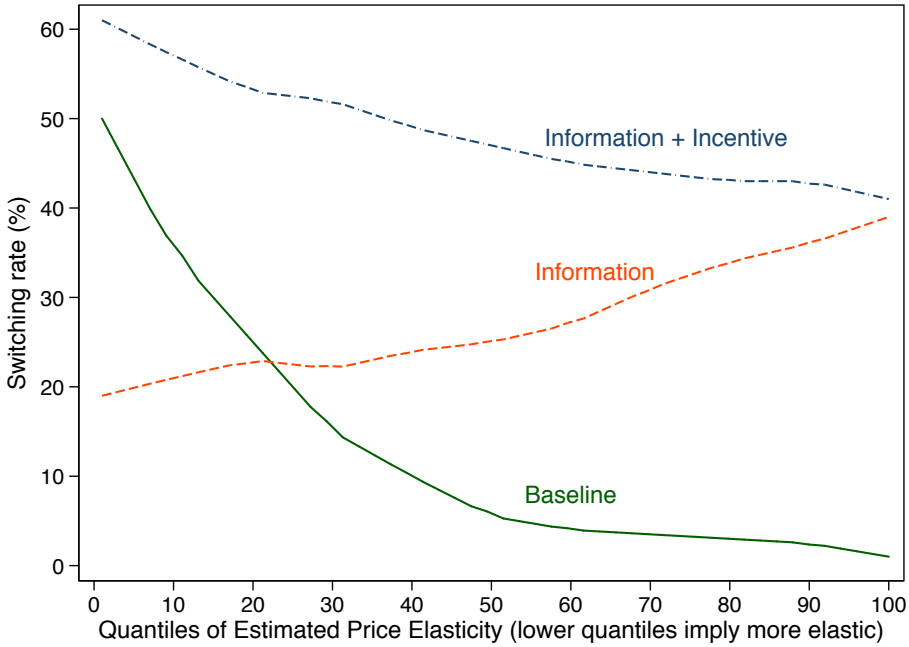
Note: Each histogram shows the distribution of customers over their expected gains from switching. The dark filled color indicates those who switched to the new price plan. The expected gains were calculated based on each customer's past usage, assuming zero price elasticity.

Figure 6: Switching Rate by Expected Gains from Switching



Note: The figure shows the percentage of customers who switched to the new price schedule for each treatment group by their expected gains from switching. The expected gains were calculated based on each customer's past usage, assuming zero price elasticity.

Figure 7: Selection on Elasticity ($e_{i,h1}$) by Treatment Groups



Note: This figure shows the switching rates over the quantiles of the price elasticity for peak hours that is estimated by the discrete-continuous choice model for each of the three treatment groups. Note that the price elasticity is negative so that smaller quantiles imply more price-elastic.

Table 1: Experimental Design

Group	Eligible to switch tariff	Information provision	Incentive for switching	Number of customers
Control (C)				697
Baseline treatment (Z_1)	✓			486
Information treatment (Z_2)	✓	✓		468
Information + incentive treatment (Z_3)	✓	✓	✓	502

Notes: This table shows difference in treatments in four groups that were randomly assigned.

Table 2: Summary Statistics by Groups

	Control (C)	Baseline treatment (Z_1)	Information treatment (Z_2)	Information + incentive (Z_3)
Household income (JPY10,000)	759.72	768.31	768.18	755.97
	(330.78)	(340.11)	(343.46)	(334.23)
Square Meters	106.36	107.64	111.44	107.51
	(29.99)	(30.49)	(31.09)	(30.64)
Number of Bedrooms	3.68	3.75	3.84	3.76
	(1.10)	(1.17)	(1.15)	(1.09)
Age of Building	14.02	14.58	15.59	15.14
	(11.41)	(11.31)	(11.68)	(11.39)
Number of room AC	3.06	3.24	3.17	3.10
	(1.39)	(1.37)	(1.33)	(1.34)
Number of TV	1.97	2.06	2.12	2.07
	(1.03)	(1.10)	(1.05)	(1.10)
Number of refrigerator	1.13	1.18	1.13	1.12
	(0.43)	(0.45)	(0.38)	(0.37)
Electricity usage (kWh/day)	13.11	13.04	13.09	13.02
	(5.56)	(5.71)	(5.58)	(5.71)

Notes: This table shows the mean and standard deviation of each variable.

Table 3: Elicitation of Risk Preference

Panel A: First set of questions to obtain q

q	Option A	Option B
0.1	\$100	10% chance of \$300, 90% chance of \$0
0.2	\$100	20% chance of \$300, 80% chance of \$0
0.3	\$100	30% chance of \$300, 70% chance of \$0
0.4	\$100	40% chance of \$300, 60% chance of \$0
0.5	\$100	50% chance of \$300, 50% chance of \$0
0.6	\$100	60% chance of \$300, 40% chance of \$0
0.7	\$100	70% chance of \$300, 30% chance of \$0
0.8	\$100	80% chance of \$300, 20% chance of \$0
0.9	\$100	90% chance of \$300, 10% chance of \$0
1	\$100	100% chance of \$300, 0% chance of \$0

Panel B: Second set of questions to obtain q'

q'	Option A	Option B
0.1	50% chance of \$300, 50% chance of \$0	10% chance of \$300, 90% chance of \$0
0.2	50% chance of \$300, 50% chance of \$0	20% chance of \$300, 80% chance of \$0
0.3	50% chance of \$300, 50% chance of \$0	30% chance of \$300, 70% chance of \$0
0.4	50% chance of \$300, 50% chance of \$0	40% chance of \$300, 60% chance of \$0
0.5	50% chance of \$300, 50% chance of \$0	50% chance of \$300, 50% chance of \$0
0.6	50% chance of \$300, 50% chance of \$0	60% chance of \$300, 40% chance of \$0
0.7	50% chance of \$300, 50% chance of \$0	70% chance of \$300, 30% chance of \$0
0.8	50% chance of \$300, 50% chance of \$0	80% chance of \$300, 20% chance of \$0
0.9	50% chance of \$300, 50% chance of \$0	90% chance of \$300, 10% chance of \$0
1	50% chance of \$300, 50% chance of \$0	100% chance of \$300, 0% chance of \$0

Notes: Subjects were asked to choose option A or option B for each question, which allows us to obtain q and q'. This method is originally developed by Callen et al. (2014).

Table 4: Risk Preferences and Switching Decisions

Dependent Variable: $Switch_i \equiv 1(i \in$ switched to the new tariff)

	What variable is used for $RiskPref_i$?		
	Risk aversion (q)	Risk aversion (q')	Certainty premium
$RiskPref_i$	-0.33	-0.51	-0.16
	(0.08)	(0.06)	(0.07)
$RiskPref_i$ x Group 2 (SB)	0.19	0.25	0.17
	(0.03)	(0.04)	(0.09)
$RiskPref_i$ x Group 3 (SB + Incentive)	0.38	0.45	0.09
	(0.03)	(0.03)	(0.08)
N	1331	1331	1331

Notes: This table shows the estimation results of equation (8). The dependent variable is the binary variable that equals to one if the consumer switched to the new tariff. The independent variables in columns 1-3 are risk parameter q, risk parameter q', and certainty premium cp. Treatment group fixed effects are included.

Table 5: Treatment Effect on the Treated (TOT)

Dependent Variable: $\ln x_{it} \equiv$ the natural log of electricity usage in a 30-minute interval

	All Treatment Days	Treatment Days with Price = 100	Treatment Days with Price = 45
Baseline (α_1)	-0.220	-0.318	-0.185
	(0.058)	(0.070)	(0.057)
Information (α_2)	-0.090	-0.164	-0.064
	(0.031)	(0.038)	(0.030)
Information + Incentive (α_3)	-0.132	-0.185	-0.114
	(0.020)	(0.025)	(0.020)
Observations	841180	435030	699154

Notes: This table shows the estimation results of equation (9). The dependent variable is the natural log of electricity usage in a 30-minute interval. All regressions include household fixed effects and time fixed effects. We cluster the standard errors at the household level to adjust for serial correlation.

Table 6: Intention to Treat (ITT)

Dependent Variable: $\ln x_{it} \equiv$ the natural log of electricity usage in a 30-minute interval

	All Treatment Days	Treatment Days with Price = 100	Treatment Days with Price = 45
Baseline (β_1)	-0.037	-0.053	-0.031
	(0.010)	(0.012)	(0.010)
Information(β_2)	-0.029	-0.051	-0.020
	(0.010)	(0.012)	(0.010)
Information + Incentive (β_3)	-0.066	-0.092	-0.057
	(0.010)	(0.012)	(0.010)
Observations	841180	435030	699154

Notes: This table shows the estimation results of equation (10). The dependent variable is the natural log of electricity usage in a 30-minute interval. All regressions include household fixed effects and time fixed effects. We cluster the standard errors at the household level to adjust for serial correlation.

Table 7: Mechanism of Selection on Elasticity: Heterogeneity in the Treatment Effect on the Treated

Panel A: Interaction with the Expected Gains from Switching (in $) and risk preferences

	(1)	(2)	(3)	(4)	(5)
Baseline	-0.220	-0.220	-0.316	-0.250	-0.262
	(0.058)	(0.058)	(0.051)	(0.056)	(0.051)
Information	-0.090	-0.103	-0.166	-0.121	-0.134
	(0.031)	(0.032)	(0.043)	(0.032)	(0.044)
Information + Incentive	-0.132	-0.132	-0.184	-0.135	-0.148
	(0.020)	(0.020)	(0.037)	(0.020)	(0.038)
Treatment x Expected Gains		0.015	0.010	0.012	0.012
		(0.006)	(0.006)	(0.006)	(0.006)
Treatment x Risk Aversion			0.085		0.021
			(0.059)		(0.062)
Treatment x Certainty Premium				0.185	0.180
				(0.042)	(0.043)
Observations	841180	761116	761116	761116	

Panel B: Interaction with the Expected Gains from Switching (by quartile) and risk preferences

	(1)	(2)	(3)	(4)
Baseline	-0.280	-0.367	-0.310	-0.322
	(0.068)	(0.061)	(0.067)	(0.062)
Information	-0.162	-0.219	-0.183	-0.196
	(0.048)	(0.054)	(0.048)	(0.055)
Information + Incentive	-0.192	-0.233	-0.193	-0.205
	(0.036)	(0.046)	(0.036)	(0.047)
Treatment x 2nd Quartile of Expected Gains	0.058	0.035	0.040	0.040
	(0.038)	(0.038)	(0.038)	(0.038)
Treatment x 3rd Quartile of Expected Gains	0.032	0.038	0.047	0.047
	(0.037)	(0.036)	(0.037)	(0.037)
Treatment x 4th Quartile of Expected Gains	0.146	0.131	0.143	0.143
	(0.037)	(0.037)	(0.037)	(0.037)
Treatment x Risk Aversion		0.081		0.019
		(0.059)		(0.061)
Treatment x Certainty Premium			0.180	0.176
			(0.042)	(0.043)
Observations	841180	761116	761116	761116

Notes: This table shows the estimation results of equation (11). The dependent variable is the natural log of electricity usage in a 30-minute interval. All regressions include household fixed effects and time fixed effects. We cluster the standard errors at the household level to adjust for serial correlation.

Table 8: Structural Estimates of the Discrete and Continuous Choice Model

	Elasticity in peak hours (ϵ_{i,h_1})	Elasticity in off-peak hours (ϵ_{i,h_0})	Information frictions (ι_i)	Inertia (δ_i)
Mean	-0.134	-0.008	74.93	67.96
	(0.051)	(0.027)	(13.94)	(8.72)
Interaction with	0.0034	0.0004	0.77	-0.69
Expected Gain ($)	(0.0012)	(0.0008)	(0.37)	(0.19)
Interaction with	0.0057	0.0007	-15.84	-9.10
Log Household Income	(0.0021)	(0.0015)	(6.55)	(3.80)
Std. Dev.	0.198	0.083	11.34	13.6
	(0.032)	(0.028)	(5.04)	(4.23)

Notes: This table shows the estimation results of the discrete continuous choice model in equation (14).