

ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

*An International Society for the Advancement of Economic
Theory in its Relation to Statistics and Mathematics*

<https://www.econometricsociety.org/>

Econometrica, Vol. 94, No. 1 (January, 2026), 225–247

CHOOSING WHO CHOOSES: SELECTION-DRIVEN TARGETING IN ENERGY REBATE PROGRAMS

TAKANORI IDA

Graduate School of Economics, Kyoto University

TAKUNORI ISHIHARA

Faculty of Economics and Business Administration, Kyoto University of Advanced Science

KOICHIRO ITO

Harris School of Public Policy, University of Chicago

DAIDO KIDO

Department of Economics, Otaru University of Commerce

TORU KITAGAWA

Department of Economics, Brown University

SHOSEI SAKAGUCHI

Faculty of Economics, University of Tokyo

SHUSAKU SASAKI

Center for Infectious Disease Education and Research, University of Osaka

The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs. No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society. For such commercial purposes contact the Office of the Econometric Society (contact information may be found at the website <http://www.econometricsociety.org> or in the back cover of *Econometrica*). This statement must be included on all copies of this Article that are made available electronically or in any other format.

CHOOSING WHO CHOOSES: SELECTION-DRIVEN TARGETING IN ENERGY REBATE PROGRAMS

TAKANORI IDA

Graduate School of Economics, Kyoto University

TAKUNORI ISHIHARA

Faculty of Economics and Business Administration, Kyoto University of Advanced Science

KOICHIRO ITO

Harris School of Public Policy, University of Chicago

DAIDO KIDO

Department of Economics, Otaru University of Commerce

TORU KITAGAWA

Department of Economics, Brown University

SHOSEI SAKAGUCHI

Faculty of Economics, University of Tokyo

SHUSAKU SASAKI

Center for Infectious Disease Education and Research, University of Osaka

We develop an optimal policy assignment rule that integrates two distinctive approaches commonly used in economics—targeting by *observables* and targeting through *self-selection*. Our method can be used with experimental or quasi-experimental data to identify who should be treated, be untreated, and self-select to achieve a policymaker's objective. Applying this method to a randomized controlled trial on a residential energy rebate program, we find that targeting that optimally exploits both observable data and self-selection outperforms conventional targeting. We use the Local Average Treatment Effect (LATE) framework (Imbens and Angrist (1994)) to investigate the mechanism in our approach. By estimating several key LATEs based on the random variation created by our experiment, we demonstrate how our method allows policymakers to identify whose self-selection would be valuable and harmful to social welfare.

KEYWORDS: Policy targeting, self-selection, empirical welfare maximization, local average treatment effect, energy rebate programs.

Takanori Ida: ida@econ.kyoto-u.ac.jp

Takunori Ishihara: ishihara.takunori@kuas.ac.jp

Koichiro Ito: ito@uchicago.edu

Daido Kido: dkido@res.otaru-uc.ac.jp

Toru Kitagawa: toru_kitagawa@brown.edu

Shosei Sakaguchi: sakaguchi@e.u-tokyo.ac.jp

Shusaku Sasaki: ssasaki.econ@cider.osaka-u.ac.jp

We would like to thank Keisuke Ito for exceptional research assistance and Judson Boomhower, Severin Borenstein, Fiona Burlig, Lucas Davis, Mark Jacobsen, Ryan Kellogg, Louis Preonas, Wolfram Schlenker, Frank Wolak, and seminar participants at UC Berkeley, the Global Energy Challenge Conference at University of Chicago, and the Center on Japanese Economy and Business (CJEB) at Columbia Business School, NBER Summer Institute Public Economics Meeting for their helpful comments. We thank the Japanese Ministry of Environment for their collaboration for this study. Kitagawa and Sakaguchi gratefully acknowledge financial support from ERC Grant 715940 and the ESRC Centre for Microdata Methods and Practice (CeMMAP) Grant RES-589-28-0001. The AEA RCT Registry number for our experiment is 6139.

1. INTRODUCTION

TARGETING has become a central question in economics and policy design. When policymakers face budget constraints, identifying those who should be treated is critical to maximizing policy impacts. Advances in machine learning and econometric methods have led to a surge in research on targeting in many policy domains, including job training programs (Kitagawa and Tetenov (2018)), social safety net programs (Finkelstein and Notowidigdo (2019), Deshpande and Li (2019)), energy efficiency programs (BKRR+ (2020)), behavioral nudges for electricity conservation (Knittel and Stolper (2021)), and dynamic electricity pricing (Ito, Ida, and Tanaka (2023)).

Economists generally consider two distinctive approaches to the design of effective targeting. The first approach is based on *observable characteristics*. In this approach, policymakers use individuals' observable data to explore optimal targeting (Kitagawa and Tetenov (2018), Athey and Wager (2021)). The second approach is based on *self-selection*. In this approach, policymakers consider individuals' self-selection as valuable information to target certain individual types (Heckman and Vytlacil (2005), Heckman (2010), APWB+ (2016), Ito, Ida, and Tanaka (2023)).

A priori, which approach is desirable for policymakers is unclear. For example, referring to the two distinctive approaches above as “planner’s decisions” and “laissez-faire,” Manski (2013) summarizes,

“The bottom line is that one should be skeptical of broad assertions that individuals are better informed than planners and hence make better decisions. Of course, skepticism of such assertions does not imply that planning is more effective than laissez-faire. Their relative merits depend on the particulars of the choice problem.”

—Charles F. Manski, *Public Policy in an Uncertain World*

A common view in the literature, reflected in this quote, is that the appropriate approach depends on the context, and therefore, researchers and policymakers need to decide which to use on a case-by-case basis. In this study, we develop an optimal policy assignment rule that systematically integrates these two distinctive approaches commonly used in economics. Consider a treatment from which the social welfare gains are heterogeneous across individuals and can be positive, negative, or zero, depending on who takes the treatment. Our comment is that policymakers can leverage both of the *observable* and *unobservable* information by considering three treatment arms, treated, untreated, and free to choose. Once these individual types are identified, policymakers can design a targeting policy that takes advantage of observed and unobserved heterogeneity in the treatment effect.

We begin by formulating this idea by characterizing a social planner’s optimal policy assignment problem following the statistical treatment choice literature (Manski (2004)). We highlight that the Local Average Treatment Effect (LATE) framework (Imbens and Angrist (1994)) can be used to investigate the mechanism in our approach. When individuals have an option to take a treatment, we can define two individual types. *Takers* are those who would take the treatment and *non-takers* are those who would not take the treatment. We demonstrate that the planner’s decision rule can be characterized by the LATEs for takers and non-takers as well as the average treatment effect (ATE), all conditional on individuals’ observable characteristics.

We then show that the optimal policy assignment rule, LATEs for takers and non-takers and the ATE can be identified and estimated by a randomized controlled trial (RCT) or a quasi-experiment with three randomly-assigned groups: an untreated group, a treated group, and a self-selection group. To estimate the optimal policy assignment, we use

the empirical welfare maximization (EWM) method developed by Kitagawa and Tetenov (2018) with policy trees (Zhou, Athey, and Wager (2023)). Further, we demonstrate that the conventional estimation strategy for the LATE (Imbens and Angrist (1994)) can be applied to the three randomly-assigned groups to estimate the LATEs for takers and non-takers.

Our theoretical framework clarifies what variation has to be generated by an RCT or quasi-experiment to estimate the optimal policy assignment. With this insight, we designed an RCT on a residential electricity rebate program and implemented a field experiment in collaboration with the Japanese Ministry of Environment. The policy goal of the rebate program is to incentivize energy conservation in peak demand hours when the marginal cost of electricity tends to be substantially higher than the time-invariant residential electricity price. In our context, the social welfare gain from this rebate program can be heterogeneous across individuals and can be positive, negative, or zero given the existence of per-household implementation cost. This implies that optimal targeting could improve the social welfare gain from this program.

We randomly assigned households to an untreated group, a treated group, and a self-selection group to generate data for our empirical analysis. Using the data from this RCT, we estimate the optimal policy assignment, the ATE, and LATEs for takers and non-takers. We then use our framework to quantify the program's social welfare gain for each of the five policies: (i) all consumers get untreated, (ii) all consumers get treated, (iii) all consumers self-select, (iv) optimal targeting without self-selection (*selection-absent targeting*), and (v) optimal targeting with self-selection (*selection-driven targeting*). Our findings suggest that the selection-driven targeting substantially improves welfare relative to the non-targeting policies and selection-absent targeting. The optimal assignment in the selection-driven targeting suggests that 24% of households in our sample should be untreated, 31% should be treated, and 45% should self-select. The welfare gain from the selection-driven targeting more than doubles compared to the selection-absent targeting.

We then use the LATE framework described above to investigate the mechanism in our optimal policy assignment. Given the random assignment in our field experiment, we are able to estimate the LATEs for takers and non-takers conditional on observables. This implies that we can estimate these LATEs for each of the three groups obtained by the optimal assignment rule. Consider households who would be assigned to the self-selection group by the optimal assignment rule. For these households, we find that the LATE for takers is positive and large, and the LATE for non-takers is negative. Hence, self-selection is useful for the planner to sort customers in this group to get treated or untreated by their choice. In contrast, these two LATEs for those who are not assigned to the self-selection group suggest that allowing self-selection for them would not be optimal because the planner can obtain higher social welfare gains by assigning them to either compulsory treatment or compulsory un-treatment.

Related literature and our contributions—Our study is related to three strands of the literature. First, many recent studies in economics have explored targeting based either on “observables” or “unobservables” through self-selection. Along with the papers cited earlier in this Introduction, recent studies on targeting solely based on individuals' observable characteristics include Johnson, Levine, and Toffel (2023), Murakami, Shimada, Ushifusa, and Ida (2022), Cagala, Glogowsky, Rincke, and Strittmatter (2021), CFMS+ (2024), Gerarden and Yang (2023), and studies on targeting based on self-selection include Dynarski, Libassi, Michelmore, and Owen (2021), Lieber, Ethan, and Lockwood (2019), Unrath (2021), Waldinger (2021). However, to the best of our knowledge, this is

the first study to build an algorithm that systematically integrates these two distinctive targeting approaches to maximize a policy's social welfare gain.¹

Second, the medical statistics literature has studied hybrid sampling designs that combine randomization and treatment choice by patients. See, for example, [JJDL+ \(2003\)](#), [Long, Little, and Lin \(2008\)](#), and references therein. In the medical literature, the sampling process used in our experiment is referred to as “a doubly randomized preference trial” ([Rücker \(1989\)](#)). An example of a clinical trial that implements a doubly randomized preference design is the Woman Take Pride study analyzed in [JJDL+ \(2003\)](#). These studies focus on assessing whether letting patients choose their own treatment can have a direct causal effect on their health status beyond the causal effect of the treatment itself. See [Knox, Yamamoto, Baum, and Berinsky \(2019\)](#) for partial identification analysis in such a context and an application to political science. Doubly randomized preference trials have received less attention in economics. [Bhattacharya \(2013\)](#) is the only study that uses double randomization between randomized control trials and planner's allocation to assess the efficiency of the planner's treatment allocations. To our knowledge, no work has analyzed doubly randomized preference trial data to integrate targeting by observable characteristics and targeting through self-selection.

Third, our econometric framework builds on the growing statistical treatment choice literature. Generally assuming discrete characteristics, earlier studies in this literature ([Manski \(2004\)](#), [Dehejia \(2005\)](#), [Hirano and Porter \(2009\)](#), [Stoye \(2009, 2012\)](#), [Chamberlain \(2011\)](#), [Tetenov \(2012\)](#), among others) formulate estimation of a treatment assignment rule as a statistical decision problem. The empirical welfare maximization approach proposed by [Kitagawa and Tetenov \(2018\)](#) estimates a treatment assignment rule by maximizing the in-sample empirical welfare criterion over a class of assignment rules. As shown in the Supplemental Appendix of [Kitagawa and Tetenov \(2018\)](#) and [Zhou, Athey, and Wager \(2023\)](#), this approach can accommodate multi-armed treatment assignment and a rich set of household characteristics, including continuous characteristics, as in our empirical application. We employ a class of tree partitions considered in [Athey and Wager \(2021\)](#) and [Zhou, Athey, and Wager \(2023\)](#) as our class of policy rules. Finally, building on the LATE framework by [Imbens and Angrist \(1994\)](#), we demonstrate that the newly-defined estimators, the LATEs for *takers* and *non-takers*, can be used to investigate the mechanism in the optimal policy assignment in the presence of self-selection. These LATE estimands can be viewed as the complier's average treatment effects under a multi-valued discrete instrument, which indexes the three arms randomly assigned in the experiment.

2. CONCEPTUAL FRAMEWORK

In this section, we present a theoretical framework of optimal policy assignment in the presence of self-selection. We begin by formulating an optimal policy assignment problem in Section 2.1. In Section 2.2, we present that the LATE framework ([Imbens and Angrist \(1994\)](#)) can be used to investigate the mechanism in our approach. In Section 2.3, we describe how to empirically estimate the optimal policy assignment and LATEs using data from an RCT and the EWM method.

¹For example, [Gerarden and Yang \(2023\)](#) and other recent studies on residential electricity demand explore targeting based on the observable characteristics of customers without using self-selection. Our study differs from these studies because our objective is to develop an algorithm that systematically integrates targeting based on self-selection and observables.

2.1. *Optimal Policy Assignment in the Presence of Self-Selection*

Consider a planner who wishes to introduce a policy intervention (program) to a population of interest. Instead of the uniform assignment over the entire population, the planner is interested in targeted assignment for heterogeneous individuals. A novel feature of our setting is that the planner can control not only who is compulsorily exposed to the program but also who is given an option to opt-in to the program. Interpreting an individual's take-up of the program as their exposure to the treatment, the planner's goal is therefore to assign each individual in the population to one of the three arms: *compulsorily treated* (indexed as T), *compulsorily untreated* (indexed as U), and *self-selection* (indexed as S). An individual assigned to T or U is exposed to or excluded from the program with no opt-out or opt-in option, whereas an individual assigned to S chooses whether to take it up by themselves. In our RCT, the treatment refers to participation in the energy rebate program, rather than assignment to the program. Hence, individuals assigned to T and U are those who are compulsorily exposed to and excluded from the rebate program, respectively. Individuals assigned to S are those who are given the choice to decide whether to participate in the program on their own.

The planner's goal is to optimize a social welfare criterion by assigning individuals to these three arms. Following the statistical treatment choice literature (Manski (2004)), we specify the planner's social welfare criterion to be the sum of individuals' welfare contributions. An individual's welfare contribution is a known function of the individual's response to being assigned to arm T , U , or S , and the per-person cost of the treatment. An individual's welfare contribution may not correspond to their utility. Hence, if an individual is assigned to S , their utility maximizing decision may not correspond to the choice that maximizes the planner's objective. For example, some individuals assigned to S participate in the energy rebate program to obtain monetary benefits but save less electricity consumption than that needed to compensate for the implementation cost of the program for the planner.

Let Y_T , Y_U , and Y_S denote the potential welfare contributions that would be realized if an individual were assigned to T , U , and S . We assume that the planner observes a pre-treatment characteristic vector for each individual $x \in \mathcal{X}$, where \mathcal{X} denotes the support of the characteristics. Depending on these observable characteristics, the planner assigns each individual to one of the three arms. We consider partitioning the characteristics space \mathcal{X} into three subspaces G_T , G_U , and G_S . We denote by $G_T \subseteq \mathcal{X}$ a set of the pre-treatment characteristics x such that any individual whose x belongs to G_T is assigned to T . Similarly, G_U and G_S denote sets of the pre-treatment characteristics x such that the individuals with $x \in G_U$ are assigned to U and individuals with $x \in G_S$ are assigned to S . Let $\mathcal{G} := \{G = (G_T, G_U, G_S) : G \text{ is a measurable partition of } \mathcal{X}\}$ be the set of feasible partitions.

We call a partition $G := (G_T, G_U, G_S)$ an *assignment policy*. G describes how individuals are assigned to arms according to their observable characteristics x . The realized welfare contribution after assignment for an individual with characteristics x is either Y_T , Y_U , or Y_S depending on $x \in G_T$, $x \in G_U$, or $x \in G_S$. Hence, their welfare contribution under the policy G can be written as $\sum_{j \in \{T, U, S\}} Y_j \cdot 1\{x \in G_j\}$. Viewing individual characteristics and their potential welfare contributions as random variables, the average welfare contribution under assignment policy G can be written as

$$\mathcal{W}(G) \equiv E \left[\sum_{j \in \{T, U, S\}} Y_j \cdot 1\{X \in G_j\} \right], \tag{1}$$

where the expectation is with respect to (Y_T, Y_U, Y_S, X) .

We define $\mathcal{W}(G)$ as our social welfare function. The social welfare function depends on the assignment policy G through the post-assignment distribution of individual welfare contributions, which can be manipulated by changing the individuals assigned to the different arms. This form of social welfare is standard in the statistical treatment choice literature. Y_j is not restricted to any specific functional form. Therefore, the planner can choose an appropriate social welfare function.

The planner's objective is to find the optimal assignment policy G^* that maximizes the social welfare $\mathcal{W}(G)$ over a set of possible assignment policies. If the planner can implement any assignment policy, this set of assignment policies corresponds to the set of measurable partitions of \mathcal{X} . G^* can be defined by

$$G^* \in \arg \max_{G \in \tilde{\mathcal{G}}} \mathcal{W}(G). \quad (2)$$

It is desirable that individuals with characteristics x be assigned to an arm that provides the largest conditional mean welfare contribution among $\{E[Y_j|x] : j \in \{T, U, S\}\}$. In the absence of a self-selection treatment arm, the planner's assignment policy is to allocate them to either T or U . The optimal choice is then determined by comparing $E[Y_T|x]$ and $E[Y_U|x]$. In other words, an optimal assignment policy exploits only heterogeneity in the average welfare contribution conditional on observable characteristics x , which can be assessed by the planner prior to assignment (e.g., Johnson, Levine, and Toffel (2023), Murakami, Shimada, Ushifusa, and Ida (2022), Cagala, Glogowsky, Rincke, and Strittmatter (2021), CFMS+ (2024), Gerarden and Yang (2023)). We use G^\dagger to denote this sub-optimal policy assignment and call it *the selection-absent targeting*.

Once individuals are permitted to self-select treatment, social welfare can be improved beyond the level attained by the selection-absent targeting. This is because an individual may possess private information, which drives or helps predict their response to the treatment, and choose whether to receive treatment based on it. Importantly, there can be significant heterogeneity in the usefulness of self-selection for the planner's objective. Individuals with some values of x choose by themselves the treatment that is optimal in terms of the social welfare. In contrast, individuals with other values of x may choose treatment that does not improve social welfare. Moreover, under some conditions x , the individuals would select the same treatment as the social planner. Thus, an optimal assignment policy that identifies who should be assigned to S along with T and U could further improve welfare. We use G^* to denote this optimal policy assignment and call it *the selection-driven targeting*. In this case, the planner allocates individuals with x to either T , U , or S by comparing $E[Y_T|x]$, $E[Y_U|x]$, and $E[Y_S|x]$.²

2.2. Using the LATE Framework to Investigate the Mechanism

In this section, we present a simple model that clarifies how the optimal assignment policy G^* assigns T , U , and S to individuals in accordance with individual observable characteristics x . Let $D_S \in \{0, 1\}$ denote the individual's take-up of treatment when assigned to S . $D_S = 1$ means that the consumer would take the treatment if she is assigned to S , and $D_S = 0$ means that she would not take the treatment if she is assigned to S . The

²This also implies that comparing the sub-optimal assignment policies (such as the assignment policy with T and U only) against the optimal assignment policy allows us to estimate the welfare cost of eliminating an arm option or options.

choice D_S may depend on both observable characteristics X and unobservable characteristics (i.e., private information). Note that the potential outcomes are the quadruple (Y_T, Y_U, Y_S, D_S) .

We define the LATEs for *takers* and *non-takers* as follows, which will be useful statistics to characterize the mechanism of optimal policy assignment.

DEFINITION 2.1—The LATEs for takers and non-takers: We define the LATE for takers by $E[Y_T - Y_U|D_S = 1]$ and the LATE for non-takers by $E[Y_T - Y_U|D_S = 0]$.³

In our RCT, the takers are individuals who voluntarily participate in the energy rebate program. The non-takers, on the other hand, are those who choose not to participate in the program when assigned to S .

Additionally, we make the following assumption. Importantly, this assumption is not required for the validity of our method (Section 2.1) and main empirical results (Section 4), but it is useful to investigate the mechanism in this section and Section 5.

ASSUMPTION 2.2: $Y_S = D_S Y_T + (1 - D_S) Y_U$.

The meaning of Assumption 2.2 is that an individual’s response to the treatment is the same irrespective of whether they self-select themselves or are assigned to it by the planner. That is, who chooses the treatment, either the individuals themselves or the planner, does not have causal impact on the individuals’ outcomes, and this can be viewed as the exclusion restriction for instrumental variables, with an indicator for assignment to the self-selection treatment corresponding to an instrumental variable.

We use $p_1(x) = P(D_S = 1|x)$ and $p_0(x) = P(D_S = 0|x)$ to denote the probability of take-up conditional on x . Under Assumption 2.2, $E[Y_j|x]$ can be decomposed by

$$E[Y_j|x] = \begin{cases} p_1(x) \cdot E[Y_T|D_S = 1, x] + p_0(x) \cdot E[Y_U|D_S = 0, x] & \text{if } j = S, \\ p_1(x) \cdot E[Y_j|D_S = 1, x] + p_0(x) \cdot E[Y_j|D_S = 0, x] & \text{if } j \in \{T, U\}. \end{cases} \quad (3)$$

We can use equation (3) to investigate how the planner ranks the three assignments (T, U, S) for individuals with x . First, consider what condition makes the planner prefer S over U . Equation (3) implies that $E[Y_S - Y_U|x] = p_1(x) \cdot E[Y_T - Y_U|D_S = 1, x]$. Assuming $p_1(x) > 0$, $E[Y_S - Y_U|x] \geq 0$ if and only if $E[Y_T - Y_U|D_S = 1, x] \geq 0$. That is, the LATE for takers has to be greater than or equal to 0. Second, consider what condition makes the planner prefer S over T . Equation (3) implies that $E[Y_S - Y_T|x] = p_0(x) \cdot E[Y_U - Y_T|D_S = 0, x]$. Assuming $p_0(x) > 0$, $E[Y_S - Y_T|x] \geq 0$ if and only if $E[Y_T - Y_U|D_S = 0, x] \leq 0$. That is, the LATE for non-takers has to be less than or equal to 0.

Finally, the condition that makes the planner prefer T over U is trivial such that $E[Y_T - Y_U|x] \geq 0$. Combining the three conditions, we can characterize the optimal assignment

³An alternative way to define $E[Y_T - Y_U|D_S = 1]$ and $E[Y_T - Y_U|D_S = 0]$ is to use the average treatment effects on the treated (ATT) and the average treatment effects on the untreated (ATU). $E[Y_T - Y_U|D_S = 1]$ is the ATT for individuals assigned to group S and $E[Y_T - Y_U|D_S = 0]$ is the ATU for individuals assigned to group S . In our context, these terms could create confusion because there is another ATT for those assigned to the compulsory treatment group (T) and another ATU for those assigned to the compulsory untreated group (U). To avoid this confusion, we use the terms defined in Definition 2.1.




FIGURE 1.—Example of optimal policy assignment G^* . *Notes:* This figure illustrates how an optimal policy rule G^* described in equation (4) partitions the two-dimensional characteristic space \mathcal{X} . Let $ATE(x) := E[W_T - W_U|x]$, $LATE_T(x) := E[W_T - W_U|D_S = 1, x]$ (i.e., the LATE for takers), and $LATE_{NT}(x) := E[W_T - W_U|D_S = 0, x]$ (i.e., the LATE for non-takers). Among the six subspaces from A to F, $ATE(x)$ is non-negative only in A, B, and C; $LATE_T(x)$ is non-negative only in B, C, and D; $LATE_{NT}(x)$ is non-negative only in A, B, and F. Therefore, according to the optimal policy characterization (4), $G_T^* = A \cup B$, $G_U^* = E \cup F$, and $G_S^* = C \cup D$.

policy G^* as defined in equation (2) that has the form $G^* = (G_T^*, G_U^*, G_S^*)$ with

$$\begin{aligned}
 G_T^* &= \{x \in \mathcal{X} : E[Y_T - Y_U|x] \geq 0 \text{ and } E[Y_T - Y_U|D_S = 0, x] > 0\}, \\
 G_U^* &= \{x \in \mathcal{X} : E[Y_T - Y_U|x] < 0 \text{ and } E[Y_T - Y_U|D_S = 1, x] < 0\}, \\
 G_S^* &= \{x \in \mathcal{X} : E[Y_T - Y_U|D_S = 1, x] \geq 0 \text{ and } E[Y_T - Y_U|D_S = 0, x] \leq 0\}.
 \end{aligned}
 \tag{4}$$

Equation (4) implies that the key statistics that characterize the optimal assignment mechanism are the ATE ($E[Y_T - Y_U|x]$), the LATE for takers ($E[Y_T - Y_U|D_S = 1, x]$), and the LATE for non-takers ($E[Y_T - Y_U|D_S = 0, x]$), all conditional on observables. Figure 1 illustrates an example of how an optimal policy G^* is characterized in the two-dimensional characteristic space \mathcal{X} .

2.3. Estimation

In this section, we describe how data from an RCT allow us to estimate the optimal policy assignment (G^*) presented in Section 2.1 and LATEs for takers and non-takers described in Section 2.2. To estimate G^* , we use the EWM method in Kitagawa and Tetenov (2018). Let the RCT data be a size n random sample of (Y_i, Z_i, X_i) , where $Z_i \in \{T, U, S\}$ is individual i 's randomly-assigned treatment arm, Y_i is their observed outcome (welfare contribution), and X_i are their observable pre-treatment characteristics. We use $\{Y_{T,i}, Y_{U,i}, Y_{S,i}\}$ to denote potential outcomes for individual i . The observed outcome Y_i is subject to $Y_i = \sum_{j \in \{T, U, S\}} Y_{j,i} \cdot 1\{Z_i = j\}$. We assume that $\{Y_{T,i}, Y_{U,i}, Y_{S,i}, X_i\}_{i=1, \dots, n}$ are independently and identically distributed as $\{Y_T, Y_U, Y_S, X\}$.

Using the RCT data and a class \mathcal{G} of policies G , the EWM method estimates an optimal policy G^* by maximizing the empirical analogue of the social welfare function over \mathcal{G} :

$$\hat{G}^* \in \arg \max_{G \in \mathcal{G}} \widehat{W}(G),$$

$$\text{where } \widehat{\mathcal{W}}(G) \equiv \frac{1}{n} \sum_{i=1}^n \sum_{j \in \{T, U, S\}} \left(\frac{Y_i \cdot 1\{Z_i = j\}}{P(Z_i = j|X_i)} \cdot 1\{X_i \in G_j\} \right), \quad (5)$$

where $\widehat{\mathcal{W}}(G)$ is an empirical welfare function of G that produces an unbiased estimate of the population social welfare $\mathcal{W}(G)$. Observations are weighted by the inverse of the propensity scores, $P(Z_i = j|X_i)$, which are known from the RCT design.

The EWM approach is model-free: It does not require any assumptions or a functional form specification for the potential outcome distributions. However, the class of policies \mathcal{G} must be specified, considering any feasibility constraints for assignment policies. If the class \mathcal{G} is too rich, the EWM solution \hat{G}^* overfits the RCT data, and the social welfare attained by the estimated policy falls. We use a class of decision trees (Breiman, Friedman, Olshen, and Stone (2017)) as \mathcal{G} because of the ease of interpretation of the decision tree-based assignment policies and the availability of partition search algorithms from the classification tree literature.

We now demonstrate that the LATE for takers ($E[Y_T - Y_U|D_S = 1, x]$) and non-takers ($E[Y_T - Y_U|D_S = 0, x]$) can be also identified and estimated by the RCT data under Assumption 2.2. Supposing that the experimental assignment Z is randomly assigned and that Assumption 2.2 holds, our identification strategy is the same as that of Imbens and Angrist (1994).⁴ We denote the observed take-up by $D \in \{0, 1\}$, which obeys $D = 1\{Z = T\} + 1\{Z = S, D_S = 1\}$. Furthermore, we suppress the dependence on x for ease of notation, although all expectations are taken conditional on x .

First, we discuss the identification and estimation of the LATE for takers. As illustrated in Section 2.2, the ITT between S and U (i.e., $E[Y_S - Y_U]$) equals to $p_1 \cdot E[Y_T - Y_U|D_S = 1]$ where $p_1 = P(D_S = 1)$. Then, the experimental variation of S and U allows us to identify the ITT and p_1 by $E[Y|Z = S] - E[Y|Z = U]$ and $P(D = 1|Z = S)$, respectively. Consequently, the LATE for takers can be identified by

$$E[Y_T - Y_U|D_S = 1] = \frac{E[Y|Z = S] - E[Y|Z = U]}{P(D = 1|Z = S)}. \quad (6)$$

This identification result is simply the application of the conventional LATE framework to experimental groups S and U . Thus, we can estimate this LATE by running the instrumental variable (IV) estimation using data from two groups ($Z \in \{S, U\}$) with the randomly-assigned Z as an instrument for take-up D .

Similarly, the ITT between T and S (i.e., $E[Y_T - Y_S]$) can be written as $p_0 \cdot E[Y_T - Y_U|D_S = 0]$ with $p_0 = P(D_S = 0)$. The experimental variation of T and S allows us to identify the ITT and p_0 by $E[Y|Z = T] - E[Y|Z = S]$ and $P(D = 0|Z = S)$. Thus, the LATE for non-takers can be identified by

$$E[Y_T - Y_U|D_S = 0] = \frac{E[Y|Z = T] - E[Y|Z = S]}{P(D = 0|Z = S)}. \quad (7)$$

This result can be regarded as the application of conventional LATE framework to experimental groups T and S . In our empirical application in Section 2.2, we estimate equations (6) and (7).

⁴In general, the identification of LATE requires monotonicity assumption. In our application, this assumption is automatically satisfied by the nature of groups. In fact, if we define $D_T \in \{0, 1\}$ and $D_U \in \{0, 1\}$ as the individual's potential take-up when assigned to T and U , it always holds that $1 \equiv D_T \geq D_S \geq D_U \equiv 0$ since non-compliance is not allowed under T and U .

Similarly to the testable implications of the instrument validity assumption for LATE models shown by [Balke and Pearl \(1997\)](#), [Imbens and Rubin \(1997\)](#), [Heckman and Vytacil \(2005\)](#), and [Kitagawa \(2015\)](#), non-negativity of the potential outcome distributions for takers and non-takers identified by Assumption 2.2 and the random assignment of Z requires the following inequalities on the distribution of observables:

$$\begin{aligned} f(y|Z = T) &\geq f(y|D = 1, Z = S) \cdot P(D = 1|Z = S), \\ f(y|Z = U) &\geq f(y|D = 0, Z = S) \cdot P(D = 0|Z = S), \end{aligned} \quad (8)$$

for all $y \in \mathbb{R}$, where $f(y|\cdot)$ denotes the probability density function of the observed outcome Y conditional on the corresponding event. The instrument validity test available in the literature such as the test of [Kitagawa \(2015\)](#) can be applied to empirically assess these inequalities and it can serve as a specification test for Assumption 2.2. We perform this test with our data in Section 5.

3. FIELD EXPERIMENT AND DATA

The framework in Section 2 highlighted that data from an RCT can be used to estimate the optimal policy assignment in the presence of self-selection. In this section, we describe how we designed and implemented such an RCT in the context of a residential energy rebate program in Japan. Section 3.1 provides an overview of the field experiment. Section 3.2 presents summary statistics and balance test.

3.1. Field Experiment

We conducted our field experiment in the summer of 2020 in collaboration with the Ministry of the Environment, Government of Japan in the Kansai (around Osaka) and Chubu (around Nagoya) regions of Japan. To include a broad set of households, we invited customers in these regions both by letter and email with a participation reward with 2000 JPY (≈ 20 USD, given $1 \text{ } \text{¢} \approx 1$ JPY in the summer of 2020). A total of 4446 customers pre-registered for the experiment. Non-residential customers, those who canceled their electricity contracts in the middle of the experiment, and those who have incomplete high-frequency electricity usage data were excluded. This left us with 3870 residential customers. That is, our experiment was an RCT for households who agreed to participate in the experiment, which is common in the literature of residential electricity demand ([Wolak \(2011\)](#), [Ito, Ida, and Tanaka \(2023\)](#)).⁵

We randomly assigned the 3870 households to one of the following three groups: an untreated group (U), a treated group (T), and a self-selection group (S).⁶

Untreated group (U): 1577 customers did not participate in the rebate program.

Treated group (T): 1486 customers participated in the rebate program.

Selection group (S): 807 customers were asked to self-select into the rebate program.

⁵Because our experiment was an RCT for households who agreed to participate in the experiment, the external validity of the sample is an important question. To investigate this point, we collected data from a random sample of 2070 customers who resided in the experimental locations but did not participate in the experiment. We find that the experimental sample has slightly higher sample averages in their monthly electricity usage, number of people at home on weekdays, self-efficacy in energy conservation, and household income.

⁶The random assignment process was designed such that $U : T : S = 2 : 2 : 1$. A relatively large number of households were assigned to the U and T groups in consideration that the data for these groups were going to be used for other studies.

The rebate program in our experiment is called the “peak-time rebate” (PTR) program (Wolak (2011)). The fundamental inefficiency in electricity markets in many countries is that residential electricity prices do not fully reflect the time-varying marginal cost of electricity. In peak hours, the time-invariant residential price tends to be too low relative to time-variant marginal cost. This creates a textbook example of short-run deadweight loss. The goal of peak-time rebate programs is to lower this deadweight loss by setting the rebate so that the price minus the rebate is equal to the marginal cost.

The objective of our PTR was to reduce residential electricity consumption in the system peak hours (between 1 pm and 5 pm) during the week of August 24 to 30, 2020. To prevent customers from manipulating their baseline usage, we did not tell them how the baseline was calculated until August. The baseline usage is each customer’s average electricity usage during the peak hours from July 1 to 31. During the treatment week (from August 24 to 30), customers who enrolled in the rebate program received a rebate that was equal to the energy conservation during the peak hours relative to the baseline (kWh) times 100 JPY per kWh. Customers who enrolled in the program were notified about the information about the treatment week, peak hours, and reward calculation procedure in the beginning of August.

Customers in the selection group (S) were asked to send an email or a prepaid postcard during the two-week period from July 31 to August 11 if they intended to participate in the rebate program. The take-up rate was 37.17%, which was rather higher than those for Critical Peak Pricing (CPP) in previous studies.⁷ As mentioned above, the PTR never makes consumers pay more, unlike the CPP treatment, which may have contributed to the higher take-up rate. At the same time, although the PTR would not make any participating household worse off financially, the take-up rate was lower than 100%, which could imply that there were non-financial reasons for a relatively low take-up, including inertia to participate in a new program.

3.2. Data and Summary Statistics

Our primary data are household-level electricity usage over a 30-minute interval. We collected these data in the pre-experimental period (from July 1 to 31, 2020) and the experimental period (from August 24 to 30, 2020). We also conducted a survey before the experiment to collect a variety of household characteristics.

Table I presents summary statistics and balance check. Columns 1, 2, and 3 present the sample averages by the randomly-assigned group with the standard deviations in brackets. Columns 4 to 6 report the difference in sample means with the standard error in parentheses. The first three variables are electricity usage (watt hour per 30-minute) in peak hours (from 1 pm to 5 pm), pre-peak hours (from 10 am to 1 pm), and post-peak hours (from 5 pm to 8 pm). The rest of the variables are from the survey. “Number of people at home” is the number of household members usually at home on weekdays. The survey also asked the self-efficacy in energy conservation using the 5-point Likert scale, in which higher scores imply higher self-efficacy. The household income is reported in 10,000 JPY. “All electric” equals 1 if a customer has an all-electric service with no natural gas service. The survey also asked about the numbers of room air conditioners, electric fans, household members, and the total living area.

⁷The take-up rate for the CPP was 20% in FWBS+ (2021) and 16% (without a take-up incentive) in Ito, Ida, and Tanaka (2023).

TABLE I
SUMMARY STATISTICS AND BALANCE CHECK.

| | Sample Mean by Group [Standard Deviation] | | | Difference in Sample Means (Standard Error) | | |
|--|--|------------------------|--------------------------|--|-----------------|-------------------|
| | Untreated ($Z = U$) | Treated ($Z = T$) | Selection ($Z = S$) | U vs. T | U vs. S | T vs. S |
| Peak hour usage (Wh) | 192 [141] | 190 [138] | 189 [134] | 2.57 (5.03) | 2.87 (5.91) | 0.29 (5.93) |
| Pre-peak hour usage (Wh) | 179 [137] | 176 [135] | 180 [142] | 3.79 (4.92) | -1.11 (6.07) | -4.89 (6.11) |
| Post-peak hour usage (Wh) | 299 [175] | 297 [171] | 293 [174] | 1.94 (6.26) | 6.02 (7.54) | 4.08 (7.56) |
| Number of people at home | 2.48 [1.24] | 2.44 [1.24] | 2.47 [1.27] | 0.04 (0.04) | 0.01 (0.05) | -0.03 (0.06) |
| Self-efficacy in energy conservation (1–5 scale) | 3.45 [0.85] | 3.46 [0.85] | 3.49 [0.83] | -0.01 (0.03) | -0.04 (0.04) | -0.02 (0.04) |
| Household income (JPY 10,000) | 645 [399] | 613 [362] | 637 [391] | 31.69 (13.75) | 8.45 (17.06) | -23.23 (16.67) |
| All electric | 0.32 [0.47] | 0.31 [0.46] | 0.30 [0.46] | 0.01 (0.02) | 0.02 (0.02) | 0.00 (0.02) |
| Number of air conditioners | 3.14 [1.69] | 3.11 [1.71] | 3.08 [1.67] | 0.03 (0.06) | 0.05 (0.07) | 0.02 (0.07) |
| Number of fans | 2.80 [1.63] | 2.73 [1.63] | 2.77 [1.56] | 0.07 (0.06) | 0.04 (0.07) | -0.04 (0.07) |
| Number of household members | 2.76 [1.27] | 2.73 [1.27] | 2.75 [1.28] | 0.04 (0.05) | 0.01 (0.06) | -0.03 (0.06) |
| Total living area (m^2) | 107 [49] | 106 [50] | 103 [46] | 1.78 (1.78) | 3.87 (2.03) | 2.09 (2.07) |

Note: Columns 1–3 present the sample mean and standard deviations in brackets for the pre-experiment consumption data and demographic variables by randomly-assigned group: untreated ($Z = U$), treated ($Z = T$), and selection ($Z = S$). Columns 4–6 show the difference in the sample means with the standard error of the difference in parentheses. The number of households are 1577 (U), 1486 (T), and 807 (S). The monetary unit is given as $1 \text{ €} \approx 1 \text{ JPY}$ in the summer of 2020.

4. OPTIMAL ASSIGNMENT POLICY AND WELFARE GAINS

In this section, we apply the framework developed in Section 2 to our experimental data. In our framework, the planner’s objective is to find the optimal policy assignment rule $G^* = (G_U^*, G_T^*, G_S^*)$ that maximizes the welfare gain $\mathcal{W}(G)$. We define $\mathcal{W}(G)$ in our empirical context in Section 4.1, describe exogenous parameters and estimation details in Section 4.2, and report the results in Section 4.3.

4.1. Construction of the Social Welfare Criterion

We use p and c to denote the price and marginal cost of electricity. In peak hours, the time-invariant residential price p tends to be too low relative to c . The goal of peak-time rebate programs is to reduce welfare loss from this economic inefficiency by setting the rebate incentive equal to c .

Consider a household that takes the rebate program. We use Q_U and Q_T to denote the potential untreated and treated outcomes of electricity consumption. We assume a locally-linear demand curve for electricity usage. Then, the short-run social welfare gain from this program can be written by $\frac{1}{2}(p - c)(Q_T - Q_U)$. Further, we consider that the

reduction in consumption creates an additional long-run social welfare gain as it saves the cost of power plant investments. We denote this long-run gain by $\delta(Q_T - Q_U)$, where δ is the price per kW in the capacity market. Finally, the participation to the rebate program incurs an implementation cost per customer by a .

Then, for each $j \in \{U, T, S\}$, the social welfare gain from the rebate program can be written by

$$\Delta Y_j := Y_j - Y_U = b \cdot (Q_j - Q_U) - a \cdot 1\{D_j = 1\}, \tag{9}$$

where Y_j is the potential outcome of social welfare for $j \in \{U, T, S\}$, Q_j is the potential outcome of electricity usage, D_j is the potential outcome of consumer's take-up of the program for $j \in \{U, T, S\}$, and $b = \frac{1}{2}(p - c) + \delta$. Note that $-b \cdot Q_U$ in equation (9) does not depend on policy assignment, and therefore, we can replace ΔY_j with $Y_j \equiv b \cdot Q_j - a \cdot 1\{D_j = 1\}$ and define a population optimal assignment policy G^* as a maximizer of the criterion of $\mathcal{W}(G) = E[Y_j \cdot 1\{X \in G_j\}]$. Using the sample, we estimate G^* by maximizing the following objective function with respect to G over a class of policies \mathcal{G} :

$$\widehat{\mathcal{W}}(G) = \frac{1}{n} \sum_{i=1}^n \sum_{j \in \{U, T, S\}} \frac{Y_i \cdot 1\{Z_i = j\}}{P(Z_i = j|X_i)} \cdot 1\{X_i \in G_j\}, \tag{10}$$

where i indicates each household in the sample, n is the sample size, $Y_i = b \cdot Q_i - a \cdot 1\{D_i = 1\}$, Q_i is the observed electricity usage for household i , $D_i = 1$ if household i is treated, and $Z_i \in \{U, T, S\}$ is the randomly-assigned group.

4.2. Estimation Details

Equation (9) includes four exogenous parameters: p , c , a , and δ . We use data from the Japanese electricity market during our experimental period to set the values for these parameters. p is the unit price of electricity. We set $p = 25$ JPY/kWh, approximately the regulated price of electricity in Japan, which is independent of the time of day. c is the marginal cost of production for electricity. We specify $c = 125$ JPY/kWh, so that the difference between p and c is equal to the rebate per kWh, which is 100 JPY. The wholesale price of electricity sometimes soars during peak hours such as summer afternoons or winter evenings, reflecting supply constraints. In the past, the wholesale price has occasionally exceeded 100 JPY/kWh in summer afternoons. Parameter a represents the administrative cost of implementing our energy saving program. This cost comprises several items, including the installation cost of the Home Energy Management System (HEMS). In 2016, the Japanese government estimated the cost of implementing a demand reduction program, including the installation cost of HEMS, to be 291.1 JPY per household per season (Ida and Ushifusa (2017)). We use this as the value of the administrative cost.

Parameter δ represents the long-term benefits of a unit reduction in energy consumption. We consider the effect of a unit reduction on the capacity market, where future supply capacity is traded between the power generation and retail sectors. In Japan, the capacity market was established in 2020, with the first auction held at that time. In that auction, the Japanese government provided a reference price 9425 JPY/kW to bidders, which we use as the value for δ .

To estimate the optimal policy G^* , we need to solve the optimization problem with the objective function in Section 4.1. To do so, we specify the policy class to be the class of decision trees of depth 6. We select five variables among candidates to be used in constructing the decision trees. The first two variables are constructed by each household's

hourly electricity usage data in the pre-experimental period: the average usage in peak hours relative to pre-peak hours and the average usage in peak hours relative to post-peak hours. The other three variables are from pre-experimental survey data: household income, the number of household members usually at home on weekdays, and a measure of the households’ self-efficacy in energy conservation. We select these variables by running two off-the-shelf machine learning algorithms, lasso (least absolute shrinkage) and random forest, with all the available covariates and assessing the importance of each variable. For lasso, we regress Q_i on all the available covariates with an l_1 -penalization term. We order variables in terms of importance by increasing the penalization parameter step-wise and checking which variables remain selected for large penalization parameter values. For random forest, we estimate the conditional average treatment effects with the causal forest algorithm of [Wager and Athey \(2018\)](#) with all available covariates. We use the frequency with which a variable is used to split nodes as a measure of its importance. These selected variables are those that appeared on the lists of important variables produced by both methods.

We use the decision tree at depth 6, and maximize the empirical welfare criterion by applying the exhaustive search algorithm of [Zhou, Athey, and Wager \(2023\)](#).⁸ An important technical detail of the EWM estimation is that the optimized empirical welfare value from the estimation will be an upwardly biased estimate of the true welfare attained by the estimated policy. This is known as the winner’s curse bias (see, e.g., [Andrews, Kitagawa, and McCloskey \(2024\)](#)), and is caused by using the same data twice: once to learn the policy and once to infer the policy’s welfare.⁹

To mitigate the winner’s curse bias in our point estimates and confidence intervals, we create an artificial test sample by running random forest regressions of the outcome onto all the covariates with cross-fitting, and generating outcome observations by the sum of regression fits and permuted regression residuals. We then estimate the optimal welfare by the welfare value in the test sample evaluated at the EWM optimal policy obtained in the original sample. One-sided $1 - \alpha$ confidence intervals are constructed by applying the standard normal approximation to t-ratio centered at the point estimate and standard errors estimated with the artificial test sample.

Specifically, we construct test data $\{Y_i^{\text{test}}, Z_i, X_i\}_{i=1}^n$ by generating $(Q_i^{\text{test}}, D_i^{\text{test}})$ in the following procedure and substituting them into $Y_i^{\text{test}} = b \cdot Q_i^{\text{test}} - a \cdot \mathbb{1}\{D_i^{\text{test}} = 1\}$. For each $z \in \{U, T, S\}$, we define $I_z := \{i : Z_i = z\}$ as the set of experimental units assigned to arm z . We randomly partition I_z into 10 equal-sized folds $I_z^{(1)}, \dots, I_z^{(10)}$, and define $I_z^{(-k)} := I_z \setminus I_z^{(k)}$, for $k = 1, \dots, 10$:

1. For each $k = 1, \dots, 10$, assign $D_i^{\text{test}} = 1$ for $i \in I_T^{(k)}$ and $D_i^{\text{test}} = 0$ for $i \in I_U^{(k)}$, respectively. Regress D on X using observations from $I_S^{(-k)}$ to obtain $\hat{P}^{(-k)}(D = 1|X, Z = S)$. For $i \in I_S^{(k)}$, sample $D_i^{\text{test}} \sim \hat{P}^{(-k)}(D = 1|X = X_i, Z = S)$.

⁸ For computational reasons, it is difficult to obtain a globally optimal tree of depth 6 that exactly maximizes the empirical welfare. To alleviate this difficulty, we employ a heuristic two-step procedure to approximate the globally optimal depth-6 tree. Specifically, we first optimize the parent tree of depth 3 that maximizes the empirical welfare in the entire sample; the resulting parent tree divides the entire sample into eight subsamples. Then, for each subsample, we search the child depth-3 tree that maximizes the empirical welfare within the subsample. We finally graft the child depth-3 trees on the parent tree to construct the depth-6 tree. In the machine learning literature, this grafted-tree approach is common when constructing tree classifiers for computational feasibility (see, e.g., Chapter 2 of [Breiman, Friedman, Olshen, and Stone \(2017\)](#) and Section 9.2 in [Hastie, Tibshirani, and Friedman \(2009\)](#)).

⁹ The estimation and inference procedures proposed by [Andrews, Kitagawa, and McCloskey \(2024\)](#) cannot be directly applied to decision tree based policies because the number of candidate policies is infinite.

2. For each $z \in \{T, U\}$,
 - (a) For each $k = 1, \dots, 10$, estimate the conditional expectation function of electricity usage given arm z and covariates X using observations in $I_z^{(-k)}$ to obtain $\hat{E}^{(-k)}[Q|X, Z = z]$, and compute residuals $\hat{\epsilon}_i = Q_i - \hat{E}^{(-k)}[Q|X = X_i, Z = z]$ for $i \in I_z^{(k)}$. We then estimate the conditional variance of the regression residuals $E[\epsilon^2|X, Z = z]$ by regressing $\hat{\epsilon}_i^2$ on X using observations from $I_z^{(-k)}$ to get $\hat{E}^{(-k)}[\epsilon^2|X, Z = z]$, and calculate $\hat{\sigma}_i = \sqrt{\hat{E}^{(-k)}[\epsilon^2|X = X_i, Z = z]}$ for $i \in I_z^{(k)}$.
 - (b) Sample $\{\tilde{\epsilon}_i\}_{i \in I_z}$ independently from the empirical distribution of the standardized residuals $\{\hat{\epsilon}_i/\hat{\sigma}_i\}_{i \in I_z}$, and compute $\epsilon_i^{\text{test}} = \tilde{\epsilon}_i \cdot \hat{\sigma}_i$ for $i \in I_z$.
 - (c) Construct $Q_i^{\text{test}} = \hat{E}^{(-k)}[Q|X = X_i, Z = z] + \epsilon_i^{\text{test}}$ for $i \in I_z^{(k)}$ and $k = 1, \dots, 10$.
3. For $i \in I_S$, we additionally include take-up status D_i in the conditioning variables of the regressions:
 - (a) For each $k = 1, \dots, 10$, we obtain a regression estimate $\hat{E}^{(-k)}[Q|X, D, Z = S]$ using observations of $I_S^{(-k)}$, and compute residuals $\hat{\epsilon}_i = Q_i - \hat{E}^{(-k)}[Q|X = X_i, D = D_i, Z = S]$ for $i \in I_S^{(k)}$. Estimate the residual conditional variance $\hat{E}^{(-k)}[\epsilon^2|X, D, Z = S]$ using the observations $i \in I_S^{(-k)}$ and let $\hat{\sigma}_i(D_i) = \sqrt{\hat{E}^{(-k)}[\epsilon^2|X = X_i, D = D_i, Z = S]}$ for $i \in I_S^{(k)}$.
 - (b) For each $d = 0, 1$, units with $D_i^{\text{test}} = d$ sample $\tilde{\epsilon}_i$ independently from the empirical distribution of $\{\hat{\epsilon}_i/\hat{\sigma}_i(D_i) : D_i = d, i \in I_S\}$ and obtain $\epsilon_i^{\text{test}} = \tilde{\epsilon}_i \cdot \hat{\sigma}_i(d)$.
 - (c) Construct $Q_i^{\text{test}} = \hat{E}^{(-k)}[Q|X = X_i, D = D_i^{\text{test}}, Z = S] + \epsilon_i^{\text{test}}$ for $i \in I_S^{(k)}$ and $k = 1, \dots, 10$.

In this procedure, we estimate the conditional expectation functions of D_i and Q_i and the conditional variances of the residuals $\hat{\epsilon}_i$ using random forests (Friedberg, Tibshirani, Athey, and Wager (2021), Wager and Athey (2018)).

With these test data, we obtain a point estimator for the maximized welfare gain $\Delta\mathcal{W}(G^*) = \mathcal{W}(G^*) - \mathcal{W}_U$ relative to the welfare level \mathcal{W}_U attained by the uniformly untreated policy by

$$\widehat{\Delta\mathcal{W}(G^*)} \equiv \frac{1}{n} \sum_{i=1}^n \sum_{j \in \{U, T, S\}} \left(\frac{Y_i^{\text{test}} \cdot 1\{Z_i = j\}}{P(Z_i = j|X_i)} \cdot 1\{X_i \in \hat{G}_j^*\} \right) - \frac{1}{n} \sum_{i=1}^n \frac{Y_i^{\text{test}} \cdot 1\{Z_i = U\}}{P(Z_i = U|X_i)}, \quad (11)$$

where \hat{G}^* is an EWM policy defined in (5) constructed upon the original sample. We form one-sided confidence intervals for the maximal welfare gain $\Delta\mathcal{W}(G^*)$ with coverage $1 - \alpha$ by $[\widehat{\Delta\mathcal{W}(G^*)} - z_{1-\alpha} \cdot \hat{\sigma}_{\mathcal{W}}/n^{1/2}, \infty]$, where $z_{1-\alpha}$ is the $(1 - \alpha)$ th quantile of the standard normal distribution and $\hat{\sigma}_{\mathcal{W}}$ is a standard deviation estimator for the summands in equation (11) with \hat{G}^* fixed.

Our approach constructs a test sample by resampling the residuals from cross-fitted regressions. In contrast to sample splitting, we do not sacrifice the sample sizes for training and test samples. At the same time, we can mitigate the winner’s bias by performing cross-fitting. We conduct an empirical Monte Carlo study (EMCS) to validate the performance of this method in Section 4.4.

4.3. Results of the Optimal Policy Assignment

We estimate the optimal policy assignment that maximizes social welfare based on equation (5) in Section 2. We compare five alternative policies: (1) assigning everyone

TABLE II
WELFARE GAINS FROM EACH POLICY.

| Policy | Welfare Gain | Share of Customers in Each Arm | | | Share of Treated Customers |
|--|------------------|--------------------------------|--------|--------|----------------------------|
| | | G_U | G_T | G_S | |
| 100% untreated | 0.0 (—) | 100.0% | 0.0% | 0.0% | 0.0% |
| 100% treated | 120.7 (98.8) | 0.0% | 100.0% | 0.0% | 100.0% |
| 100% self-selection | 180.6 (112.1) | 0.0% | 0.0% | 100.0% | 37.2% |
| Selection-absent targeting (G^\dagger) | 186.4 (66.9) | 47.6% | 52.4% | 0.0% | 52.4% |
| Selection-driven targeting (G^*) | 477.0 (87.2) | 23.9% | 31.4% | 44.7% | 48.8% |

Note: This table summarizes characteristics of three benchmark policies (100% untreated, 100% treated, and 100% self-selection), selection-absent targeting (G^\dagger), and selection-driven targeting (G^*). The column titled “Welfare Gain” shows the estimated ITT of welfare gain in JPY per household per season, with its standard error in parentheses. The monetary unit is given as $1 \text{ €} \approx 1 \text{ JPY}$ in the summer of 2020.

to U , (2) assigning everyone to T , (3) assigning everyone to S , (4) the selection-absent targeting G^\dagger , and (5) the selection-driven targeting G^* .

In Table II, we present the welfare performances of three benchmark policies without targeting (100% U , 100% T , and 100% S) estimated by the sample averages in the original sample, and those of the sub-optimal and optimal targeting policies (G^\dagger and G^*) estimated with the test sample method shown in the previous section. For each policy, we estimate the ITT of the welfare gain in JPY per household per season. We find that the 100% T policy induces a welfare gain of 120.7 per consumer, but the effect is not statistically significant. The 100% S policy results in a welfare gain by 180.6 per consumer and is marginally significant at a p -value of 0.107. These results suggest that without targeting, we cannot reject that the policy’s net welfare gain can be zero.

Our policy intervention induces both cost (from the implementation cost) and benefit (from the energy conservation), and therefore, the net welfare gain from a consumer can be positive, negative, or zero. This implies that we could increase the policy performance by targeting policies, G^\dagger and G^* . The results in Table II suggest that the selection-absent targeting (G^\dagger) attains a welfare gain by 186.4 per consumer. Our algorithm identifies that 52.4% of consumers should be treated, and 47.6% of them should be untreated. Furthermore, we find that the selection-driven targeting (G^*) results in a welfare gain of 477.0 per consumer. With this policy, our algorithm identifies that 31.4% of consumers should be treated, 23.9% of them should be untreated, and 44.7% of them should self-select.

In Table III, we statistically compare the welfare gains between the alternative policies. The results imply that the selection-driven targeting (G^*) statistically and economically outperforms other policies. Compared to the selection-absent targeting (G^\dagger), it generates an additional welfare gain by 290.6, which makes its welfare gain more than double the one obtained by the selection-absent targeting.

Table IV presents the covariates distribution by the optimal policy assignment group $G^* = (G_U^*, G_T^*, G_S^*)$. Columns 1, 2, and 3 show the mean and standard deviation by group, and columns 4, 5, and 6 show the difference between the means and their standard errors. For example, the means of household income indicate that higher-income households are

TABLE III
COMPARISONS OF ALTERNATIVE POLICIES.

| | Difference in Welfare Gains | p -Value |
|---|-----------------------------|------------|
| 100% self-selection vs. 100% treated | 59.9 (110.0) | 0.293 |
| Selection-absent targeting (G^\dagger) vs. 100% treated | 65.7 (113.9) | 0.282 |
| Selection-absent targeting (G^\dagger) vs. 100% self-selection | 5.8 (126.9) | 0.482 |
| Selection-driven targeting (G^*) vs. 100% treated | 356.4 (125.7) | 0.002 |
| Selection-driven targeting (G^*) vs. 100% self-selection | 296.5 (133.6) | 0.013 |
| Selection-driven targeting (G^*) vs. selection-absent targeting (G^\dagger) | 290.6 (78.9) | 0.000 |

Note: This table compares welfare gains from each policy. For each row, the column “Difference in Welfare Gains” shows the estimated welfare gain of the policy on the left-hand side (W_L) relative to the policy on the right-hand side (W_R) in JPY per household per season, with its standard error in parentheses. The column “ p -value” gives the p -value for the null hypothesis: $H_0 : W_L \leq W_R$. The monetary unit is given as $1 \text{ €} \approx 1 \text{ JPY}$ in the summer of 2020.

more likely to be assigned to U rather than T or S . Similarly, the means of self-efficacy in energy conservation suggest that households with lower efficacy in energy conservation are more likely to be assigned to U .

4.4. Monte Carlo Simulation

As described in Section 4.2, we perform resampling of regression residuals with cross-fitting to address the winner’s curse bias in the estimation of welfare. As we do not have an analytical claim on the bias and coverage properties, we conduct an empirical Monte Carlo study (EMCS) in this section to present evidence for valid performance.

To make the EMCS’s data generating process closely tailored to our experimental data, we use the approach developed by [Athey, Imbens, Metzger, and Munro \(2024\)](#). We begin by estimating the conditional expectation of the observed electricity consumption Q_i given the treatment take-up $D_i \in \{0, 1\}$, the experimental arm $Z_i \in \{U, T, S\}$, all covariates X_i , and the conditional expectation of D_i on Z_i and X_i . We apply random forest to each of these regressions. In addition, we apply the conditional Wasserstein-GAN proposed by [Athey, Imbens, Metzger, and Munro \(2024\)](#) to learn the conditional distribution $\hat{F}_{X|Z}$ of the covariates given the experimental arm and the conditional distribution $\hat{F}_{\epsilon|D,Z,X}$ of the electricity consumption residuals $\epsilon = Q - \hat{E}[Q|D, Z, X]$ given (D, Z, X) .¹⁰

We then draw 500 Monte Carlo data sets $\{(Y_i, Z_i, X_i)\}_{i=1}^n$ of size $n = 3870$ by the following steps: (i) randomly draw the experimental arm Z_i to match the sample proportions in the original data; (ii) draw the covariate X_i from $\hat{F}_{X|Z_i}$; (iii) draw the treatment take-up D_i from a Bernoulli distribution with parameter $\hat{E}[D|Z_i, X_i]$; (iv) draw the error term ϵ_i

¹⁰This conditional distribution $\hat{F}_{\epsilon|D,Z,X}$ is not guaranteed to have its mean zero, since the residuals constructed from the initial nonparametric regression are not guaranteed to have conditional mean zero. To address this issue, we symmetrize its distribution by drawing a residual from $\hat{F}_{\epsilon|D,Z,X}$ and multiply $+1$ or -1 randomly with an equal chance.

TABLE IV
COVARIATE DISTRIBUTION BY OPTIMALLY ASSIGNED GROUP G^* .

| | Sample Mean by Group [Standard Deviation] | | | Difference in Sample Means (Standard Error) | | |
|--|--|----------------|----------------|--|---------------------|---------------------|
| | G_U^* | G_T^* | G_S^* | G_U^* vs. G_T^* | G_U^* vs. G_S^* | G_T^* vs. G_S^* |
| Peak hour usage (Wh) | 203 [146] | 180 [136] | 191 [135] | 23.03 (6.18) | 11.98 (5.79) | -11.05 (5.08) |
| Pre-peak hour usage (Wh) | 198 [150] | 167 [133] | 175 [132] | 30.56 (6.23) | 23.08 (5.86) | -7.48 (4.97) |
| Post-peak hour usage (Wh) | 329 [176] | 255 [176] | 310 [164] | 73.22 (7.67) | 18.82 (7.00) | -54.40 (6.41) |
| Number of people at home | 2.87 [1.34] | 2.27 [1.32] | 2.38 [1.08] | 0.60 (0.06) | 0.48 (0.05) | -0.11 (0.05) |
| Self-efficacy in energy conservation (1-5 scale) | 3.30 [1.02] | 3.49 [0.82] | 3.53 [0.75] | -0.19 (0.04) | -0.23 (0.04) | -0.04 (0.03) |
| Household income (JPY 10,000) | 787 [433] | 597 [397] | 572 [318] | 190.12 (18.23) | 215.11 (16.15) | 25.00 (13.73) |
| All electric | 0.36 [0.48] | 0.25 [0.43] | 0.33 [0.47] | 0.11 (0.02) | 0.03 (0.02) | -0.08 (0.02) |
| Number of air conditioners | 3.41 [1.72] | 2.82 [1.66] | 3.16 [1.67] | 0.58 (0.07) | 0.24 (0.07) | -0.34 (0.06) |
| Number of fans | 2.99 [1.75] | 2.58 [1.57] | 2.78 [1.55] | 0.41 (0.07) | 0.20 (0.07) | -0.21 (0.06) |
| Number of household members | 3.17 [1.31] | 2.54 [1.36] | 2.67 [1.14] | 0.63 (0.06) | 0.50 (0.05) | -0.13 (0.05) |
| Total living area (m^2) | 115 [48] | 97 [49] | 107 [47] | 18.25 (2.11) | 8.68 (1.94) | -9.57 (1.81) |

Note: This table shows the covariate distribution by group based on the optimal policy assignment G^* .

from the symmetrized distribution of $\epsilon|D_i, Z_i, X_i; (v)$ set the electricity consumption Q_i by $Q_i = \hat{E}[Q|D_i, Z_i, X_i] + \epsilon_i$ and the welfare contribution Y_i by $Y_i = b \cdot Q_i - a \cdot 1\{D_i = 1\}$.

Using these Monte Carlo samples, we evaluate the bias and coverage of our estimator and confidence intervals for the maximal welfare within the class of decision trees of depth 6.¹¹ In Table V, we compare three methods in terms of the bias and standard errors of the estimator and the coverage of one-sided confidence intervals. The first method is a naive approach that regards the in-sample optimized empirical welfare contrast as an estimator for the population maximal welfare gain and forming CIs under the assumption of asymptotic normality of the point estimate. The second method is our estimator $\Delta\widehat{\mathcal{W}}(G^*)$ and the associated confidence intervals based on resampling of the cross-fitted residuals. The third method is sample splitting, in which we randomly split the sample into training and test subsamples of equal size. We use the training subsample to estimate an optimal policy and the test subsample to estimate and infer the welfare gain.

¹¹Since the exact computation of the population optimal policy is difficult, we approximate it by $\arg \max_{G=(G_U, G_T, G_S)} \sum_{i=1}^{10,000} \sum_{z \in \{U, T, S\}} \hat{E}[Y|Z=z, X_i^{\text{policy}}] \cdot 1\{X_i^{\text{policy}} \in G_z\}$, where the maximization of G is over the class of tree partitions of depth 6, X_i^{policy} is the covariates that are used for targeting, and $\hat{E}[Y|Z=z, X_i^{\text{policy}}]$, $z \in \{U, T, S\}$ is the regression estimate of $\hat{E}[Y|Z=z, X_i]$ specified for EMCS data generating process on X_i^{policy} . We then draw covariate data $\{X_i\}$ of size 1,000,000, and evaluate the welfare at the optimal policy based on $\hat{E}[Y|Z=z, X_i]$, $z \in \{U, T, S\}$.

TABLE V
MONTE CARLO SIMULATION RESULTS.

| | Naive | | | Resampling with Cross-Fitting | | | Sample Splitting | | |
|-----------------------|-------|------|----------|-------------------------------|------|----------|------------------|-------|----------|
| | Bias | S.E. | Coverage | Bias | S.E. | Coverage | Bias | S.E. | Coverage |
| S vs. T | 1.0 | 74.8 | 0.94 | 1.0 | 74.8 | 0.94 | 3.6 | 105.3 | 0.94 |
| G^\dagger vs. T | 459.3 | 41.3 | 0.00 | 29.1 | 49.9 | 0.97 | -77.8 | 63.7 | 1.00 |
| G^\dagger vs. S | 458.3 | 67.6 | 0.00 | 28.1 | 80.1 | 0.92 | -81.4 | 103.7 | 1.00 |
| G^* vs. T | 573.1 | 51.2 | 0.00 | -8.0 | 73.1 | 0.98 | -126.2 | 82.9 | 1.00 |
| G^* vs. S | 572.1 | 50.2 | 0.00 | -9.0 | 70.2 | 0.99 | -129.8 | 85.5 | 1.00 |
| G^* vs. G^\dagger | 113.8 | 42.9 | 0.25 | -37.0 | 60.9 | 0.98 | -48.4 | 82.5 | 0.98 |

Note: T is 100% treated, S is 100% self-selection, G^\dagger is the selection-absent targeting, and G^* is the selection-driven targeting. In sample splitting, 50% of the sample is used for training. In resampling with cross-fitting, we perform inference based on test samples constructed by resampled residuals with cross-fitting. See the method shown in Section 4.2.

Our main finding is that the resampling with cross-fitting method outperforms the other two methods in terms of the bias correction, coverage, and precision. The naive method is subject to severe upward bias and under-coverage due to winner’s curse bias. The sample splitting method meets the desired coverage while the point estimates are biased downward since insufficient training sample size forces limited learning of the optimal policy. It also sacrifices the precision of the estimator. The resampling with cross-fitting method provides smaller bias, better coverage, and precision.

5. USING THE LATE FRAMEWORK TO UNCOVER THE MECHANISM

As presented in Section 2.2, an advantage of our research design is that we can identify both the LATE for *takers* ($E[Y_T - Y_U|D_S = 1]$) and the LATE for *non-takers* ($E[Y_T - Y_U|D_S = 0]$). In this section, we demonstrate that these two LATEs can be used to examine the mechanism in the selection-driven targeting.

As shown in equation (6) in Section 2.2, we can use the conventional LATE framework by [Imbens and Angrist \(1994\)](#) to demonstrate that $E[Y_T - Y_U|D_S = 1] = \frac{E[Y|Z=S] - E[Y|Z=U]}{P(D=1|Z=S)}$, where $Z = \{S, U\}$ is randomly assigned in our RCT, $Z = S$ is the selection group, $Z = U$ is the untreated group, and D is the observed treatment take-up for those who were assigned to $Z = S$. The numerator of the right-hand side of the equation is the difference in the ITTs between groups S and U , and the denominator is the take-up rate in groups S . Therefore, the sample analogue of this equation can be estimated from our experimental data.¹² A unique feature of our research design is that we have a randomly-assigned compulsory treatment group ($Z = T$) along with groups $Z = \{S, U\}$. As presented in equation (6) in Section 2.2, we can use two groups $Z = \{S, T\}$ to estimate the LATE for non-takers by $E[Y_T - Y_U|D_S = 0] = \frac{E[Y|Z=T] - E[Y|Z=S]}{P(D=0|Z=S)}$.

¹²Equation (6) shows that the LATE for takers is equivalent to the LATE for compliers when we consider two groups with a binary instrument $Z = \{U, S\}$. This implies that we can use the conventional IV estimation to estimate equation (6) under the regular assumptions for identifying the LATE. In particular, a key assumption is the exclusion restriction in Assumption 2.2..

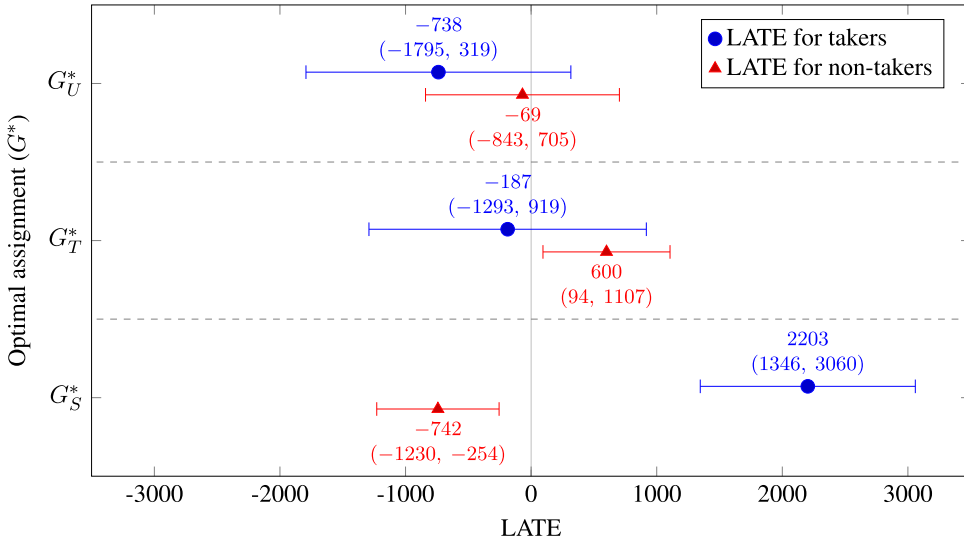


FIGURE 2.—Mechanism behind the algorithm: The LATEs for takers and non-takers. *Notes:* This figure shows the estimation results in Section 5. For each of the three groups in the optimal assignment ($x \in G_U^*$, $x \in G_T^*$, $x \in G_S^*$), we estimate the LATE for *takers* ($E[Y_T - Y_U | D_S = 1]$) and the LATE for *non-takers* ($E[Y_T - Y_U | D_S = 0]$) to investigate the mechanism in the optimal assignment. We show the point estimates with the 95% confidence intervals. For example, for those who are assigned to the selection group (G_S^*), the LATE for takers is 2203, and the LATE for non-takers is -742 . This implies that self-selection is a useful tool for this group to let customers sort into the treatment choice that is in line with the planner’s objective. The point estimates and confidence intervals are calculated with test data constructed by the method delineated in Section 4.2. The monetary unit is given as $1 \text{ €} \approx 1 \text{ JPY}$ in the summer of 2020.

The LATEs for takers and non-takers can be estimated conditional on X because the randomization of $Z = (U, T, S)$ holds given X . This implies that we can estimate these LATEs by customer types based on X . Consider the optimal assignment rule with the selection-driven targeting policy $G^* = (G_U^*, G_T^*, G_S^*)$. This policy divides customers into three groups based on their observables: those who should be untreated ($X \in G_U^*$), those who should be treated ($X \in G_T^*$), and those who should self-select ($X \in G_S^*$).

In Figure 2, we estimate equations (6) and (7) for these three groups, G_U^* , G_T^* , and G_S^* . For those who are assigned to the selection group (G_S^*), the LATE for takers is 2203 and the LATE for non-takers is -742 . This implies that self-selection is a useful tool for this group to let customers sort into the treatment choice that is in line with the planner’s objective. By contrast, if we allow self-selection for customers in G_U^* , it is likely to decrease welfare because the LATE for takers is -738 . Similarly, if we allow self-selection in G_T^* , it is likely to lower welfare because self-selection would make the non-takers untreated even though their LATE is positive and large at 600. Therefore, the LATE for takers and non-takers presented in Figure 2 highlights how our algorithm chooses who should get treated, untreated, and choose to get treated by themselves.¹³

¹³To empirically assess Assumption 2.2, we perform the test developed by Kitagawa (2015) for the null hypothesis of the inequalities (8). We find that the p -value of this test is 1.000 and insensitive to various choices of tuning parameters, which provides supporting evidence for Assumption 2.2 with our data.

6. CONCLUSION

We develop an optimal policy assignment rule that systematically integrates two distinctive approaches commonly used in the literature—targeting by “observables” and targeting through “self-selection.” Our method identifies those who should be treated, should be untreated, and should self-select into a treatment to maximize a policy’s social welfare gain. We show that targeting that leverages information on both observables and self-selection outperforms conventional targeting. Finally, we use the LATE framework (Imbens and Angrist (1994)) to uncover the mechanism in our approach. We introduce new estimators, the LATEs for *takers* and *non-takers*, to demonstrate how our method identifies whose self-selection is useful and harmful for the planner to maximize social welfare.

REFERENCES

- [APWB+] Alatas, Vivi, Ririn Purnamasari, Matthew Wai-Poi, Abhijit Banerjee, Benjamin A. Olken, and Rema Hanna (2016), “Self-Targeting: Evidence From a Field Experiment in Indonesia.” *Journal of Political Economy*, 124, 371–427. [0226]
- Andrews, Isaiah S., Toru Kitagawa, and Adam McCloskey (2024), “Inference on Winners.” *The Quarterly Journal of Economics*, 139, 305–358. [0238]
- Athey, Susan and Stefan Wager (2021), “Efficient Policy Learning With Observational Data.” *Econometrica*, 89, 133–161. [0226,0228]
- Athey, Susan, Guido W. Imbens, Jonas Metzger, and Evan Munro (2024), “Using Wasserstein Generative Adversarial Networks for the Design of Monte Carlo Simulations.” *Journal of Econometrics*, 240, 105076. [0241]
- Balke, Alexander and Judea Pearl (1997), “Bounds on Treatment Effects From Studies With Imperfect Compliance.” *Journal of the American Statistical Association*, 92, 1171–1176. [0234]
- Bhattacharya, Debopam (2013), “Evaluating Treatment Protocols by Combining Experimental and Observational Data.” *Journal of Econometrics*, 173, 160–174. [0228]
- Breiman, Leo, Jerome H. Friedman, Richard A. Olshen, and Charles J. Stone (2017), *Classification and Regression Trees*. Routledge. [0233,0238]
- [BKRR+] Burlig, Fiona, Christopher Knittel, David Rapson, Mar Reguant, and Catherine Wolfram (2020), “Machine Learning From Schools About Energy Efficiency.” *Journal of the Association of Environmental and Resource Economists*, 7, 1181–1217. [0226]
- Cağala, Tobias, Ulrich Glogowsky, Johannes Rincke, and Anthony Strittmatter (2021), “Optimal Targeting in Fundraising: A Causal Machine-Learning Approach.” arXiv:2103.10251. arXiv preprint. [0227,0230]
- Chamberlain, Gary (2011), “Bayesian Aspects of Treatment Choice.” In *The Oxford Handbook of Bayesian Econometrics* (John Geweke, Gary Koop, and Herman van Dijk, eds.), 11–39, Oxford University Press. [0228]
- [CFMS+] Christensen, Peter, Paul Francisco, Erica Myers, Hansen Shao, and Mateus Souza (2024), “Energy Efficiency Can Deliver for Climate Policy: Evidence From Machine Learning-Based Targeting.” *Journal of Public Economics*, 234, 105098. [0227,0230]
- Dehejia (2005), “Program Evaluation as a Decision Problem.” *Journal of Econometrics*, 125, 141–173. [0228]
- Deshpande, Manasi and Yue Li (2019), “Who Is Screened out? Application Costs and the Targeting of Disability Programs.” *American Economic Journal: Economic Policy*, 11, 213–248. [0226]
- Dynarski, Susan, CJ Libassi, Katherine Micheltore, and Stephanie Owen (2021), “Closing the Gap: The Effect of a Targeted, Tuition-Free Promise on College Choices of High-Achieving, Low-Income Students.” *American Economic Review*, 111, 1721–1756. [0227]
- Finkelstein, Amy and Matthew J. Notowidigdo (2019), “Take-up and Targeting: Experimental Evidence From SNAP.” *Quarterly Journal of Economics*, 134, 1505–1556. [0226]
- [FWBS+] Fowlie, Meredith, Catherine Wolfram, Patrick Baylis, C. Anna Spurlock, Annika Todd-Blick, and Peter Cappers (2021), “Default Effects and Follow-on Behaviour: Evidence From an Electricity Pricing Program.” *The Review of Economic Studies*, 88, 2886–2934. [0235]
- Friedberg, Rina, Julie Tibshirani, Susan Athey, and Stefan Wager (2021), “Local Linear Forests.” *Journal of Computational and Graphical Statistics*, 30, 503–517. [0239]
- Gerarden, Todd D. and Muxi Yang (2023), “Using Targeting to Optimize Program Design: Evidence From an Energy Conservation Experiment.” *Journal of the Association of Environmental and Resource Economists*, 10, 687–716. [0227,0228,0230]

- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2009), *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2 edition. Springer Series in Statistics. Springer, New York. [0238]
- Heckman, James J. (2010), “Building Bridges Between Structural and Program Evaluation Approaches to Evaluating Policy.” *Journal of Economic Literature*, 48, 356–398. [0226]
- Heckman, James J. and Edward Vytlacil (2005), “Structural Equations, Treatment Effects, and Econometric Policy Evaluation.” *Econometrica*, 73, 669–738. [0226,0234]
- Hirano, Keisuke and Jack R. Porter (2009), “Asymptotics for Statistical Treatment Rules.” *Econometrica*, 77, 1683–1701. [0228]
- Ida, Takanori and Yoshiaki Ushifusa (2017), “Cost-Benefit Analysis of Price-Based Residential Demand Response.” *Proceedings of the Japan Joint Automatic Control Conference*, 60, 304–307. [0237]
- Imbens, Guido W. and Joshua D. Angrist (1994), “Identification and Estimation of Local Average Treatment Effects.” *Econometrica*, 62, 467–475. [0226-0228,0233,0243,0245]
- Imbens, Guido W. and Donald B. Rubin (1997), “Estimating Outcome Distributions for Compliers in Instrumental Variables Models.” *The Review of Economic Studies*, 64, 555–574. [0234]
- Ito, Koichiro, Takanori Ida, and Makoto Tanaka (2023), “Selection on Welfare Gains: Experimental Evidence From Electricity Plan Choice.” *American Economic Review*, 113, 2937–2973. [0226,0234,0235]
- [JJDL+] Janevic, Mary R., Nancy K. Janz, Julia A. Dodge, Xihong Lin, Wenqin Pan, Brandy R. Sinco, and Noreen M. Clark (2003), “The Role of Choice in Health Education Intervention Trials: A Review and Case Study.” *Social Science and Medicine*, 56, 1581–1594. [0228]
- Johnson, Matthew S., David I. Levine, and Michael W. Toffel (2023), “Improving Regulatory Effectiveness Through Better Targeting: Evidence From Osha.” *American Economic Journal: Applied Economics*, 15, 30–67. [0227,0230]
- Kitagawa, Toru (2015), “A Test for Instrument Validity.” *Econometrica*, 83, 2043–2063. [0234,0244]
- Kitagawa, Toru and Aleksey Tetenov (2018), “Who Should Be Treated? Empirical Welfare Maximization Methods for Treatment Choice.” *Econometrica*, 86, 591–616. [0226-0228,0232]
- Knittel, Christopher R. and Samuel Stolper (2021), “Machine Learning About Treatment Effect Heterogeneity: The Case of Household Energy Use.” *AEA Papers and Proceedings*, 111, 440–444. [0226]
- Knox, Dean, Tepei Yamamoto, Matthew A. Baum, and Adam J. Berinsky (2019), “Design, Identification, and Sensitivity Analysis for Patient Preference Trials.” *Journal of the American Statistical Association*, 114, 1532–1546. [0228]
- Lieber, Ethan M.J., and Lee M. Lockwood (2019), “Targeting With in-Kind Transfers: Evidence From Medicaid Home Care.” *American Economic Review*, 109, 1461–1485. [0227]
- Long, Qi, Roderick Little, and Xihong Lin (2008), “Causal Inference in Hybrid Intervention Trials Involving Treatment Choice.” *Journal of the American Statistical Association*, 103, 474–484. [0228]
- Manski, Charles F. (2004), “Statistical Treatment Rules for Heterogeneous Populations.” *Econometrica*, 72, 1221–1246. [0226,0228,0229]
- Manski, Charles F. (2013), *Public Policy in an Uncertain World*. Harvard University Press, Cambridge, MA. [0226]
- Murakami, Kayo, Hideki Shimada, Yoshiaki Ushifusa, and Takanori Ida (2022), “Heterogeneous Treatment Effects of Nudge and Rebate: Causal Machine Learning in a Field Experiment on Electricity Conservation.” *International Economic Review*, 63, 1779–1803. [0227,0230]
- Rücker, Gerta (1989), “A Two-Stage Trial Design for Testing Treatment, Self-Selection, and Treatment Preference Effects.” *Statistics in Medicine*, 8, 477–485. [0228]
- Stoye, Jörg (2009), “Minimax Regret Treatment Choice With Finite Samples.” *Journal of Econometrics*, 151, 70–81. [0228]
- Stoye, Jörg (2012), “Minimax Regret Treatment Choice With Covariates or With Limited Validity of Experiments.” *Journal of Econometrics*, 166, 138–156. [0228]
- Tetenov, Aleksey (2012), “Statistical Treatment Choice Based on Asymmetric Minimax Regret Criteria.” *Journal of Econometrics*, 166, 157–165. [0228]
- Unrath, Matthew (2021), “Targeting, Screening, and Retention: Evidence From California’s Food Stamps Program.” California Policy Labo Working paper, No2021-1. <https://capolicylab.org/wp-content/uploads/2021/02/CalFresh-Working-Paper.pdf>. [0227]
- Wager, Stefan and Susan Athey (2018), “Estimation and Inference of Heterogeneous Treatment Effects Using Random Forests.” *Journal of the American Statistical Association*, 113, 1228–1242. [0238,0239]
- Waldinger, Daniel (2021), “Targeting in-Kind Transfers Through Market Design: A Revealed Preference Analysis of Public Housing Allocation.” *American Economic Review*, 111, 2660–2696. [0227]
- Wolak, Frank A. (2011), “Do Residential Customers Respond to Hourly Prices? Evidence From a Dynamic Pricing Experiment.” *The American Economic Review*, 101, 83–87. [0234,0235]

Zhou, Zhengyuan, Susan Athey, and Stefan Wager (2023), "Offline Multi-Action Policy Learning: Generalization and Optimization." *Operations Research*, 71, 148–183. [0227,0228,0238]

Co-editor Guido W. Imbens handled this manuscript.

Manuscript received 19 September, 2022; final version accepted 26 September, 2025; available online 6 October, 2025.

The replication package for this paper is available at <https://doi.org/10.5281/zenodo.17074824>. The authors were granted an exemption to publish parts of their data because either access to these data is restricted or the authors do not have the right to republish them. Therefore, the replication package only includes the codes and the parts of the data that are not subject to the exemption. However, the authors provided the Journal with (or assisted the Journal to obtain) temporary access to the restricted data. The Journal checked the provided and restricted data and the codes for their ability to reproduce the results in the paper and approved online appendices. Given the highly demanding nature of the algorithms, the reproducibility checks were run on a simplified version of the code, which is also available in the replication package.

All authors assume responsibility for all aspects of the paper.