

Selection on Welfare Gains: Experimental Evidence from Electricity Plan Choice*

Koichiro Ito Takanori Ida Makoto Tanaka
University of Chicago and NBER Kyoto University GRIPS

Abstract

We study a problem in which policymakers need to screen self-selected individuals by unobserved heterogeneity in social welfare gains from a policy intervention. In our framework, the marginal treatment effects and marginal treatment responses arise as key statistics to characterize social welfare. We apply this framework to a randomized field experiment on electricity plan choice. Consumers were offered welfare-improving dynamic pricing with randomly assigned take-up incentives. We find that price-elastic consumers—who generate larger welfare gains—are more likely to self-select. Our counterfactual simulations quantify the optimal take-up incentives that exploit observed and unobserved heterogeneity in selection and welfare gains.

Keywords: Selection, Roy Model, Electricity, Marginal Treatment Effect, Sufficient Statistics

*Ito: Harris School of Public Policy, University of Chicago, 1307 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 would like to thank Andrew Smith and Hideki Shimada for excellent research assistance, and Naoki Aizawa, Vivek Bhattacharya, Dan Black, Severin Borenstein, Meghan Busse, Raj Chetty, Steve Cicala, Lucas Davis, Manasi Deshpande, Amy Finkelstein, Meredith Fowlie, Peter Ganong, Joshua Gottlieb, Tatsuo Hatta, Peter Hull, Gaston Illanes, Tetsuya Kaji, Ryan Kellogg, Erin Mansur, Neale Mahoney, Magne Mogstad, Sam Norris, Matt Notowidigdo, Robert Porter, Steve Puller, Mar Reguant, Yuya Sasaki, Hitoshi Shigeoka, Charles Sprenger, Kensuke Teshima, Alexander Torgovitsky, Chris Walters, Catherine Wolfram, Frank Wolak, Matthew Zaragoza-Watkins, and seminar participants at UC Berkeley Energy Camp, Summer Juku at Stanford, University of Colorado Boulder, Arizona State, Midwest Energy Fest, NBER Summer Institute, NBER Japan Meeting, University of Pittsburg, National University of Singapore, Duke, UBC, Vanderbilt, Chicago Booth, Northwestern, Carleton University, Kellogg School of Management, University of Wisconsin at Madison and ASSA Annual Meeting for their 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 for this study. We thank the New Energy Promotion Council for financial support.

Selection is a key phenomenon in economic policies because voluntary take-up is a ubiquitous feature of policy design. For example, social safety net programs in the United States, such as food stamps and disability insurance programs, require eligible individuals to apply for benefits voluntarily (Finkelstein and Notowidigdo, 2019; Deshpande and Li, 2019). A similar design is widely used in other policy domains, including worker training programs (LaLonde, 1986), childcare programs (Cornelissen et al., 2018), medicare payment programs (Einav et al., 2022), energy efficiency rebates (Borenstein and Davis, 2016; Allcott and Greenstone, 2017), and insurance and utility service plans (Einav et al., 2013; Handel, 2013; Hortaçsu, Madanizadeh and Puller, 2017; Fowlie et al., forthcoming).

In the presence of selection, a policy’s social welfare gains depend on what we call *selection on welfare gains*. Selection is governed by each individual’s private gain from the policy. This private gain itself is an important outcome, but what matters to policymakers is the social welfare gain that can be generated by each individual’s take-up and behavioral response to the policy intervention. If this welfare gain is heterogeneous among individuals, the connection between private selection and social welfare gains plays an important role in policy design.

We formalize this idea by developing a framework that connects heterogeneity in private gains and heterogeneity in social welfare gains. Our framework builds on the principles of welfare economics and the generalized Roy model developed by Eisenhauer, Heckman and Vytlačil (2015). An individual selects into treatment based on a selection equation, which is a function of observables and unobservables. The Roy model allows us to characterize heterogeneity in potential outcomes using the marginal treatment effect (MTE) and the marginal treatment responses (MTRs), as shown by Mogstad, Santos and Torgovitsky (2018). We combine this framework with the principles of welfare economics to demonstrate that heterogeneity in social welfare gains can be expressed by a function of the selection equation and the MTRs. With this social welfare function in hand, we can examine a policy’s social welfare gain in the presence of selection.

An advantage of this framework is that it clarifies the parameters that need to be estimated through empirical analysis. With this insight, we designed a randomized controlled trial (RCT) in the field to study electricity plan choice, where the policy goal was to generate social welfare gains from the adoption of dynamic electricity pricing. Similar to nearly all households in the United States and the rest of the world, consumers in our experiment had a non-dynamic electricity price

(i.e., the price did not vary by hour). This commonly used pricing is socially inefficient because the marginal cost of electricity tends to vary substantially by hour (Joskow and Wolfram, 2012). To address this inefficiency, we offered customers a dynamic pricing plan. In addition, we offered them a randomly assigned financial incentive for take-up. We use this random variation in take-up incentive as an instrument to estimate the MTE and MTRs based on a method developed by Brinch, Mogstad and Wiswall (2017).

We begin by estimating the selection equation. Before the field experiment, we collected demographic information and solicited household-level risk preferences based on the method developed by Callen et al. (2014). We also collected historical hourly usage data at the household level. These data allowed us to calculate each consumer’s expected saving from dynamic pricing, that is, the financial saving from dynamic pricing under the assumption of no behavioral response to price changes. We find that selection is strongly related to the expected savings from dynamic pricing. This provides empirical evidence of “selection on the level” that has been documented in the healthcare markets (Einav et al., 2013). We also find that selection is positively associated with the take-up incentive (Z) and years of schooling, but negatively related to risk aversion, certainty premium, and employment.

The propensity score that is obtained from the selection equation plays a central role in estimating the MTE and MTRs. We estimate the causal effect of price changes on consumption (i.e., the price elasticity of demand). The advantage of the MTE is that it allows us to estimate the relationship between the average treatment effects (ATE) and selection. We use the MTE to test “selection on the slope” (Einav et al., 2013), that is, whether more-elastic consumers are more likely to adopt dynamic pricing. We find strong evidence of selection on the slope for demand in peak hours. The estimated MTE function also indicates that price elasticity diminishes to zero as we approach consumers less likely to adopt dynamic pricing. As the social welfare gain from dynamic pricing is directly tied to the slope of demand, our result implies that the marginal social welfare gain is likely to diminish as we approach consumers less likely to adopt.

We investigate this question by plugging our empirical estimates of the selection equation and MTRs into the social welfare function developed in our framework. This approach allows us to conduct three welfare analyses. First, we quantify the social welfare gains from the policies implemented in our field experiment. We find that the take-up incentive provided in our experiment

increased social welfare.

Second, we conduct counterfactual policy simulations to find the optimal take-up incentive. We find that the welfare gain is increasing in the take-up incentive (Z) up to a certain level, but the relationship is concave for two reasons. First, the estimated MTE function suggests that the treatment effect (i.e., the behavioral responses to the price change) diminishes as we increase the take-up incentive to induce consumers who would not adopt dynamic pricing in the absence of a high take-up incentive. Second, these consumers, who need a larger take-up incentive, are those who have larger unobserved disutility from adopting dynamic pricing. Therefore, the marginal welfare gain from increasing the take-up incentive diminishes to zero.

We show that this is empirically the case and estimate the optimal level of the take-up incentive for three scenarios. The first is the optimal uniform take-up incentive ($Z = z^*$), where policymakers cannot differentiate the incentive by observables. The second one is the optimal differentiated take-up incentives ($Z = z^*(x)$), under which policymakers can differentiate the incentive by observables. Recall that our welfare function is a function of the selection equation and the MTRs. Therefore, the differentiated take-up incentive can exploit variation in observables in both the selection equation and the treatment effects. Finally, we consider such a differentiated take-up incentive based on a restricted set of observables ($Z = z^\dagger(x)$). For example, consumer-level consumption data are typically readily available from the billing systems of utility companies, while consumer-level demographic information can be relatively more difficult to obtain. Thus, we consider a targeting policy based only on historical electricity usage data at the consumer level.

We compare the welfare gains from five policies: i) $Z = 0$, ii) $Z = 60$ (in USD), iii) $Z = z^*$, iv) $Z = z^*(x)$, and v) $Z = z^\dagger(x)$. The welfare gains from the first two policies are calculated directly from the variation in the experiment. Those from the rest of the policies are based on counterfactual policy simulations. We find that $Z = 60$ increases welfare relative to $Z = 0$ and that there are additional welfare gains by setting Z at the optimal level. This welfare gain can be further enhanced by exploiting variation in observables and using the optimal differentiated take-up incentives ($Z = z^*(x)$). Consistent with a theoretical prediction, the differentiated take-up policy with a restricted set of observables ($Z = z^\dagger(x)$) improves welfare compared to the uniform take-up incentive, but the welfare improvement is lower than the case with the differentiated take-up incentive with the full set of observables.

Finally, we use our framework to discuss implications for a mandatory take-up policy. The vast majority of countries rely on voluntary take-up policies for the adoption of dynamic electricity pricing for residential customers because a mandate is politically infeasible. However, it is still useful to discuss what conditions could make a mandate more welfare-enhancing than the counterfactual policies discussed above. A key unknown effect of a mandate is how it changes a consumer’s disutility from adopting dynamic pricing. Since this effect is unknown from our experiment and previous studies, we examine two possibilities. First, we show that if a mandate does not change such disutility, our counterfactual simulation results imply that it is inferior to the optimal take-up incentive policies because additional take-up beyond the optimal level would reduce welfare. Second, if a mandate can reduce disutility for adopting dynamic pricing, we can use our framework to calculate how much such a benefit has to be to ensure a mandate policy is more welfare-enhancing than other policies. Our calculation indicates that a mandate can be superior to other counterfactual policies if its net benefit, including the political cost of implementing it, is more than \$26 per customer per year.

Related literature and our contributions—First, our framework builds on the literature on the MTE and policy-relevant treatment effects (Heckman and Vytlačil, 2001, 2005; Heckman, 2010). The underlying concept in this literature is motivated by Marschak (1953), which states that for many policy analyses, it may not be necessary to identify fully specified structural models. Instead, researchers may be able to address policy questions by estimating a set of relevant structural parameters derived from economic models. Recently, Mogstad, Santos and Torgovitsky (2018) extends this concept by introducing a so-called *target parameter*. A target parameter is any function of MTRs, and therefore, encompasses the MTE. We show that the social welfare function in our study emerges as a target parameter. Our counterfactual policy simulations highlight that this tool can be useful for analyzing not only the private benefits from a policy intervention but also its implications for social welfare.¹

Second, our study is also related to sufficient-statistics approaches (Chetty, 2009; Kleven, 2021), which is another strand of the literature influenced by Marschak (1953). In this stream of the literature, researchers use economic theory to derive a social welfare function with parameters that

¹The spirit of our welfare analysis based on the MTE is similar to Kline and Walters (2016), which evaluates policy counterfactuals of the Head Start program using the MTE.

can be estimated from experimental or quasi-experimental variation in data. Typical sufficient-statistics approaches do not consider selection. In the absence of selection, the *average* treatment effect (ATE) is often sufficient to characterize a policy’s impact on social welfare, even if there is unobserved heterogeneity in treatment effects (Chetty, 2009; Kleven, 2021).² However, in practice, policies often have selection and unobserved treatment heterogeneity. With these two phenomena, the key information researchers need to know is how selection and unobserved treatment heterogeneity are related and how this relationship is linked to a policy’s welfare gains. We show that these two phenomena can be incorporated into the sufficient-statistics framework by characterizing a social welfare gain as a function of the selection equation and MTRs. In this regard, our study contributes to the recent literature that aims to connect welfare analysis to treatment effect heterogeneity (Andrews and Miller, 2013; Kline and Walters, 2016; Finkelstein and Hendren, 2020).

Third, our study is closely related to the recent literature on policy take-up and targeting (Finkelstein and Notowidigdo, 2019; Deshpande and Li, 2019). These studies show that optimal targeting requires knowledge of how selection in take-up is related to individual characteristics. Our study contributes to this body of the literature by showing that policymakers need additional key information—how selection in take-up is related to heterogeneity in the treatment effects and the resulting social welfare gains. To investigate this point, we need to estimate both selection and the MTRs from a policy intervention. Our experimental design allows us to conduct such an analysis, and we show that the connection between selection and heterogeneity in social welfare gains plays a key role in targeting.³

Another closely related stream of the literature is research on selection outside the electricity industry, such as in healthcare sectors. Einav et al. (2013) is the first study to decompose selection in the context of health insurance into two key elements: selection on the *level* and selection on the *slope*. Since then, both types of selection have been recognized as key phenomena in the study of healthcare markets (Einav et al., 2022). Our study highlights that this concept also plays a key role in energy markets. A notable distinction is the implication of selection on the slope. In healthcare markets, selection on the slope is generally considered to be undesirable for social welfare because

²In Chetty (2009), Kleven (2021), and others, treatment effects are already conditioned on observables, and, therefore, the focus is unobserved heterogeneity in treatment effects.

³Conceptually, our targeting strategy is similar to that of Kitagawa and Tetenov (2018), in which targeting is based on social welfare gains from a policy intervention rather than individual characteristics per se. Kitagawa and Tetenov (2018) focus on mandatory policy assignments without selection.

it exaggerates “selection on moral hazard” (Einav et al., 2013). However, such selection can be welfare-enhancing in corrective policies such as the adoption of dynamic electricity pricing. This is because the social planner benefits more from consumers with more-elastic demand in this situation.

Finally, our analysis provides important implications for ongoing discussions in energy policy. Households in the United States and many countries now have smart meters at home that record hourly consumption and, therefore, make dynamic pricing technologically feasible (Joskow, 2012). In fact, recent studies show that there are substantial efficiency gains if residential consumers adopt dynamic pricing (Wolak, 2011; Ito, Ida and Tanaka, 2018). However, the mandate of dynamic pricing is politically infeasible for most countries, and therefore, policymakers need to design a policy with voluntary take-up. Our result suggests that it can be possible to improve welfare by incentivizing consumers to adopt dynamic pricing. However, there is also likely to be a diminishing return to such policies. As we show in our welfare analysis, consumers who are less likely to adopt dynamic pricing are likely to be relatively price-inelastic consumers, who may not generate large social welfare gains from the adoption of dynamic pricing.

1 Background

In this section, we describe fundamental inefficiency in electricity markets, why dynamic pricing can mitigate this problem, and why selection is the key factor to policy design.

1.1 Fundamental Inefficiency in Electricity Markets

Inefficient retail pricing has long been a central issue in electricity markets (Joskow, 2012). In the United States, nearly all households pay electricity prices that do not reflect the marginal cost of electricity (FERC, 2011). The cost of electricity is time-varying because of substantial fluctuations in demand and a lack of storability (Borenstein, 2002). Therefore, a welfare-improving electricity tariff would be dynamic pricing that reflects the time-varying marginal cost.

However, typical residential electricity prices are time-invariant. The wedge between the marginal cost and marginal price creates a textbook example of deadweight loss. Compared with the efficient level, households consume in excess when marginal costs are higher than retail rates—typically in peak demand hours—and too little when marginal costs are lower than retail rates—typically in

off-peak hours. This phenomenon is not unique to the United States and is ubiquitous in most countries, except for a few nations in Europe.

[Figure 1 about here]

Figure 1 illustrates this inefficiency using hypothetical price schedules and demand curves. Consider flat pricing in which the marginal price of electricity does not vary over time and dynamic pricing in which the marginal price reflects the time-varying marginal cost of electricity. The social welfare gain from the adoption of dynamic pricing is Harberger’s triangle in the figure. Importantly, the size of the welfare gain depends on the slope of demand. When a more-elastic consumer adopts dynamic pricing, the welfare gain is $A + B$ in peak hours and $C + D$ in off-peak hours. The welfare gain from a less-elastic consumer is B in peak hours and C in off-peak hours. Thus, the more-elastic consumer would produce a larger welfare gain from adopting dynamic pricing.

1.2 Who Selects into Dynamic Pricing?

The heterogeneity in welfare gains presented in Figure 1 plays a key role in policy design because residential dynamic electricity pricing is usually offered through voluntary take-up. With voluntary take-up, a key question is “what types of consumers are likely to select into dynamic pricing?” Prior to providing a formal model in the next section, we describe potential selection mechanisms.

[Figure 2 about here]

The first potential selection mechanism is *selection on the slope* (Einav et al., 2013). Figure 2 describes this theoretical prediction, showing consumers’ private gains from adopting dynamic pricing with two hypothetical demand curves. When the *more-elastic* consumer adopts dynamic pricing, the private gain (the change in consumer surplus in this figure) is $-E$ in peak hours and $G + H$ in off-peak hours. The private gain for the *less-elastic* consumer is $-(E + F)$ in peak hours and G in off-peak hours. Thus, the more-elastic consumer would have a larger private gain from selecting into dynamic pricing in terms of the consumer surplus.

The second selection mechanism is *selection on the level* (Einav et al., 2013). Consider two consumers who have the same price elasticity, and hence, the same selection on the slope. Suppose that one consumer uses more electricity in off-peak hours, whereas the other consumer uses more

electricity in peak hours. The private gain from dynamic pricing is larger for the consumer who uses more in off-peak hours because dynamic pricing increases the peak price and lowers the off-peak price. Such consumers are called the *structural winners* of dynamic pricing. In Figure 2, selection on the level is related to the *intercepts* of the demand curves, whereas selection on the slope is related to the slopes of the demand curves. As we see in Section 3, a typical dynamic electricity tariff creates a wide distribution of structural winners and losers because of the variation in electricity usage in peak and off-peak hours. Therefore, selection on the level can be another key selection mechanism in our context.

The third mechanism is through the *subjective cost of treatment* (Eisenhauer, Heckman and Vytlacil, 2015). Consider two consumers with identical electricity demand, including their slope and intercept. Then, they have the same selection on the slope and level. However, their selection can still differ for at least two reasons. First, they may have differences in inertia or switching costs to change their electricity plans. Second, they may have differences in unobserved plan preferences that cannot be explained by observable plan characteristics. In the terminology of Eisenhauer, Heckman and Vytlacil (2015), both elements are part of the subjective cost of treatment in the selection equation, as described in the next section. Our framework below incorporates this subjective cost of treatment.⁴

1.3 Link Between Private and Social Welfare Gains

Figures 1 and 2 provide a few insights into the link between the private and social welfare gains from dynamic pricing.⁵ First, both of the private and social welfare gains are increasing in (the absolute value of) demand elasticity. This implies that consumers' selection on the slope is welfare-enhancing in this setting. This is notably different from other settings studied in the literature including healthcare markets, in which selection on the slope is generally considered undesirable for social welfare because it exaggerates "selection on moral hazard" (Einav et al., 2013).

Second, selection on the level matters to consumers but does not have a *direct* effect on social

⁴In the generalized Roy model in Eisenhauer, Heckman and Vytlacil (2015), the subjective cost of treatment is defined as part of a consumer's (negative) surplus, and therefore, relevant to welfare. We follow this approach because it is unlikely the case that policymakers in our context can ignore costs for consumers in their social welfare functions. In the context of healthcare markets, Handel (2013) provides detailed discussions about welfare-relevance of different types of inertia in healthcare plan choice.

⁵We provide a more formal description about this point based on our model in Section 2.3.

welfare. As we described above, structural winners are more likely to adopt dynamic pricing because their private gains are larger than structural losers. However, if a structural winner and a structural loser happen to have identical slopes in their demand curves, the social welfare gains are the same between them even though their private gains are different.

Third, the second point does not necessarily imply that selection on the level does not matter to social welfare. A potential indirect effect exists if selection on the level is correlated with selection on the slope. For example, if structural losers are more-elastic than structural winners, the social planner has a reason to encourage structural losers to adopt dynamic pricing, and vice versa.

Therefore, a consumer's selection decision is not necessarily in line with the social planner's objective in our context. The relationship between the private and social welfare gains depends on heterogeneity in the demand curves, including its slopes and intercepts, and how this heterogeneity is related to consumers' selection. We describe this point more formally in our model in Section 2 and empirically estimate this heterogeneity and its relation to selection in Section 4.

2 Conceptual Framework

In this section, we formalize the selection mechanisms and social planner's problem discussed in Section 1. Our framework has two goals. First, we use welfare economics and the generalized Roy model to connect heterogeneity in private gains and heterogeneity in social welfare gains from a policy. Second, we show that, in our context, the social welfare function in the presence of selection can be written as a *target parameter*, as defined by (Mogstad, Santos and Torgovitsky, 2018). This framework clarifies the parameters that have to be estimated empirically to conduct the welfare analysis.

Figure 3 provides an overview of our framework. The social planner can provide a financial take-up incentive (Z) for consumers to adopt dynamic pricing. Consumers self-select based on the selection equation in Section 2.1. In Section 2.2, we show that the social welfare gains can be written as a function of the selection equation and MTRs. We provide a visual summary of the key results of this framework in Section 2.3 and explain how this framework relates to our empirical analysis in Section 2.4.

[Figure 3 about here]

2.1 Consumer’s Problem

Consider an electricity consumer who has a default electricity price plan ($j = 0$) and an option to adopt a dynamic price plan ($j = 1$). A binary variable $D = \{0, 1\}$ equals one if the consumer chooses dynamic pricing. We use Y_t to denote the consumer’s hourly electricity consumption in hour $t = 1, \dots, T$. The relationship between the observed outcome (Y_t) and potential outcomes ($Y_{t,j}$) for $j = \{0, 1\}$ is

$$Y_t = DY_{t,1} + (1 - D)Y_{t,0}, \tag{1}$$

where $Y_{t,j} = \mu_{t,j}(X) + U_{t,j}$ for $j = \{0, 1\}$ in which $\mu_{t,j}(X)$ is an unspecified function of observables X and $U_{t,j}$ are the unobservables with $E[U_{t,j}|X = x] = 0$.

The consumer obtains indirect utility S_j from choice j and selects into dynamic pricing if $S_1 - S_0 > 0$. Following Heckman and Vytlacil (1999, 2005) and Mogstad, Santos and Torgovitsky (2018), we assume that the net surplus $S_1 - S_0$ is weakly separable between observables and unobservables:

$$S_1 - S_0 = \nu(X, Z) - V. \tag{2}$$

$\nu(X, Z)$ is the observable part of the selection equation with a flexible function $\nu(\cdot)$, observables X , and an instrument Z that affects selection. V is the unobserved disutility for selecting $j = 1$. We make the following assumptions.

Assumption 1. $D = 1$ if $S_1 - S_0 > 0$.

Assumption 2. $(U_{t,0}, U_{t,1}, V) \perp\!\!\!\perp Z|X$, where $\perp\!\!\!\perp$ denotes conditional independence.

Assumption 3. V is continuously distributed, conditional on X .

Assumption 1 implies that consumers self-select based on the selection equation. Assumption 2 implies that conditional on X , instrument Z is independent of the potential outcomes and unobservables in the selection equation. Assumption 3 requires the distribution of V to be continuous, but it does not impose a particular distributional assumption. We use $F_V(\cdot)$ to denote the cumulative distribution function of V . Then equation (2) implies that the propensity score is $P(X, Z) \equiv \Pr[D = 1|X, Z] = F_V(\nu(X, Z))$.⁶

⁶As noted in Brinch, Mogstad and Wiswall (2017), a traditional approach to estimating equations (1) and (2) is to assume a certain parametric joint distribution of the unobservables (U_0, U_1, V) (e.g., see Björklund and Moffitt

The key variable we define is $U \equiv F_V(V) \in [0, 1]$. Since V is the unobserved *disutility* for choosing dynamic pricing, U tells us the quantiles of consumer types in terms of their unobserved preferences. Consumers with lower U are more likely to adopt dynamic pricing for unobservable reasons, whereas those with higher U are less likely to adopt dynamic pricing for unobservable reasons.

With this setup, we can define the **Marginal Treatment Effect (MTE)** following Heckman and Vytlacil (1999, 2005):

$$Y_t^{MTE}(x, p) = E[Y_{t,1} - Y_{t,0} | X = x, U = p]. \quad (3)$$

One interpretation of the MTE is that it is a function of the ATE in observables (X) and unobserved disutility for adopting dynamic pricing (U). The MTE is useful for our research question because it allows a direct test of selection on the slope (Einav et al., 2013) in terms both of the observables and unobservables. First, consider consumer types whose observables X make them more likely to adopt dynamic pricing through the selection equation in (2). The link between the MTE and X allows us to test if these observable types have larger or smaller behavioral responses to treatment (i.e., changes in electricity consumption in response to changes in electricity price) than others.

Second, consider consumer types whose unobservables U make them more likely to self-select, conditional on observables X . In other words, their values of U —the unobserved disutility for adopting dynamic pricing—are lower than others. The link between the MTE and U allows us to test if these unobservable types have larger or smaller behavioral responses to treatment than others. The advantage of this test is that it does not have to impose distributional assumptions on the unobservables, except for Assumptions 1, 2, and 3.

Other useful estimands for our study are the **Marginal Treatment Responses (MTRs)** (1987)). By contrast, the approach we use—developed by Heckman and Vytlacil (1999, 2005); Brinch, Mogstad and Wiswall (2017); Mogstad, Santos and Torgovitsky (2018) and others—does not require assumptions on the joint distribution of these unobservables. We adopt this approach because it requires weaker assumptions about the unobservables to estimate the primitives necessary for our welfare analysis. However, the traditional approach has its own advantages, including its ability to identify a richer set of primitives, which can be valuable for other research questions.

introduced by [Mogstad, Santos and Torgovitsky \(2018\)](#):

$$\begin{aligned} m_{t,0}(x, p) &= E[Y_{t,0}|X = x, U = p] \\ m_{t,1}(x, p) &= E[Y_{t,1}|X = x, U = p]. \end{aligned} \tag{4}$$

The MTRs are simply two components of the MTE. There are at least two advantages of working with the MTRs. First, $m_{t,0}(x, p)$ provides the average untreated outcome for consumer types $U = p$ and $X = x$. Therefore, $m_{t,0}(x, p)$ provides important information on selection in terms of the untreated potential outcome ([Einav, Finkelstein and Schrimpf, 2010](#)). Second, this allows us to work with policy-relevant target parameters that can be any function of $m_{t,0}$ and $m_{t,1}$, including, but not only, functions of the MTE ($= m_{t,1}(x, p) - m_{t,0}(x, p)$). This is important in our study because we show in [Section 2.2](#) that our social welfare function can be written as a function of $m_{t,0}$ and $m_{t,1}$ but cannot be written as a function of $m_{t,1} - m_{t,0}$.

2.2 Social Planner's Problem

We now consider the social planner's problem and its key connection to the consumer's problem. As shown in [Figure 1](#), the adoption of dynamic pricing generates a social welfare gain equal to the deadweight loss of time-invariant pricing. This welfare gain can be expressed by potential outcomes for social welfare, consumer surplus, and producer surplus as follows:

$$W_1 - W_0 = (S_1 - S_0) + (PS_1 - PS_0), \tag{5}$$

where (W_j, S_j, PS_j) are potential outcomes for social welfare (W), consumer's indirect utility (S), and producer surplus (PS) for the choice of price plan $j = \{0, 1\}$.

This allows us to define the MTE of the social welfare gain for consumer types $U = p$ and $X = x$:

$$\begin{aligned} W^{MTE}(x, z, p) &= E[W_1 - W_0|X = x, Z = z, U = p] \\ &= S^{MTE}(x, z, p) + PS^{MTE}(x, z, p), \end{aligned} \tag{6}$$

where X is observables, Z is a take-up incentive, and U is unobserved disutility for adopting dynamic

pricing. Take-up incentive Z is a transfer between consumers and producers; therefore, although it affects a consumer's selection decision, the transfer will be canceled out in the calculation of social welfare.⁷ We decompose the observable part of the consumer's indirect utility from adopting dynamic pricing by a common marginal utility from Z and a flexible function that can depend on X and Z by $\nu(X, Z) = Z + \tilde{\nu}(X, Z)$. Then, the selection equation (2) implies that $S^{MTE}(x, z, p) = E[S_1 - S_0 | X = x, Z = z, U = p] = \tilde{\nu}(x, z) + z - F_V^{-1}(p)$. Using the definition of producer surplus, W^{MTE} can be written as⁸

$$W^{MTE}(x, z, p) = E[W_1 - W_0 | X = x, Z = z, U = p] \\ = \tilde{\nu}(x, z) - F_V^{-1}(p) + \sum_t \left[(\tau_{t,1} - c_t) \cdot m_{t,1}(x, p) - (\tau_{t,0} - c_t) \cdot m_{t,0}(x, p) \right], \quad (7)$$

where $\tau_{t,j}$ and $c_{t,j}$ are the marginal price and marginal cost of electricity at time t for price plan j , respectively.⁹ This key result implies that the MTE of the social welfare gain ($W^{MTE}(x, z, p)$) can be expressed by a function of the parameters in the selection equation and the MTRs: $m_{t,0}(x, p)$ and $m_{t,1}(x, p)$.

The social planner's primary objective is to maximize the welfare gain per capita (\bar{W}). In the absence of selection, this is simply the policy's ATE on the welfare gain. However, in the presence of selection, the welfare gain per capita becomes:

$$\bar{W}(x, z) \equiv \int_0^1 D \cdot W^{MTE}(x, z, p) dp = \int_0^{P(x,z)} W^{MTE}(x, z, p) dp, \quad (8)$$

where $P(x, z)$ is the propensity score, and the last equality comes from the fact that $D = 1$ for $0 \leq p < P(x, z)$ and $D = 0$ for $P(x, z) \leq p \leq 1$. One interpretation of this formula is that it is the intention-to-treat (ITT) on the welfare gain that compares the intention-to treat outcome of the policy $Z = z$ (i.e., incomplete take-up) against the absence of the dynamic pricing policy (no

⁷In a situation where such a transfer comes with a cost of public funds, our model can be extended to include a marginal cost of public funds in the welfare calculation.

⁸Note that $PS_1 - PS_0 = \sum_t [(\tau_{t,1} - c_t) \cdot Y_{t,1} - (\tau_{t,0} - c_t) \cdot Y_{t,0}] - z$. Therefore, $PS^{MTE}(x, z, p) \equiv E[PS_1 - PS_0 | X = x, Z = z, U = p] = \sum_t [(\tau_{t,1} - c_t) \cdot m_{t,1}(x, p) - (\tau_{t,0} - c_t) \cdot m_{t,0}(x, p)] - z$.

⁹In this model, we consider a constant marginal cost c_t given time t . This is consistent with our empirical context because the change in consumption from customers in our experiment was unlikely to change the very short-run marginal cost of electricity. For different applications, the model can be easily modified to incorporate the case in which the marginal cost at time t is a step function of electricity demand at time t .

take-up).

Two additional statistics are policy relevant. The first statistic is the average treatment effect on the treated (ATT) on the welfare gain: $W^{ATT}(x, z) \equiv \overline{W}(x, z)/P(x, z)$. This calculates the policy's ATE on the welfare gain for those who self-select given the value of Z .

Another policy-relevant statistic is the marginal welfare gain with respect to the policy instrument ($d\overline{W}(x, z)/dz$), which measures how much marginal welfare gain can be obtained by increasing the value of policy instrument Z . This parameter is often the primary interest in the sufficient statistics literature. Chetty (2009) and Kleven (2021) show that even if a policy's treatment effect is heterogenous among individuals, the ATE is often sufficient to describe the policy's marginal welfare gain in the absence of selection. The following equation shows that in the presence of selection, the marginal welfare gain becomes a function of $W^{MTE}(x, z, p)$ and the propensity score:

$$\begin{aligned} \frac{d\overline{W}(x, z)}{dz} &= \frac{d}{dz} \int_0^{P(x, z)} W^{MTE}(x, z, p) dp \\ &= \underbrace{\frac{dP(x, z)}{dz}}_{\text{Size of marginal consumers}} \cdot \underbrace{W^{MTE}(x, z, P(x, z))}_{\text{Welfare gain from marginal consumers}}. \end{aligned} \quad (9)$$

The second equality comes from the Leibniz integral rule. Equation (9) can be seen as a generalization of the sufficient statistics approach in the presence of the selection. This equation implies that the MTE of the welfare gain ($W^{MTE}(x, z, p)$) and the propensity score ($P(x, z)$) from the selection equation become key statistics to characterize the marginal welfare gain with respect to policy instrument Z when there is selection into policy take-up.

Equation (8) also implies that the socially optimal Z can be obtained by estimating propensity score $P(x, z)$ and the MTE of the welfare gain $W^{MTE}(x, z, p)$. When the social planner can differentiate Z based on observables X , the planner can find the optimal differentiated take-up incentive $z^*(x)$ by

$$z^*(x) = \operatorname{argmax}_{z(x)} \int_0^{P(x, z)} W^{MTE}(x, z, p) dp. \quad (10)$$

When the planner cannot differentiate Z by X , the planner can find the optimal uniform take-up

incentive z^* by

$$z^* = \underset{z}{\operatorname{argmax}} \int_X \int_0^{P(x,z)} W^{MTE}(x, z, p) dp dF_X \quad (11)$$

where F_X is the distribution of X .

2.3 Link Between Selection and Social Welfare Gains

Figure 4 summarizes the key results of our framework by illustrating the connection between a consumer's selection problem and the social planner's problem. The horizontal axis is $U \equiv F_V(V) \in [0, 1]$, which is the quantiles of unobserved consumer heterogeneity in terms of their disutility for adopting dynamic pricing. Conditional on observables X , consumers with lower values of U (i.e., the left-hand side of the horizontal axis) are more likely to adopt dynamic pricing and vice versa.

[Figure 4 about here]

Panel A illustrates a hypothetical example of the selection equation, $S_1 - S_0 = \nu(x, z) - V$, for given $X = x$ and $Z = z(x)$. The selection equation is downward sloping in this figure because U is the quantiles of V . That is, given $X = x$ and $Z = z(x)$, consumers with lower values of U have lower values of V and thus they are more likely to adopt dynamic pricing. More precisely, given take-up incentive $Z = z(x)$, consumers with $0 \leq U < P(x, z)$ adopt dynamic pricing because $S_1 - S_0 > 0$. By contrast, consumers with $U \geq P(x, z)$ do not adopt it because $S_1 - S_0 \leq 0$.

Panel B presents a hypothetical function of $W^{MTE}(x, z, p)$ as a function of $U = p$. The shape of $W^{MTE}(x, z, p)$ in this figure is hypothetical, and the function that we empirically estimate can be upward-sloping, downward-sloping, or non-monotonic in U . The welfare gain per capita $\bar{W}(x, z)$ is the integral of $W^{MTE}(x, z, p)$ in $0 \leq U < P(x, z)$. Therefore, $\bar{W}(x, z)$ equals K in the figure.

In this example, take-up incentive $Z = z(x)$ is not at the optimal level. The planner can set a higher incentive $Z = z^*(x)$ based on equation (10). This optimal incentive moves the selection equation to the right ($S_1 - S_0 = \nu(x, z^*(x)) - V$). Then, consumers with $0 \leq U < P(x, z)$ adopt dynamic pricing, and $\bar{W}(x, z)$ can be increased to $K + L$. In other words, if the social planner has knowledge of the selection equation and $W^{MTE}(x, z, p)$, it is possible to find the optimal level of Z by connecting the consumer's problem (Panel A) and the social planner's problem (Panel B). This visual interpretation is a summary of the key results in equations (9) to (11).

2.4 What Parameters Need to be Estimated?

Equations (9) to (11) clarify the parameters that need to be estimated in our empirical analysis to conduct the welfare analysis. First, we need to estimate the selection equation (2) to estimate $\nu(x) - F_V^{-1}(p)$, which also provides an estimate of propensity score $P(x, z)$. Second, we need to estimate the two MTRs ($m_{t,1}(x, p)$ and $m_{t,0}(x, p)$), which provide an estimate of the MTE. Using these estimates, we can construct an empirical estimate of $\bar{W}(x, z)$ in equation (8), calculate $d\bar{W}(x, z)/dz$, and derive the optimal level of the take-up incentive.

With this insight, we designed our field experiment to generate the variation necessary to estimate these parameters. In particular, to estimate the MTE and MTRs, we need an instrumental variable that affects selection decisions but not the potential outcomes of electricity usage ($Y_{t,0}, Y_{t,1}$). In our experiment, we created an instrument by randomly assigning a financial take-up incentive. We explain the details of experiential design in the next section.

3 Experimental Design and Data

We conducted our field experiment in the city of Yokohama, Japan with the Japanese Ministry of Economy, Trade and Industry, the city of Yokohama, Tokyo Electric Power Company, Toshiba Corporation, and Panasonic Corporation. We collected household-level hourly electricity usage data in 2013 to calculate each customer’s expected saving from adopting dynamic pricing and ran the experiment in 2014 and 2015.

Although our experiment was an RCT, participating households were not a random sample of the population, as was the case in previous experiments on electricity pricing (Wolak, 2006, 2011; Faruqui and Sergici, 2011). The city of Yokohama used online and offline public advertisements to recruit participating households. To recruit as broad ranges of households as possible, the city provided free installations of an advanced electricity meter and in-home display as well as a participation reward of JPY 20,000 (\approx \$200) for all participating households. The city was able to recruit 3,293 households. We excluded students, customers who had self-generation devices, and those without access to the Internet. This process left us with 2,153 households. We use 970 customers for the current study and used the rest of the customers for another field experiment (Ito, Ida and Tanaka, 2017). For these customers, we collected hourly electricity usage data at the customer level

as the primary outcome variable for this study.

This sample selection is important when considering external validity. For this reason, we compare observable characteristics between households in our experiment and a random sample of households in the city in the online appendix. We do not find statistically significant differences in the observables between these two groups, including electricity usage and demographics. However, it is still important to recognize that our experimental sample can differ from the population in their unobservables.

3.1 Electricity Plan Choice

All the customers in our experiment were eligible to switch from their default electricity tariff to a dynamic pricing tariff. As shown in Figure 5, under the default flat pricing, the price was 26 cents per kWh regardless of the time of use. Under the dynamic pricing plan, the price in off-peak hours was reduced to 21 cents per kWh, and the price in peak hours was 45 cents per kWh on regular days and 100 cents per kWh on critical peak days.

[Figure 5 about here]

Peak hours were between 1 pm and 4 pm in summer and between 5 pm and 8 pm in winter corresponding to the system peak hours in the Japanese electricity system. Peak-hour prices were mostly 45 cents per kWh except for critical peak days, defined as those days for which the previous day’s weather forecast predicted particularly high system demand relative to system supply. All customers received day-ahead and same-day notices about their prices via the in-home display and text notifications.¹⁰

To generate the variation to estimate the MTE and MTRs, we randomly assigned a financial take-up incentive. Consumers in the incentivized group were notified that they were going to receive an upfront incentive of JPY 6,000 (\approx \$60) upon the take-up of dynamic pricing. As summarized in Table 1, the two groups differed only by the receipt of this take-up incentive.

[Table 1 about here]

¹⁰There were fourteen critical peak days in each of the summer and winter.

3.2 Data and Descriptive Statistics

Using advanced electricity meters often called “smart meters,” we collected hourly electricity usage data at the household level in the pre-experimental period in 2013 and experimental period in 2014 and 2015. The primary intervention (i.e., dynamic pricing) started on July 8, 2014 and continued until September 30 in the summer. The winter intervention began on December 1 and ended on January 31, 2015.

In addition, we collected household demographics and elicited risk preferences using the method developed by [Callen et al. \(2014\)](#). The elicitation was based on two series of questionnaires, which is included in the online appendix. In the first task, a customer made a series of choices between a relatively safe option (option A) and a relatively risky option (option B). Where the customer switched from preferring option A to option B was used to determine the customer’s risk aversion. In the second task, we added uncertainty into option A, but kept it less risky than option B to obtain another risk aversion parameter. The advantage of this method is that we can obtain two risk preference parameters, from which we can calculate the certainty premium as well as risk aversion.

Risk preference can be important in our setting for a few reasons. First, risk-averse consumers can be less likely to adopt dynamic pricing. Second, it is ambiguous if risk-averse consumers have larger or smaller responses to dynamic pricing once they self-select. Third, the interaction of these two phenomena could matter to social welfare. For example, if risk-averse consumers are less likely to adopt dynamic pricing but have larger responses to dynamic pricing (i.e., have more elastic demand), then the social planner may want to incentivize them to adopt dynamic pricing because the social welfare gain from their take-up can be large.¹¹

Table 2 presents the summary statistics of the demographics and pre-experiment consumption data. A comparison across the groups indicates statistical balance in the observables.

[Table 2 about here]

A key variable obtained from the pre-experimental data was the expected saving from dynamic pricing. We used each household’s hourly usage data in the pre-experimental period to calculate a counterfactual payment as if the customer was on dynamic pricing rather than default flat pricing. Using this value, we calculated each customer’s expected annual saving from dynamic pricing based

¹¹This is indeed what we find in our empirical analysis in Section 4.

on the customer’s past usage data. When we notified customers of their expected savings, we explicitly told them that this value was an *expected* saving because it was based on historical usage data as well as the (conservative) assumption of no behavioral response to changes in price. This information was provided to all the customers in this experiment.

[Figure 6 about here]

Figure 6 shows the distribution of customers regarding their expected savings from dynamic pricing. The distribution implies that even with no behavioral response to price changes (zero price elasticity assumption), about half of customers can lower their payment simply by switching to the dynamic tariff. This segment of customers are called structural winners (Borenstein, 2013). Our description in Section 1.2 implies that theoretically, customers in this segment may be more likely to adopt dynamic pricing through *selection on level*. Customers at the other end of the scale are called structural losers. These customers use more electricity in peak hours than in off-peak hours, meaning that they are likely to lose money from dynamic pricing unless they are price-elastic and can re-optimize their usage in response to changes in price.

3.3 Timeline of the Experiment

The experiment consisted of two stages. The first stage was each customer’s decision to adopt dynamic pricing. In June 2014, customers were notified that they were eligible to switch from default flat pricing to dynamic pricing. Customers received a graphical explanation of the two electricity tariffs, as presented in the online appendix. In addition, we explicitly presented each customer’s annual expected saving from dynamic pricing based on his/her historical hourly usage.

The second stage of the experiment was the implementation of dynamic pricing for customers who decided to switch. Customers who switched and those who stayed with the default plan received the same type of daily information, including the information about their usage and price, via the in-home display and text messages.

4 Empirical Analysis

In this section, we use the data from our RCT to estimate the selection equation, MTE, and MTRs, which are key parameters for the welfare analysis in Section 5.

4.1 Descriptive Evidence on Selection

We first present descriptive evidence on selection. Figure 7 shows that 31% of the baseline group and 48% of the incentive group, which received a \$60 take-up incentive, adopted dynamic pricing. A key question is whether consumers self-selected based on their expected savings from dynamic pricing, that is, whether there is evidence of selection on the *level* (Einav et al., 2013).

[Figure 7 about here]

We test this prediction in Figure 8. It shows the histogram of the expected savings from dynamic pricing and frequency of take-up. The baseline group's figure shows evidence of selection on the level. Consumers with positive expected savings were more likely to adopt. Moreover, the take-up incentive group's figure indicates that the \$60 take-up incentive had different effects on take-up between structural winners (i.e., customers with positive expected savings) and structural losers (i.e., customers with negative expected savings). This result is consistent with the fact that the \$60 take-up incentive is likely to be pivotal for structural losers, whereas it may not be pivotal for structural winners, who would adopt dynamic pricing without the take-up incentive.

[Figure 8 about here]

Figure 9 summarizes this different effect of the take-up incentive by presenting the take-up rate of each group over the expected savings. The upward trend for the baseline group suggests evidence of selection on the level. For the take-up incentive group, we observe a horizontal shift to the left compared with the baseline group. The shift is approximately \$60, which is consistent with the fact that they received a \$60 take-up incentive.

[Figure 9 about here]

The figure also shows that the effect of the take-up incentive ($Z = \$60$) on take-up (i.e., the vertical distance between the two lines) is substantially different across customers with different expected savings. This variation is important for our estimation of the MTE and MTRs in the next section. In econometrics terminology, this figure implies that although instrument Z is discrete, the effect of Z on the take-up of treatment D is substantially different across different values of observables X , which in this case are the expected savings from dynamic pricing. We apply this

variation to the estimation method developed by [Brinch, Mogstad and Wiswall \(2017\)](#) to estimate the flexible functional forms of the MTE and MTRs in Section 4.3.

With these pieces of descriptive evidence on selection in mind, we provide a formal statistical analysis of the selection equation in the next subsection.

4.2 Selection Equation

Recall that in equation (2), we model the selection equation by $D = 1$ if $S_1 - S_0 = \nu(X, Z) - V > 0$, where $S_1 - S_0$ is the net surplus of adopting dynamic pricing, and $\nu(X, Z)$ is the observable part of the selection equation with a flexible function $\nu(\cdot)$, observables X , and instrument Z . V is the unobserved disutility for adopting dynamic pricing with distribution $F_V(\cdot)$.

Table 3 shows the estimation results of the selection equation. We use the logit estimation for our main analysis and provide results with other models in the online appendix.¹² To make the results interpretable, we show the marginal effects at the means of the covariates. For example, the estimate for the take-up incentive (0.003) implies that the average marginal effect of the take-up incentive (in USD) is a 0.3 percentage point increase in the take-up rate for customers with the mean values of the covariates.

[Table 3 about here]

In column 1, we estimate the selection equation with the take-up incentive and the expected savings from dynamic pricing without including other household characteristics. The positive and statistically significant partial effect of the expected savings provides evidence of selection on the *level*. Consumers with higher expected savings from dynamic pricing were more likely to self-select.

In column 2, we include more covariates that we described in Section 3.2. The partial effect of the take-up incentive does not change between columns because it was randomly assigned in our experiment.¹³ The risk aversion, certainty premium, and employment status are negatively associated with selection. Years of schooling and expected savings from dynamic pricing are positively

¹²We also use probit and a semi-parametric estimation developed by [Gallant and Nychka \(1987\)](#) to examine the robustness of our results with respect to the choice of the distributional assumption for V . As shown in the online appendix, we find that the results are nearly identical between the three models in our case, and therefore, we follow [Brinch, Mogstad and Wiswall \(2017\)](#) to use logit for our main results.

¹³One of the advantages of having the randomly-assigned take-up incentive is that it allows us to convert the estimates of the selection equation to dollar values. We can use the coefficient estimate of the take-up incentive to convert the estimated indirect utility from dynamic pricing into dollar values.

associated with selection. In column 3, we make the functional form of $\nu(X, Z)$ more flexible by including the interactions of the covariates.¹⁴ In column 4, we make the functional form of $\nu(X, Z)$ flexible further by including a non-parametric controls for expected savings (we include dummy variables for each ventile) rather than a linear control. Column 4 indicates that the association of the take-up and covariates is robust to the inclusion of this flexible control for the expected savings.

4.3 Marginal Treatment Effects and Marginal Treatment Responses

In this section, we use methods developed by [Brinch, Mogstad and Wiswall \(2017\)](#) to estimate the MTE and MTRs, which are key inputs to our framework presented in Section 2. Recall that we model potential outcomes for electricity usage at time t for tariff choice $j = \{0, 1\}$ by $Y_{t,j} = \mu_{t,j}(X) + U_{t,j}$. Then, the MTRs can be written by,

$$\begin{aligned} m_{t,j}(x, p) &= E[Y_{t,j}|X = x, U = p] = \mu_{t,j}(x) + E[U_{t,j}|U = p, X = x] \\ &= \mu_{t,j}(x) + k_{t,j}(p, x), \end{aligned} \tag{12}$$

where $k_{t,j}(p) = E[U_{t,j}|U = p, X = x]$ for $j = \{0, 1\}$. This implies that the MTE is,

$$Y_t^{MTE}(x, p) = m_{t,1}(x, p) - m_{t,0}(x, p) = \mu_t(x) + k_t(p, x), \tag{13}$$

where $\mu_t(x) = \mu_{t,1}(x) - \mu_{t,0}(x)$ and $k_t(p, x) = k_{t,1}(p, x) - k_{t,0}(p, x)$.

We estimate $m_{t,0}(x, p)$ and $m_{t,1}(x, p)$ separately to obtain $Y_t^{MTE}(x, p)$.¹⁵ A key concept that connects equations (12) and the data comes from the following conditional means. Consider the conditional means of $Y_{t,j}$ given $X = x$, $P(X, Z) = p$, and $D = j$ for $j = \{0, 1\}$. The sample analogues of these conditional means can be obtained from data because we observe $Y_{t,j}$, X , $P(X, Z)$ and D .

¹⁴We demean covariates so that the non-interaction effects are easier to compare between columns.

¹⁵Estimating $m_{t,0}(x, p)$ and $m_{t,1}(x, p)$ separately, as opposed to estimating $m_{t,1}(x, p) - m_{t,0}(x, p)$ directly; this is called "a separate estimation approach" by [Heckman and Vytlacil \(2007\)](#) and [Brinch, Mogstad and Wiswall \(2017\)](#).

Mathematically, these conditional means can be written by,

$$\begin{aligned} E[Y_{t,1}|X = x, P(X, Z) = p, D = 1] &= \mu_{t,1}(x) + E[U_{t,1}|U < p, X = x] \\ &= \mu_{t,1}(x) + K_{t,1}(p, x) \end{aligned} \quad (14)$$

$$\begin{aligned} E[Y_{t,0}|X = x, P(X, Z) = p, D = 0] &= \mu_{t,0}(x) + E[U_{t,0}|U \geq p, X = x] \\ &= \mu_{t,0}(x) + K_{t,0}(p, x) \end{aligned} \quad (15)$$

where $K_{t,1}(p, x) = E[U_{t,1}|U < p, X = x]$ and $K_{t,0}(p, x) = E[U_{t,0}|U \geq p, X = x]$. We can take the derivatives of $K_{t,1}$ and $K_{t,0}$ with respect to p and rearrange them to obtain $k_{t,1} = (\partial K_{t,1}/\partial p)p + K_{t,1}$ and $k_{t,0} = -(\partial K_{t,0}/\partial p)(1 - p) + K_{t,0}$.¹⁶

There are two approaches to estimate $m_{t,j}(x, p)$ with a binary instrument. The first approach estimates $m_{t,j}(x, p)$ as a linear function in p . An advantage of this approach is that it does not require an additional assumption on $m_{t,j}(x, p)$ as long as it is linear in p . Conversely, its limitation is that the validity of the linearity assumption is generally unknown. Moreover, as we show below, a linear function of $m_{t,j}(x, p)$ is estimated off of two points of $U \in [0, 1]$. This means that we need to rely on a linear extrapolation if we want to recover $m_{t,j}(x, p)$ for a wider range of U . To address these limitations, [Brinch, Mogstad and Wiswall \(2017\)](#) propose a second approach, which uses a separability assumption between X and U to be able to estimate $m_{t,j}(x, p)$ as a flexible nonlinear function in p . We provide results from both approaches below.

4.3.1 Linear Marginal Treatment Effects

Since [Brinch, Mogstad and Wiswall \(2017\)](#) provide detailed explanations of the estimation procedure, we focus on a brief discussion of the procedure and its relevance to our context. Suppose that $k_{t,j}$ is linear in p . Then, this implies that $K_{t,j}$ is also linear in p .¹⁷ Suppose that we have a randomly-assigned binary instrument Z . For each of the two values of Z , we observe $Y_{t,1}$ for consumers who adopted dynamic pricing ($j = 1$) and $Y_{t,0}$ for those who did not adopt dynamic pricing ($j = 0$). Thus, for each $j = \{0, 1\}$, our data provide the sample analogues for two points of $E[Y_{t,j}|X = x, P(X, Z) = p, D = j]$ conditional on X . As $\mu_{t,j} + K_{t,j}$ is linear in p , these two points

¹⁶Note that $K_{t,1} = E[U_{t,1}|U < p, X = x] = \frac{1}{p} \int_0^p E[U_{t,1}|U = u, X = x] du$ and hence $\partial K_{t,1}/\partial p = \frac{1}{p}(k_{t,1} - K_{t,1})$. $K_{t,0}(p, x) = E[U_{t,0}|U \geq p, X = x] = \frac{1}{1-p} \int_p^1 E[U_{t,0}|U = u, X = x] du$ and hence $\partial K_{t,0}/\partial p = \frac{1}{1-p}(K_{t,0} - k_{t,0})$.

¹⁷This is because $k_{t,1} = (\partial K_{t,1}/\partial p)p + K_{t,1}$ and $k_{t,0} = -(\partial K_{t,0}/\partial p)(1 - p) + K_{t,0}$.

can identify the slope and intercept of $\mu_{t,j} + K_{t,j}$.

Panel A in Table 4 shows the results of this estimation with bootstrapped standard errors in parentheses. We estimate $\mu_{t,j} + K_{t,j}$ conditional on X and report the estimates for the average values of X . Recall that consumers with $D = 1$ had a higher peak-hour price than consumers with $D = 0$. Therefore, we expect the treatment effects to be negative (i.e., reductions in electricity consumption). The positive slopes of Y^{MTE} in the summer and winter provide empirical evidence of selection on the *slope* through unobserved heterogeneity—consumers with lower values of $U = p$ (i.e., those who are more likely to adopt dynamic pricing for unobserved reasons) have larger treatment effects (i.e., more reductions in electricity usage in response to the price increase). The treatment effect diminishes as we move to customers who have higher values of $U = p$, namely, those who are less likely to adopt dynamic pricing for unobserved reasons.

[Table 4 about here]

Additional key results presented in Panel A are the MTRs for the untreated and treated potential outcomes (i.e., $E[Y_0|U]$ and $E[Y_1|U]$). We find that $E[Y_0|U]$ is downward-sloping. This implies that electricity usage without dynamic pricing (Y_0) is *higher* for consumers with lower values of U (i.e., those who are more likely to adopt dynamic pricing for unobserved reasons). In addition, $E[Y_1|U]$ is upward-sloping, suggesting that electricity usage with dynamic pricing (Y_1) is *lower* for consumers with lower values of U . These two elements make the MTE upward-sloping.¹⁸

[Figure 10 about here]

In Panel B of Table 4, we estimate the local average treatment effect (LATE) of the take-up incentive. Brinch, Mogstad and Wiswall (2017) show that the linear MTE with a binary instrument is equal to the LATE at the middle point of the two propensity scores, $p_0 = E[D|Z = 0]$ and $p_1 = E[D|Z = 1]$. That is, we can obtain the LATE by $MTE[(p_0 + p_1)/2]$, which is presented in the Table. The LATE provides the ATE for the compliers—that is, those who were induced to take up the dynamic pricing by the randomly-assigned take-up incentive.

¹⁸Another way to interpret this result is to consider the always-takers, compliers, and never-takers in relation to take-up D and instrument Z (the take-up incentive). The downward-sloping $E[Y_0|U]$ implies that $E[Y_0]$ is higher for the always-takers than the compliers and higher for the compliers than the never-takers. Similarly, the upward-sloping $E[Y_1|U]$ implies that $E[Y_1]$ is lower for the always-takers than the compliers and lower for the compliers than the never-takers.

4.3.2 Semi-parametric Estimation with the Separability Assumption

As we described above, the linear MTE approach has two main limitations. To address them, [Brinch, Mogstad and Wiswall \(2017\)](#) propose an alternative estimation strategy with a separability assumption between X and U . We assume that $m_{t,j}(x, p)$ is additively separable in X and U so that it can be expressed by $m_{t,j}(x, p) = \mu_{t,j}(x) + E[U_{t,j}|U = p] = \mu_{t,j}(x) + k_{t,j}(p)$. That is, $m_{t,j}(x, p)$ can still depend on both of X and U , but the relationships with X and U are additively separable.

The separability assumption allows us to estimate $m_{t,j}(x, p)$ as a flexible nonlinear function with a binary instrument for the following reason. Suppose that the effects of the discrete instrument (Z) on treatment take-up (D) are heterogeneous among individuals with different values of X . [Figure 9](#) presents an empirical example of this variation. Our take-up incentive (Z) had different effects on D for structural winners (i.e., those who had higher values of X) and structural losers (i.e., those who had lower values of X). This implies that the experimental variation of Z created many values of $P(X, Z) = p$ in the range of $U \in [0, 1]$. With the separability assumption, we can exploit this variation in $P(X, Z) = p$ to estimate $m_{t,j}(x, p)$ semi-parametrically over a wide range of U that is covered by the common support of $P(X, Z|D = 1)$ and $P(X, Z|D = 0)$. This allows us to recover $m_{t,j}(x, p)$ for a wide range of U without relying on a linear extrapolation, which we presented in [Section 4.3.1](#).¹⁹

Before we proceed to our estimation, it is important to summarize the advantages and limitations of the separability assumption in [equation \(12\)](#). As we mentioned above, the separability assumption allows us to estimate $m_{t,j}(x, p)$ as a semi-parametric function of $U = p$ over a wide range of U . In addition, even with the separability assumption, $m_{t,j}(x, p)$ is still a flexible function of observables and unobservables because it does not require any restriction on the functional forms of $\mu_{t,j}(x)$ and $k_{t,j}(p)$. The restriction imposed by the separability assumption is that $\mu_{t,j}(x)$ and $k_{t,j}(p)$ enter $m_{t,j}(x, p)$ in an additively separable way. This means that $m_{t,j}(x, p)$ can depend on p and x in a flexible way but not depend on the interaction of x and p . In [section 4.3.3](#), we provide a test for the validity of this assumption.

Our semi-parametric estimation follows the two-step approach of [Brinch, Mogstad and Wiswall \(2017\)](#). First, we obtain estimates for propensity score $\hat{p} = \hat{P}(X, Z)$ from the selection equation

¹⁹As we describe below, this range is $U = [0.03, 0.89]$ in our case.

in Table 3. Second, we estimate $\mu_{t,0}(X)$, $\mu_{t,1}(X)$, K_0 , and K_1 using the double residual regression method of [Robinson \(1988\)](#), as modified by [Heckman, Ichimura and Todd \(1997\)](#). This method first estimates the relationship between potential outcomes and covariates (i.e., $\mu_{t,0}(X)$ and $\mu_{t,1}(X)$) and then estimates $K_{t,0}$ and $K_{t,1}$ semi-parametrically by running local quadratic regression of $Y_{t,1} - \hat{\mu}_{t,1}(X)$ and $Y_{t,0} - \hat{\mu}_{t,0}(X)$ on \hat{p} . Based on these estimates, we can obtain $k_{t,j}$, $m_{t,j}(x, p)$, and $Y_t^{MTE}(x, p)$ as we described above. As noted by [Brinch, Mogstad and Wiswall \(2017\)](#), this method allows semi-parametric estimation in the range of U that has sufficient overlap of the propensity scores from the treated and untreated customers. In our case, this overlap range is $U \in [0.03, 0.89]$.

This procedure allows us to jointly estimate the observable part $\mu_t(x)$ and unobservable part $k_t(p)$ of the MTE in equation in (13). Below, we investigate both types of heterogeneity in the estimated MTE—heterogeneity explained by the observables $\mu_t(x)$ and heterogeneity explained by the unobservables $k_t(p)$.

Table 5 shows the estimates of $\mu_t(x)$ for peak hours, where treated customers faced an increase in their electricity price. The negative treatment effect thus implies a reduction in usage in response to this higher price. The positive coefficient on the expected savings implies that consumers with larger expected savings (i.e., structural winners) had *lower* reductions in usage in response to a price increase. Recall that the selection equation in Table 3 and Figure 8 shows evidence of selection on the *level*—customers with larger expected savings were more likely to adopt dynamic pricing. However, the result in Table 5 implies that structural winners are likely to reduce electricity usage to a lower extent, suggesting that the welfare gains from structural winners are smaller than those from structural losers. This finding implies that encouraging structural losers to adopt dynamic pricing (e.g., by providing relatively larger take-up incentives) is likely to be welfare-improving.

[Table 5 about here]

In addition to expected savings, similar phenomena can be observed for risk aversion and certainty premiums. Table 3 suggests that customers with higher risk aversion and certainty premiums are less likely to adopt dynamic pricing. However, Table 5 indicates that these customers produce a larger reduction in peak-hour usage, leading to larger welfare gains. Therefore, customers with higher risk aversion and certainty premiums are observable types that may benefit from relatively large take-up incentives when the take-up incentive can be differentiated by the observables.

In contrast, selection on years of schooling is likely to be in line with the social planner’s objective. Table 3 indicates that customers with more years of schooling are more likely to adopt dynamic pricing, and Table 5 suggests that these customers would reduce peak-hour usage to a larger degree. This finding suggests positive selection on the *slope* based on these observables. Customers with these observable types such as highly educated consumers are more price-elastic and also more likely to self-select into dynamic pricing.

The question is whether there is selection on the slope through *unobserved* heterogeneity, conditional on selection through the observables. We presented evidence from the linear MTE approach in the previous section. The semi-parametric MTEs shown in Figure 10 confirm that this is consistent with evidence from the semi-parametric estimation. Panels A and B show the results for peak-hour electricity usage in the summer and winter, respectively. These figures show evidence of selection on the slope through unobserved heterogeneity—consumers with lower U (i.e., those who are more likely to adopt dynamic pricing for unobserved reasons) have larger treatment effects (i.e., more reductions in electricity usage in response to the price increase).²⁰

[Figure 10 about here]

In sum, the selection on the slope found in this section provides two key implications for social welfare and take-up incentives. First, there is a chance that a take-up incentive can improve welfare if consumers’ self-selection based on observables and unobservables results in take-up below the socially optimal level. Second, however, the marginal return from such an incentive is likely to diminish because consumers with higher U are likely to produce a smaller social welfare gain. In addition, consumers with higher U induce a higher utility loss when they are pushed into take-up because the definition of U implies that their unobserved disutility from adopting dynamic pricing is likely to be larger. This suggests that at some level of take-up incentive, the marginal net welfare gain can reach zero. With this intuition, we investigate the welfare implications more formally in Section 5 by combining the empirical results in this section and framework in Section 2.

²⁰We conduct the same analysis for off-peak hours. In off-peak hours, those who selected into dynamic pricing had a small decrease in price (a change from 26 cents to 21 cents per kWh). In theory, they could thus increase their usage in off-peak hours, and this increase would be welfare-improving. However, we do not find statistically significant evidence of such behavior in our experiment—the treatment effects in the off-peak hours are not statistically significant from zero. This finding is consistent with those of previous studies (Wolak, 2011). A possible reason is that the change in off-peak price was small relative to the large change in peak-hour price; thus, consumers may not have strongly reacted to such a small price change in our experiment.

4.3.3 Validity of the Separability Assumption

Before we proceed to the welfare analysis in Section 5, we investigate the validity of the separability assumption with our data. Recall that we estimated the local quadratic regression of $Y_{t,j} - \mu_{t,j}(X)$ on $\hat{p} = \hat{P}(X, Z)$ for $j = \{0, 1\}$ in Section 4.3. The separability assumption implies that the residuals from this estimation should be uncorrelated with the interactions of the propensity score \hat{p} and observables X .²¹

This can be tested with our data. For summer and winter, we regress the residuals from the local quadratic regression estimated in Section 4.3 on the interactions of \hat{p} and X and test the statistical significance of the coefficients. An obvious limitation of this approach is that failing to reject these null hypotheses does not necessarily guarantee that the separability assumption is valid. It does, however, provide a test for the violation of the separability assumption. In the online appendix, we show the results of this estimation. For both of the treated and untreated outcomes in the summer and winter, we find that the coefficients for the interaction terms between the propensity score and observables are statistically insignificant. These results provide supporting empirical evidence for the separability assumption in our estimation.

5 Welfare Analysis

In this section, we examine the welfare implications of various policies based on the framework developed in Section 2 and the parameter estimates from the selection equation and the semi-parametric estimation of the MTE and MTRs in Section 4.

5.1 Welfare Gains with Respect to Take-Up Incentive (Z)

We first explore how a policy's social welfare gain per capita $\bar{W}(x, z)$ changes with respect to take-up incentive Z . To perform this calculation, we plug the parameter estimates from our empirical analysis into equation (8). Our framework allows $\bar{W}(x, z)$ to vary between individuals with different values of observables X . For illustration purposes, we use given values of X to show the following figures.

Figure 11 shows the welfare gain per capita with respect to the take-up incentives. The figure

²¹We would like to thank anonymous referees for suggesting this testable implication.

shows that providing a positive take-up incentive improves welfare compared with a policy with a zero take-up incentive. In addition, the welfare gain per capita is concave in the take-up incentives—it increases to a certain value of Z and then starts to decline.

This concavity occurs for two primary reasons. The first reason is the shape of the MTE estimated in the previous section. A higher Z induces more consumers to adopt dynamic pricing. However, marginal consumers with higher Z do not provide as large treatment effects (i.e., changes in electricity usage in response to changes in price) as infra-marginal consumers who would adopt dynamic pricing with a lower Z . The second reason is heterogeneity in the unobserved disutility for adopting dynamic pricing, which is term V in the selection equation. The selection equation implies that as we increase Z , consumers with a higher V (i.e., those with higher disutility from adoption) would adopt dynamic pricing. This mechanism also contributes to the diminishing return of Z in the figure.

[Figure 11 about here]

Figure 12 also presents this point, showing the *marginal* welfare gain per capita with respect to the take-up incentive $\left(\frac{d\bar{W}(x,z)}{dz}\right)$. The marginal welfare gain is high when Z is near zero but diminishes and eventually becomes negative as the benefit of marginal consumers becomes smaller than the cost.

[Figure 12 about here]

5.2 Welfare Comparison Between Counterfactual Policies

Table 6 compares the welfare implications of five policies—the two policies from the RCT ($Z = 0$ and $Z = 60$) and three counterfactual policies. The first counterfactual is a policy with an optimal uniform take-up incentive ($Z = z^*$) that maximizes the welfare gain per capita $\bar{W}(x, z)$ under the constraint that a take-up incentive cannot be differentiated by observables X . The second counterfactual is a policy with optimal differentiated take-up incentives ($Z = z^*(x)$) that maximize $\bar{W}(x, z)$ with take-up incentives that can vary by observables X .

The outcome variables of interest are 1) the take-up rate, 2) the ATE on the treated on the welfare gain (W^{ATT}), and 3) the intention-to-treat on the welfare gain (\bar{W}). In general, the ultimate policy goal is to maximize \bar{W} because it measures the policy’s overall welfare gain per capita. W^{ATT} is also

an informative outcome variable for understanding the welfare gain per capita among consumers who adopt dynamic pricing in different policy scenarios.

[Table 6 about here]

Results in rows 1 and 2 in Panel A of Table 6 are obtained using the experimental variation in our RCT, whereas the results in rows 3, 4, and 5 are obtained by counterfactual policy simulations based on the model described in Section 2 and parameter estimates from Section 4. Providing a take-up incentive ($Z = 60$) increases the take-up rate from 31% to 48%. Simultaneously, providing such an incentive reduces W^{ATT} because the marginal welfare gain decreases in Z , as described in Section 5.1. Overall, providing $Z = 60$ increases \bar{W} from \$18.0 to \$23.2 per consumer per year. Row 3 investigates the case with the optimal uniform take-up incentive ($Z = z^*$). The optimal uniform incentive increases the take-up rate to 43.9%. Compared with the case with $Z = 60$, $Z = z^*$ improves \bar{W} , but the additional welfare gain is not large.

Row 4 shows the policy outcomes with the optimal differentiated take-up incentive ($Z = z^*(x)$). Compared with the case with $Z = z^*$, $Z = z^*(x)$ further improves \bar{W} to \$33.1 because the differentiated take-up incentives can target consumer types that generate high W^{ATT} based on observables X . For example, suppose that consumers with certain values of observables have high W^{MTE} for a wider range of U than other consumers. Then, an optimal differentiated incentive can induce these consumers with a higher subsidy specific to this observable type. Another example is consumers with certain values of the observables who have high W^{MTE} but a low propensity score. If incentives can be differentiated by the observables, policymakers can capture these customers by increasing the incentive for this observable type. This point is mathematically seen in equation (10) in Section 2. The optimal incentive formula exploits both the propensity score from the selection equation ($P(x, z)$) and the MTE of the welfare gain ($W^{MTE}(x, z, p)$). The results in this table confirm that the optimal subsidies indeed improve welfare empirically with our data.

Finally, we consider a restricted targeting policy. In practice, utility companies of electricity, natural gas, or water have consumer-level consumption data in their billing systems, whereas they may not have consumer-level information on demographics and risk preferences. In our last counterfactual simulation, we consider a targeting policy based only on the historical electricity usage data at the consumer level. With this subset of observables, we can still use the expected gains from

dynamic pricing for obtaining the optimal differentiated subsidy $Z = z^\dagger(x)$ but cannot use other observables. Our finding in the final row of Table 6 suggests that the restricted targeting policy still improves the outcome compared with the non-targeting policy, but the benefit of targeting is less than the optimal targeting based on $Z = z^*(x)$.

In Panel B, we calculate the difference in \bar{W} and difference in W^{ATT} between policies with the bootstrapped standard errors in parentheses. For example, row 1 suggests that relative to $Z = 0$, the policy with $Z = 60$ increases \bar{W} by 5.6 but decrease W^{ATT} by -6.3. This is because a higher take-up incentive attracts more customers, but the marginal welfare gains from these additional customers are smaller than those who adopt dynamic pricing without the take-up incentive. Panel B also shows that with targeting (both with $Z = z^*(x)$ and $Z = z^\dagger(x)$), policymakers can target customers with higher W^{ATT} , which results in higher \bar{W} .

5.3 Welfare Comparison with a Mandate

As discussed in the introduction, the vast majority of countries rely on voluntary take-up policies for the adoption of dynamic electricity pricing because a mandate is politically infeasible. However, it is still useful to discuss what conditions could make a mandate more welfare-enhancing than the counterfactual policies discussed above.

While our study cannot comprehensively analyze a mandate, we can use our model to provide a discussion below, which would be useful for comparing a mandate and counterfactual policies considered in the previous section.

To analyze a mandate, a key question is how it would change consumers' utility, particularly their disutility for adopting dynamic pricing. For example, consumers may dislike adopting dynamic pricing because they have an unobserved preference for non-dynamic pricing relative to dynamic pricing. Another example is that they may have inertia or switching costs for moving from default pricing to dynamic pricing. An important question is how a mandate affects these factors.

We can consider two possibilities. The first is that a mandate does not change or may increase the costs of these factors. In this case, consumers have the same or higher disutility from these factors as they have under voluntary take-up policies. In this case, our results suggest that a mandate is likely to be inferior to optimal take-up incentive policies because additional take-up beyond the optimal take-up in Table 6 lowers the welfare gain (\bar{W}). This is because the social welfare gain from

the marginal consumer is lower than the welfare loss from the consumer’s disutility for adopting dynamic pricing.

The second possibility is that a mandate may lower consumers’ disutility from adopting dynamic pricing. For example, a mandate may be able to lower switching costs. For this possibility, we can use our counterfactual simulation to examine how much such a benefit has to be to make a mandate policy more welfare-enhancing than other policies. In our counterfactual simulation, if we increase the take-up rate to 100%, the welfare gain per capita (\overline{W})—note that this calculation still includes consumers’ disutility from adoption—becomes \$5.5 per year per consumer. This implies that a mandate can be more welfare-improving than the take-up incentive policy if it can reduce consumers’ disutility by more than \$27.6 ($= 33.1 - 5.5$) per customer per year. This result implies that if the net benefit of a mandate (the economic benefit minus the political cost) exceeds this number, a mandate might be preferred over voluntary take-up policies.

6 Conclusion and Directions for Further Research

In this study, we investigate a problem in which policymakers need to screen self-selected individuals by unobserved heterogeneity in social welfare gains from a policy intervention. In our framework, the marginal treatment effects and marginal treatment responses arise as key statistics to characterize social welfare. We apply this framework to a randomized field experiment on electricity plan choice. In the experiment, consumers were offered welfare-improving dynamic pricing with randomly assigned financial take-up incentives. We find that price-elastic consumers, who generate larger social welfare gains, are more likely to self-select. Finally, we use counterfactual simulations to quantify the optimal take-up incentives by exploiting observed and unobserved heterogeneity in selection and welfare gains.

We want to describe several key empirical issues that were not fully addressed in our study and potential directions for further research. First, the targeting policy considered in this study is static. When the targeting policy is implemented repeatedly, policymakers need to consider the dynamic implications of such policies. For example, if consumers have knowledge about a policymaker’s targeting function, this may create an incentive to manipulate observables in the current period to obtain a subsidy in the future. This could make a dynamically-optimal targeting policy different

from a statically-optimal targeting policy.

Second, our experimental analysis abstracts from a general equilibrium effect that could occur when the policy is implemented on a large scale. If substantial part of the population adopt dynamic pricing and reduce peak-hour electricity usage, it lowers the marginal cost of peak-hour electricity, making the benefit of additional take-up of dynamic pricing diminish as more and more customers adopt it. These issues are discussed in [Joskow and Wolfram \(2012\)](#) and analytically studied by [Borenstein \(2013\)](#).

Third, our results indicate that the optimal take-up rate is not 100% in our context because some consumers have large disutility from adopting dynamic pricing. This result suggests that it would be informative to understand about the sources of the disutility as well as what policy could alleviate it. For example, if this disutility comes from the aversion to bill volatility, a policy instrument that reduces bill volatility could increase take-up and social welfare.²² Conversely, if this disutility comes from physical/psychological switching costs, a policy instrument that makes switching easier could increase take-up and social welfare.

Fourth, our results are based on a randomized controlled trial that lasted for one year. Therefore, the external validity is an important issue when one considers a policy that could continue for many years. For example, if consumers can adopt dynamic pricing for a longer time period, their expected gains or losses from dynamic pricing can be larger than those in our experimental period. This could make the subjective cost of adoption relatively smaller than the expected gains or losses and, therefore, could make the marginal effect of take-up incentives smaller. Although our experiment was not able to create variation in the duration of experimental periods, such an experimental design would be valuable to quantify this important point.

References

Allcott, Hunt, and Michael Greenstone. 2017. “Measuring the Welfare Effects of Residential Energy Efficiency Programs.” NBER Working Paper 23386.

Andrews, Isaiah, and Conrad Miller. 2013. “Optimal social insurance with heterogeneity.” *Unpublished manuscript*.

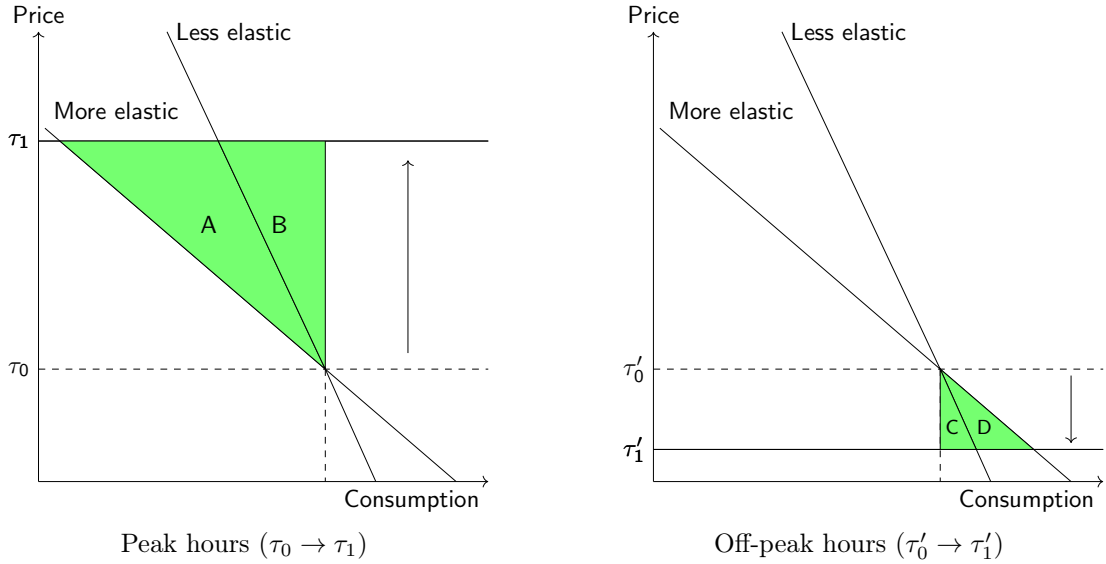
²²In fact, some utility companies offer a payment plan that smooths seasonal variability across a year. Utilities could also offer a dynamic pricing plan that provides a non-dynamic price for infra-marginal usage and a dynamic price for marginal usage.

- Björklund, Anders, and Robert Moffitt.** 1987. “The Estimation of Wage Gains and Welfare Gains in Self-Selection Models.” *The Review of Economics and Statistics*, 69(1): 42–49.
- Borenstein, Severin.** 2002. “The Trouble With Electricity Markets: Understanding California’s Restructuring Disaster.” *Journal of Economic Perspectives*, 16(1): 191–211.
- Borenstein, Severin.** 2013. “Effective and Equitable Adoption of Opt-In Residential Dynamic Electricity Pricing.” *Review of Industrial Organization*, 42(2): 127–160.
- Borenstein, Severin, and Lucas W. Davis.** 2016. “The Distributional Effects of US Clean Energy Tax Credits.” *Tax Policy and the Economy*, 30(1): 191–234.
- Brinch, Christian N., Magne Mogstad, and Matthew Wiswall.** 2017. “Beyond LATE with a Discrete Instrument.” *Journal of Political Economy*, 125(4): 985–1039.
- 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.
- Chetty, Raj.** 2009. “Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods.” *Annual Review of Economics*, 1(1): 451–488.
- Cornelissen, Thomas, Christian Dustmann, Anna Raute, and Uta Schönberg.** 2018. “Who Benefits from Universal Child Care? Estimating Marginal Returns to Early Child Care Attendance.” *Journal of Political Economy*, 126(6): 2356–2409.
- Deshpande, Manasi, and Yue Li.** 2019. “Who Is Screened Out? Application Costs and the Targeting of Disability Programs.” *American Economic Journal: Economic Policy*, 11(4): 213–248.
- Einav, Liran, Amy Finkelstein, and Paul Schrimpf.** 2010. “Optimal Mandates and the Welfare Cost of Asymmetric Information: Evidence From the U.K. Annuity Market.” *Econometrica*, 78(3): 1031–1092.
- Einav, Liran, Amy Finkelstein, Stephen Ryan, Paul Schrimpf, and Mark R. Cullen.** 2013. “Selection on Moral Hazard in Health Insurance.” *American Economic Review*, 103(1): 178–219.
- Einav, Liran, Amy Finkelstein, Yunan Ji, and Neale Mahoney.** 2022. “Voluntary Regulation: Evidence from Medicare Payment Reform.” *Quarterly Journal of Economics*, 137(1): 565–618.
- Eisenhauer, Philipp, James J. Heckman, and Edward Vytlacil.** 2015. “The Generalized Roy Model and the Cost-Benefit Analysis of Social Programs.” *Journal of Political Economy*, 123(2): 413–443.
- 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.
- FERC.** 2011. “2010 Assessment of Demand Response and Advanced Metering.” *Staff Report*.
- Finkelstein, Amy, and Matthew J Notowidigdo.** 2019. “Take-up and Targeting: Experimental Evidence from SNAP.” *Quarterly Journal of Economics*, 134(3): 1505–1556.

- Finkelstein, Amy, and Nathaniel Hendren.** 2020. “Welfare Analysis Meets Causal Inference.” *Journal of Economic Perspectives*, 34(4): 146–167.
- Fowle, Meredith, Catherine Wolfram, Anna Spurlock, Annika Todd, Patrick Baylis, and Peter Cappers.** forthcoming. “Default Bias and Follow-On Behavior: Evidence from an Electricity Pricing Program.” *Review of Economic Studies*.
- Gallant, A., and Douglas W. Nychka.** 1987. “Semi-nonparametric Maximum Likelihood Estimation.” *Econometrica*, 55(2): 363–90.
- Handel, Benjamin R.** 2013. “Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts.” *American Economic Review*, 103(7): 2643–82.
- Heckman, James J.** 2010. “Building Bridges between Structural and Program Evaluation Approaches to Evaluating Policy.” *Journal of Economic Literature*, 48(2): 356–398.
- Heckman, James J., and Edward Vytlacil.** 1999. “Local Instrumental Variables and latent Variable Models for Identifying and Bounding Treatment Effects.” *Proceedings of the National Academy of Sciences of the United States of America*, 96(8): 4730–4734.
- Heckman, James J., and Edward Vytlacil.** 2001. “Policy-Relevant Treatment Effects.” *The American Economic Review*, 91(2): 107–111.
- Heckman, James J., and Edward Vytlacil.** 2005. “Structural Equations, Treatment Effects, and Econometric Policy Evaluation.” *Econometrica*, 73(3): 669–738.
- Heckman, James J., and Edward Vytlacil.** 2007. “Econometric evaluation of social programs, part II: Using the marginal treatment effect to organize alternative econometric estimators to evaluate social programs, and to forecast their effects in new environments.” *Handbook of econometrics*, 6: 4875–5143. Publisher: Elsevier.
- Heckman, James J., Hidehiko Ichimura, and Petra E. Todd.** 1997. “How Details Make a Difference: Semiparametric Estimation of the Partially Linear Regression Model.” *Manuscript*.
- Hortaçsu, Ali, Seyed Ali Madanizadeh, and Steven L. Puller.** 2017. “Power to Choose? An Analysis of Consumer Inertia in the Residential Electricity Market.” *American Economic Journal: Economic Policy*, 9(4): 192–226.
- Ito, Koichiro, Takanori Ida, and Makoto Tanaka.** 2017. “Information Frictions, Inertia, and Selection on Elasticity: A Field Experiment on Electricity Tariff Choice.” *Working Paper*.
- Ito, Koichiro, Takanori Ida, and Makoto Tanaka.** 2018. “Moral Suasion and Economic Incentives: Field Experimental Evidence from Energy Demand.” *American Economic Journal: Economic Policy*, 10(1): 240–67.
- 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.
- Kitagawa, Toru, and Aleksey Tetenov.** 2018. “Who Should Be Treated? Empirical Welfare Maximization Methods for Treatment Choice.” *Econometrica*, 86(2): 591–616.


- Kleven, Henrik J.** 2021. "Sufficient Statistics Revisited." *Annual Review of Economics*, 13: 515–538.
- Kline, Patrick, and Christopher R. Walters.** 2016. "Evaluating Public Programs with Close Substitutes: The Case of Head Start." *The Quarterly Journal of Economics*, 131(4): 1795–1848.
- LaLonde, Robert J.** 1986. "Evaluating the Econometric Evaluations of Training Programs with Experimental Data." *The American Economic Review*, 76(4): 604–620.
- Marschak, Jacob.** 1953. *Economic Measurements for Policy and Prediction*. Cowles Comm. Res. Econ., Monogr. No. 14. New York: Wiley.
- Mogstad, Magne, Andres Santos, and Alexander Torgovitsky.** 2018. "Using Instrumental Variables for Inference About Policy Relevant Treatment Parameters." *Econometrica*, 86(5): 1589–1619.
- Robinson, P. M.** 1988. "Root-N-Consistent Semiparametric Regression." *Econometrica*, 56(4): 931–954.
- 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: Social Welfare Gains from Dynamic Pricing



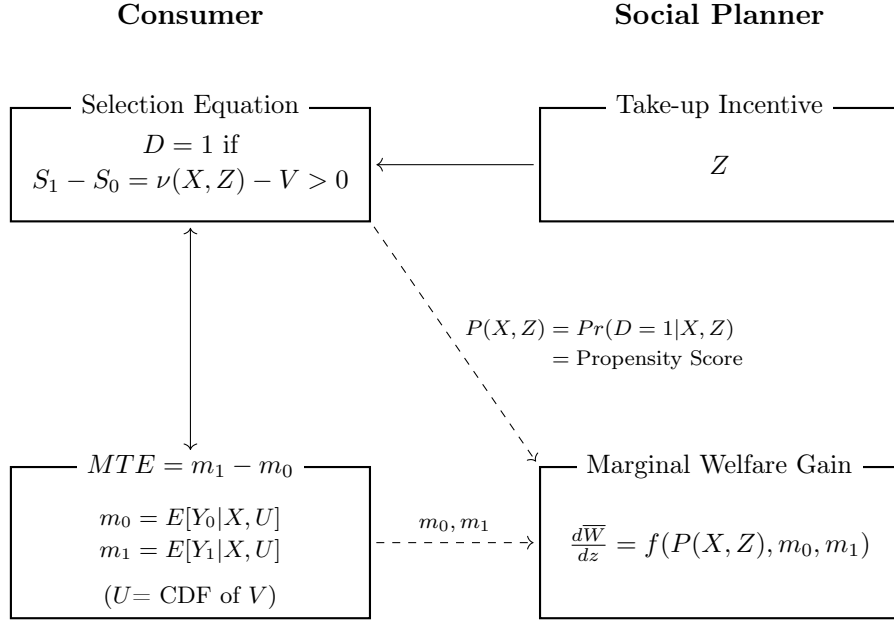
Note: These figures show welfare gains from dynamic pricing for two hypothetical demand curves. Consider a change from flat pricing (τ_0, τ'_0) to dynamic pricing (τ_1, τ'_1), which reflects the time-varying marginal cost of electricity. When the more-elastic consumer adopts dynamic pricing, the welfare gain (i.e., the reduction in deadweight loss) is $A + B$ in peak hours and $C + D$ in off-peak hours. The welfare gain from the less-elastic consumer is B in peak hours and C in off-peak hours. Thus, the more-elastic consumer produces a larger welfare gain upon take-up. These figures do not include the inertia/switching costs from adopting dynamic pricing, which we incorporate in our model and estimation.

Figure 2: Consumers' Private Gains from Dynamic Pricing



Note: These figures show consumers' private gains from dynamic pricing for two hypothetical demand curves. When the more-elastic consumer adopts dynamic pricing, the private gain (i.e., the change in the consumer surplus) is $-E$ in peak hours and $G + H$ in off-peak hours. The private gain for the less-elastic consumer is $-(E + F)$ in peak hours and G in off-peak hours. Thus, the more-elastic consumer has a larger private gain upon take-up. These figures do not include inertia/switching cost from adopting dynamic pricing, which we incorporate in our model and estimation.


Figure 3: Conceptual Framework: Overview



Note: This figure visualizes the conceptual framework in Section 2. The social planner provides a financial take-up incentive Z for consumers to switch from inefficient time-invariant pricing to efficient dynamic pricing. A consumer's switching decision is described by the Roy model with a selection equation and potential outcomes. In Section 2.1, we use this framework to derive the MTE and MTRs (m_0 and m_1 in the diagram). We then show the connection between the consumer's problem and social planner's problem in Section 2.2. A key insight is that the social planner's objective function (i.e., the welfare gain from a policy) can be written as a function of the selection equation and the MTRs.


Figure 4: Conceptual Framework: Consumer's Selection and Social Planner's Welfare Gains

A) Consumer: Selection Equation: $D = 1\{\nu(x, z) - V > 0\}$



Consumer type $U \in [0, 1]$
(unobserved disutility for adopting dynamic pricing)


B) Social Planner: Marginal Welfare Gains (W^{MTE})



Consumer type $U \in [0, 1]$
(unobserved disutility for adopting dynamic pricing)

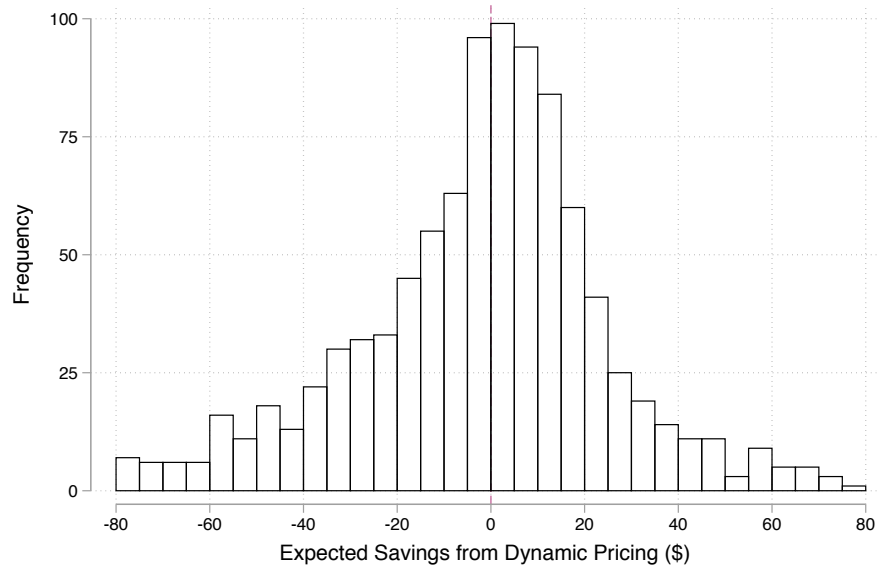
Note: These figures visualize the link between a consumer's selection and the social planner's welfare gains from the consumer. See the text in Section 2.3. Panel A illustrates a consumer's hypothetical selection equation $S_1 - S_0 = \nu(x, z) - V$ over unobserved consumer type U . It is downward-sloping because U is defined as the disutility for adopting dynamic pricing. With the initial level of $Z = z(x)$, consumers with $U < P(x, z)$ self-select, where $P(x, z)$ is the propensity score with observables $X = x$ and the take-up incentive $Z = z(x)$. The take-up increases to $P(x, z^*(x))$ if the take-up incentive is increased to $Z = z^*(x)$. Panel B illustrates a hypothetical function of $W^{MTE}(x, z, p)$. In this example, the welfare gain per capita, $\bar{W}(x, z)$, equals K . This can be further increased to $K + L$ by increasing the take-up incentive to $Z = z^*(x)$. The actual slope of W^{MTE} can be upward, downward, or flat, and we estimate it empirically in Section 4.

Figure 5: Electricity Price Plans



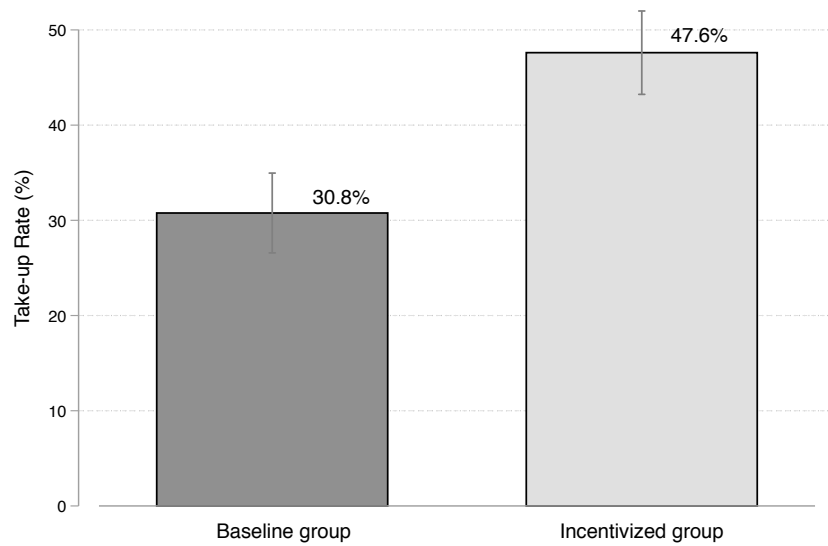
Note: This figure shows the electricity tariffs for the default flat price plan and new dynamic pricing. The dynamic pricing plan charges 21 cents/kWh in off-peak hours and either 45 or 100 cents/kWh in peak hours depending on the day. Peak hours are between 1 pm and 4 pm in summer and between 5 pm and 8 pm in winter.

Figure 6: Expected Savings from Dynamic Pricing



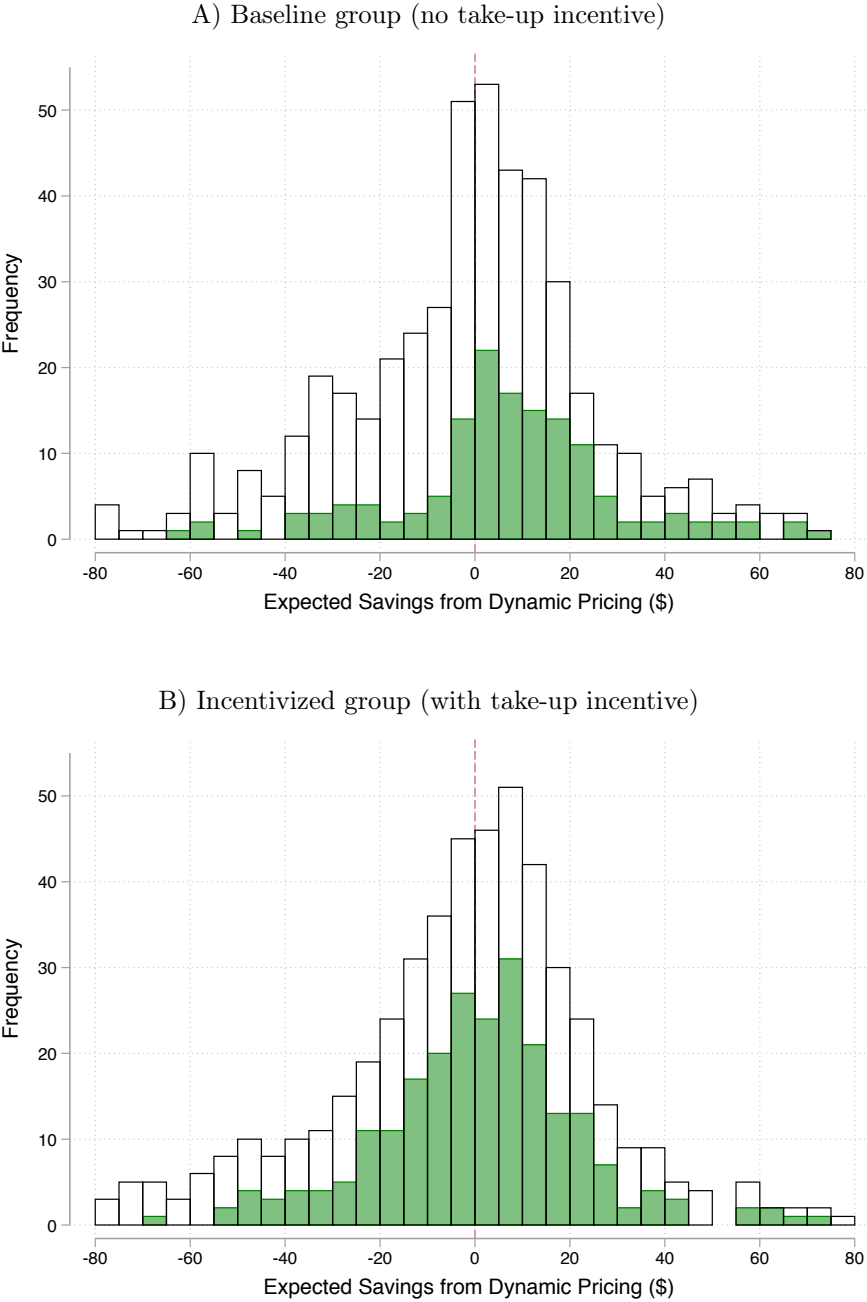
Note: This figure shows the histogram of the expected savings from dynamic pricing. The expected savings are calculated based on each customer's past usage, assuming no behavioral response to changes in price.

Figure 7: Take-Up Rates of Dynamic Pricing



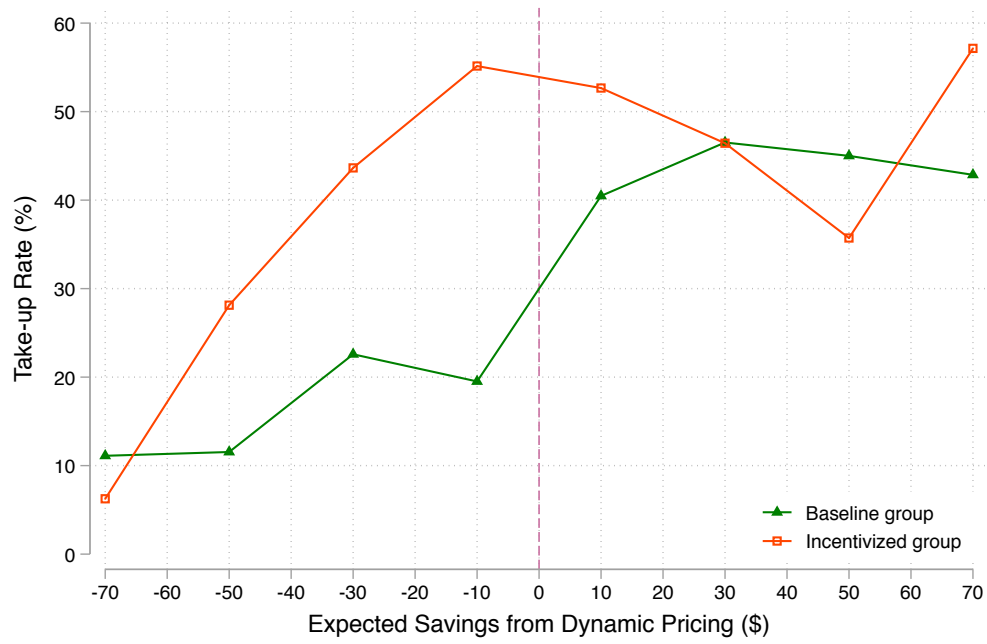
Note: The figure shows the take-up rates for the baseline group and incentivized group that received the take-up incentive. The bars are 95% confidence intervals.

Figure 8: Selection on the *Level*: Take-Up Rates and Expected Savings from Dynamic Pricing



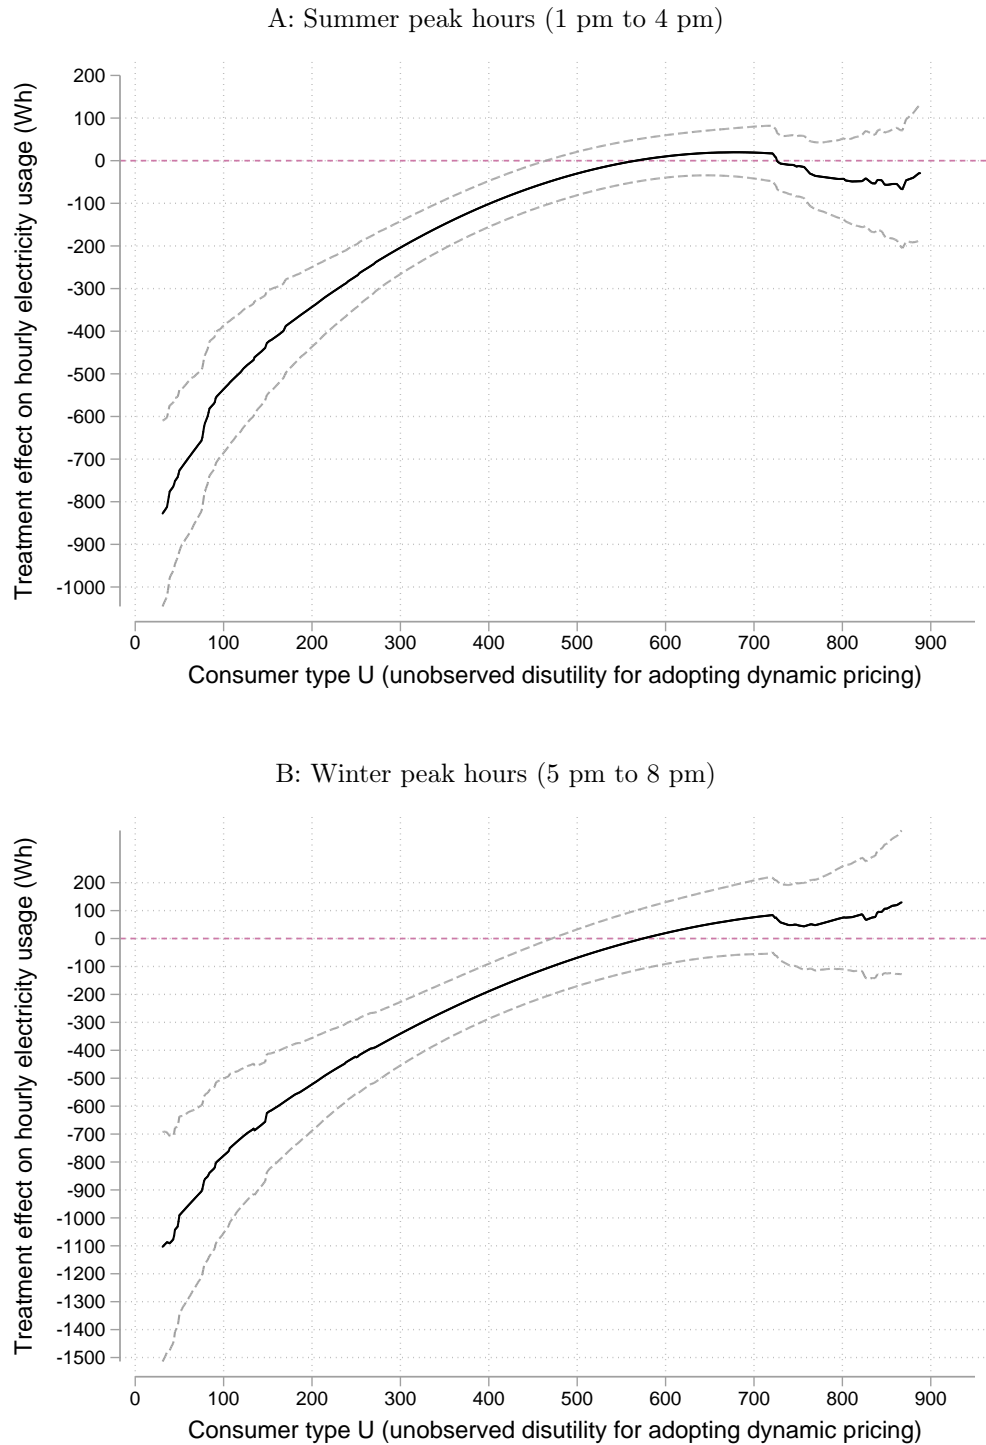
Note: These histograms show the distributions of customers in terms of their expected savings from dynamic pricing. The dark color indicates those who self-selected into dynamic pricing. The expected savings were calculated based on each customer’s past usage, assuming no behavioral response to changes in price. Panel A indicates that structural winners, whose expected savings were larger, are more likely to self-select. Panel B indicates that the take-up incentive increased take-ups of structural losers, whose expected savings are small or negative.

Figure 9: Heterogeneous Effects of the Take-Up Incentive on Take-up Rates




Note: This figure shows the take-up rate for each group by the expected savings from dynamic pricing. The take-up incentive caused different effects on the take-up rates of structural winners and losers.

Figure 10: Selection on the *Slope*: Marginal Treatment Effects on Electricity Usage (Y^{MTE})




Note: This figure shows $Y_t^{MTE}(x, p)$ for peak hours in the summer and winter based on the semi-parametric MTEs estimated in Section 4.3.2. We allow the MTE to vary by X , and in these figures, we show the result for the average values of X . Consumers whose U is lower (i.e., those more likely to adopt dynamic pricing for unobserved reasons) have larger reductions in electricity usage in response to the increase in peak-hour price. The sample size of hourly electricity data in peak hours is 294,810 in the summer and 330,672 in the winter. The dashed lines show the 95 percent confidence intervals based on bootstrapped standard errors clustered at the customer level.

Figure 11: Welfare Gain per Capita: $\bar{W}(x, z)$



Note: This figure shows the welfare gain per capita $\bar{W}(x, z)$ with respect to different values of take-up incentive $Z = z$. We allow the $\bar{W}(x, z)$ to vary by X , and in this figure, we show the result for a given value of X .

Figure 12: Marginal Welfare Gain per Capita: $\frac{d\bar{W}}{dz}$



Note: This figure shows the marginal welfare gain per capita with respect to take-up incentive Z : $d\bar{W}(x, z)/dz$. We allow $d\bar{W}(x, z)/dz$ to vary by X , and in this figure, we show the result for a given value of X .

Table 1: Experimental Design

Group	Eligible to adopt dynamic pricing	Information provision	Take-up incentive
Baseline group	✓	✓	
Incentivized group	✓	✓	✓

Table 2: Summary Statistics

	Baseline group	Incentivized group
Household income (JPY10,000)	742.31 (296.29)	749.80 (311.25)
Years of schooling	14.84 (2.22)	14.62 (2.30)
Employed	0.84 (0.37)	0.86 (0.35)
Risk aversion	0.66 (0.28)	0.63 (0.30)
Certainty premium	0.04 (0.24)	0.06 (0.26)
Square meters	99.82 (33.20)	100.91 (33.43)
Number of bedrooms	3.70 (1.10)	3.81 (0.99)
Age of building	12.71 (12.29)	11.63 (11.15)
Number of room AC	3.18 (1.25)	3.13 (1.25)
Number of TV	2.08 (1.02)	2.05 (0.94)
Number of refrigerator	1.16 (0.41)	1.12 (0.38)
Electricity usage (kWh/day) in the pre-experimental period	13.17 (5.82)	13.18 (6.13)
Expected savings from dynamic pricing (USD/year)	-2.00 (29.73)	-1.86 (34.96)

Notes: This table shows the summary statistics of the two groups. Standard deviations are in parentheses. The sample size of customers is 468 for the baseline group and 502 for the incentivized group.

Table 3: Selection Equation

	Marginal effects on $\Pr[D_i = 1(\text{household } i \text{ selected into dynamic pricing})]$			
	(1)	(2)	(3)	(4)
Take-up incentive (USD)	0.0029 (0.0005)	0.0030 (0.0005)	0.0031 (0.0006)	0.0033 (0.0006)
Expected savings (USD)	0.0019 (0.0005)	0.0021 (0.0005)	0.0026 (0.0006)	
Risk aversion		-0.2045 (0.0811)	-0.2430 (0.0862)	-0.2606 (0.0876)
Certainty premium		-0.3130 (0.0963)	-0.3355 (0.1015)	-0.3432 (0.1027)
Years of schooling		0.0210 (0.0075)	0.0160 (0.0079)	0.0115 (0.0080)
Employed		-0.1308 (0.0496)	-0.0541 (0.0638)	-0.0818 (0.0646)
Income (100,000 USD)		0.0527 (0.0593)	0.0541 (0.0645)	0.0690 (0.0650)
Covariates interacted with each other	No	No	Yes	Yes
Non-parametric controls for expected savings	No	No	No	Yes
Log likelihood	-628.8	-613.1	-592.7	-579.6

Notes: This table shows the estimation results of the selection equation in equation (2). We show the marginal effects at the means of the covariates. The sample size is 970. We use the delta method to obtain standard errors and report them in parentheses.

Table 4: Linear Marginal Treatment Effects and Local Average Treatment Effects

Panel A: Linear MTE

	Summer		Winter	
	Intercept	Slope	Intercept	Slope
$\mu_1 + K_1 = E[Y_1 U < p]$	338.3 (5.6)	103.1 (13.6)	465.2 (5.3)	387.3 (25.3)
$\mu_0 + K_0 = E[Y_0 U \geq p]$	560.5 (20.8)	-200.9 (22.6)	756.6 (25.1)	-162.7 (29.7)
$\mu_1 + k_1 = E[Y_1 U = p]$	338.3 (5.6)	51.6 (6.8)	465.2 (5.3)	193.6 (12.6)
$\mu_0 + k_0 = E[Y_0 U = p]$	460.1 (10.6)	-100.5 (11.3)	675.2 (11.1)	-81.4 (14.9)
Linear $Y^{MTE} = E[Y_1 - Y_0 U = p]$	-222.2 (17.3)	304.0 (28.6)	-291.4 (23.6)	550.0 (44.9)

Panel B: LATE (Local average treatment effects of the take-up incentive)

	Summer	Winter
LATE from linear MTE: Y^{MTE} with $p = (0.308 + 0.476)/2$	-103.2 (7.9)	-76.1 (9.4)

Notes: This table shows the estimation results described in Section 4.3.1. The outcome variable is hourly electricity usage (in Wh) in peak hours. Treated customers had higher prices in peak hours and therefore the negative treatment effects imply reductions in electricity usage in response to increases in price. The sample size of hourly electricity data in peak hours is 294,810 in the summer and 330,672 in the winter. We compute bootstrapped standard errors clustered at the customer level by bootstrapping the entire estimation process, including the propensity score estimation and MTE estimation.

Table 5: Observables in the Marginal Treatment Effect

Estimates of $\mu_t(x)$ in the MTE functions in equation (13)

	Summer		Winter	
	(1)	(2)	(3)	(4)
Risk aversion	-42.96 (11.39)	-39.74 (11.11)	-73.12 (17.58)	-66.05 (16.85)
Certainty premium	-45.25 (12.80)	-50.87 (12.80)	-89.42 (24.78)	-98.60 (23.69)
Expected savings (USD/year)	1.13 (0.08)	1.13 (0.08)	1.85 (0.16)	1.79 (0.16)
Employed	-16.26 (8.10)	-50.25 (8.81)	49.27 (10.89)	-7.97 (12.03)
Years of schooling	-7.54 (0.87)	-8.93 (0.87)	-9.25 (1.62)	-11.74 (1.66)
Income (JPY10,000)		54.76 (7.17)		102.86 (11.64)

Notes: This table shows the estimates of $\mu_t(x)$ in equation (13)—the observable part of the MTE function—for electricity usage in peak hours. The outcome variable is hourly electricity usage (in Wh) in peak hours. Treated customers had higher prices in peak hours and therefore the negative treatment effects imply reductions in electricity usage in response to increases in price. The sample size of hourly electricity data in peak hours is 294,810 in the summer and 330,672 in the winter. We compute bootstrapped standard errors clustered at the customer level by bootstrapping the entire estimation process, including the propensity score estimation and MTE estimation.

Table 6: Welfare Comparison Between Counterfactual Policies

Panel A: Welfare gain from each policy

Policy	Evidence from	Targeting	Take-up	Welfare gain: W^{ATT} (\$/year/consumer)	Welfare gain: \bar{W} (\$/year/consumer)
$Z = 0$	RCT	No	30.8%	60.2	18.0
$Z = 60$	RCT	No	47.6%	47.9	23.2
$Z = z^*$	Counterfactual	No	43.9%	53.8	23.7
$Z = z^*(x)$	Counterfactual	Based on X	44.1%	75.1	33.1
$Z = z^\dagger(x)$	Counterfactual	Based on a subset of X	42.0%	66.9	28.1

Panel B: Welfare comparison

	Difference in \bar{W}	Difference in W^{ATT}
$Z = z^*$ vs. $Z = 0$	5.6 (0.89)	-6.3 (5.81)
$Z = z^*(x)$ vs. $Z = 0$	15.1 (1.12)	14.9 (5.18)
$Z = z^*(x)$ vs. $Z = z^*$	9.5 (0.42)	21.3 (0.96)
$Z = z^*(x)$ vs. $Z = z^\dagger(x)$	5.0 (0.40)	8.2 (1.57)

Notes: This table compares the welfare gains from five policies. The first two rows are derived from the results of our RCT. The rest of the rows are from counterfactual simulations based on the model presented in Section 2 and results in Section 4. $Z = z^*$ is the optimal uniform take-up incentive that does not depend on observables X , whereas $Z = z^*(x)$ is the optimal differentiated take-up incentive that is allowed to vary by observables X , and $Z = z^\dagger(x)$ is the optimal differentiated take-up incentive when only a subset of X can be used. We compute bootstrapped standard errors clustered at the customer level by bootstrapping the entire estimation including the propensity score estimation, MTE estimation, and welfare calculation.