Comments on Machine Learning Who to Nudge

Athey, Keleher & Spiess

Dean Eckles
Sloan School of Management
Institute for Data, Systems, and Society
MIT

Review

- Analysis of large field experiments — nice to have a 2018 replication:
  - “students who received the treatment interventions were on average 6.4 ± 0.6 (2017) and 12.1 ± 0.7 (2018) percentage points more likely to submit their FAFSA forms by the priority deadline, increasing early filing rates from 37% to 43% and 38% to 50%”
- Use causal forests (and other techniques) to estimate treatment effects for students with different characteristics (HTEs, CATEs)
- Learn and evaluate targeting policies for this intervention

Outline

1. Challenges in heterogeneous treatment effect estimation
2. Learning a targeting policy vs.
   learning heterogeneous treatment effects
3. Theory and intuitions about heterogeneous treatment effects
4. Designing the first experiment

Learning HTEs is hard

- Even in a very well-powered experiment for overall effects, we often find ourselves struggling to precisely and credibly estimate heterogeneous treatment effects
  - Here we have z-statistics >10 for the ATEs, but things get harder for CATEs
  - In practice, large sample sizes often accompanied by high dimensional covariates

- This paper documents challenges in calibrated estimation of CATEs

Learning policies vs. learning HTEs

- When learning HTEs (CATEs), the focus is often on point estimates and then we evaluate these by looking at, e.g., MSE and hopefully other measures of fit (as this paper does!) like calibration
  - In a narrow (statistical decision theory) sense, point estimates are a decision — but what really are we deciding based on them?
- Learning a *policy* is just learning a decision rule for assigning units to treatments
  - This is often essentially a classification problem (if the actions are discrete) with loss from misclassification depending on the true treatment effects
  - Assigning unit i, which should be in control, to treatment results in a loss of $-\tau_i$

Learning policies vs. learning HTEs

- Learning a *policy* is just learning a decision rule for assigning units to treatments
  - This is often essentially a classification problem (if the actions are discrete) with loss from misclassification depending on the true treatment effects
  - Assigning unit i, which should be in control, to treatment results in a loss of $-\tau_i$

A *policy* π is a way of making these choices about actions. It maps from characteristics to actions, i.e., $\pi : \mathbb{X} \to \Delta(\mathbb{A})$

Then targeting is a matter of finding a good (or the best) policy within some set of (perhaps simple) possible policies π, i.e.,

$$\pi^* = \text{argmax}_\pi V(\pi)$$

where $V(\pi) = \mathbf{E}_\pi[Y_i(A_i)]$

Comparing targeting policies

- From a single experiment, one can evaluate many targeting policies — even if each is not specified in advance:
  - Look at cases where observed randomized treatment matches what the policy of interest would have done
- Many policies agree on many cases, so there is lots of cancelation

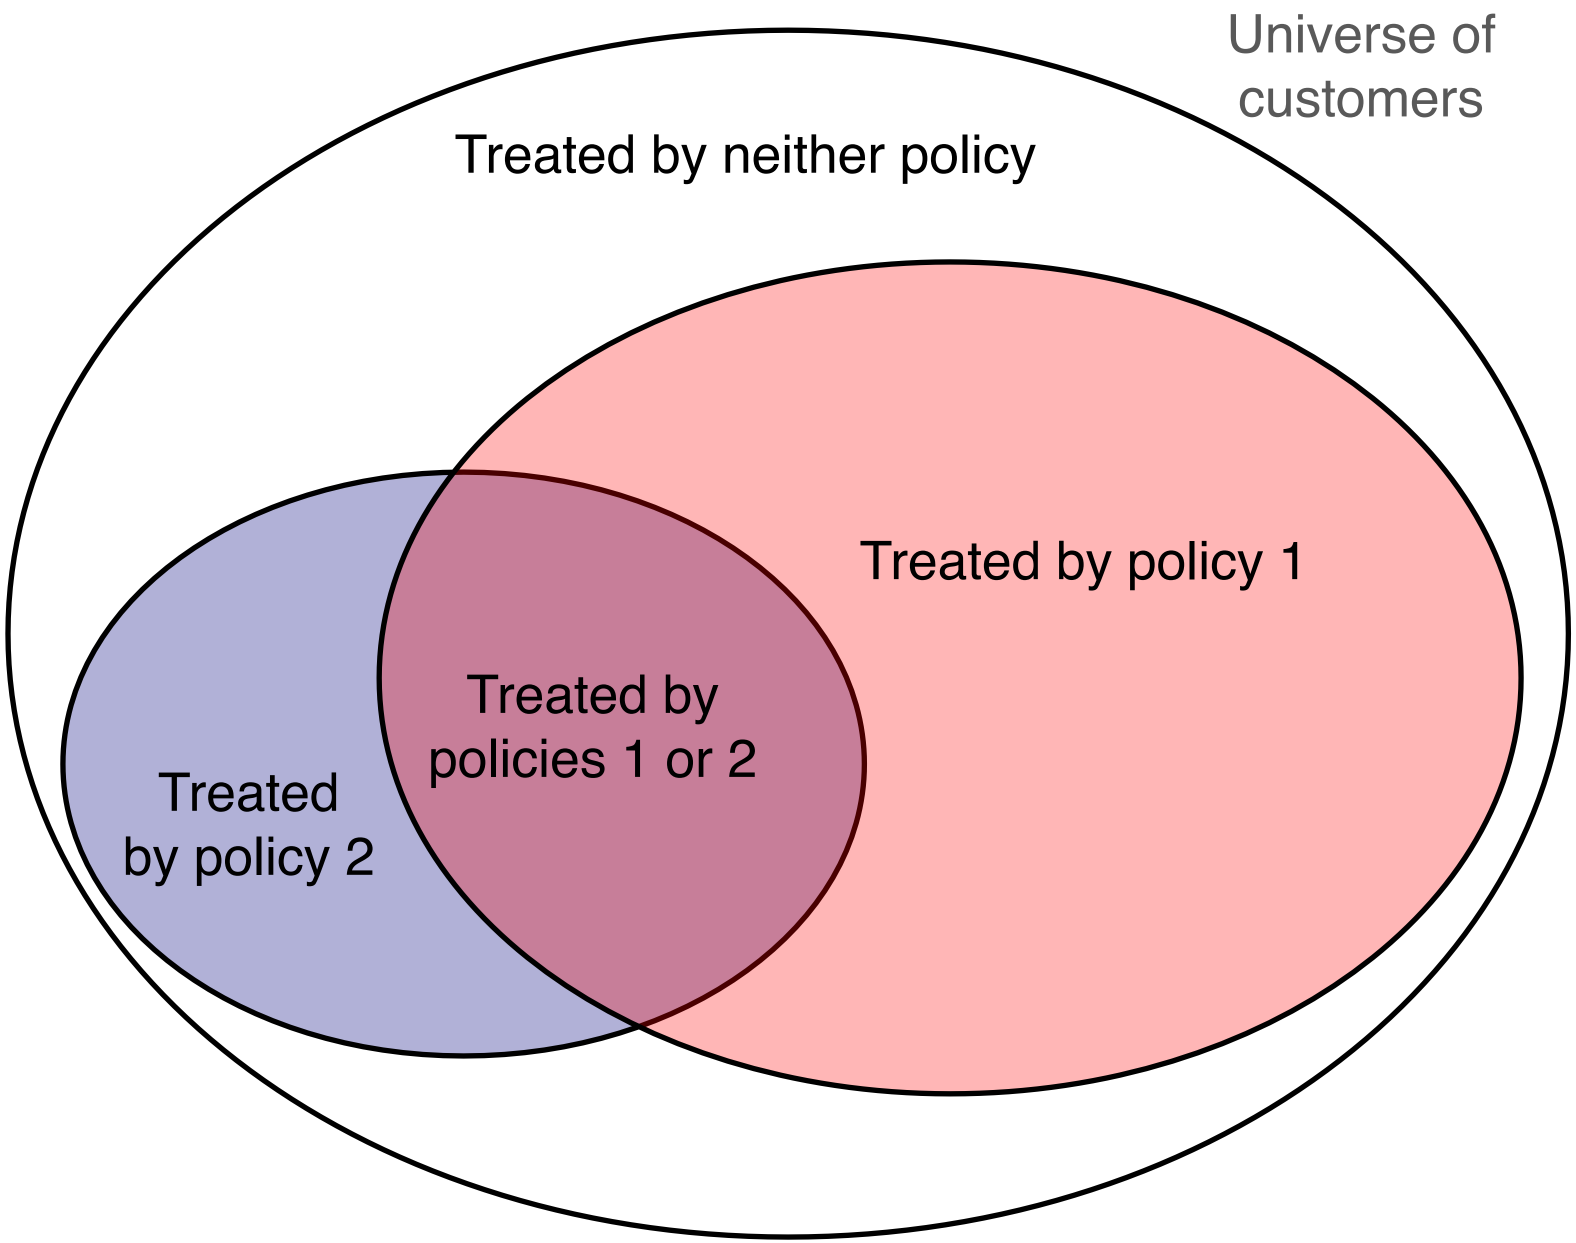

Comparing targeting policies

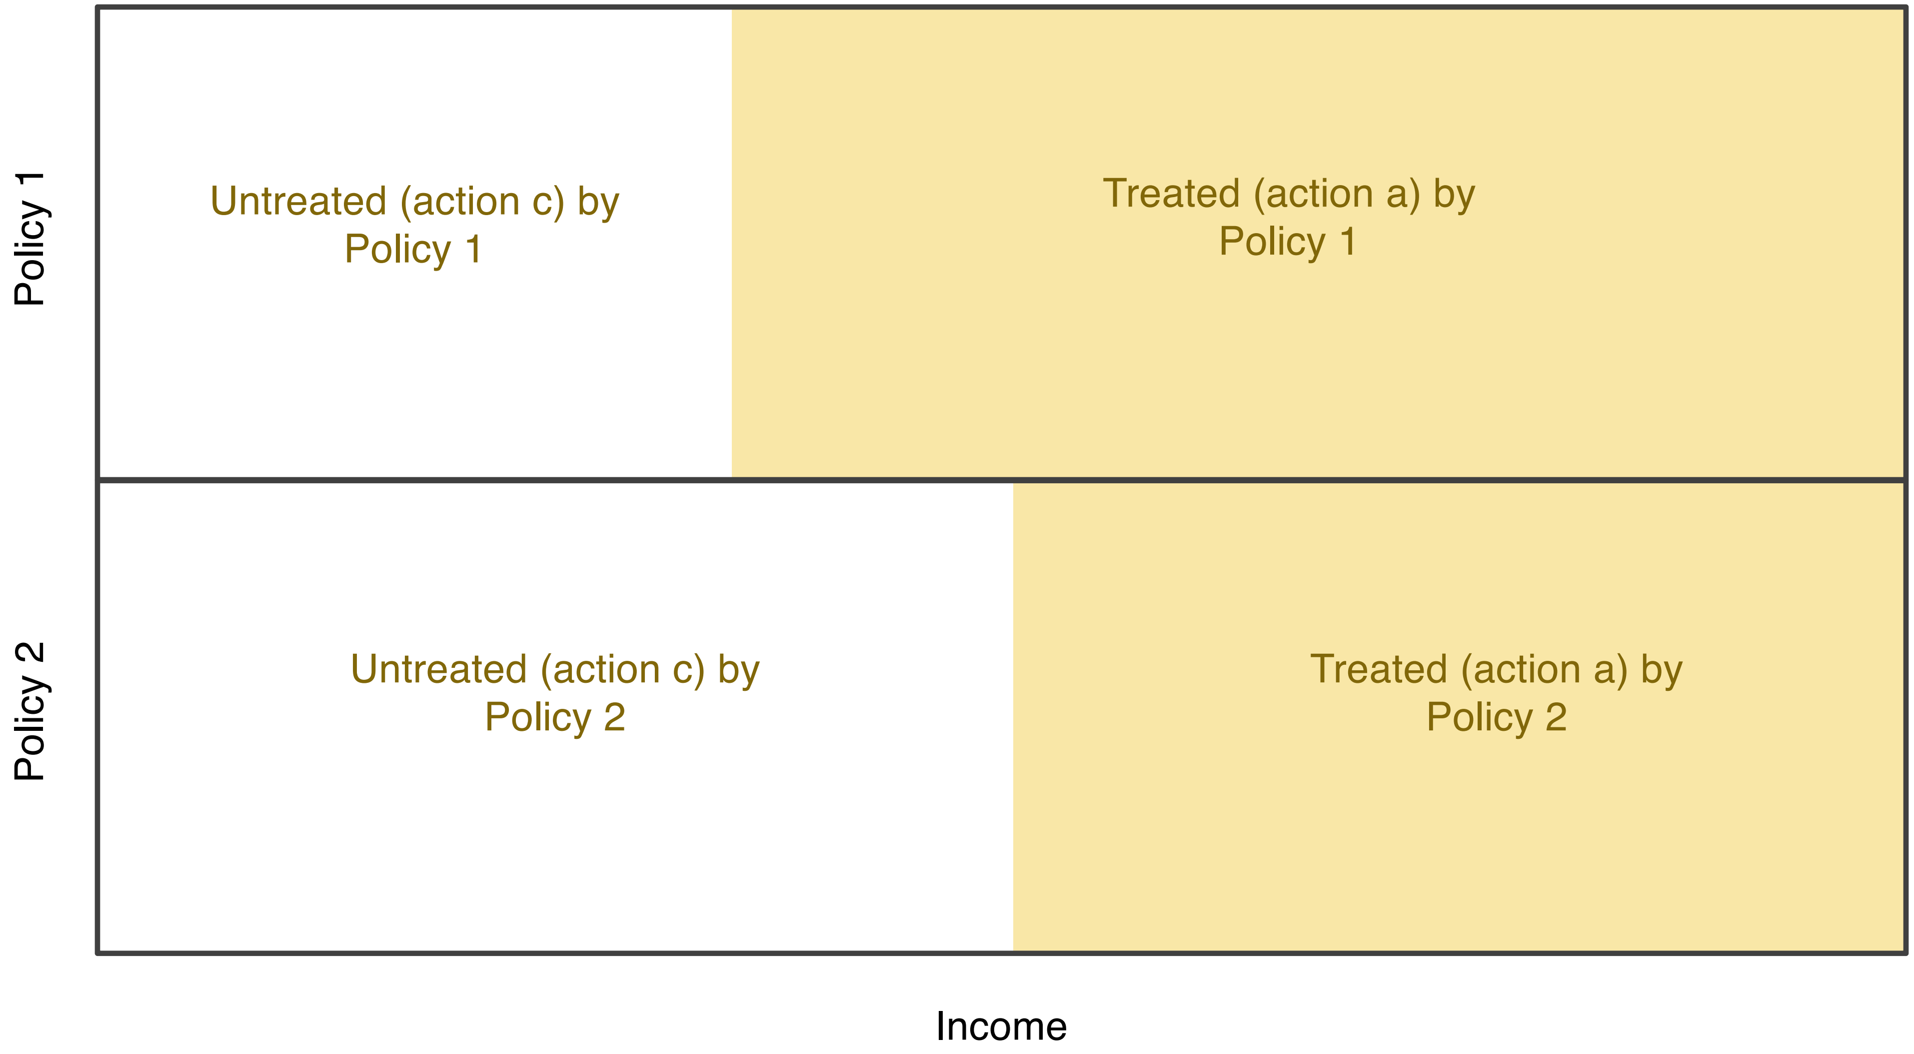

Comparing targeting policies

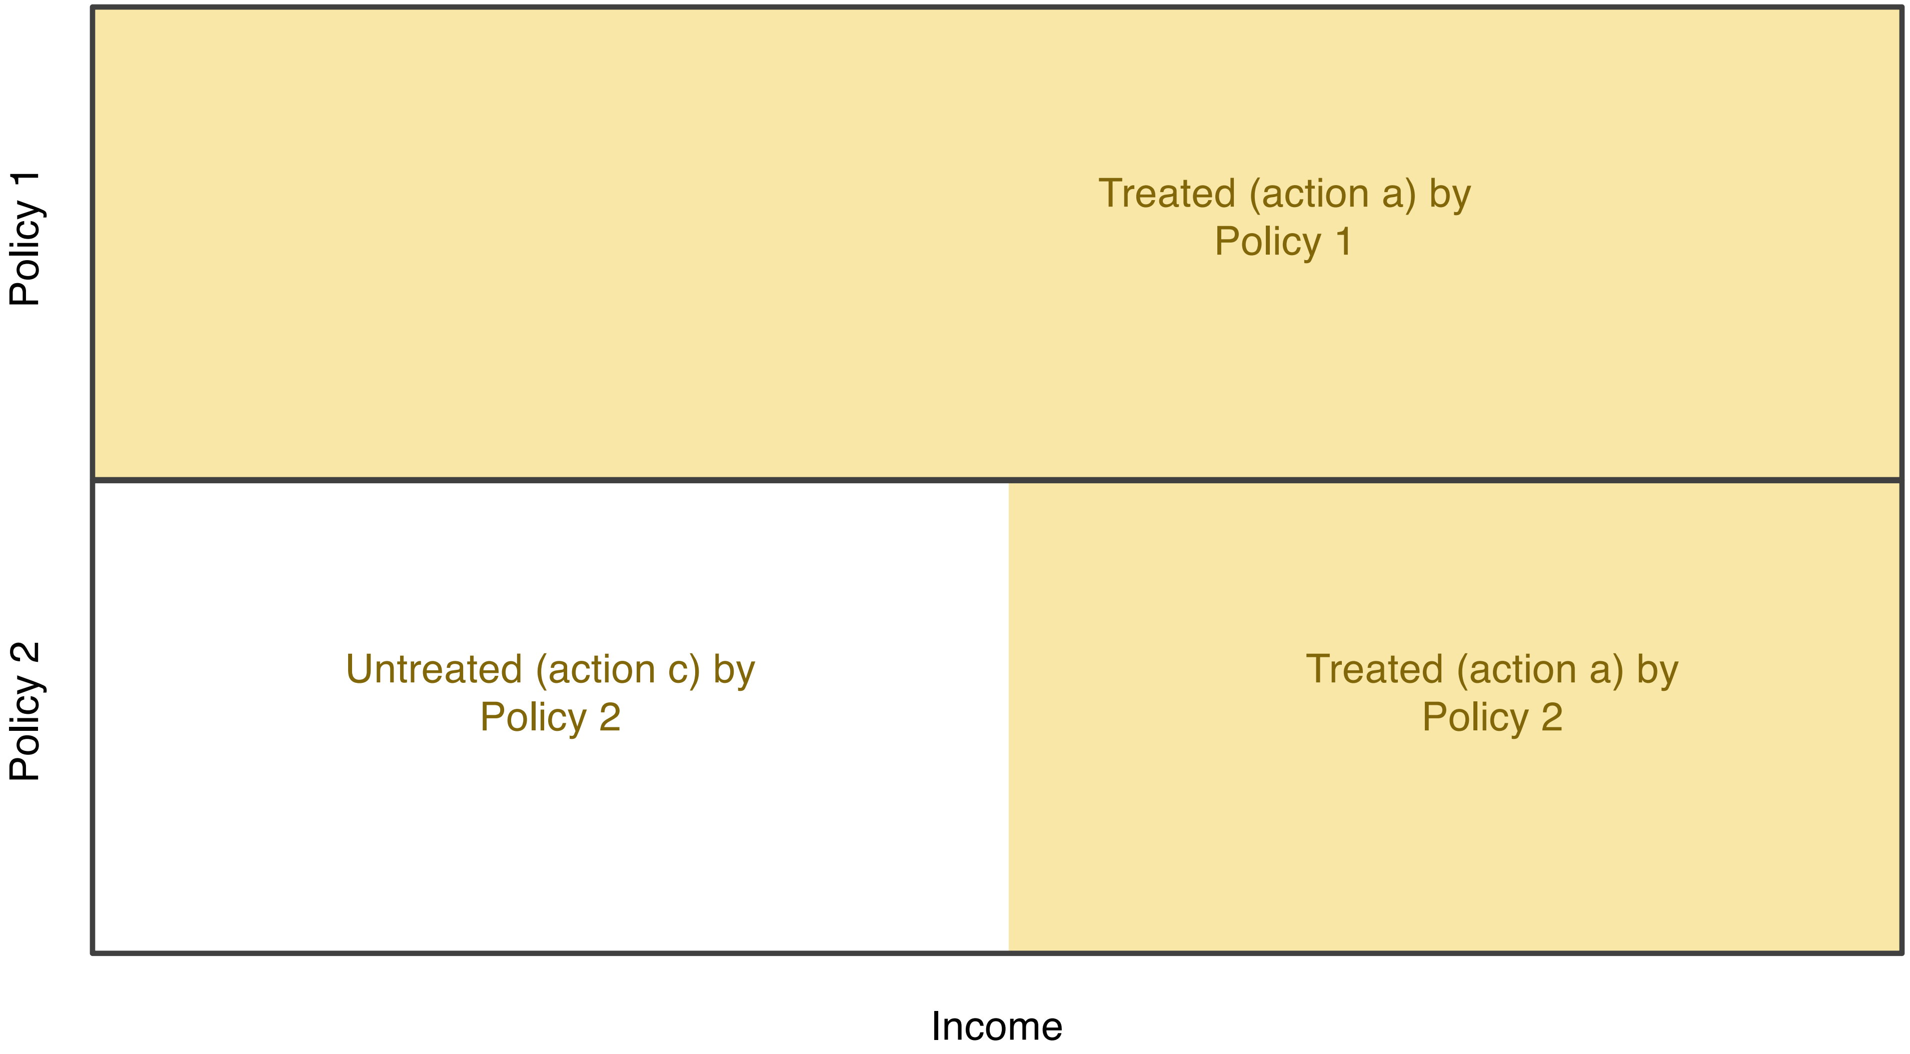

Comparing targeting policies

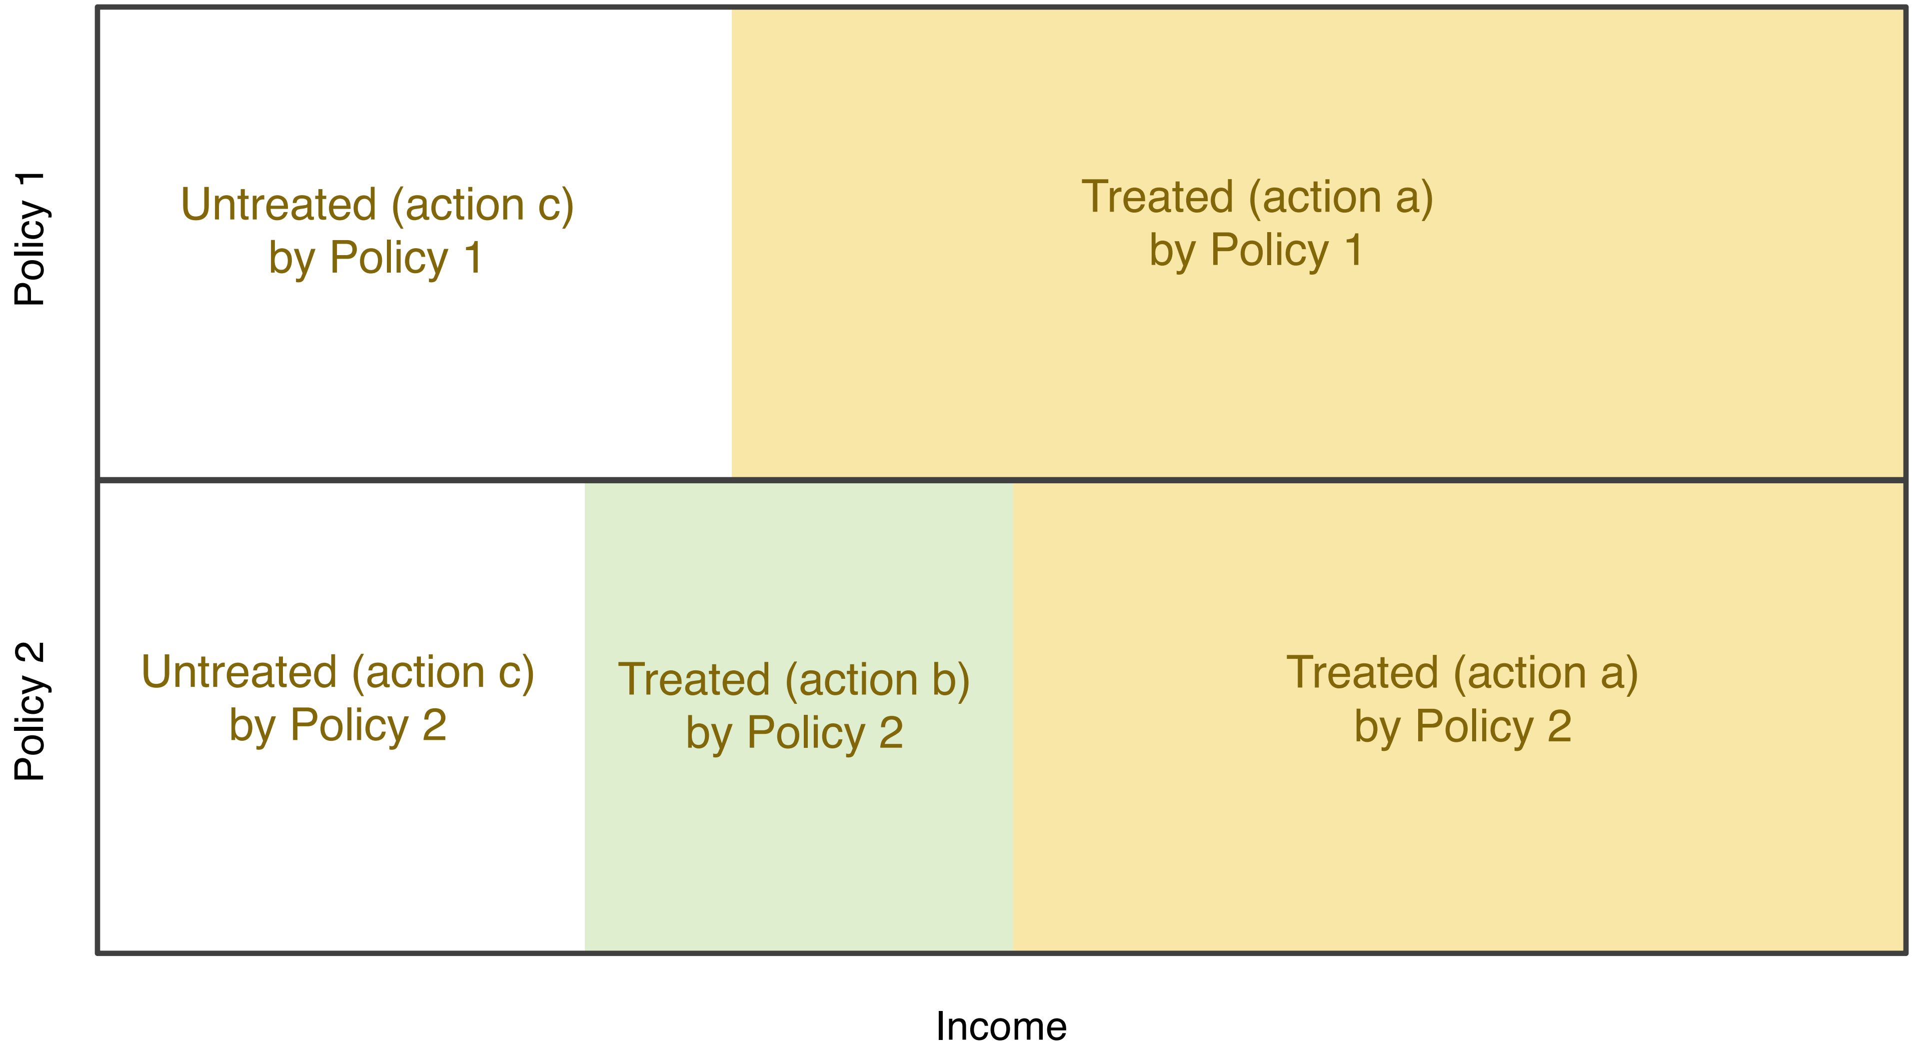

Segment Label	Recommended Action		Sample Size		
	Policy 1	**Policy 2**	Action *a*	Action *b*	Action *c*
Segment_aa	*a*	*a*	595	2,562	824
Segment_ab	*a*	*b*	13,038	11,495	16,086
Segment_ac	*a*	*c*	20,229	8,148	16,164
Segment_ba	*b*	*a*	15,602	12,847	18,784
Segment_bb	*b*	*b*	198,416	215,824	195,885
Segment_bc	*b*	*c*	51,421	45,211	65,311
Segment_ca	*c*	*a*	1,239	2,098	70
Segment_cb	*c*	*b*	19,551	19,972	16,215
Segment_cc	*c*	*c*	37,562	42,616	41,597
	Total		357,653	360,773	370,936

Where do policy 1 and policy 2 disagree?

Simester, Duncan, Artem Timoshenko, and Spyros I. Zoumpoulis. "Efficiently evaluating targeting policies: Improving on champion vs. challenger experiments." Management Science 66.8 (2020): 3412-3424.

Targeting promotional discounts

Post-stratified by what the targeted policy recommends

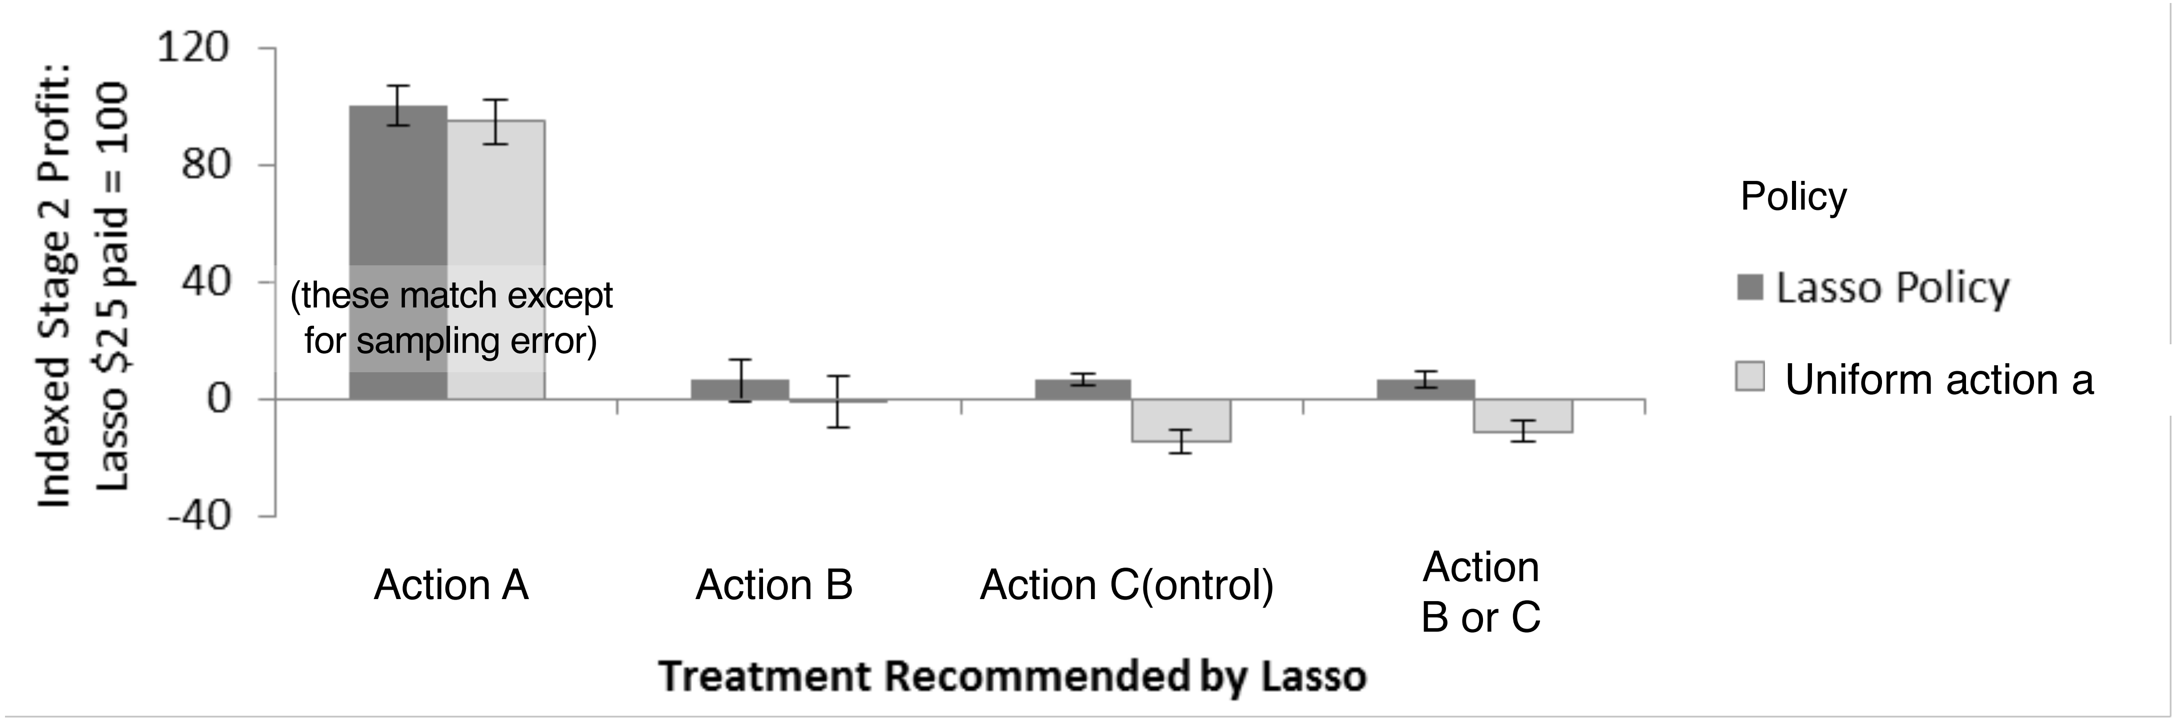

Simester, Duncan, Artem Timoshenko, and Spyros I. Zoumpoulis. "Efficiently evaluating targeting policies: Improving on champion vs. challenger experiments." Management Science 66.8 (2020): 3412-3424.

Targeting retail catalogs

Results of randomizing which customers get catalogs

- Substantial ATE (2.6 in 2015, 2.4 in 2106) on sales, which is larger than cost/margin = 2.003
- Maybe we should just treat everyone?

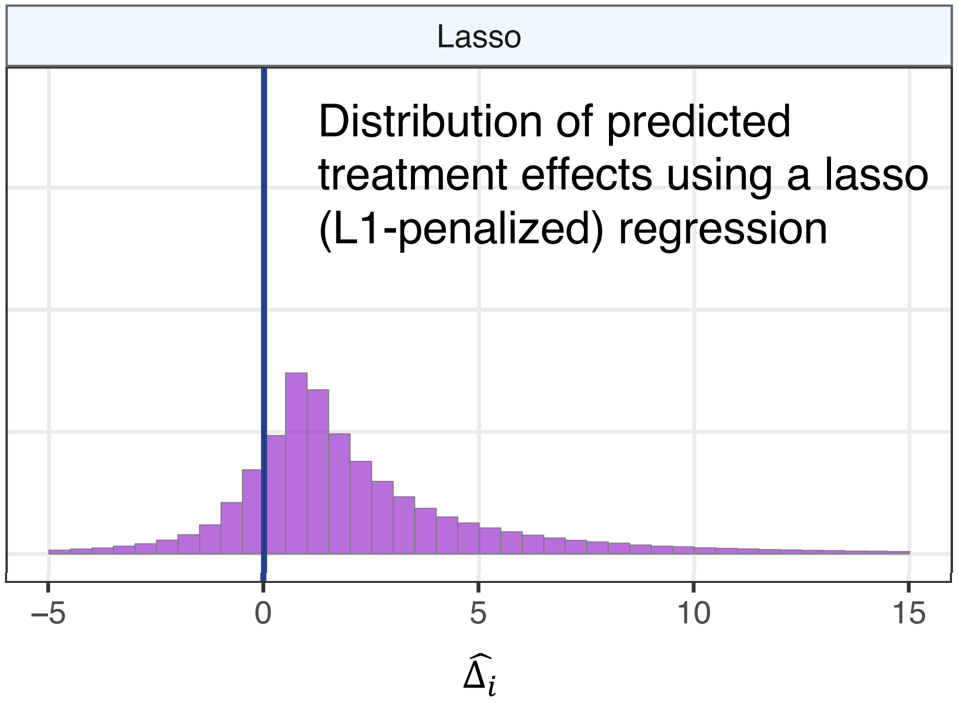

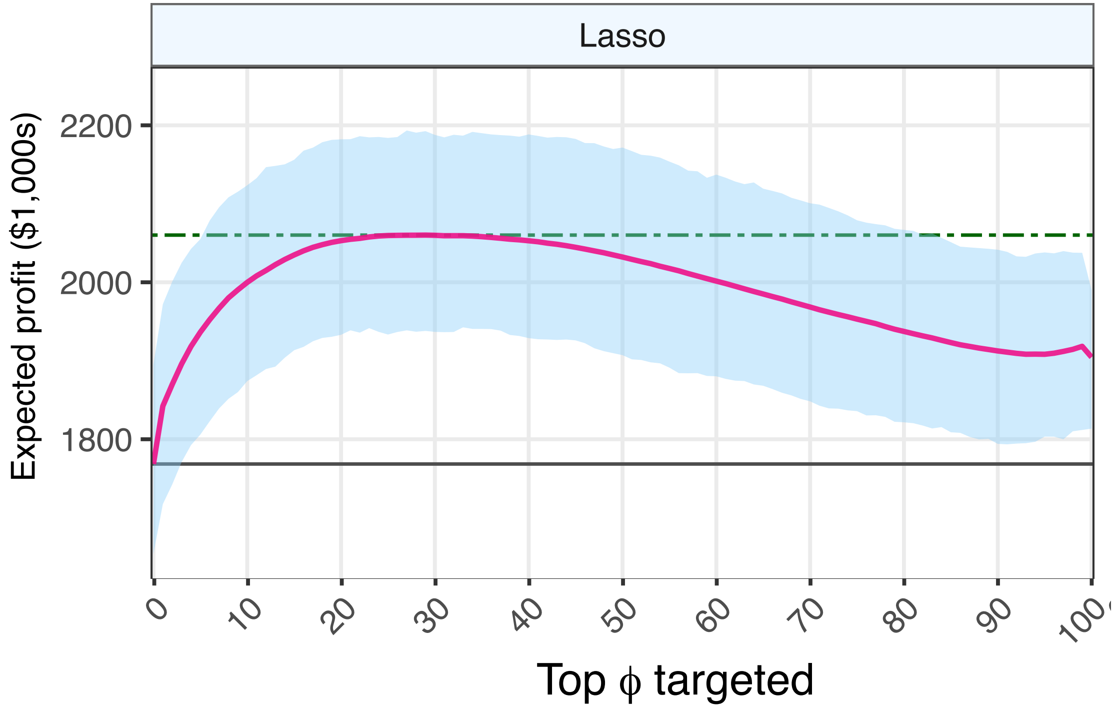

Hitsch, G., & Misra, Sanjog. (2018). Heterogeneous Treatment Effects and Optimal Targeting Policy Evaluation. Working paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3111957

Targeting and budget constraints

Costs of treatment vary from person to person

- In addition to each person having a (possibly different) treatment effect Δ_i, they have some cost C_i
- So far we've been able to just treat everyone with $\Delta_i > C_i$ (or just incorporate costs into Δ_i already)

But we can't spend more than some budget $b

- This might arise from firm-wide constraints or because finance is (reasonably) skeptical about how marketing is spending
- So we want to find who we should target — while limited by this budget

Luedtke, A. R., & van der Laan, M. J. (2016). Optimal individualized treatments in resource-limited settings. *The International Journal of Biostatistics*, *12*(1), 283-303.
Sun, H., Du, S., Wager, S. (2021). Treatment allocation under uncertain costs. https://arxiv.org/abs/2103.11066

Targeting and budget constraints

Rank people by ROI, targeting up to budget

- Rank people by Δ_i / C_i
  - (This is just ROI + 1)
- Keep adding until we reach our budget
- This is a version of "knapsack problem"
- Can use fancier methods to estimate this ratio

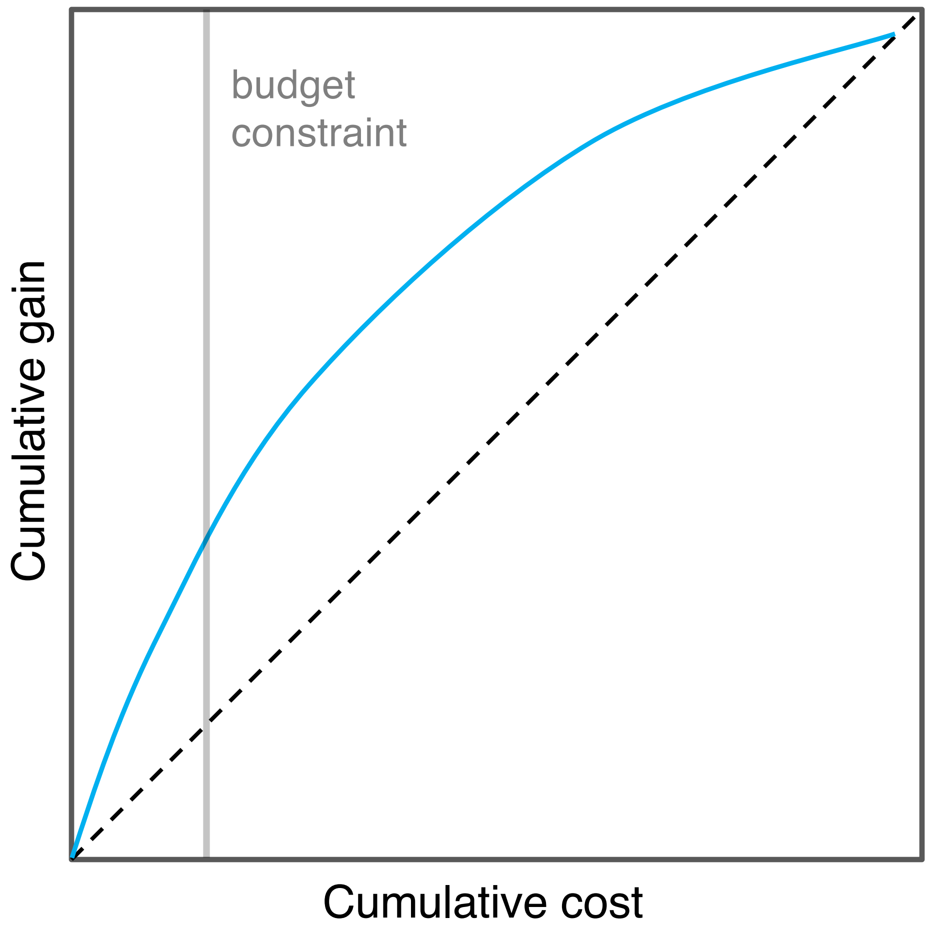

Luedtke, A. R., & van der Laan, M. J. (2016). Optimal individualized treatments in resource-limited settings. *The International Journal of Biostatistics*, *12*(1), 283-303.
Sun, H., Du, S., Wager, S. (2021). Treatment allocation under uncertain costs. https://arxiv.org/abs/2103.11066

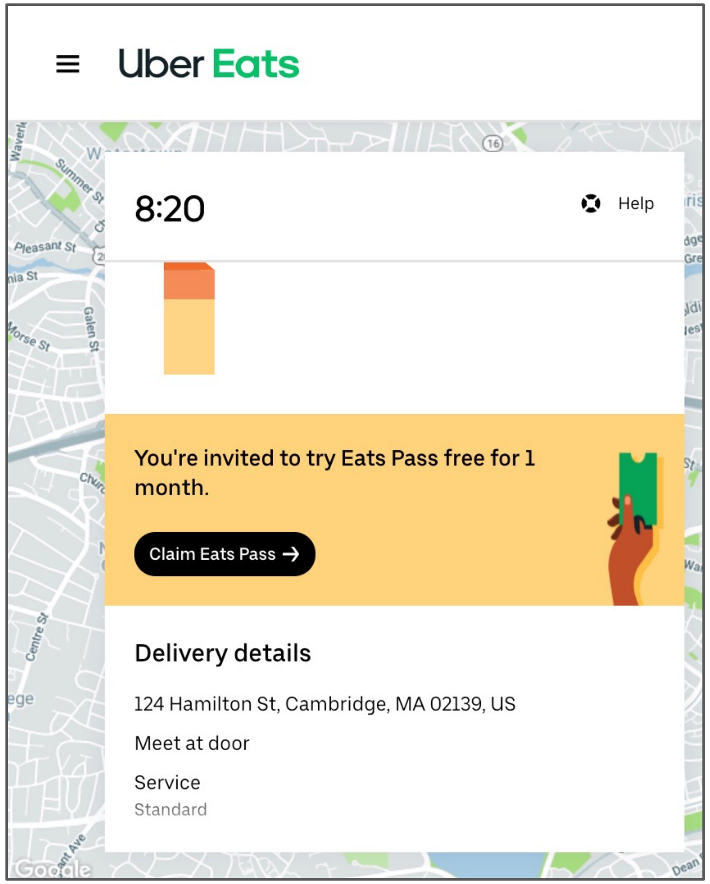
Uber Eats
8:20
Help
You're invited to try Eats Pass free for 1 month.
Claim Eats Pass
Delivery details
124 Hamilton St, Cambridge, MA 02139, US
Meet at door
Service
Standard

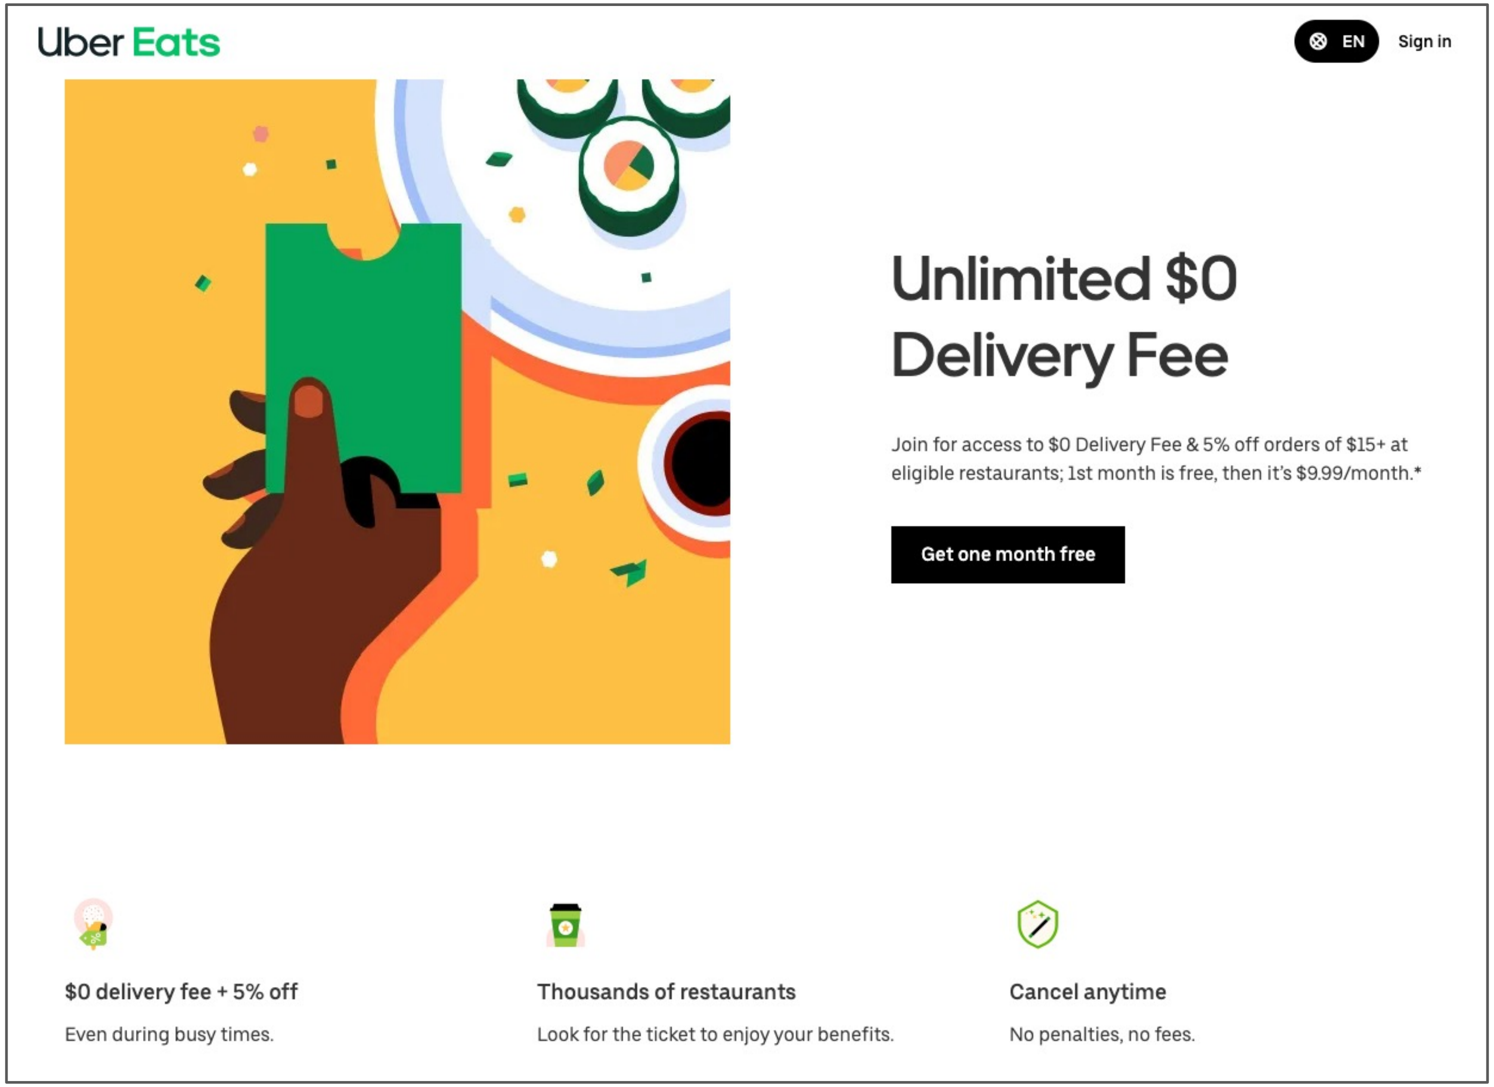
Uber Eats
EN
Sign in
Unlimited $0 Delivery Fee
Join for access to $0 Delivery Fee & 5% off orders of $15+ at eligible restaurants; 1st month is free, then it's $9.99/month.*
Get one month free
$0 delivery fee + 5% off
Even during busy times.
Thousands of restaurants
Look for the ticket to enjoy your benefits.
Cancel anytime
No penalties, no fees.

Targeting a promotion

A free "Uber Pass"

- Have 39 covariates X_i
- Training: 50k; Test: 500k users

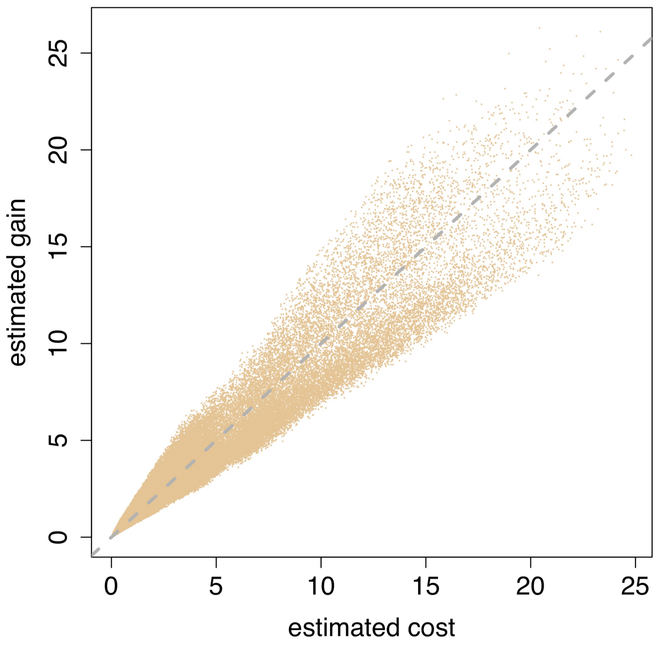

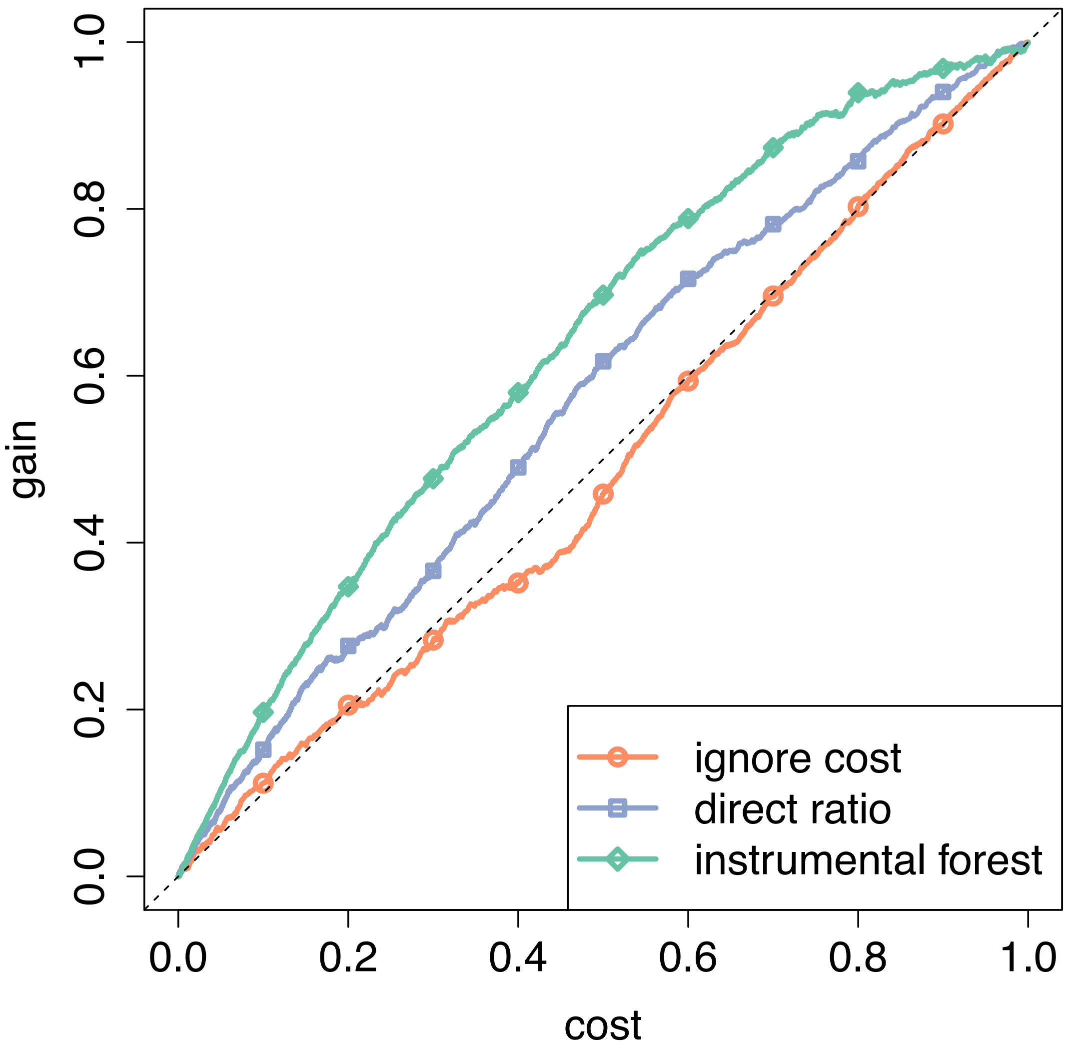

Sun, H., Du, S., Wager, S. (2021). Treatment allocation under uncertain costs. https://arxiv.org/abs/2103.11066

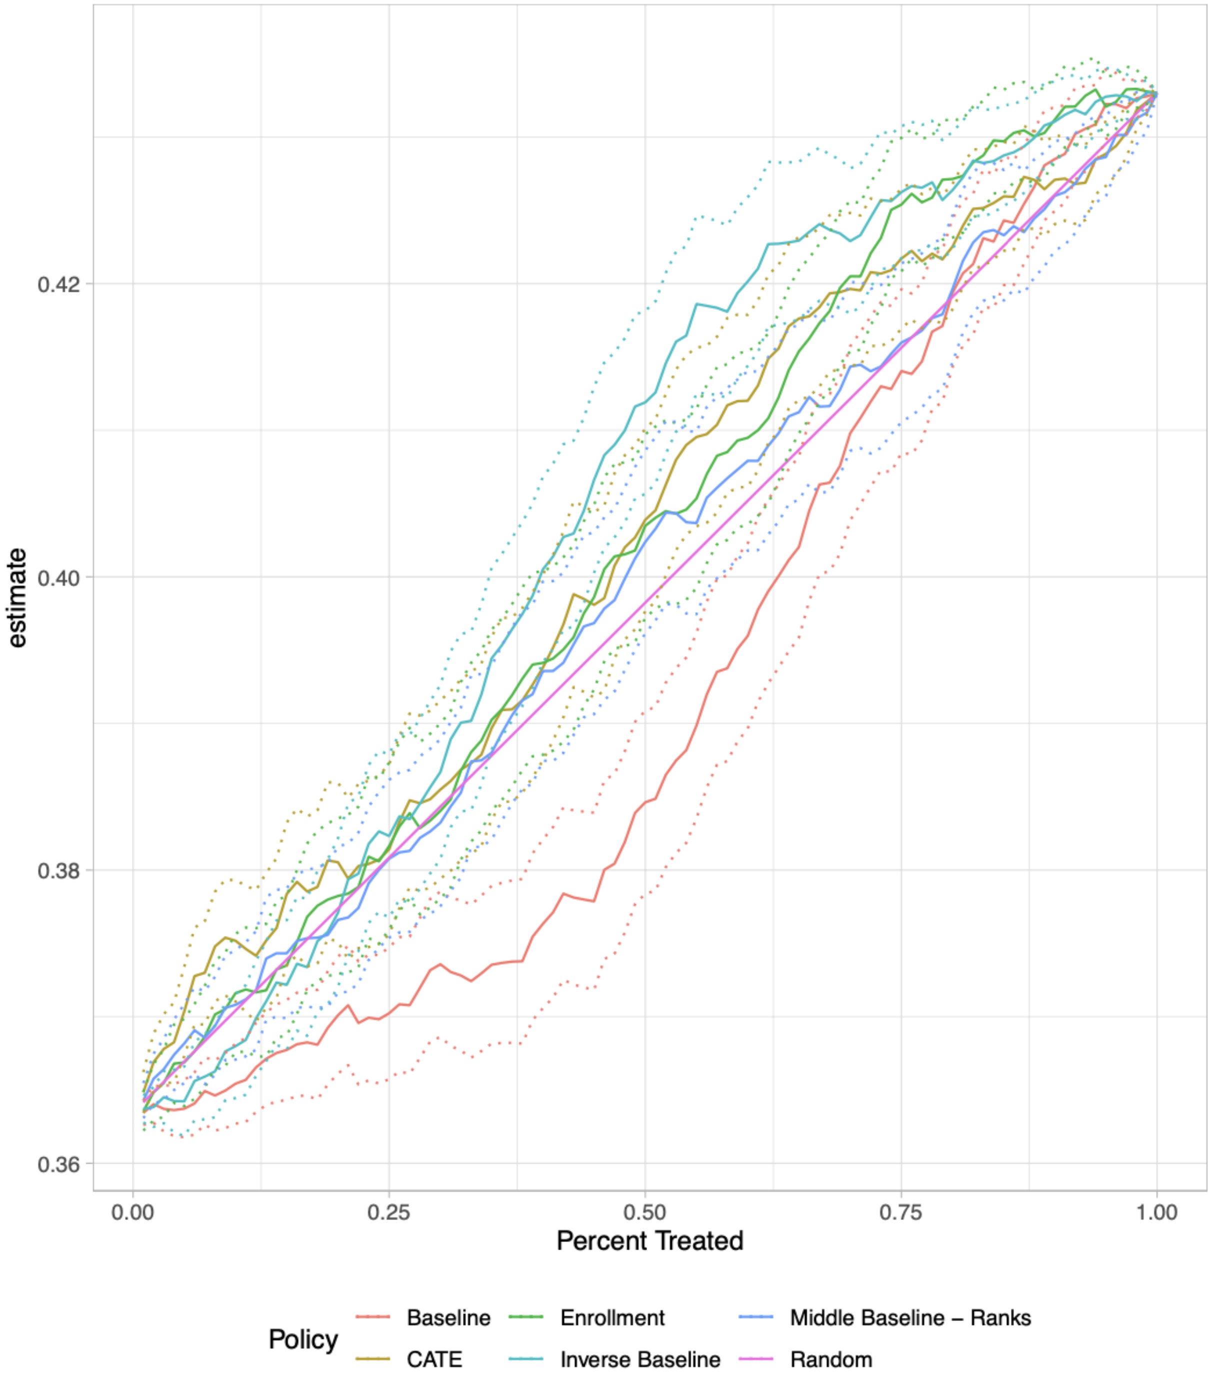

Do we have budget constraints in this setting?

Given that we don't find groups with negative CATEs, can we just treat everyone?

Some vague intuitions about HTEs

- Say our outcome Y is a binary choice

1. If treatment increased utility of choice, seems like effects should be largest on people with baseline expected utility near 0 (if people are random utility maximizers)
2. If treatment works for the same fraction of all people, but can't get people who already take choice 1 to do anything different, then effects are largest for groups with small $E[Y(0) \mid X = x]$
3. If the treatment has substantial costs or might have negative effects, what then?

HTE distributions

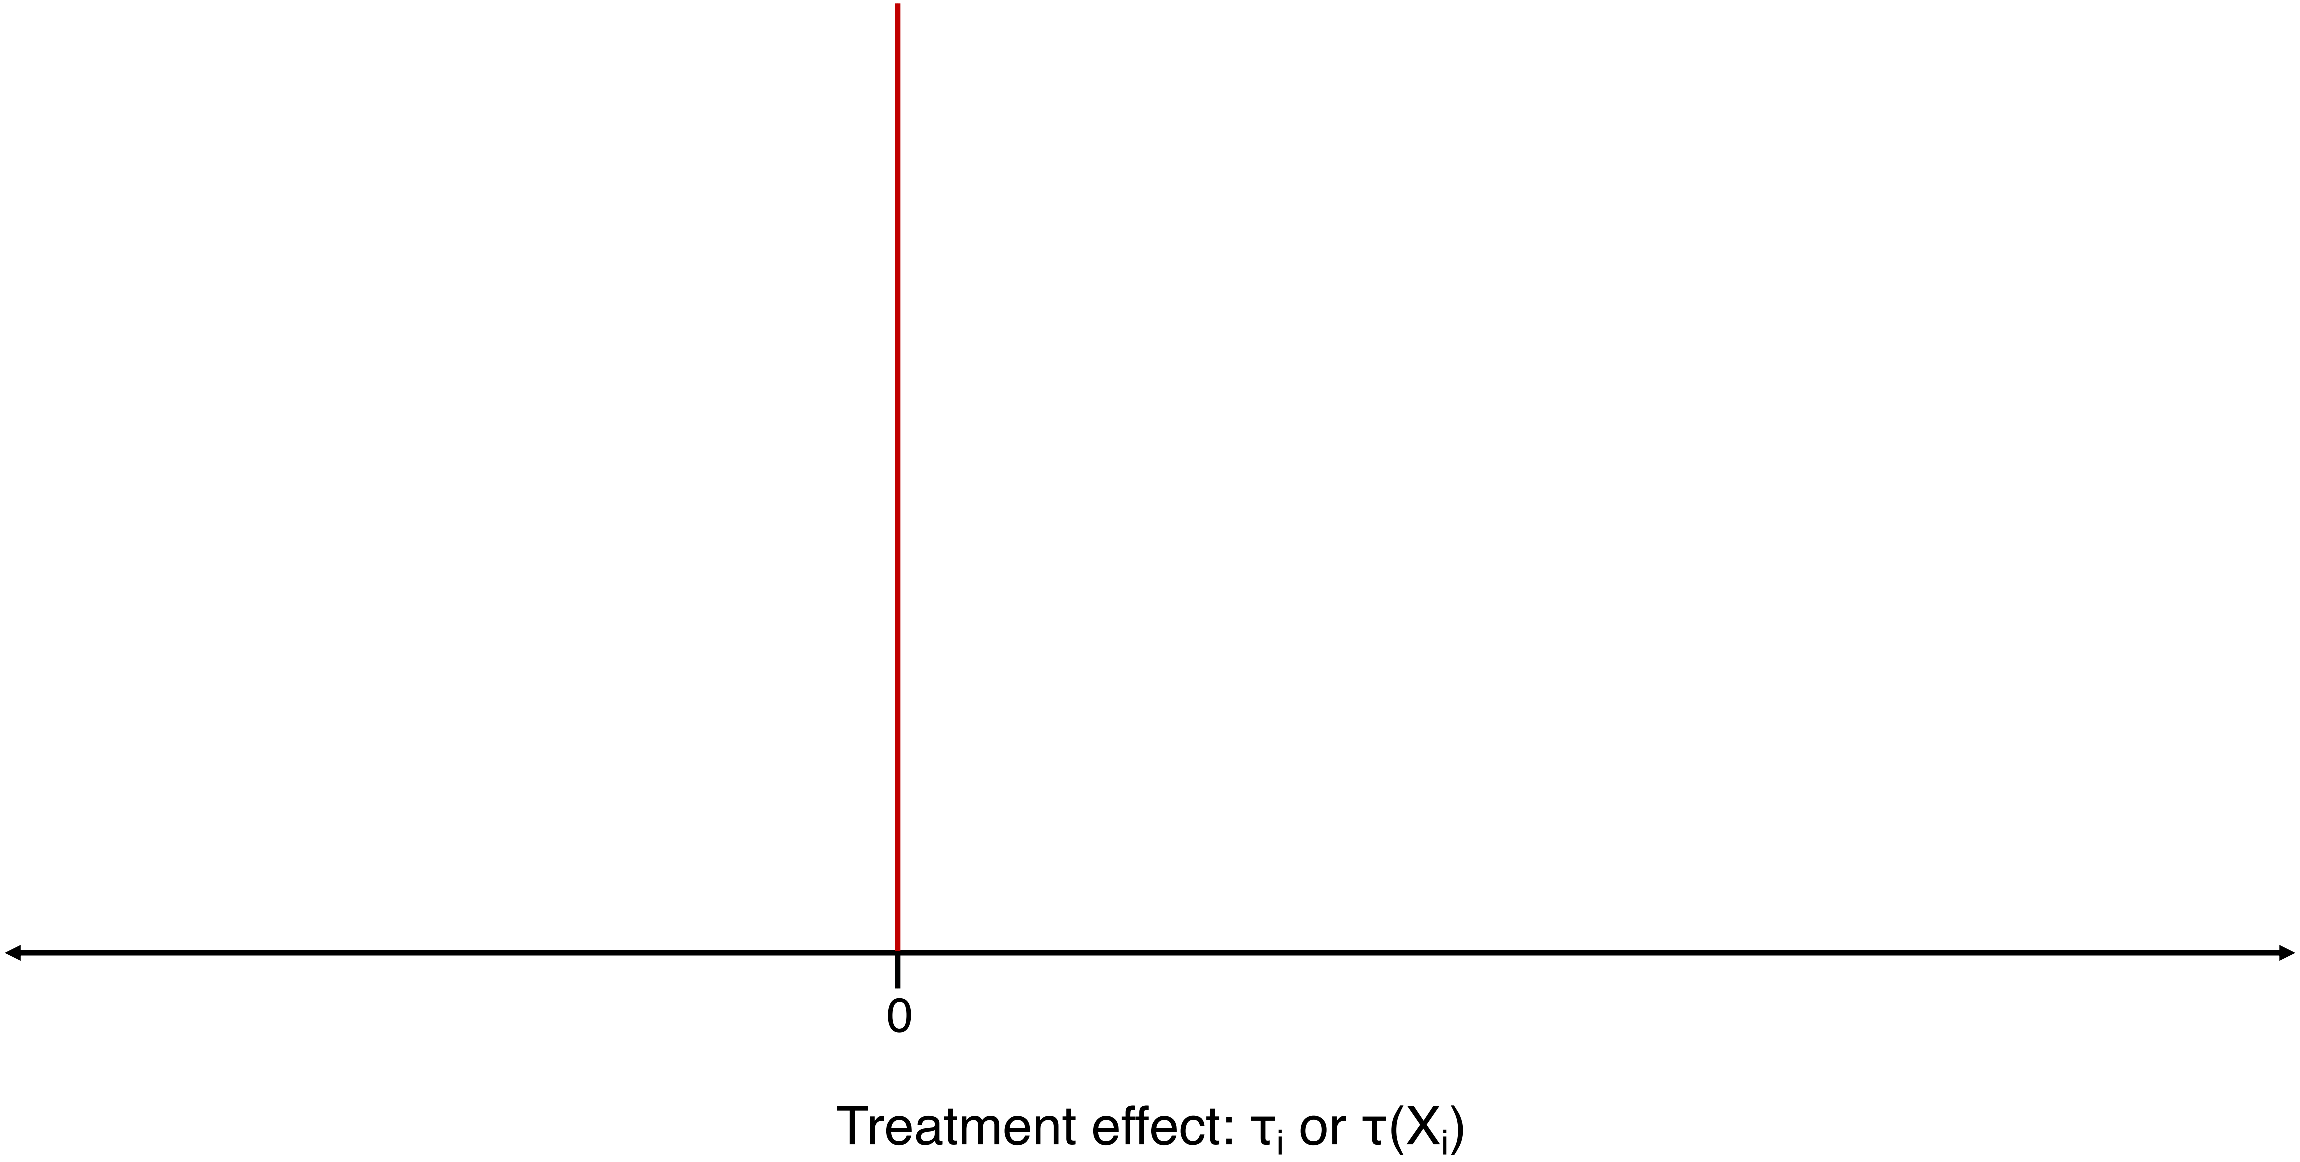

DO YOUR PART...
VOTE IN THIS ELECTION.

AFTER ALL, VOTING IS A MATTER OF PUBLIC RECORD.

YOUR NEIGHBORS NEED YOU. APPLY TO VOTE BY MAIL TODAY!

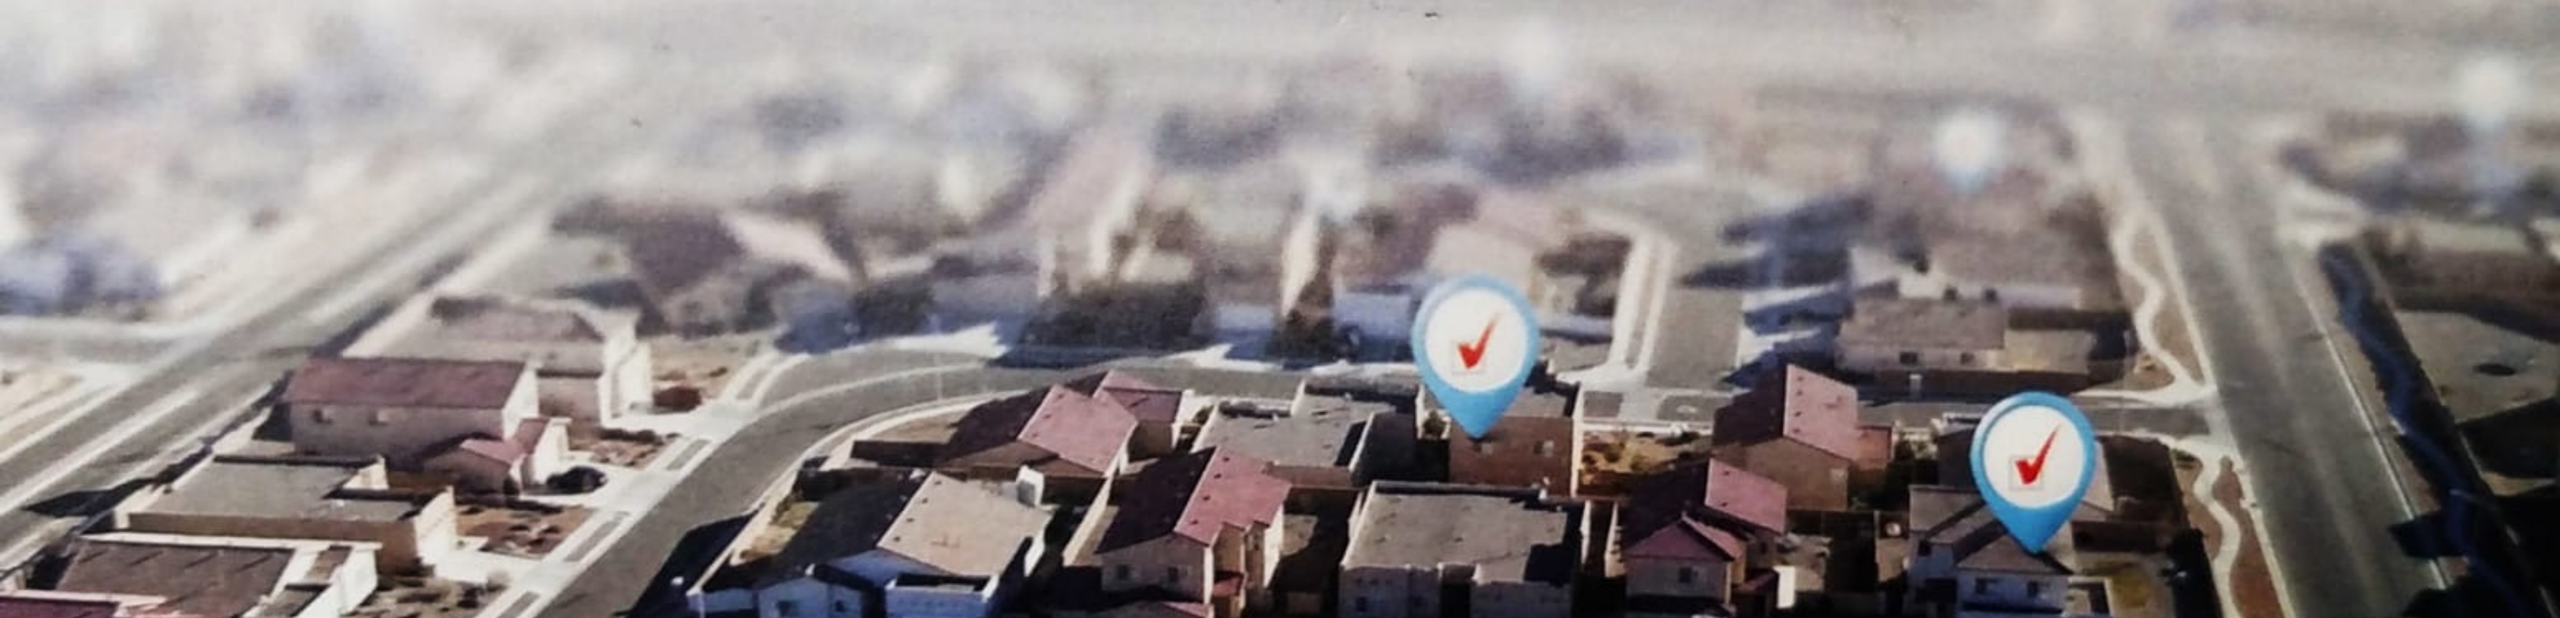

WHEN THE DEMOCRATS WIN THE ELECTION AND YOU DIDN'T DO YOUR PART TO STOP IT...

YOUR NEIGHBORS WILL KNOW.

WHAT ARE YOU WAITING FOR?

APPLY TO VOTE BY MAIL TODAY!

Paid for by the Republican Party Of New Mexico.
Not Authorized By Any Candidate Or Candidates Committee.

Non Profit Org
US Postage
PAID
Republican Party
Of New Mexico

Printed by Majority Strategies

*************AUTO**SCH 5-DIGIT 87112
13946 4/85/1

[redacted]

Targeting by predicted churn rates

- Fit a model predicting churn, getting a churn risk score per customer
- What if we targeted the top churners?

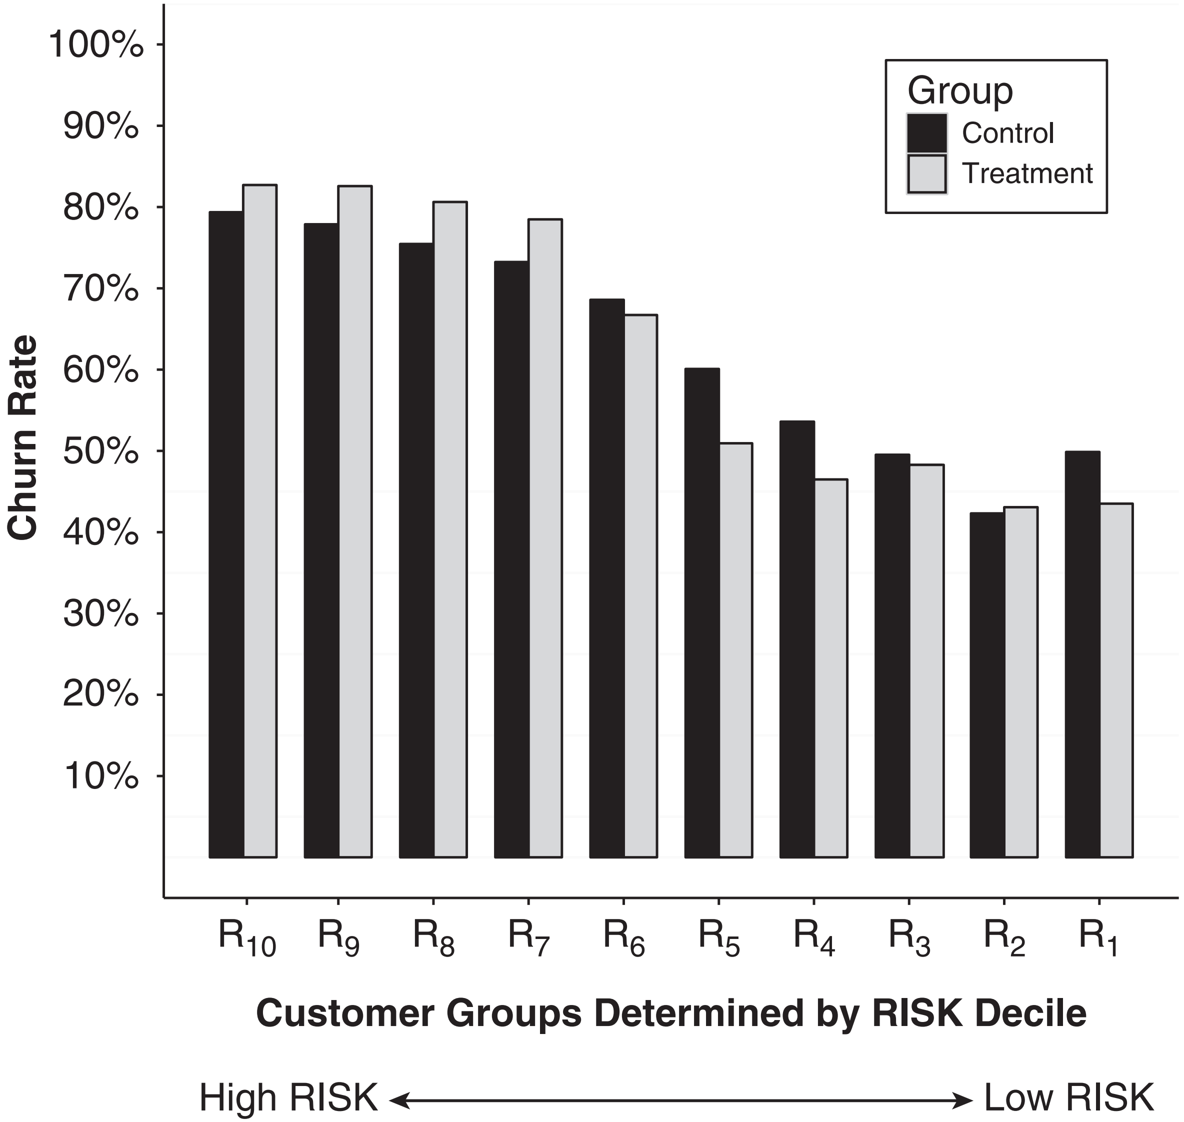

Ascarza (2018) "Retention Futility". Study 2: Membership organization in North America

Targeting by predicted churn rates

- Fit a model predicting churn, getting a churn risk score per customer
- What if we targeted the top churners?
- Can work quite poorly, but maybe this is not typical

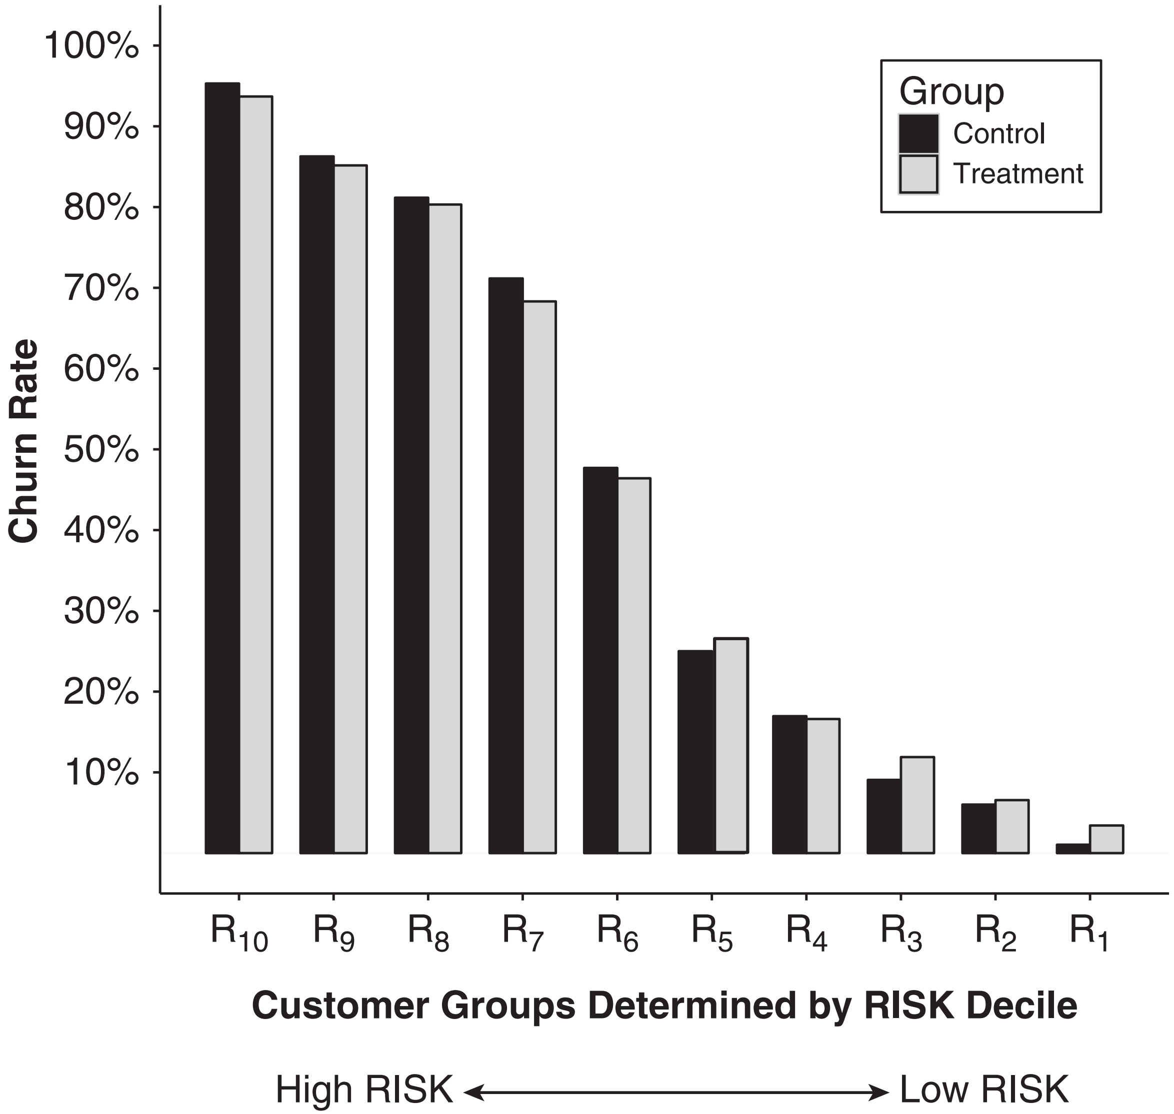

Ascarza (2018) “Retention Futility”. Study 1: Wireless service provider in the Middle East

Intuition: Large effects for "fence sitters"

Vaccine **norms info: narrow and *[broad]* conditions**

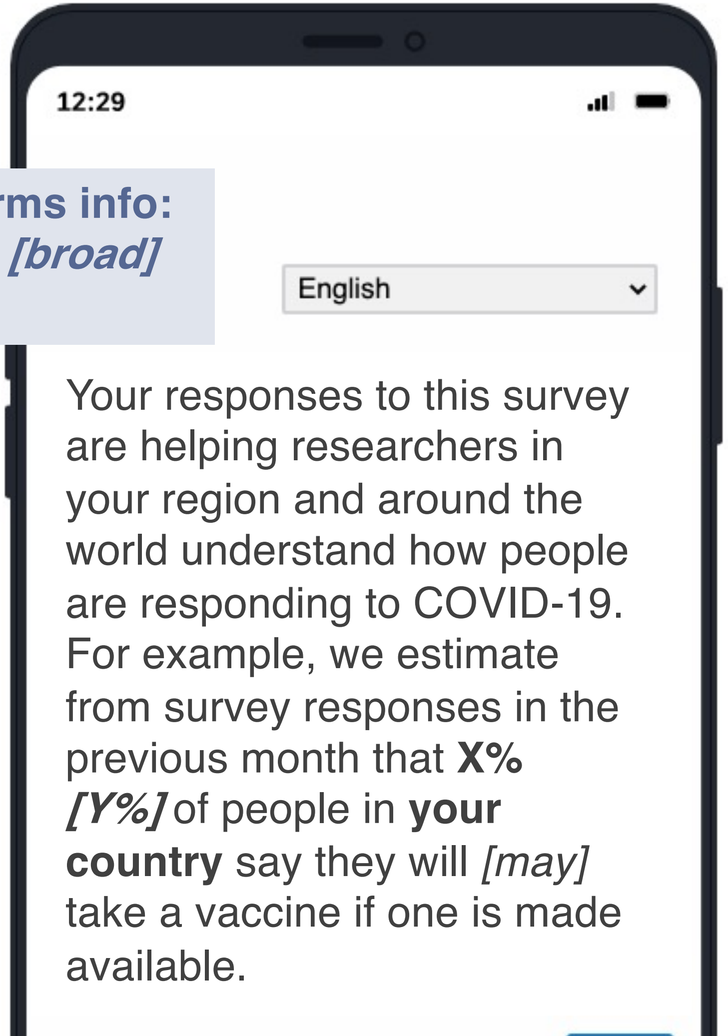

Data from a pre-registered, randomized experiment (N=484,239) embedded in a survey in 23 countries

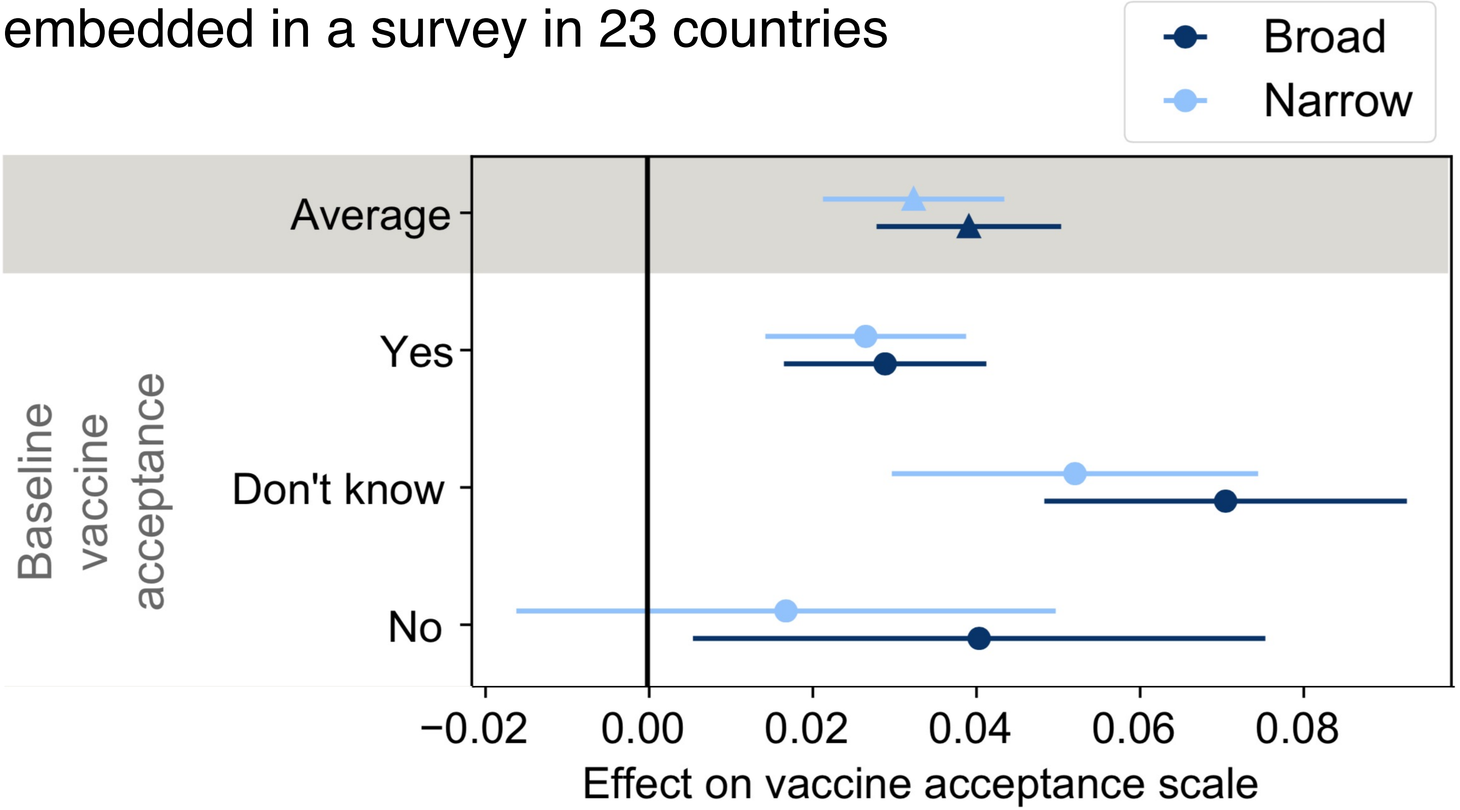

Moehring, A., Collis, A., Garimella, K., Rahimian, M., Aral, S., & Eckles, D. (2021). Surfacing norms to increase vaccine acceptance. https://doi.org/10.31234/osf.io/srv6t

The Boston Globe

Thank you.

Dear Jessie,

Thank you for being a loyal Globe.com subscriber. As a Boston Globe reader, your support allows us to continue producing award-winning journalism.

To say thank you, we are reducing your subscription rate to $4.99/week for the next 8 weeks. This discounted rate will be automatically applied to your next two invoices.

Go to Globe.com →

Using (often informal) prior about HTEs for experimental design

- Treating people with low (or high, or middle) probability of outcome with higher probability, but allow everyone to be either treated or not with some probability
  - Can be interpreted as reflecting our prior belief about the imperfect relationship between predicted Y(0) and treatment effects
  - If we're right, lowers regret compared with uniform experiment — and gives us more precision to learn good targeting policies