Choosing Who Chooses:

Selection-Driven Targeting in Energy Rebate Programs

Takanori Ida [1] Takunori Ishihara [2] Koichiro Ito [3] Daido Kido [1]
Toru Kitagawa [4] Shosei Sakaguchi [5] Shusaku Sasaki [6]

[1]Kyoto University [2]Kyoto University of Advanced Science [3]University of Chicago
[4]Brown University [5]University College London [6]Tohoku Gakuin University

July 25, 2023

Targeting has become a central interest in policy design

- Many policies are costly. Budgets are limited.
- How to maximize a policy's impact given a limited budget?
- Policymakers could target individuals who generate large welfare gains

- Examples:
  - Job training program (Kitagawa and Tetenov, 2018)
  - SNAP (Finkelstein and Notowidigdo, 2019)
  - Disability program (Deshpande and Li, 2019)
  - Energy efficiency (Burlig, Knittel, Rapson, Reguant, and Wolfram, 2020)
  - Behavioral nudge (Knittel and Stolper, 2019)
  - Dynamic electricity pricing (Ito, Ida, and Tanaka, forthcoming)

Two competing views: Observables vs. Self-selection

1. Targeting based on observables
    - Policymakers can use observable information (X) to target individuals
    - Machine learning techniques can be used to find "who should be treated"
    - e.g. Kitagawa and Tetenov (2018), Athey and Wager (2021)

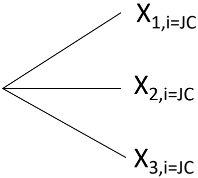

Two competing views: Observables vs. Self-selection

1. Targeting based on observables
   - Policymakers can use observable information (X) to target individuals
   - Machine learning techniques can be used to find "who should be treated"
   - e.g. Kitagawa and Tetenov (2018), Athey and Wager (2021)

2. Targeting through self-selection
   - Individuals' self-selection may have important (unobservable) information
   - Policymakers could take advantage of this self-selection
   - e.g., Alatas, Purnamasari, Wai-Poi, Banerjee, Olken, and Hanna (2016), Ito, Ida, and Tanaka (forthcoming)

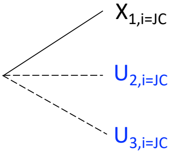

A priori, it is not clear which approach works better

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.

from “Public Policy in an Uncertain World”
by Charles F. Manski (2013)

Our idea: Build an algorithm that takes advantage of both

1. Setting: A costly treatment that could generate a social welfare gain
   - Field experiment: A peak-hour rebate program for energy conservation
   - Benefit: A reduction in DWL if a participant actually conserves energy
   - Cost: Implementation cost per participating household

2. Use an RCT & the Empirical Welfare Maximization (EWM) to identify
   - Consumer type $x \in X$ who should be "treated"
   - Consumer type $x \in X$ who should be "untreated"
   - Consumer type $x \in X$ who should "decide by themselves"

3. Test hypotheses: What policy rule can maximize policy impacts?
   - Targeting based on observables vs. targeting through self-selection
   - Optimal mix of these two approaches

Road map of the talk

1. Introduction
2. Conceptual Framework
3. Field Experiment and Data
4. Optimal Assignment Policy and Welfare Gains
5. Mechanism Behind the Optimal Policy Assignment
6. Welfare Maximization with Redistribution
7. Conclusion

Conceptual Framework

Setup

- Consider a net social welfare gain (W) from a costly treatment
- Define three potential outcomes (heterogeneous across individuals)
  - W_T: a potential outcome if an individual is treated
  - W_U: a potential outcome if an individual is untreated
  - W_S: a potential outcome if an individual self-selects
- Which outcome is the best for a social planner? It is ambiguous.
  - W_U can be best for those who generate gain lower than cost
  - W_S can be useful if self-selection works in line with a planner's goal
  - W_S can be worse if self-selection is adverse to a planner's goal
- Challenge: Potential outcomes are unobservable
- Can we use an RCT to identify best assignment for each type $x \in X$?

Optimal targeting

- The planer assigns each individual to one of three arms "T", "U", "S" depending on the observed information $x \in \mathcal{X}$.
- $G_j \subseteq \mathcal{X}$ ($j = T, U, S$): A set of x such that any individual with $x \in G_j$ is assigned the arm j.
- Targeting policy $G \equiv (G_T, G_U, G_S)$, a partition of $\mathcal{X}$.
- The average welfare contribution under a targeting policy G:

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

- **Goal:** Find an optimal policy G^* that maximizes the average welfare contribution.

Empirical Welfare Maximization (EWM) method

- RCT (or quasi-experimental) data: $\{(W_i, Z_i, X_i) : i = 1, \ldots, n\}$ where $Z_i \in \{T, U, S\}$ is randomly assigned.

- With random assignment, the empirical analogue of $\mathcal{W}(G)$ is

$$\widehat{\mathcal{W}}(G) = \frac{1}{n} \sum_{i=1}^{n} \sum_{j \in \{T,C,S\}} \left(\frac{W_i \cdot 1\{Z_i = j\}}{P(Z_i = j \mid X_i)} \cdot 1\{X_i \in G_j\} \right).$$

- EWM method (Kitagawa and Tetenov, 2018): With a pre-specified class of feasible policies $\mathcal{G}$, estimate the optimal policy over $\mathcal{G}$ by

$$\widehat{G} \in \arg\max_{G \in \mathcal{G}} \widehat{\mathcal{W}}(G).$$

- We use a class of policy trees (Zhou, Athey, Wager, 2022) for $\mathcal{G}$.

Policy tree

Depth 2 policy tree

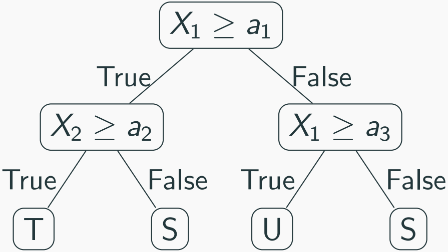

Partition of $\mathcal{X}$ by a depth 2 policy tree

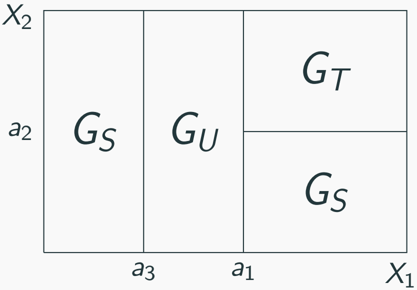

Mechanism behind the optimal targeting

- The LATE framework (Imbens and Angrist, 1994) can be used to uncover the mechanism behind our approach

- We define the LATEs for "takers" and "non-takers"
  - $D_S \in \{0, 1\}$ is an individual's treatment take-up when they self-select.
  - (W_T, W_U) are the treated and untreated potential outcomes.
  - CLATE for takers: $E[W_T - W_U | D_S = 1, x]$.
  - CLATE for non-takers: $E[W_T - W_U | D_S = 0, x]$.

- Suppose that the **exclusion restriction** holds:

$$W_S = W_T \cdot 1\{D_S = 1\} + W_U \cdot 1\{D_S = 0\}$$

- Note: Exclusion restriction is required to estimate the optimal policy but helpful to investigate the mechanism

Mechanism behind the optimal targeting

- For individuals with type x,

$$E[W_T \mid x] \geq E[W_U \mid x] \Leftrightarrow \underbrace{E[W_T - W_U \mid x]}_{CATE(x)} \geq 0;$$

$$E[W_S \mid x] \geq E[W_U \mid x] \Leftrightarrow \underbrace{E[W_T - W_U \mid D_S = 1, x]}_{CLATE_{taker}(x)} \geq 0.$$

$$E[W_S \mid x] \geq E[W_T \mid x] \Leftrightarrow \underbrace{E[W_T - W_U \mid D_S = 0, x]}_{CLATE_{non\text{-}taker}(x)} \leq 0;$$

- Using these properties, we can show:
  - T is best $\Leftrightarrow$ $CATE(x) \geq 0$ and $CLATE_{non\text{-}taker}(x) \geq 0$;
  - U is best $\Leftrightarrow$ $CATE(x) \leq 0$ and $CLATE_{taker}(x) \leq 0$;
  - S is best $\Leftrightarrow$ $CLATE_{taker}(x) \geq 0$ and $CLATE_{non\text{-}taker}(x) \leq 0$.

- We use our RCT data to estimate these LATEs to show this relationship in our empirical analysis

Road map of the talk

1. Introduction
2. Conceptual Framework
3. Field Experiment and Data
4. Optimal Assignment Policy and Welfare Gains
5. Mechanism Behind the Optimal Policy Assignment
6. Welfare Maximization with Redistribution
7. Conclusion

Field Experiment and Data

Field experiment

1. Treatment: A peak-hour rebate program for residential electricity use
   - Partner: Japanese Ministry of the Environment
   - Peak-hour: 1 pm to 5 pm in critical peak days in summer 2020
   - Baseline: Average hourly usage in the same hours before experiment
   - Customers were unaware of baseline until experiment began
   - All customers were on "non-dynamic retail prices"
   - Rebate = \$1/kWh conservation ≈ peak-hour wholesale price
   - Implementation cost per consumer = 291.1 JPY (≈ cents)
   - Welfare gain = a reduction in DWL − implementation cost

2. Experimental sample: 3,870 households in Japan
   - Not a random sample of population
   - Recruitment by mail and email

3. Randomization:
   - Control: 1,577, Treatment: 1,486, Opt-in: 807

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)	107.29 [48.57]	105.51 [49.61]	103.42 [46.14]	1.78 (1.78)	3.87 (2.03)	2.09 (2.07)

Impacts on ln(peak-hour usage): ITT analysis

		Peak hour usage − Pre-peak hour usage (in pre-experiment)	
	All	Low	High
100% Treatment	-0.097	-0.108	-0.079
	(0.021)	(0.028)	(0.031)
100% Opt-in	-0.052	-0.022	-0.073
	(0.027)	(0.034)	(0.041)
Observations	1,176,480	588,240	588,240
p-value (T = O)	0.088	0.013	0.880
Opt-in rate	37.17	36.92	37.44

- Note that this is the impact on consumption, not the net welfare gain
- Suggests there is important heterogeneity by consumer type $x \in X$
- Standard errors are clustered at the customer level

More evidence on heterogeneity by X: ITT analysis

	Number of people at home in peak hours		Interested in energy conservation	
	Low	High	Low	High
100% Treatment	-0.096	-0.098	-0.134	-0.057
	(0.027)	(0.034)	(0.028)	(0.031)
100% Opt-in	-0.022	-0.094	-0.036	-0.072
	(0.034)	(0.042)	(0.035)	(0.040)
p-value (T = O)	0.020	0.934	0.004	0.715
Opt-in rate	37.64	36.57	33.83	40.55

- Suggests there is important heterogeneity by consumer type $x \in X$
- We will exploit this variation & EWM method to estimate optimal policy rule

Road map of the talk

1. Introduction
2. Conceptual Framework
3. Field Experiment and Data
4. Optimal Assignment Policy and Welfare Gains
5. Mechanism Behind the Optimal Policy Assignment
6. Welfare Maximization with Redistribution
7. Conclusion

Optimal Assignment Policy and Welfare Gains

Welfare gain

- Y_j ($j \in \{T, U, S\}$): Potential outcome of household's peak-hour electricity consumption (kWh) in experimental-period
- $D_j \in \{0, 1\}$ ($j \in \{T, U, S\}$): Household's potential choice if the arm j is assigned. (Note: $D_T = 1$ and $D_U = 0$ w.p.1.)
- Household's potential welfare contribution:

$$W_j \equiv \underbrace{b}_{\text{benefit}} \times \underbrace{(Y_U - Y_j)}_{\text{electricity conservation}} - \underbrace{c}_{\text{cost}} \times 1\{D_j = 1\}.$$

  - $j \in \{T, U, S\}$
  - b: social welfare gain from a unit reduction in energy use
  - c: implementation cost of the program.

Optimal targeting policy

- Social Welfare of a targeting policy $G = (G_T, G_U, G_S)$:

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

- **Goal:** Find an optimal policy G^* that maximizes $\mathcal{W}(\cdot)$.

- Observables (X): electricity consumption in the pre-experimental period; household income; number of people at home, self-efficacy in energy conservation (scale 1-5)

- We apply EWM method with policy trees (depth 6): With a class of decision trees $\mathcal{G}$, estimate the optimal policy over $\mathcal{G}$ by

$$\widehat{G} \in \arg\max_{G\in\mathcal{G}} \widehat{\mathcal{W}}(G).$$

We compare welfare gains form five policies

1. 100% untreated (baseline)
   - Everyone is assigned to U
2. 100% treated
   - Everyone is assigned to T
3. 100% self-selection
   - Everyone is assigned to S
4. Selection-absent targeting ($\hat{G}^{\dagger}$)
   - Optimal assignment of (U, T)
5. Selection-driven targeting ($\hat{G}^{*}$)
   - Optimal assignment of (U, T, S)

Welfare Gains from Each Policy

Policy	Welfare Gain	Share of customers in each arm		
		G_U	G_T	G_S
100% untreated	0 (—)	100%	0.0%	0.0%
100% treated	120.7 (98.8)	0.0%	100%	0.0%
100% self-selection	180.6 (112.1)	0.0%	0.0%	100%
Selection-absent targeting ($\hat{G}^{\dagger}$)	387.8 (55.7)	47.6%	52.4%	0.0%
Selection-driven targeting ($\hat{G}^{*}$)	553.7 (68.0)	23.9%	31.4%	44.7%

- Selection-absent targeting assigns 47.6% to U and 52.4% to T
- Selection-driven targeting assigns 23.9% to U, 31.4% to T, and 44.7% to S

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 ($\hat{G}^{\dagger}$) vs. 100% treated	267.2 (99.7)	0.004
Selection-absent targeting ($\hat{G}^{\dagger}$) vs. 100% self-selection	207.3 (116.9)	0.038
Selection-driven targeting ($\hat{G}^{*}$) vs. 100% treated	433.0 (106.8)	0.000
Selection-driven targeting ($\hat{G}^{*}$) vs. 100% self-selection	373.1 (113.3)	0.000
Selection-driven targeting ($\hat{G}^{*}$) vs. Selection-absent targeting ($\hat{G}^{\dagger}$)	165.8 (61.1)	0.003

- Both targeting policies improve welfare compared to non-targeting policies
- $\hat{G}^{*}$ further improves welfare compared to $\hat{G}^{\dagger}$

	Sample mean by group [standard deviation]			Difference in sample means (standard error)		
	$\hat{G}_U^*$	$\hat{G}_T^*$	$\hat{G}_S^*$	$\hat{G}_U^*$ vs. $\hat{G}_T^*$	$\hat{G}_U^*$ vs. $\hat{G}_S^*$	$\hat{G}_T^*$ vs. $\hat{G}_S^*$
Peak hour usage (Wh)	203	180	191	23.03	11.98	−11.05
	[146]	[136]	[135]	(6.18)	(5.79)	(5.08)
Pre-peak hour usage (Wh)	198	167	175	30.56	23.08	−7.48
	[150]	[133]	[132]	(6.23)	(5.86)	(4.97)
Post-peak hour usage (Wh)	329	255	310	73.22	18.82	−54.40
	[176]	[176]	[164]	(7.67)	(7.00)	(6.41)
Number of people at home	2.87	2.27	2.38	0.60	0.48	−0.11
	[1.34]	[1.32]	[1.08]	(0.06)	(0.05)	(0.05)
Self-efficacy in energy	3.30	3.49	3.53	−0.19	−0.23	−0.04
conservation (1-5 scale)	[1.02]	[0.82]	[0.75]	(0.04)	(0.04)	(0.03)
Household income	787	597	572	190.12	215.11	25.00
(JPY 10,000)	[433]	[397]	[318]	(18.23)	(16.15)	(13.73)
All electric	0.36	0.25	0.33	0.11	0.03	−0.08
	[0.48]	[0.43]	[0.47]	(0.02)	(0.02)	(0.02)
Number of air conditioners	3.41	2.82	3.16	0.58	0.24	−0.34
	[1.72]	[1.66]	[1.67]	(0.07)	(0.07)	(0.06)
Number of fans	2.99	2.58	2.78	0.41	0.20	−0.21
	[1.75]	[1.57]	[1.55]	(0.07)	(0.07)	(0.06)
Number of household members	3.17	2.54	2.67	0.63	0.50	−0.13
	[1.31]	[1.36]	[1.14]	(0.06)	(0.05)	(0.05)
Total living area (m^2)	115.41	97.16	106.73	18.25	8.68	−9.57
	[47.77]	[49.13]	[47.37]	(2.11)	(1.94)	(1.81)

Road map of the talk

1. Introduction
2. Conceptual Framework
3. Field Experiment and Data
4. Optimal Assignment Policy and Welfare Gains
5. Mechanism Behind the Optimal Policy Assignment
6. Welfare Maximization with Redistribution
7. Conclusion

Mechanism Behind the Optimal Policy Assignment

Mechanism behind the optimal assingnment

- How does our algorithm improve the social welfare gains?
- We highlight that the LATE framework (Imbens and Angrist, 1994) can be used to investigate the mechanism
- We define the LATEs for "takers" and "non-takers"
  - $D_S \in \{0, 1\}$ is an individual's treatment take-up when they self-select
  - (W_T, W_U) are the treated and untreated potential outcomes
  - LATE for takers: $E[W_T - W_U | D_S = 1]$
  - LATE for non-takers: $E[W_T - W_U | D_S = 0]$

Mechanism behind the optimal assignment

- Two key LATEs
  - LATE for takers: $E[W_T - W_U | D_S = 1]$
  - LATE for non-takers: $E[W_T - W_U | D_S = 0]$

- Random variation in $(U,\ T,\ S)$ allows us to estimate both LATEs
  - Use data from groups (U,S) to estimate the LATE for takers
  - Use data from groups (T,S) to estimate the LATE for non-takers

Mechanism behind the optimal assingnment

- Two key LATEs
  - LATE for takers: $E[W_T - W_U | D_S = 1]$
  - LATE for non-takers: $E[W_T - W_U | D_S = 0]$

- Random variation in $(U,\ T,\ S)$ allows us to estimate both LATEs
  - Use data from groups (U,S) to estimate the LATE for takers
  - Use data from groups (T,S) to estimate the LATE for non-takers
  - Recall that we have random (U, T, S) in each group $(\hat{G}^*_U,\ \hat{G}^*_T, \hat{G}^*_S)$
    $\rightarrow$ We can estimates these LATEs in each group $(\hat{G}^*_U,\ \hat{G}^*_T, \hat{G}^*_S)$

Mechanism behind the optimal assingnment

- Two key LATEs
  - LATE for takers: $E[W_T - W_U | D_S = 1]$
  - LATE for non-takers: $E[W_T - W_U | D_S = 0]$

- Random variation in $(U,\ T,\ S)$ allows us to estimate both LATEs
  - Use data from groups $(U{,}S)$ to estimate the LATE for takers
  - Use data from groups $(T{,}S)$ to estimate the LATE for non-takers
  - Recall that we have random (U, T, S) in each group $(\hat{G}^*_U,\ \hat{G}^*_T, \hat{G}^*_S)$
    $\rightarrow$ We can estimates these LATEs in each group $(\hat{G}^*_U,\ \hat{G}^*_T, \hat{G}^*_S)$

- These two LATEs play key roles in the optimal policy assignment
  - In theory section, we show this point theoretically
  - In empirical section, we use our RCT data to demonstrate it (next page)

Mechanism: LATEs for takers and non-takers

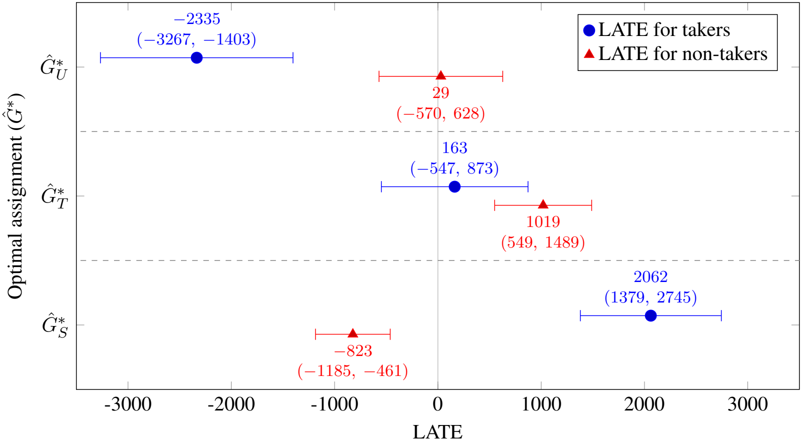

- $\hat{G}_S^*$ (individuals whose optimal assignment is S)
  - LATE for takers is positive & LATE for non-takers is negative
  - Self-selection improves social welfare

Mechanism: LATEs for takers and non-takers

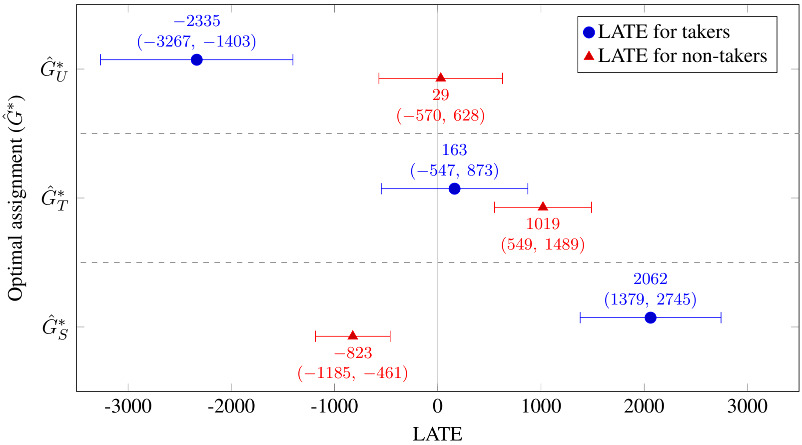

- $\hat{G}^*_T$ (individuals whose optimal assignment is T)
  - Both of the LATEs for takers non-takers are positive
  - Everyone should be treated & self-selection would lower social welfare

Mechanism: LATEs for takers and non-takers

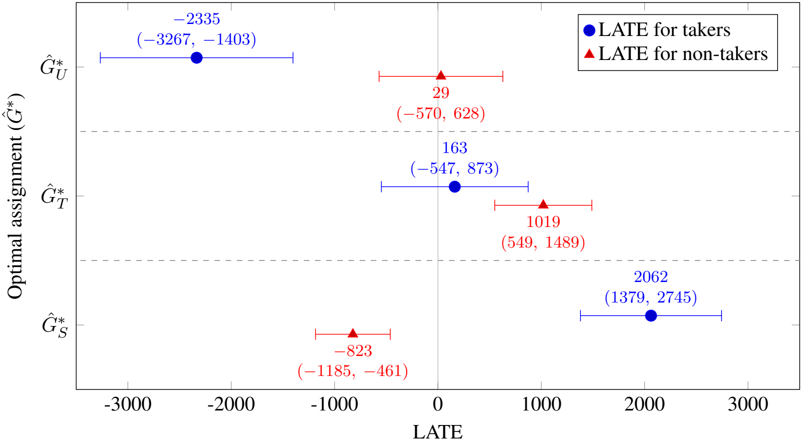

- $\hat{G}^*_U$ (individuals whose optimal assignment is U)
  - LATE for takers is negative and LATE for non-takers is near zero
  - Everyone should be untreated & self-selection would lower social welfare

Mechanism: Counterfactual ITTs

	Consumer types based on the optimal assignment rule $\hat{G}^*$		
	$\hat{G}^*_U$	$\hat{G}^*_T$	$\hat{G}^*_S$
Counterfactual ITT (if assigned to U)	0 (——)	0 (——)	0 (——)
Counterfactual ITT (if assigned to T)	−905.4 (157.8)	662.5 (131.4)	257.8 (117.5)
Counterfactual ITT (if assigned to S)	−923.0 (166.7)	67.9 (150.9)	772.5 (119.7)

- We have randomly-generated (U, T, S) within each group $(\hat{G}^*_U,\ \hat{G}^*_T, \hat{G}^*_S)$
- This variation allows us to estimate counterfactual ITTs for each group
- Result: The counterfactual ITTs confirm the optimality of the assignment

Road map of the talk

1. Introduction
2. Conceptual Framework
3. Field Experiment and Data
4. Optimal Assignment Policy and Welfare Gains
5. Mechanism Behind the Optimal Policy Assignment
6. Welfare Maximization with Redistribution
7. Conclusion

Welfare Maximization with Redistribution

What about equity/re-distributions?

- Our main result uses the utilitarian social welfare function
  - It maximizes efficiency gains but ignores equity
- Is there an equity-efficiency tarde-off?
  - We can look at re-distributional consequence of the policy
  - This table suggests higher-income households would receive more rebates

	Efficiency gain	Average rebate by the quartiles of household income			
		[0%,25%]	(25%,50%]	(50%,75%]	(75%,100%]
Utilitarian ($\nu = 0$)	553.7 (68.0)	72.8 (10.3)	93.7 (12.9)	144.1 (19.1)	148.9 (18.9)

Welfare maximization with redistribution

- Our method is not restricted to a conventional utilitarian framework
  - Any welfare function most appropriate for a policy goal can be used
- Consider a welfare function that balances the equity-efficiency trade-off
  - This function is introduced by Saez (2002)
  - Used by Allcott, Lockwood, and Taubinsky (2019) and Lockwood (2020)
  - Weigh each household's welfare by Pareto weight $w = h^{-\nu}$
  - h is household income
  - ν represents a policymaker's preference for redistribution.
  - $\nu = \infty$ corresponds to the Rawlsian criterion
  - $\nu = 0$ corresponds to utilitarianism.

Welfare maximization with redistribution

	Efficiency gain	Average rebate by the quartiles of household income			
		[0%,25%]	(25%,50%]	(50%,75%]	(75%,100%]
Utilitarian ($\nu = 0$)	553.7	72.8	93.7	144.1	148.9
	(68.0)	(10.3)	(12.9)	(19.1)	(18.9)
With a redistribution goal ($\nu = 1$)	431.2	77.0	132.3	140.2	116.5
	(69.2)	(13.0)	(17.4)	(18.1)	(17.6)
With a redistribution goal ($\nu = 2$)	366.1	105.2	115.7	109.9	119.2
	(69.3)	(14.8)	(16.5)	(16.2)	(20.8)

- We apply our method to a welfare function with different values of ν
- Higher ν indeed improves equity at the cost of sacrificing efficiency
- Policymakers can choose the appropriate level of ν based on this trade-off

Conclusion

Summary

1. Setting: A costly treatment that could generate a social welfare gain
   - Field experiment: A peak-hour rebate program for energy conservation
   - Cost: Implementation cost per participating household
   - Benefit: A social welfare gain if a participant actually conserves energy

2. Our method can be used with RCT or quasi-experimental data to find
   - Consumer type $x \in X$ who should be "treated"
   - Consumer type $x \in X$ who should be "untreated"
   - Consumer type $x \in X$ who should "self-select"

3. Key results:
   - Selection-driven targeting outperforms conventional targeting
   - We show that the LATE framework can be use to reveal the mechanism
   - Our method allows policymakers to identify whose self-selection would be valuable and harmful to social welfare

Thank you!

Koichiro Ito (ito@uchicago.edu)

Appendix Slides

Optimization

- Solving the EWM with the class of depth 6 decision trees is **computationally infeasible**.
- Use the following Two-step Procedure (does not achieve the global optimal):
    1. Search for the best decision tree of depth 3 with the three arms (T,U,S).
    2. For each leaf node in the depth 3 decision tree, search the best decision tree of depth 3 with the three arms (T,U,S).

Estimation and inference on welfare gain $\mathcal{W}(\widehat{G})$

- We want to evaluate the welfare gain $\mathcal{W}(\widehat{G})$ of $\widehat{G}$.
- **Caution:** $\widehat{\mathcal{W}}(\widehat{G})$ is an upwardly biased estimate of $\mathcal{W}(\widehat{G})$ (*winner's bias*).
- We make an artificial test data set for unbiased estimation and inference of $\mathcal{W}(\widehat{G})$.

Source of the bias

- The outcomes Y_j ($j \in \{T, U, S\}$) can be decomposed into **essential terms** and **noises**:

$$Y_j = \underbrace{E[Y_j \mid X]}_{\text{essential term}} + \underbrace{\epsilon_j}_{\text{noise}}.$$

- A learning algorithm inevitably responds to the noise (i.e., overfitting to the data) $\Rightarrow \widehat{\mathcal{W}}(\widehat{G})$ is upwardly biased.
- One possible solution is to split the whole data to **training data** (to estimate the policy) and **test data** (to estimate the welfare performance) $\Rightarrow$ Inefficient!
- Idea of our solution: Replace the noise in the training data with a second independent data to make an artificial test data.

Artificial test data

- Generate artificial test data $\{(Y_i^{test}, D_i^{test}, Z_i, X_i) : i = 1, \ldots, n\}$ as follows:

- $\{Y_i^{test} : i = 1, \ldots, n\}$ is generated as follows: For $j \in \{T, U, S\}$ and samples $i \in I_j := \{i : Z_i = j\}$,
  1. Estimate $E[Y_j \mid X]$ (by e.g., random forest) and calculate residuals $\hat{\epsilon}_i = Y_i - \widehat{E}[Y_{j,i} \mid X_i]$.
  2. Randomly sample $\{\hat{\epsilon}_i^{test} : i \in I_j\}$ from $\{\hat{\epsilon}_i\}_{i \in I_j}$ with replacement.
  3. Construct $Y_i^{test} = \widehat{E}[Y_{j,i} \mid X_i] + \hat{\epsilon}_i^{test}$ for each $i \in I_j$.

- $\{D_i^{test} : i = 1, \ldots, n\}$ is generated as follows: For $i \in I_S$,
  1. Estimate $P(D_{S,i} = 1 \mid X_i)$.
  2. Sample $\{D_i^{test} : i \in I_S\}$ according to $D_i^{test} \sim \hat{P}(D_{S,i} = 1 \mid X_i)$.
- Use this test data to estimate $\mathcal{W}(\widehat{G})$ and do its inference.