Dynamic Targeting:

Experimental Evidence from Energy Rebate Programs

Takanori Ida [1] Takunori Ishihara [2] Koichiro Ito [3] Daido Kido [4]
Toru Kitagawa [5] Shosei Sakaguchi [6] Shusaku Sasaki [7]

[1]Kyoto University [2]Kyoto University of Advanced Science [3]University of Chicago
[4]Otaru University of Commerce [5]Brown University [6]University of Tokyo [7]Osaka University

June 26, 2024

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)
  - Transfer program in development (Alatas, Purnamasari, Wai-Poi, Banerjee, Olken, Hanna, 2016)
  - Energy efficiency (Burlig, Knittel, Rapson, Reguant, and Wolfram, 2020)
  - Behavioral nudge (Knittel and Stolper, 2019)
  - Electricity pricing (Ito, Ida, and Tanaka, 2023)
  - Selection-driven targeting (Ida, Ishihara, Ito, Kido, Kitagawa, Sakagushi, Sasaki, 2023)

The literature has been focusing on “static” targeting

- However, many economic policies involve dynamics
  - Individuals often receive policy interventions repeatedly
    - Job training programs (Lechner, 2009; Rodríguez et al., 2022)
    - Unemployment insurance programs (Meyer, 1995; Kolsrud et al., 2018)
    - Healthcare programs (Luckett et al., 2019)
    - Educational interventions (Ding and Lehrer, 2010).

- How should we think about dynamic targeting?

- Consider two-period interventions with a binary treatment
  - $d_t =$ (T,U) is treatment assignment at time $t = 1, 2$
  - How can we think about dynamically-optimal targeting for $t = 1, 2$?

How should we think about dynamic targeting?

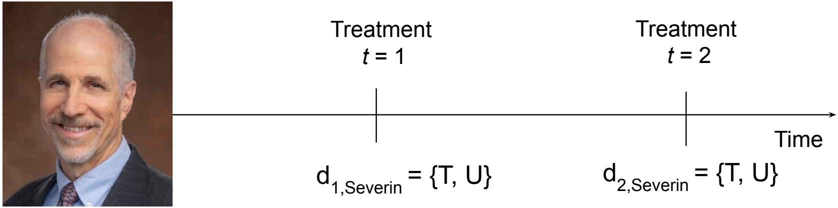

An example question: "Should Severin get treated at $t = 1$?"

How should we think about dynamic targeting?

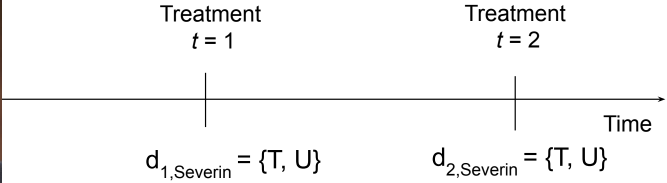

An example question: "Should Severin get treated at $t = 1$?"

1. Yes if welfare gain from his treatment at $t = 1$ is large (static reason)

How should we think about dynamic targeting?

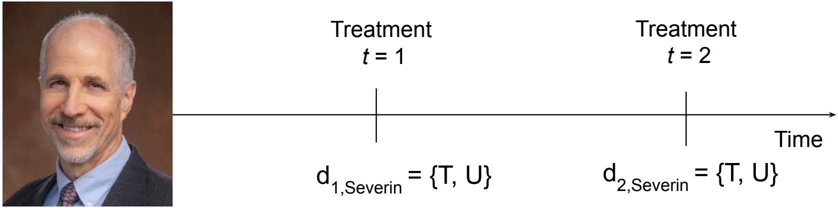

An example question: "Should Severin get treated at $t = 1$?"

1. Yes if welfare gain from his treatment at $t = 1$ is large (static reason)
2. Yes if he has a learning effect
   - Experiencing treatment at $t = 1$ enhances treatment response at $t = 2$

How should we think about dynamic targeting?

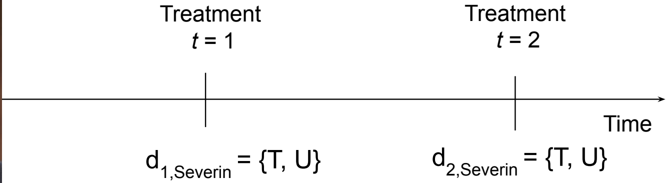

An example question: "Should Severin get treated at $t = 1$?"

1. Yes if welfare gain from his treatment at $t = 1$ is large (static reason)
2. Yes if he has a learning effect
   - Experiencing treatment at $t = 1$ enhances treatment response at $t = 2$
3. Yes if he has a screening effect
   - How he responds to treatment at $t = 1$ helps us to identify his optimal assignment at $t = 2$

We theoretically and empirically study dynamic targeting

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
   - Goal: Find dynamically-optimal targeting for a multi-period intervention
2. Use an RCT & the Empirical Welfare Maximization (EWM) to identify
   - Who should be treated & when they should be treated
3. Test hypotheses for several possible mechanisms
   - Learning (or fatigue) effects
   - Habit formation effects
   - Screening effects

Road map of the talk

1. Introduction
2. Conceptual Framework
3. Estimation Method
4. Field Experiment and Data
5. Welfare Gains from Dynamic Targeting
6. Mechanism Behind the Dynamic Targeting
7. Conclusion

Conceptual Framework

Setup

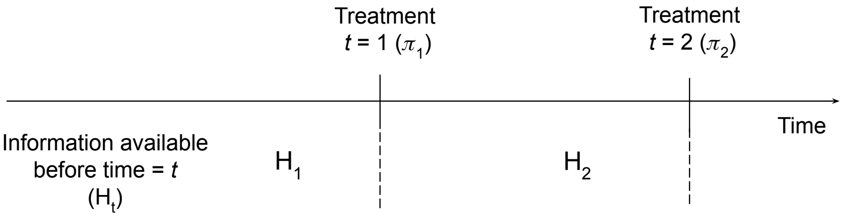

- Consider a two-period model (can be extended to more periods)
  - Treatment is binary: $d_t = \{T, U\}$ in $t = 1, 2$
  - Potential outcome of welfare in time 1: $Y_1(d_1)$
  - Potential outcome of welfare in time 2: $Y_2(d_1, d_2)$

Setup

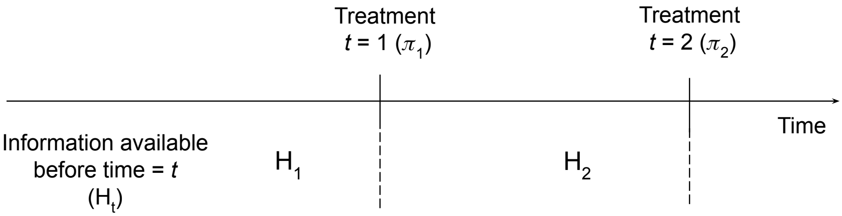

- Consider a two-period model (can be extended to more periods)
  - Treatment is binary: $d_t = \{T, U\}$ in $t = 1, 2$
  - Potential outcome of welfare in time 1: $Y_1(d_1)$
  - Potential outcome of welfare in time 2: $Y_2(d_1, d_2)$
- Planner considers targeting policy π based on observable data
  - Information available before time t: $H_t \in \mathcal{H}_t$
  - Targeting policy $\pi_t : \mathcal{H}_t \to \{T, U\}$

What information is available for the planner?

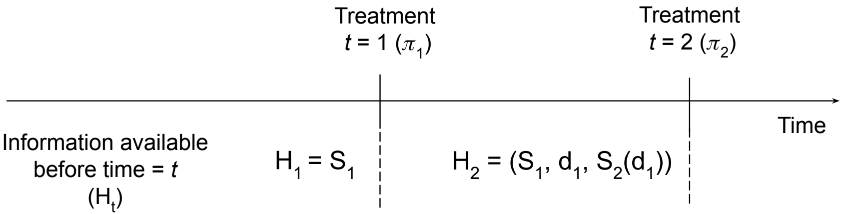

- Information available before $t = 1$
  - S_1: data from the pre-period (e.g., demographics, past energy use)

What information is available for the planner?

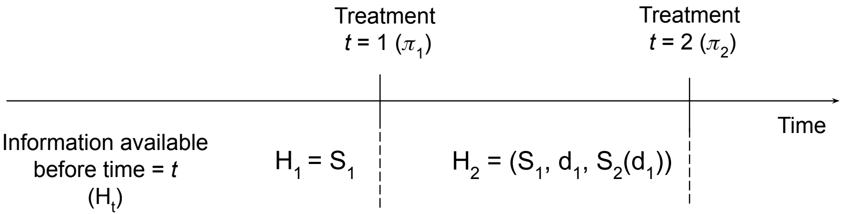

- Information available before $t = 1$
  - S_1: data from the pre-period (e.g., demographics, past energy use)

- Information available after $t = 1$
  - S_1: data from the pre-period (e.g., demographics, past energy use)
  - d_1: treatment assignment at $t = 1$
  - $S_2(d_1)$: data available after $t = 1$ (e.g., each consumer's response to d_1, i.e., how their electricity usage responded to d_1)

Dynamic targeting

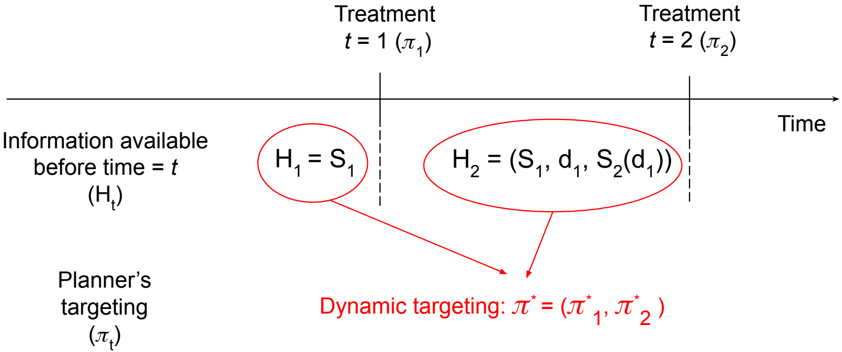

- Planner exploits both H_1 and H_2 to design targeting π
  - We allow $S_2(d_1)$ to be endogenous to d_1
  - d_1 not only affects $Y_1(d_1)$ but also affects $S_2(d_1)$

Dynamic targeting

- The optimal dynamic targeting π^* is obtained by:

$$\begin{aligned} \max_{\pi} \; W(\pi) \equiv \; & E\Big[Y_1(d_1) + Y_2(d_1, d_2)\Big], \\ \text{s.t. } d_1 = \; & \pi_1(H_1) \in \{T, U\}, \\ d_2 = \; & \pi_2(H_2(d_1)) \in \{T, U\}. \end{aligned}$$

  - Y_1 and Y_2: welfare gains in time 1 and 2
  - $d_1, d_2 \in \{T, U\}$: treatment assignment in time 1 and 2
  - H_1, H_2: information available before time 1 and 2

Three reasons why dynamic targeting could matter

1. Learning & habituation effects on $Y_2(d_1, d_2)$
   - Learning effects if $Y_2(T, T) - Y_2(U, T) > 0$
   - Habituation (fatigue) effects if $Y_2(T, T) - Y_2(U, T) < 0$

Three reasons why dynamic targeting could matter

1. Learning & habituation effects on $Y_2(d_1, d_2)$
   - Learning effects if $Y_2(T, T) - Y_2(U, T) > 0$
   - Habituation (fatigue) effects if $Y_2(T, T) - Y_2(U, T) < 0$

2. Habit formation effects on $Y_2(d_1, d_2)$
   - Habit formation effects if $Y_2(T, U) - Y_2(U, U) > 0$

Three reasons why dynamic targeting could matter

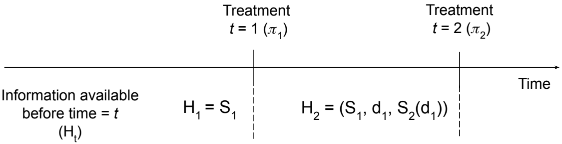

3. Screening effects of $d_1 = T$ on $Y_2(d_1, d_2)$
   - Information H_2 depends on d_1 (treatment in $t = 1$)
   - For example, suppose $H_2(d_1 = T)$ is more informative to predict treatment heterogeneity in $t = 2$ than $H_2(d_1 = U)$
   - In this case, assigning $d_1 = T$ is beneficial (screening effects), even though it could come at the cost of not-maximizing welfare in $t = 1$

Three reasons why dynamic targeting could matter

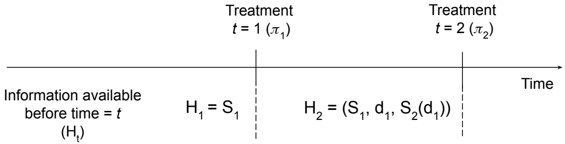

3. Screening effects of $d_1 = T$ on $Y_2(d_1, d_2)$
   - Information H_2 depends on d_1 (treatment in $t = 1$)
   - For example, suppose $H_2(d_1 = T)$ is more informative to predict treatment heterogeneity in $t = 2$ than $H_2(d_1 = U)$
   - In this case, assigning $d_1 = T$ is beneficial (screening effects), even though it could come at the cost of not-maximizing welfare in $t = 1$

   - Screening effect of $d_1 = T$ on $Y_2(d_1, d_2)$:

$$Y_2\Big(T, \pi_2^*(H_2(T))\Big) - Y_2\Big(T, \pi_2^*(H_2(U))\Big)$$

     - Both terms have $d_1 = T$
     - $d_2 = \pi_2^*(H_2(T))$ on the left and $d_2 = \pi_2^*(H_2(U))$ on the right

(Supplemental) The same logic applies to $d_1 = U$

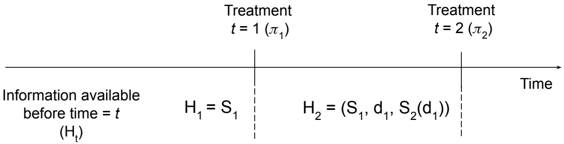

3. Screening effects of $d_1 = U$ on $Y_2(d_1, d_2)$
   - The same logic can apply to $d_1 = U$
   - Suppose $H_2(d_1 = U)$ is more informative to predict treatment heterogeneity in $t = 2$ than $H_2(d_1 = T)$
   - In this case, assigning $d_1 = U$ is beneficial (screening effects), even though it could come at the cost of not-maximizing welfare in $t = 1$

   - Screening effect of $d_1 = U$ on $Y_2(d_1, d_2)$:

$$Y_2\Big(U, \pi_2^*(H_2(U))\Big) - Y_2\Big(U, \pi_2^*(H_2(T))\Big)$$

     - Both terms have $d_1 = U$
     - $d_2 = \pi_2^*(H_2(U))$ on the left and $d_2 = \pi_2^*(H_2(T))$ on the right

Decomposition of gains from dynamic targeting

- We derive a formula that decomposes dynamic targeting's welfare gain
- Welfare gains from dynamic targeting is the sum of the followings:
  - Treatment effect in $t = 1$
  - Treatment effect in $t = 2$
  - Learning effects
  - Habit formation effect
  - Screening effect
- We develop a method to empirically estimate each of these components

Decomposition Theorem

For any dynamic targeting policy $\pi = (\pi_1, \pi_2)$,

$$
\begin{aligned}
W(\pi) - W(U,U) = {} & \underbrace{E[Y_1(T) - Y_1(U) | \pi_1(H_1) = T]}_{\text{Treatment effect on the treated in } t=1} \cdot P_1 \\
& + \underbrace{E[Y_2(U,T) - Y_2(U,U) | \pi_2(H_2(U)) = T]}_{\text{Treatment effect on the treated in } t=2} \cdot P_2 \\
& + \underbrace{E[Y_2(T,U) - Y_2(U,U) | \pi_1^*(H_1) = T, \pi_2^*(H_2(T)) = U]}_{\text{Habit formation effect for those assigned to } (T,U)} \cdot P_3 \\
& + \underbrace{E[Y_2(T,T) - Y_2(U,T) | \pi_1^*(H_1) = T, \pi_2^*(H_2(T)) = T]}_{\text{Learning effect for those assigned to } (T,T)} \cdot P_4 \\
& + \underbrace{E[Y_2(T, \pi_2^*(H_2(T))) - Y_2(T, \pi_2^*(H_2(U))) | \pi_1^*(H_1) = T]}_{\text{Screening effect for those assigned to } T \text{ in } t=1} \cdot P_5,
\end{aligned}
$$

where $\{P_k : k = 1, \ldots, 5\}$ are probabilities of the conditioning events in the conditional expectations that P_k's are multiplied to.

Estimation Method

Empirical Welfare Maximization (EWM) method

- RCT (or quasi-experimental) data: $\{(Y_{it}, Z_{it}, H_{it}) : t = 1, 2\}$ where $Z_{it} \in \{T, U\}$ at $t = 1, 2$ is randomly assigned.

- With random assignment, the empirical analogue of $W(\pi)$ is

$$\widehat{W}(\pi) = \frac{1}{n}\sum_{i=1}^{n}\left(\frac{Y_{i1}\cdot 1\{Z_{i1}=\pi_1(H_{i1})\}}{P(Z_1 = Z_{i1} \mid H_1 = H_{i1})}\right.$$
$$\left. + \frac{Y_{i2}\cdot 1\{Z_{i1}=\pi_1(H_{i1}), Z_{i2}=\pi_2(H_{i2})\}}{P(Z_1 = Z_{i1} \mid H_1 = H_{i1})\cdot P(Z_2 = Z_{i2} \mid H_2 = H_{i2})}\right).$$

- We use a class of policy trees (Zhou, Athey, Wager, 2022) for Π.

Road map of the talk

1. Introduction
2. Conceptual Framework
3. Estimation Method
4. Field Experiment and Data
5. Welfare Gains from Dynamic Targeting
6. Mechanism Behind the Dynamic Targeting
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 ($t = 1$)
   - Peak-hour: 5 pm to 9 pm in critical peak days in winter 2020 ($t = 2$)
   - 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. Randomize 2,400 residential customers into four groups
   - $(Z_1, Z_2) = (U, U)$: 625
   - $(Z_1, Z_2) = (U, T)$: 606
   - $(Z_1, Z_2) = (T, U)$: 581
   - $(Z_1, Z_2) = (T, T)$: 588

Summary statistics and balance check

	Sample mean by group [standard deviation]			
	(U, U)	(U, T)	(T, U)	(T, T)
Peak hour usage (2020 summer, Wh)	201 [145]	200 [136]	196 [136]	198 [136]
Pre-peak hour usage (2020 summer, Wh)	189 [143]	184 [130]	183 [137]	182 [130]
Post-peak hour usage (2020 summer, Wh)	311 [175]	311 [171]	308 [164]	305 [163]
Peak hour usage (2020 winter, Wh)	311 [194]	309 [170]	304 [179]	306 [170]
Pre-peak hour usage (2020 winter, Wh)	171 [117]	171 [102]	169 [112]	166 [102]
Post-peak hour usage (2020 winter, Wh)	287 [198]	295 [198]	280 [203]	287 [192]
Number of people at home (1 PM - 5 PM)	1.31 [1.04]	1.32 [0.96]	1.31 [1.04]	1.34 [1.01]
Number of people at home (5 PM - 9 PM)	2.57 [1.29]	2.48 [1.20]	2.47 [1.23]	2.51 [1.20]
Self-efficacy in energy conservation (1-5 scale)	3.44 [0.84]	3.44 [0.86]	3.47 [0.86]	3.44 [0.82]
Household income (JPY 10,000)	651 [400]	639 [387]	614 [393]	606 [333]

Road map of the talk

1. Introduction
2. Conceptual Framework
3. Estimation Method
4. Field Experiment and Data
5. Welfare Gains from Dynamic Targeting
6. Mechanism Behind the Dynamic Targeting
7. Conclusion

Welfare Gains from Dynamic Targeting

Welfare gain

- $Q_t(T)$ and $Q_t(U)$: Potential outcomes of household's peak-hour electricity consumption (kWh) in $t = 1, 2$

- Household's potential welfare contribution:

$$Y_t(T) \equiv \underbrace{b}_{\text{benefit}} \times \underbrace{\Big(Q_t(U) - Q_t(T)\Big)}_{\text{electricity conservation}} - \underbrace{c}_{\text{cost}}.$$

  - b: marginal social welfare gain from a unit reduction in energy use
  - c: implementation cost of the program.

We compare welfare gains form several policies

1. Non-targeting policies
   - Everyone is assigned to $(U, U) \rightarrow$ baseline
   - Everyone is assigned to (T, U)
   - Everyone is assigned to (U, T)
   - Everyone is assigned to (T, T)

2. Static targeting I
   - Planner uses only H_1 (pre-intervention information)
   - Assignment cannot change over time: (U, U) or (T, T)

3. Static targeting II
   - Planner uses only H_1 (pre-intervention information)
   - Assignment can change over time: (U, U), (T, U) , (U, T), (T, T)

4. Dynamic targeting
   - Planner uses H_1 and H_2 to allocate (U, U), (T, U) , (U, T), (T, T)
   - Planner solves dynamic optimization

Welfare Gains from Each Policy

Policy	Welfare gain	Share of customers in each arm			
		(U,U)	(T,U)	(U,T)	(T,T)
100% (U,U)	0.0 (0.0)	100.0%	0.0%	0.0%	0.0%
100% (T,U)	311.8 (378.4)	0.0%	100.0%	0.0%	0.0%
100% (U,T)	470.8 (457.5)	0.0%	0.0%	100.0%	0.0%
100% (T,T)	463.9 (452.2)	0.0%	0.0%	0.0%	100.0%
Static targeting I ($\pi^*_{S(I)}$)	770.6 (283.7)	45.6%	0.0%	0.0%	54.4%
Static targeting II ($\pi^*_{S(II)}$)	845.3 (348.9)	3.1%	31.3%	41.5%	24.0%
Dynamic targeting (π^*)	1684.3 (303.1)	19.5%	22.9%	25.6%	32.0%

- Both static and dynamic targeting improves welfare
- Dynamic targeting can double the welfare gain compared to static targeting

Comparisons of Alternative Policies

	Difference in welfare gains	p-value
Dynamic targeting (π^*) vs. 100% (T, U)	1365.7 (309.1)	0.000
Dynamic targeting (π^*) vs. 100% (U, T)	1546.5 (328.7)	0.000
Dynamic targeting (π^*) vs. 100% (T, T)	1397.7 (319.8)	0.000
Dynamic targeting (π^*) vs. Static targeting I ($\pi^*_{S(I)}$)	913.8 (269.2)	0.000
Dynamic targeting (π^*) vs. Static targeting II ($\pi^*_{S(II)}$)	839.1 (287.8)	0.002

- Welfare improvement from dynamic targeting is statistically significant

Road map of the talk

1. Introduction
2. Conceptual Framework
3. Estimation Method
4. Field Experiment and Data
5. Welfare Gains from Dynamic Targeting
6. Mechanism Behind the Dynamic Targeting
7. Conclusion

Mechanism Behind the Dynamic Targeting

Decomposition Theorem

For any dynamic targeting policy $\pi = (\pi_1, \pi_2)$,

$$\begin{aligned}
W(\pi) - W(U,U) = & \underbrace{E[Y_1(T) - Y_1(U) \mid \pi_1(H_1) = T]}_{\text{Treatment effect on the treated in } t=1} \cdot P_1 \\
& + \underbrace{E[Y_2(U,T) - Y_2(U,U) \mid \pi_2(H_2(U)) = T]}_{\text{Treatment effect on the treated in } t=2} \cdot P_2 \\
& + \underbrace{E[Y_2(T,U) - Y_2(U,U) \mid \pi_1^*(H_1) = T, \pi_2^*(H_2(T)) = U]}_{\text{Habit formation effect for those assigned to } (T,\,U)} \cdot P_3 \\
& + \underbrace{E[Y_2(T,T) - Y_2(U,T) \mid \pi_1^*(H_1) = T, \pi_2^*(H_2(T)) = T]}_{\text{Learning effect for those assigned to } (T,\,T)} \cdot P_4 \\
& + \underbrace{E[Y_2(T, \pi_2^*(H_2(T))) - Y_2(T, \pi_2^*(H_2(U))) \mid \pi_1^*(H_1) = T]}_{\text{Screening effect for those assigned to } T \text{ in } t=1} \cdot P_5,
\end{aligned}$$

where $\{P_k : k = 1, \ldots, 5\}$ are probabilities of the conditioning events in the conditional expectations that P_k's are multiplied to.

Decomposition of gains from dynamic targeting

- We derive a formula that decomposes dynamic targeting's welfare gain:

$$
\begin{aligned}
& W(\pi^*) - W(U, U) \\
= \; & \text{Treatment effect on the treated in } t = 1 \\
+ \; & \text{Treatment effect on the treated in } t = 2 \\
+ \; & \text{Habit formation effect for those assigned to } (T, U) \\
+ \; & \text{Learning effect for those assigned to } (T, T) \\
+ \; & \text{Screening effect for those assigned to } T \text{ in } t = 1
\end{aligned}
$$

Decomposition of gains from dynamic targeting

	Welfre contribution
1st-stage treatment effect	214.3
	(103.0)
2nd-stage treatment effect	563.5
	(198.3)
Habit formation effect	287.4
	(184.4)
Learning effect	186.4
	(128.8)
Screening effect	361.5
	(98.0)
Total effect	1613.1
	(397.8)

Conclusion

We theoretically and empirically study dynamic targeting

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
   - Goal: Find dynamically-optimal targeting for a multi-period intervention

2. Use an RCT & the Empirical Welfare Maximization (EWM) to identify
   - Who should be treated & when they should be treated

3. Test hypotheses for several possible mechanisms
   - Learning (or fatigue) effects
   - Habit formation effects
   - Screening effects

Thank you!

Koichiro Ito (ito@uchicago.edu)

Appendix Slides