

Using Targeting to Optimize Program Design: Evidence from an Energy Conservation Experiment

Todd Gerarden and Muxi Yang

Cornell University

IAEE, 2021

Research Objectives

- | We study optimal assignment of home energy reports (HERs), a.k.a. “Opower”
 - | HERs have been implemented by at least 85 utilities and reached at least 6.2 million households
 - | Estimated annual cost of \$1.2 billion if scaled up nationally

- | Research objectives:
 - | Use a policy learning algorithm
 - | Search for simple treatment assignment rules that maximize the program’s effects
 - | Provide empirical evidence on the potential gains

Behavioral Intervention: *Home Energy Reports*

Behavioral Intervention: *Home Energy Reports*



Behavioral Intervention: *Home Energy Reports*

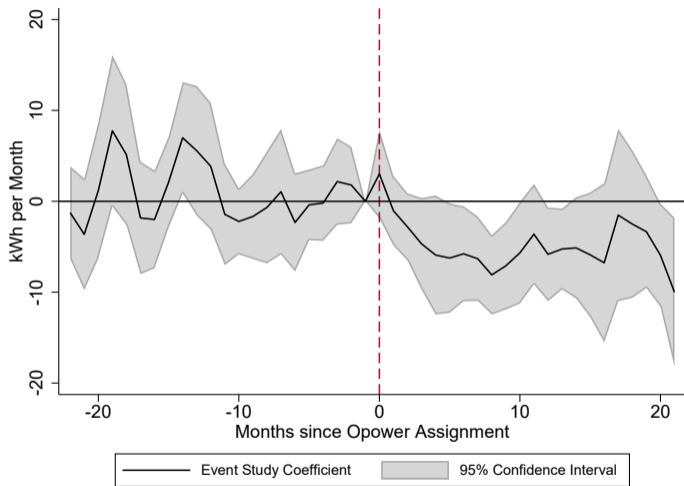


Data

Data on residential electricity accounts provided by a utility in a Northeastern state:

- | Monthly electricity consumption: 2014-2018
- | Opower program participation: multiple RCTs in 2013-2018
- | Household demographics in 2015
 - | Income bin
 - | Number of household members
 - | Marital status of head of household
 - | Square footage of unit, building
 - | Year of construction
- | All data are at the account level, with 50k households in the estimation sample

Behavioral Interventions Reduce Electricity Consumption ON AVERAGE



Heterogeneous Treatment Effects by Household Characteristics

			Dependent Variable: Electricity Usage in kWh			
			Baseline Usage	House Size	Income	House Year Built
Opower	Post	Below Median	-2.48 (2.36)	-3.15 (3.43)	-1.57 (3.64)	-6.41 (3.71)
Opower	Post	Above Median	-12.33 (6.19)	-8.53 (3.75)	-9.27 (3.59)	-5.08 (3.52)
Demographics			Yes	Yes	Yes	Yes
Wave	year-month	category FE	Yes	Yes	Yes	Yes
p-value, test of equal coefficients			0.14	0.29	0.13	0.80
N			2,186,105	2,186,105	2,186,105	2,186,105

* p<0.10, ** p<0.05, *** p<0.01

Using Targeted Treatment to Maximize Gains

Objective: use a policy learning algorithm to empirically select a rule that uses pre-treatment characteristics to assign treatment in order to maximize gains in the target population

The value of the rule is proportional to:

$$V(\pi) = E[Y_1(x_2) + Y_0(x_2)]$$

Objective Function

Assuming unconfoundedness, equivalence of distributions for the target and sampled populations, and overlap for propensity scores in the sampled population,

$$V(\hat{\tau}) = E[Y_1 | X_2] - E[Y_0 | X_2]$$

can be rewritten as

$$V(\hat{\tau}) = E[Y_0] + E \left[\underbrace{\frac{YD}{e(X)} - \frac{Y(1-D)}{1-e(X)}}_{\text{value gain relative to } E[Y_0]} \mid X_2 \right]$$

Empirical Welfare Maximization (EWM)

Optimal utilitarian treatment rule maximizes welfare gain relative to $E[Y_0]$:

$$2 \arg \max_{2\pi} E \left[\frac{YD}{e(X)} - \frac{Y(1-D)}{1-e(X)} \right] \quad (X2)$$

Empirical Welfare Maximization (EWM)

Optimal utilitarian treatment rule maximizes welfare gain relative to $E[Y_0]$:

$$2 \arg \max_{\Pi} E \left[\frac{YD}{e(X)} - \frac{Y(1-D)}{1-e(X)} \right] \quad (X2)$$

Idea of Empirical Welfare Maximization is to solve the sample analog:

$$\hat{EWM} 2 \arg \max_{\Pi} \frac{1}{N} \sum_{i=1}^N \left[\frac{Y_i D_i}{e(X_i)} - \frac{Y_i(1-D_i)}{1-e(X_i)} \right] \quad (X_i2)$$

within a constrained class of candidate rules Π

Applying EWM to Home Energy Reports

HER goal is to reduce energy usage or energy cost, so we solve:

$$\min_{\Pi} \frac{1}{N} \sum_{i=1}^N \frac{Y_i D_i}{e(X_i)} - \frac{Y_i(1 - D_i)}{1 - e(X_i)} \quad (X_i^2)$$

Two policy classes Π :


1. Quadrant rules:

$$\Pi_Q = \{X : f s_1(X_1 - t_1) > 0 \ \& \ s_2(X_2 - t_2) > 0\}; \quad s_1; s_2 \in \{1, -1\}; t_1; t_2 \in \mathbb{R}$$

2. Linear rules with cubic terms:

$$\Pi_{LES}^3 = \{X : (0 + \alpha_1 X_1 + \alpha_2 X_1^2 + \alpha_3 X_1^3 + \alpha_4 X_2 > 0)\}; \quad \alpha_i; \alpha_j \in \mathbb{R}$$

EWM Rules



Other Rules


Targeting based on EWM Rules Achieves Significant Energy Savings

Relative to the RCT

Treatment rule	Variables	Share treated %	Δ EWM v. RCT kWh/ hh-month
EWM-Quadrant	Income, baseline usage	92	-1.88 (-4.19,0.44)
EWM-Quadrant	Size, baseline usage	92	-2.3 (-4.74,0.14)
EWM-Quadrant	House age, baseline usage	78	-1.83 (-4.37,0.71)
EWM-cubic	Income, baseline usage	69	-3.8 (-6.62,-0.98)
EWM-cubic	Size, baseline usage	92	-3.2 (-5.75,-0.64)
EWM-cubic	House age, baseline usage	51	-2.92 (-5.68,-0.16)

Targeting based on EWM Rules Achieves Significant Cost Savings

Relative to the RCT, with Energy Conservation Valued at the Retail Electricity Rate



Functions of Pre-Treat Usage also Achieve Significant Cost Savings

Relative to the RCT, with Energy Conservation Valued at the Retail Electricity Rate

Properties of Empirical Welfare Maximization

Kitagawa and Tetenov (2018) provide conditions under which:

- | Average social welfare from EWM rules converges to the maximum obtainable welfare within ϵ at optimal rate
- | EWM can find constrained-optimal policy without estimating all causal effects (i.e., $\hat{\mu}(x)$ $\approx \mu(x)$)

EWM integrates economic decision problem and statistical inference

comparison

Conclusion

- | We use a policy learning algorithm to provide empirical evidence on the potential gains from targeted assignment of home energy reports
- | Targeting using transparent and easily implemented treatment rules yields significant energy and cost savings relative to actual treatment assignment
- | The gains from targeting based on electricity consumption alone further underscores the practical value of simple treatment rules

Thank You!

Appendix

Objective Function

Assuming unconfoundedness, equivalence of distributions for the target and sampled populations, and overlap for propensity scores in the sampled population,

$$V(\hat{\tau}) = E[Y_1 \cdot 1_{(X \in \mathcal{X}_2)} + Y_0 \cdot 1_{(X \notin \mathcal{X}_2)}]$$

can be rewritten as

$$V(\hat{\tau}) = E[Y_0] + E\left[\underbrace{\frac{YD}{e(X)} - \frac{Y(1-D)}{1-e(X)}}_{\text{IPW estimator of ATE (from universal treatment)}} \cdot 1_{(X \in \mathcal{X}_2)}\right]$$

Applying EWM to Home Energy Reports: Data

HER goal is to reduce energy consumption, so we solve:

$$\min_{\beta} \frac{1}{N} \sum_{i=1}^N \frac{Y_i D_i}{e(X_i)} - \frac{Y_i (1 - D_i)}{1 - e(X_i)} \quad \mathbf{1}_{(X_i \in \mathcal{X}_2)}$$

- | Expected energy or cost savings Y_i
- | Treatment status from RCT D_i
- | Pre-treatment characteristics X_i
- | Propensity scores $e(X_i)$: wave-specific treatment shares

Inference

We take two complementary approaches to inference:

- | Bootstrap confidence intervals for EWM rules per Kitagawa and Tetenov (2018)
 - | For each bootstrap sample, search for rules that lead to the biggest difference in savings between original and bootstrap samples, which produces conservative CIs
- | To demonstrate the practical value of our approach, we also bootstrap CIs for the savings from applying specific EWM rules we estimate relative to the RCT

Related Literature

- | Average treatment effect of home energy reports:
 - | Allcott (2011); Allcott and Rogers (2014); Ayres et al. (2013); many others
- | Optimizing treatment rules for home energy reports:
 - | Allcott and Kessler (2019)
 - | Knittel and Stolper (2019)
- | Policy learning in economics and statistics:
 - | Kitagawa and Tetenov (2018)
 - | Athey and Wager (2021)

Applying EWM Rules Across Waves

- | Program designers only have access to historical data
- | To mimic this, we study the performance of the EWM method when using past waves to derive treatment rules for future waves
- | EWM method can be extended to this case by reweighting:

$$\min_2 \frac{1}{N} \sum_{i=1}^N \frac{Y_i D_i}{e(X_i)} - \frac{Y_i (1 - D_i)}{1 - e(X_i)} \frac{P_X^T(x)}{P_X^S(x)} \mathbb{1}_{(X_i \in \mathcal{X})}$$

where $\frac{P_X^T(x)}{P_X^S(x)}$ is the density ratio of the marginal distributions \mathcal{X} for the sample (past) and target (future) populations

Steps to Apply EWM Rules Across Waves

1. Estimate $\frac{P_X^T(x)}{P_X^S(x)}$ nonparametrically by taking the ratio of sample shares within bins
2. Reweight data: $\frac{Y_i D_i}{e(X_i)} \quad \frac{Y_i (1 - D_i)}{1 - e(X_i)} \quad \frac{P_X^T(x)}{P_X^S(x)}$
3. Use reweighted data to estimate the EWM treatment rules as before
4. Use experimental data from the target wave to evaluate the performance of the rules relative to the actual RCT 'ex-post'

Applying EWM Rules Across Waves Outperforms the RCT

EWM Cubic Rules

Target wave	Sample wave	Pre-treatment characteristics used for targeting	Energy changes kWh/hh-month	Private cost changes \$/hh-month	Social cost changes \$/hh-month
6	3	Income and mean usage	-0.62	-0.07	-0.23
		House size and mean usage	-0.15	0.08	-0.06
		House age and mean usage	-1.83	-0.34	-0.25
7	6	Income and mean usage	-0.68	0.01	-0.13
		House size and mean usage	-2.92	-0.48	-0.40
		House age and mean usage	-0.04	-0.08	-0.41

How do EWM Rules Compare with other Types of Treatment Rules?

The unconstrained plug-in rule:

$$\hat{\tau}_{\text{plug in}} = \mathbb{1}\{x : \hat{\tau}(x) \geq 0\}$$

where $\hat{\tau}(x)$ is a conditional average treatment effect (CATE) estimator

- | Approach:
 - | Estimate CATEs for each household (e.g., via ML)
 - | Use CATEs to calculate cost savings for each household
 - | Treat all households with cost savings

- | Allcott and Kessler (2019), Knittel and Stolper (2019) use this approach

Comparison to Previous Studies: Methods

- | Prior work uses plug-in rules
 - | Use of ML for model selection searches over a large set of candidate rules
 - | Statistical performance hinges on efficient estimation of CATEs

- | We use empirical welfare maximization
 - | Find that simple and transparent rules perform well
 - | Desirable statistical properties: average social welfare converges to the maximum obtainable welfare within
 - | Integrate economic decision problem with statistical inference

Comparison to Previous Studies: Results

- | Allcott and Kessler (2019):
 - | Approach: plug-in rule to treat all households with above median predicted savings
 - | Result: targeting increases gas conservation by 85% relative to the original RCT
 - | We find targeting increases electricity savings by 50 - 105% based on point estimates

- | Knittel and Stolper (2019)
 - | Approach: plug-in rule to treat all households with positive net benefit
 - | Result: targeting yields a private cost reduction of \$1.17/hh-month and a social cost reduction of \$0.26/hh-month
 - | We find smaller reductions in private cost, but larger reductions in social cost

EWM rules

Objective: Cost Minimization with Electricity Consumption Valued at the Retail Electricity Rate

[Return](#)

The Home Energy Reports Program was Implemented in \Waves"

Opower wave	Month/Year	Number of electric accounts assigned into Opower	Number of electric accounts treated
1	03/2013	183,789	166,911
2	04/2013	19,838	17,943
3	03/2014	43,435	36,759
4	08/2014	42,069	38,174
5	10/2015	0	0
6	08/2016	25,974	12,992
7	03/2017	44,372	31,199
8	02/2018	31,534	21,688

Billing
Data
Time
Frame

Raw Data: Average Electricity Consumption by Wave and Treatment

Treatment was Randomized within Waves, but not across Waves

Covariate Balance

Balance Test: Wave 3

	Control	Treatment	Difference	t-statistic
12-month pre-treatment consumption (kWh)	650	647	3.56	0.46
Income (\$)	72,487	72,786	-299	-0.37
Number of household members	2.56	2.55	.00624	0.20
Building size (ft ²)	3,681	3,744	-63.6	-0.51
Unit size (ft ²)	1,761	1,778	-17.1	-1.04
House Year Built	1,951	1,951	.025	0.04
Married	.575	.565	.0108	1.19

Covariate Balance

Balance Test: Wave 6

	Control	Treatment	Difference	t-statistic
12-month pre-treatment consumption (kWh)	451	445	6.29	1.02
Income (\$)	64,085	64,641	-556	-0.69
Number of household members	2.02	1.98	.032	1.16
Building size (ft ²)	5,439	5,371	67.6	0.29
Unit size (ft ²)	1,794	1,793	.536	0.03
House Year Built	1,948	1,949	-.983	-1.43
Married	.432	.428	.00414	0.44

Balance Test: Wave 7

	Control	Treatment	Difference	t-statistic
12-month pre-treatment consumption (kWh)	479	487	-7.66	-1.13
Income (\$)	53,653	54,389	-735	-1.02
Number of household members	1.88	1.89	-.0157	-0.65
Building size (ft ²)	5,660	5,778	-118	-0.43
Unit size (ft ²)	2,186	2,201	-14.6	-0.70
House Year Built	1,937	1,937	.539	0.85
Married	.322	.33	-.00736	-0.90

Pooled Sample Summary Statistics

	mean	sd	min	max
Monthly electricity consumption (kWh)	505	384	0	2,705
Income (\$)	66,104	43,985	5,000	150,000
Number of household members	2.29	1.62	1	8
Building size (ft ²)	4,655	11,146	262	232,146
Unit size (ft ²)	1,886	1,053	210	5,000
House Year Built	1,947	36	1,850	2,013
Married	.482	.5	0	1
Observations	2186105			

Number of Households and Average Consumption by Opower Wave

Wave-Specific Event Study Plots

$$\text{kWh}_{iwt} = \sum_{r=23}^{r=2} \text{Opower}_{ft}^r + \sum_{r=0}^{r=22} \text{Opower}_{ft}^r + X_i + \text{wt} + \epsilon_{iwt}$$

Equivalence between IPW Estimator and Difference in Means Estimator

Under conditional independence, the inverse probability weighting estimator is a difference in means estimator.

$$\begin{aligned} \text{ATE} &= E[Y_i | D_i = 1; X_i] - E[Y_i | D_i = 0; X_i] \\ &= E \left[\frac{D_i Y_i}{e(X_i)} \right] - E \left[\frac{(1 - D_i) Y_i}{1 - e(X_i)} \right] \end{aligned}$$

$$\begin{aligned} E \left[\frac{D_i Y_i}{e(X_i)} \right] &= E \left[E \left[\frac{D_i Y_i}{e(X_i)} \middle| X_i \right] \right] \\ &= E \left[\frac{1}{e(X)} E[D_i Y_i | X_i] \right] \\ &= E \left[\frac{E[Y_i(1) | X_i] E_D[D_i | X_i]}{e(X_i)} \right] \\ &= E [E[Y_i(1) | X_i]] \\ &= E[Y_i(1)] \end{aligned}$$

Targeting based on EWM Rules Achieves Significant Energy Savings

Treatment rule	Variables	Share treated %	Net energy changes kWh/hh-month	EWM v. RCT kWh/hh-month
Actual RCT	Scaled ATT	72	-3.63 (-10.50,0.43)	-
EWM-Quadrant	Income, baseline usage	92	-5.50 (-11.69,0.68)	-1.88 (-4.19,0.44)
EWM-Quadrant	Size, baseline usage	92	-5.93 (-11.58,-0.28)	-2.30 (-4.74,0.14)
EWM-Quadrant	House age, baseline usage	78	-5.46 (-11.38,0.47)	-1.83 (-4.37,0.71)
EWM-cubic	Income, baseline usage	69	-7.43 (-14.55,-0.31)	-3.80 (-6.62,-0.98)
EWM-cubic	Size, baseline usage	92	-6.83 (-14.25,0.60)	-3.20 (-5.75,-0.64)
EWM-cubic	House age, baseline usage	51	-6.55 (-13.51,0.41)	-2.92 (-5.68,-0.16)

Targeting based on EWM Rules Achieves Significant Cost Savings

Energy Conservation Valued at the Retail Electricity Rate

Treatment rule	Variables	Share treated %	Net cost changes \$/hh-month	EWM v. RCT \$/hh-month
Actual RCT	Scaled ATT	72	-0.09 (-0.99,0.74)	- -
EWM-Quadrant	Income, baseline usage	13	-0.65 (-1.61,0.31)	-0.56 (-1.06,-0.06)
EWM-Quadrant	Size, baseline usage	27	-0.79 (-1.83,0.26)	-0.70 (-1.05,-0.34)
EWM-Quadrant	House age, baseline usage	21	-0.66 (-1.73,0.41)	-0.57 (-1.02,-0.12)
EWM-cubic	Income, baseline usage	15	-0.81 (-2.07,0.46)	-0.72 (-1.22,-0.21)
EWM-cubic	Size, baseline usage	28	-0.76 (-2.04,0.53)	-0.67 (-1.15,-0.19)
EWM-cubic	House age, baseline usage	26	-0.82 (-2.05,0.41)	-0.73 (-1.32,-0.14)


Targeting based on EWM Rules Achieves Significant Economic Benefits

Energy Conservation Valued at the Short-Run Average SMC of Electricity


Treatment rule	Variables	Share treated %	Net cost changes \$/ hh-month	EWM v. RCT \$/ hh-month
Actual RCT	Scaled ATT	72	0.32 (0.09,0.78)	.
EWM-Quadrant	Income, baseline usage	13	-0.18 (-0.58,0.23)	-0.49 (-0.68,-0.3)
EWM-Quadrant	Size, baseline usage	22	-0.17 (-0.57,0.23)	-0.49 (-0.63,-0.34)
EWM-Quadrant	House age, baseline usage	4	-0.15 (-0.53,0.22)	-0.47 (-0.67,-0.27)
EWM-cubic	Income, baseline usage	5	-0.24 (-0.71,0.24)	-0.55 (-0.9,-0.21)
EWM-cubic	Size, baseline usage	12	-0.19 (-0.63,0.25)	-0.51 (-0.87,-0.14)
EWM-cubic	House age, baseline usage	17	-0.22 (-0.65,0.22)	-0.53 (-0.94,-0.12)

Targeting based on EWM Rules Achieves Significant Energy Savings

Relative to the RCT



Targeting based on EWM Rules Achieves Significant Economic Benefits Relative to the RCT, with Energy Conservation Valued at the Short-Run Average SMC of Electricity



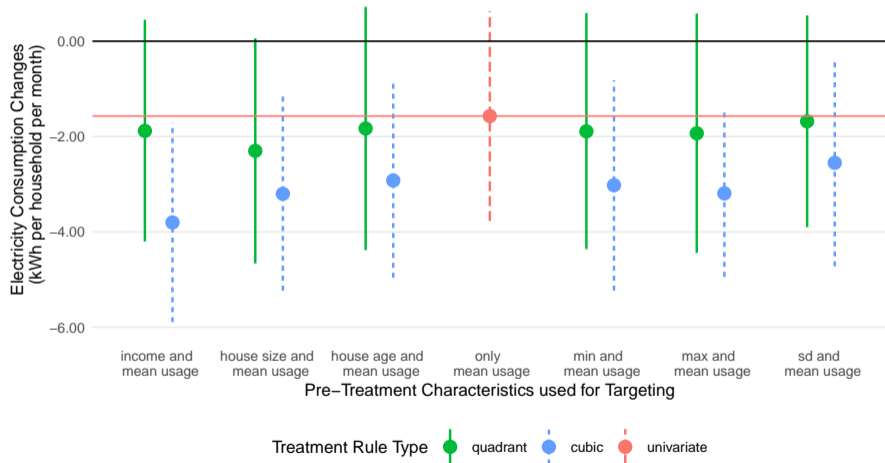
Targeting based on Pre-Treatment Consumption Alone Performs Well

Energy Conservation Valued at the Retail Electricity Rate

Treatment rule	Variables	Share treated %	Net cost changes \$/ hh-month	Δ EWM v. RCT \$/ hh-month
Actual RCT	Scaled ATT	72	-0.09 (-0.99,0.74)	.
EWM-one-dimension	Mean baseline usage	13	-0.65 (-1.54,0.24)	-0.56 (-0.88,-0.24)
EWM-Quadrant	Min baseline usage, mean baseline usage	28	-0.76 (-1.91,0.4)	-0.66 (-1.12,-0.21)
EWM-Quadrant	Max baseline usage, mean baseline usage	29	-0.73 (-1.74,0.28)	-0.64 (-0.99,-0.28)
EWM-Quadrant	Sd baseline usage, mean baseline usage	29	-0.66 (-1.81,0.49)	-0.57 (-0.89,-0.24)
EWM-cubic	Min baseline usage, mean baseline usage	24	-0.92 (-2.13,0.29)	-0.83 (-1.41,-0.25)
EWM-cubic	Max baseline usage, mean baseline usage	37	-0.91 (-2.11,0.28)	-0.82 (-1.43,-0.21)
EWM-cubic	Sd baseline usage, mean baseline usage	20	-0.79 (-1.99,0.41)	-0.7 (-1.21,-0.19)


Targeting based on Pre-Treatment Consumption Alone Performs Well

Relative to the RCT in terms of energy conservation



Targeting based on Pre-Treatment Consumption Alone Performs Well

Relative to the RCT, with Energy Conservation Valued at the Social Marginal Cost



Optimal Rules based on Pre-Treatment Consumption Alone

