A PERSONALIZED BDM MECHANISM FOR EFFICIENT MARKET INTERVENTION EXPERIMENTS

Imanol Arrieta Ibarra
joint work with Johan Ugander

Stanford University

OBJECTIVES

- Introduce a subsidized product or service to a new market.
- Estimate its causal benefits on a desired population.
- Estimate the demand for the product.
- Be cost-efficient.
- Personalization allows more efficient experiments to be run without sacrificing causal interpretation.

MOTIVATION

- Providing Water Filters to a low-income population in Ghana. (Berry, Fischer and Gutieras (2015))
- Introducing new agricultural product. (Kremer, Duflo and Robinson (2011))
- Introducing new health products in developing countries.

DEMAND ESTIMATION

- Dynamic Pricing.
- Take-it-or-leave-it (TIOLI).
- Willingness to pay: maximum price at which someone is willing to pay in order to acquire a product.
  - Vickrey auctions among the population.
  - Becker-DeGroot-Marschak mechanism (BDM).

ESTIMATING CAUSAL EFFECTS

- Randomized Control Trials (A/B) testing.
- Observational Studies (Offline Policy Evaluation).
- Becker-DeGroot-Marschak mechanism (BDM).

BDM MECHANISM (SECOND PRICE AUCTION AGAINST RANDOM BIDDER)

1. Offer a product (a water filter) to user U.
2. Ask U for her willingness to pay: The maximum amount she would pay to get the product.
3. Draw a price at random P from an interval (0,T), where T is some fixed number *.
4. If the random price P is below user U's reported willingness to pay, she gets the product and pays P. Otherwise she does not get the product.

* This is what we'll personalize.

CAUSALITY (AS NEYMAN-RUBIN CAUSAL MODEL)

- Measure an outcome variable Y that takes values Y(1) under treatment and Y(0) under control.
- We are interested in the difference Y(1)-Y(0)
- In reality we get to only observe one of the two potential outcomes.

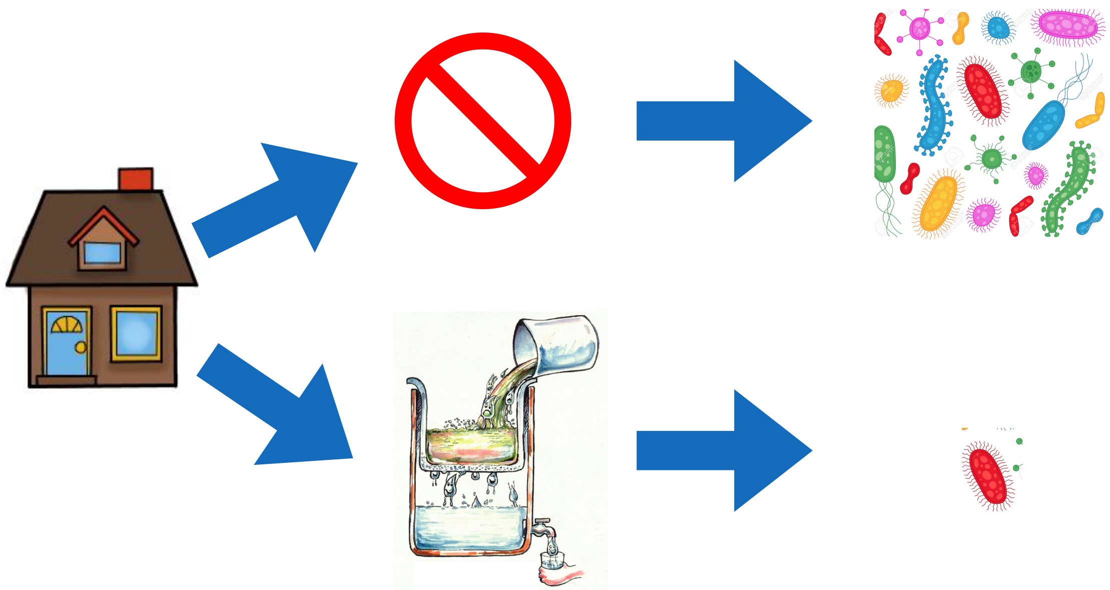

AVERAGE TREATMENT EFFECT

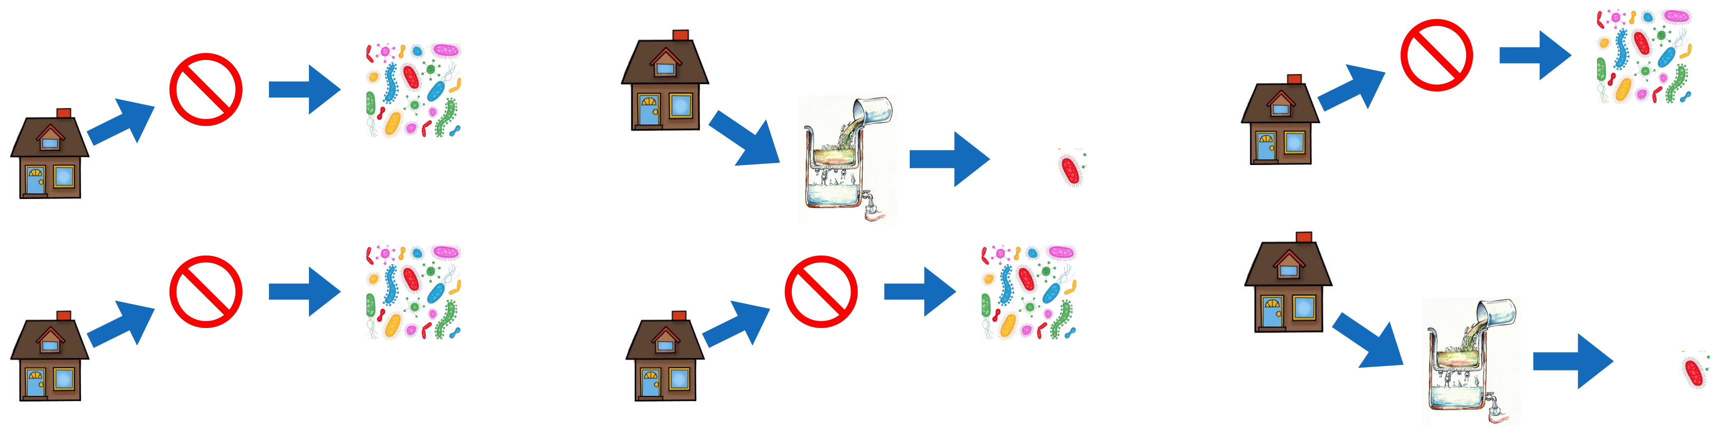

- Estimate the Average Treatment Effect for a given population: E[Y(1)-Y(0)].
- Can use difference in means estimator:

$$\hat{\tau} = \frac{1}{|G_T|} \sum_{i \in G_T} Y_i - \frac{1}{|G_C|} \sum_{i \in G_C} Y_i$$

DIFFERENT PROBABILITIES OF ASSIGNMENT

- If the probability of treatment assignment is different for each unit we can de-bias the estimates by using the Hajek estimator:

$$\hat{\tau_{Hajek}} = \frac{\sum \frac{1}{p_i} Y_{obs} W_i}{\sum \frac{1}{p_i} W_i} - \frac{\sum \frac{1}{1 - p_i} Y_{obs} (1 - W_i)}{\sum \frac{1}{1 - p_i} (1 - W_i)}$$

- Where Wi is a binary variable representing treatment assignment.

STRATIFICATION

- At each level of willingness to pay, the assignment to treatment and control is random:

(1) Choose M points from the WTP domain.
(2) For each point $W\bar{T}P_j$ select all users whose willingness to pay is close to $W\bar{T}P_j$:

$$|W\bar{T}P - WTP_i| < \epsilon$$

.

(3) For each j, compute the Horvitz-Thompson estimator for those users: $\hat{\tau}^j_{HT}$.
(4) Compute $\hat{\tau} = \sum_{j=1}^{N_j} \frac{N_j}{N} \hat{\tau}^j_{HT}$ as the post-stratified weighted average of the HT estimators.

DEMAND ESTIMATION

- Dynamic Pricing.
- Take-it-or-leave-it (TIOLI).
- Willingness to pay: maximum price at which someone is willing to pay in order to acquire a product.
  - Vickrey auctions among the population.
  - Becker-DeGroot-Marschak mechanism (BDM).

ESTIMATING CAUSAL EFFECTS

- Randomized Control Trials (A/B) testing.
- Observational Studies (Offline Policy Evaluation).
- Becker-DeGroot-Marschak mechanism (BDM).

DEMAND ESTIMATION

- Having elicited the users' willingness to pay, we can count the number of users which were willing to buy the product at each price point.

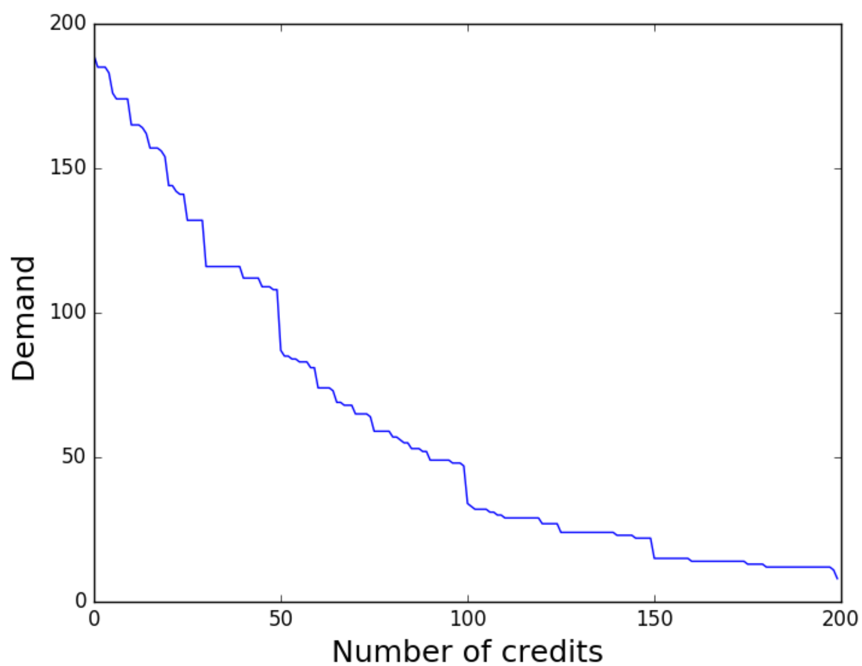

ESTIMATING CAUSAL EFFECTS

- As in Berry, Fischer and Gutieras (2015).
- Two sources of randomness:
  - Conditional on willingness to pay, treatment is random.
  - Conditional on willingness to pay and being treated, price is random.

PERSONALIZATION

- Reduce unnecessary costs for researchers by minimizing potential subsidies.
- Reduce variance in estimations by allowing better balance at each level of willingness to pay.
- Maintain incentive compatibility to elicit correct valuations.

PERSONALIZED BDM MECHANISM

1. Offer a product with cost C to a subject with Xi observable characteristics.
2. Draw a price ϕ from $F_{\Phi|X_i}$ without showing it to the subject
3. Ask the subject to report her willingness to pay w_i which comes from a population $W_i|X_i$
4. If $\phi < w$ the subject gets the intervention and pays ϕ (the lower price). Otherwise there is no exchange.

THREE ALTERNATIVES FOR $\Phi \mid X_i$

1. Make $\Phi \mid X_i$ = E[W|Xi].

$$F_{\Phi|X_i}(w) = \mathbb{I}\{w > E[W|X_i]\}$$

2. Run an RCT (Randomized Control Trial):

.

$$F_{\Phi|X_i}(w) = \frac{1}{2} + \frac{1}{2}\mathbb{I}\{w > C\}$$

3. Draw prices from a personalized uniform distribution:

$$F_{\Phi|X_i}(w) = \frac{\epsilon}{2} + (1-\epsilon)\frac{w\mathbb{I}\{a \le w \le b\}}{b-a} + \frac{\epsilon}{2}\mathbb{I}\{w \ge C\}$$

where a and b are chosen based on W|Xi .

DRAW PRICES FROM A PERSONALIZED UNIFORM DISTRIBUTION

$$F_{\Phi|X_i}(w) = \frac{\epsilon}{2} + (1-\epsilon)\frac{w\mathbb{I}\{a \le w \le b\}}{b-a} + \frac{\epsilon}{2}\mathbb{I}\{w \ge C\}$$

- BDM is a special case where :

$$\forall i \quad \epsilon = 0, \quad a = 0, \quad b = C$$

- We'll refer to as PBDM the case where

$$a(X_i) = F^{-1}_{W|X_i}\left(\frac{\delta}{2}\right), \quad b(X_i) = F^{-1}_{W|X_i}\left(1 - \frac{\delta}{2}\right)$$

PBDM

$$F_{\Phi|X_i}(w) = \frac{\epsilon}{2} + (1-\epsilon)\frac{w\mathbb{I}\{a \le w \le b\}}{b-a} + \frac{\epsilon}{2}\mathbb{I}\{w \ge C\}$$

$$a(X_i) = F^{-1}_{W|X_i}\left(\frac{\delta}{2}\right), \quad b(X_i) = F^{-1}_{W|X_i}\left(1 - \frac{\delta}{2}\right)$$

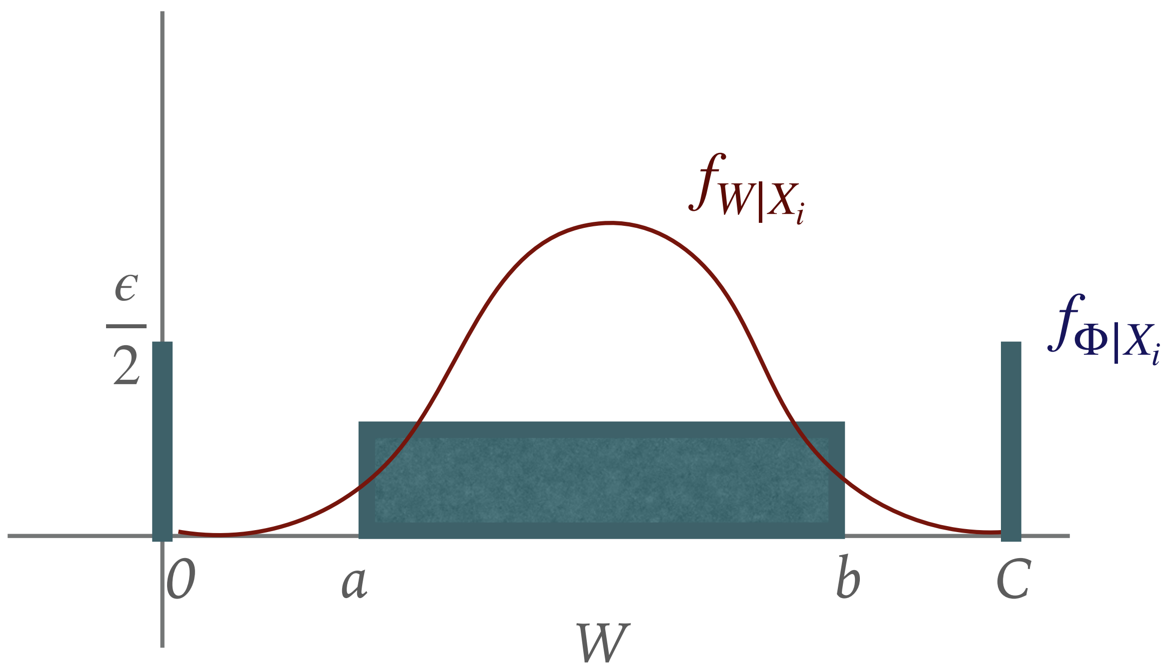

PBDM

$$F_{\Phi|X_i}(w) = \frac{\epsilon}{2} + (1-\epsilon)\frac{w\mathbb{I}\{a \le w \le b\}}{b-a} + \frac{\epsilon}{2}\mathbb{I}\{w \ge C\}$$

$$a(X_i) = F^{-1}_{W|X_i}\left(\frac{\delta}{2}\right), \quad b(X_i) = F^{-1}_{W|X_i}\left(1 - \frac{\delta}{2}\right)$$

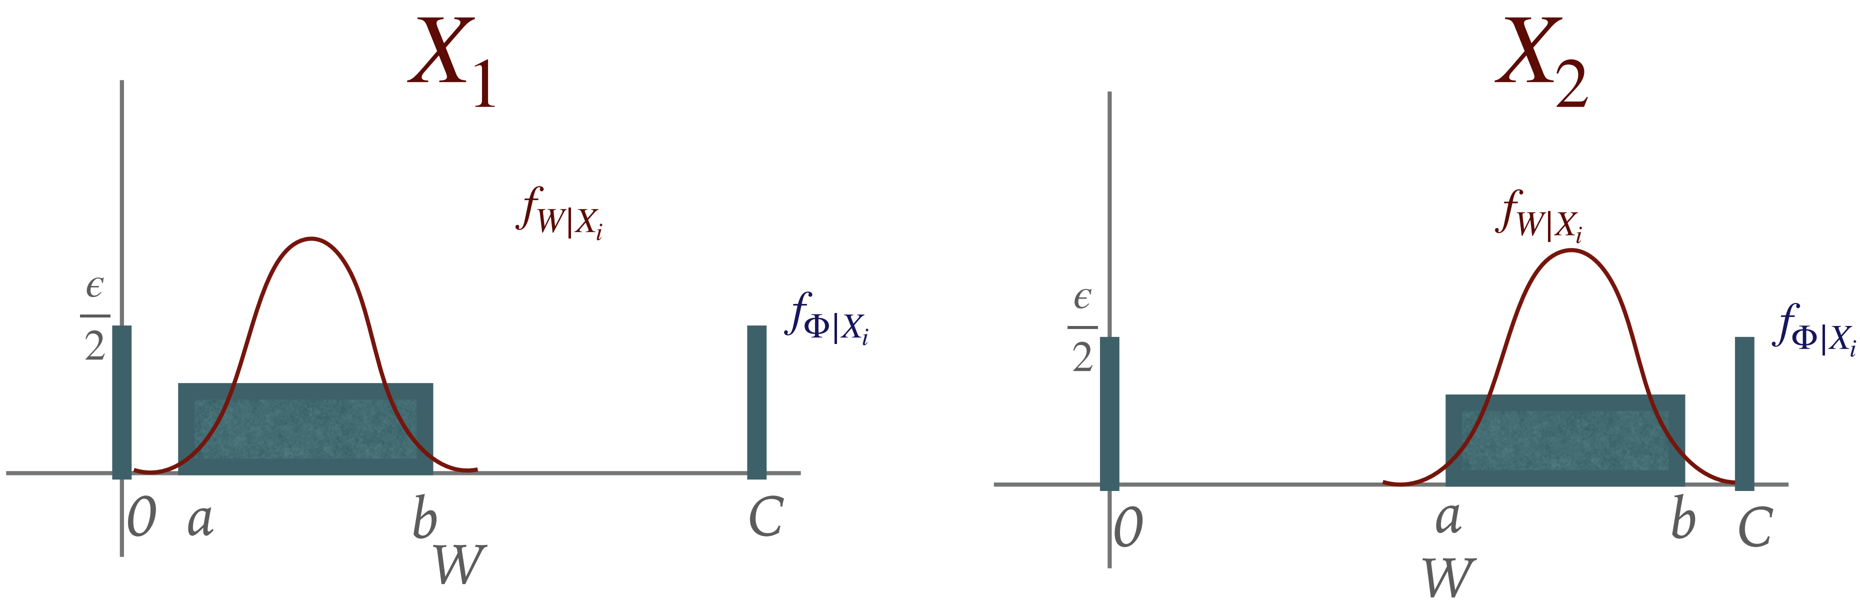

ESTIMATOR VARIANCE

- Under Fisher's null (Y(1)=Y(0)= α) we get that the H-T variance is:

$$Var(\hat{\tau_{HT}}) = \frac{1}{N^2}\left(\sum_{i=1}^{N} \frac{1 - F_{\Phi|X_i}(W_i)}{F_{\Phi|X_i}(W_i)} Y_i(1)^2 + \frac{F_{\Phi|X_i}(W_i)}{1 - F_{\Phi|X_i}(W_i)} Y_i(0)^2\right)$$

$$= \frac{1}{N^2}\left(\sum_{i=1}^{N} \alpha_i \left(\frac{1 - F_{\Phi|X_i}(W_i)}{F_{\Phi|X_i}(W_i)} + \frac{F_{\Phi|X_i}(W_i)}{1 - F_{\Phi|X_i}(W_i)}\right)\right)$$

- Minimized when:

$$\forall i \quad F_{\Phi|X_i}(W_i) = \frac{1}{2}$$

ESTIMATOR VARIANCE

1. Make $\Phi | X_i$ = E[W|Xi].

$$Var(\hat{\tau_{HT}}) = \infty$$

2. Run an RCT (Randomized Control Trial) minimizes the variance:

$$Var(\hat{\tau_{HT}}) = 2\bar{\alpha} \; .$$

3. Draw prices from a personalized uniform distribution:

$$2\bar{\alpha} < Var(\hat{\tau_{HT}}) < \infty$$

EXPECTED VARIANCE

- Using a Taylor expansion around the mean and taking a first degree approximation.

$$E[Var(\hat{\tau_{HT}})\,|\,X_1,\ldots,X_N] \approx \frac{1}{N^2}\sum_{i=1}^{N}\left(\frac{1-F_{\Phi|X_i}(E[W_i\,|\,X_i])}{F_{\Phi|X_i}(E[W_i\,|\,X_i])}+\frac{F_{\Phi|X_i}(E[W_i\,|\,X_i])}{1-F_{\Phi|X_i}(E[W_i\,|\,X_i])}\right)$$

- which gets minimized when:

$$F_{\Phi|X_i}(E[W_i\,|\,X_i]) = \frac{1}{2}$$

BUDGET REGRET

- We define budget regret as:

$$BR(\Phi, W) = (W - \Phi)\mathbf{I}\{\Phi < W\}$$

- and expected budget regret as:

$$br(F_{\Phi|X}) = E_{X,W}[E_{F_{\Phi|X}}[BR(\Phi, W)]]$$

- Every time we assign someone to treatment we incur some regret derived from having been able to treat that subject with a lower subsidy had we known their true willingness to pay.

BUDGET REGRET

1. Make $\Phi | X_i$ = E[W|Xi].

$$br(F_{\Phi|X}) = E\left[\frac{W}{2}\right] - \theta, \quad 0 < \theta < E\left[\frac{W}{2}\right]$$

2. Run an RCT (Randomized Control Trial) maximizes budget regret: .

$$br(F_{\Phi|X}) = E[W]$$

3. Draw prices from a personalized uniform distribution:

$$br(F_{\Phi|X}) = E\left[\frac{W - \hat{a}}{2}\right]$$

INCENTIVE COMPATIBILITY

1. Make $\Phi | X_i$= E[W|Xi].
   - Subject is indifferent after convergence.

2. Run an RCT (Randomized Control Trial) minimizes the variance:
   - Subject indifferent amongst valuations .

3. Draw prices from a personalized uniform distribution:
   - Incentive compatible with probability higher than $1 - \delta$

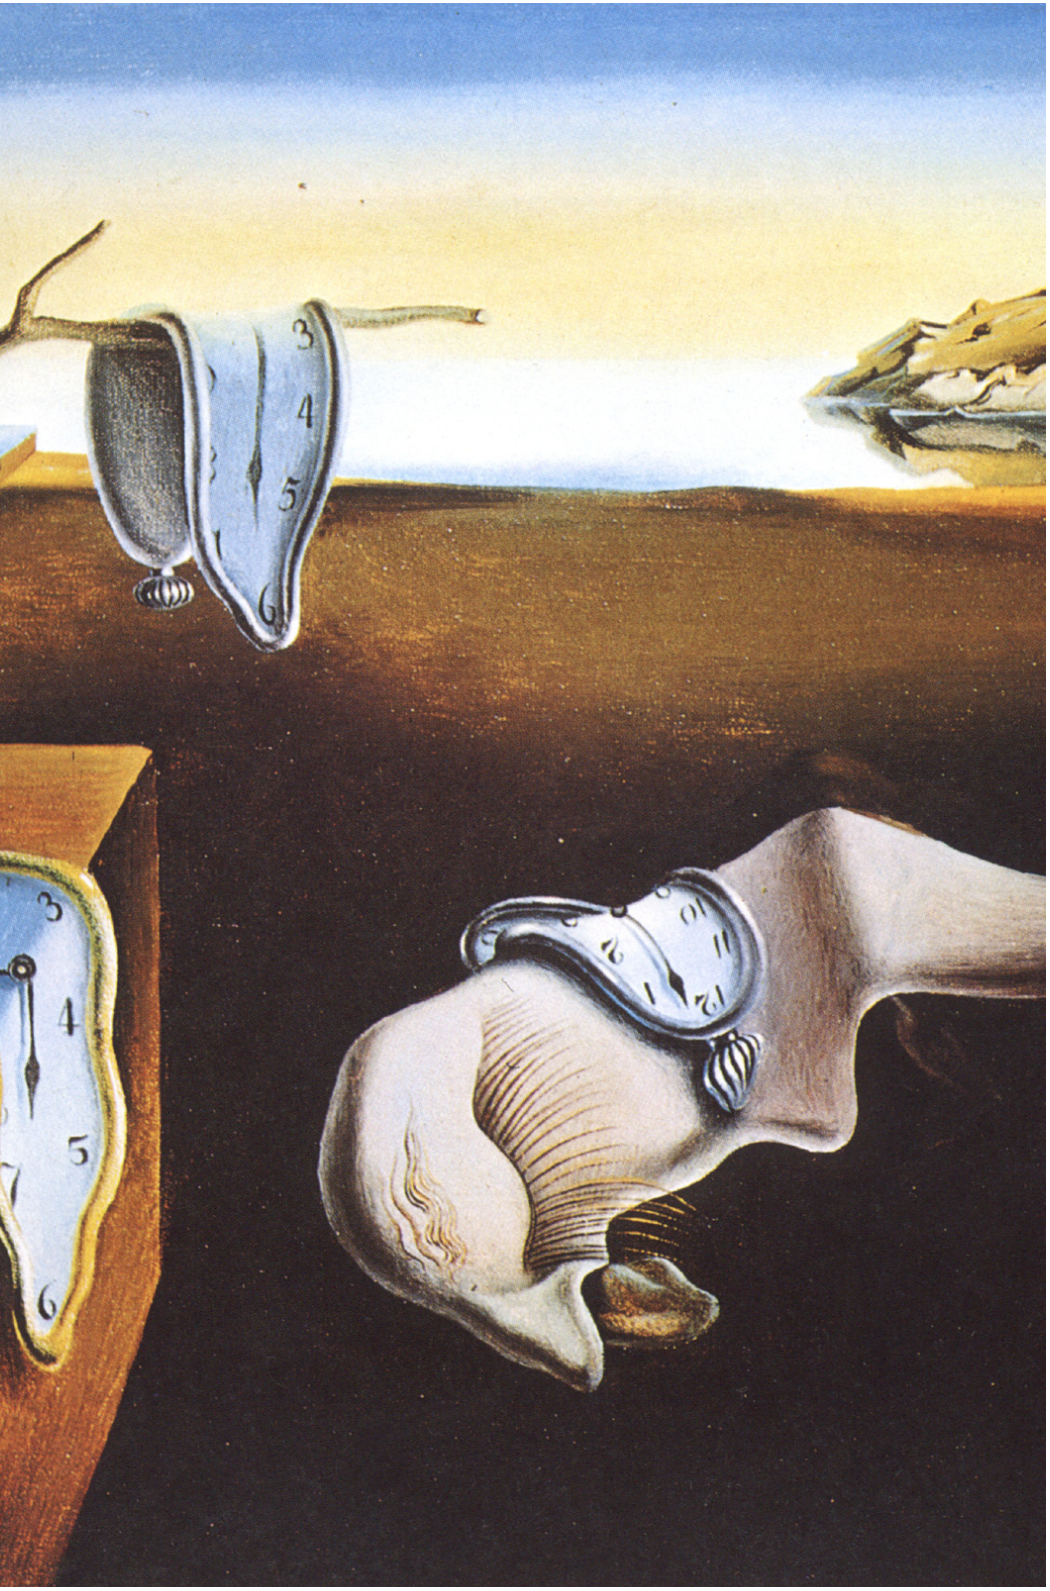

TIME CONSTRAINTS ON MECHANICAL TURK

- Understand MT workers performance under time constraints.
- Measure how performance changes conditional on how much workers value not being constrained.
- Evaluate performance on turkers who paid not to be timed conditional on what they were paid.
- Used STAN to estimate the distribution of $W|X$

CONTEXT (DEMOGRAPHICS)

Demographic Survey

Age

Household Size

Country of Residence

Select Country

Gender

Select Gender

Yearly Income in US dollars

Select Income

Marital Status

Select Status

Operating System you use

Select OS

CONTEXT (RISK AVERSION)

Risk Questionaire

For the next 7 questions indicate from 1 (Strongly Agree) to 7 (Strongly Disagree), how you feel about the statements.

I do not feel comfortable about taking chances.

Select Option

I prefer situations that have foreseeable outcomes.

Select Option

Before I make a decision, I like to be absolutely sure how things will turn out.

Select Option

I feel nervous when I have to make decisions in uncertain situations.

Select Option

I avoid situations that have uncertain outcomes.

Select Option

I feel comfortable improvising in new situations.

Select Option

The next two questions are hypothetical scenarios about your risk preferences.

Select the option you would prefer in the following hypothetical situations:

Select Option

Select the option you would prefer in the following hypothetical situations:

Select Option

PBDM

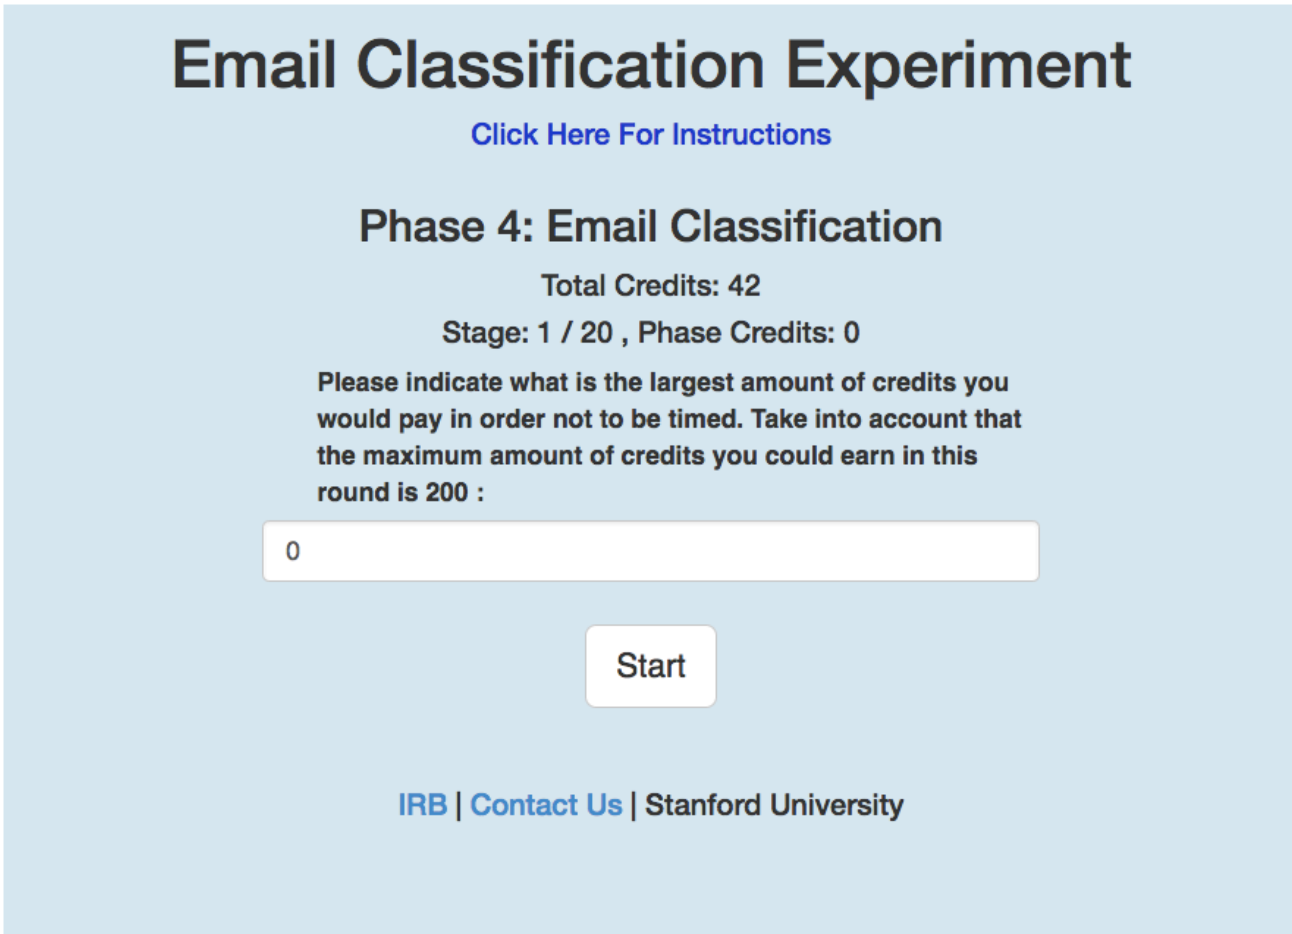
Email Classification Experiment
Click Here For Instructions
Phase 4: Email Classification
Total Credits: 42
Stage: 1 / 20 , Phase Credits: 0
Please indicate what is the largest amount of credits you would pay in order not to be timed. Take into account that the maximum amount of credits you could earn in this round is 200 :
0
Start
IRB | Contact Us | Stanford University

PBDM

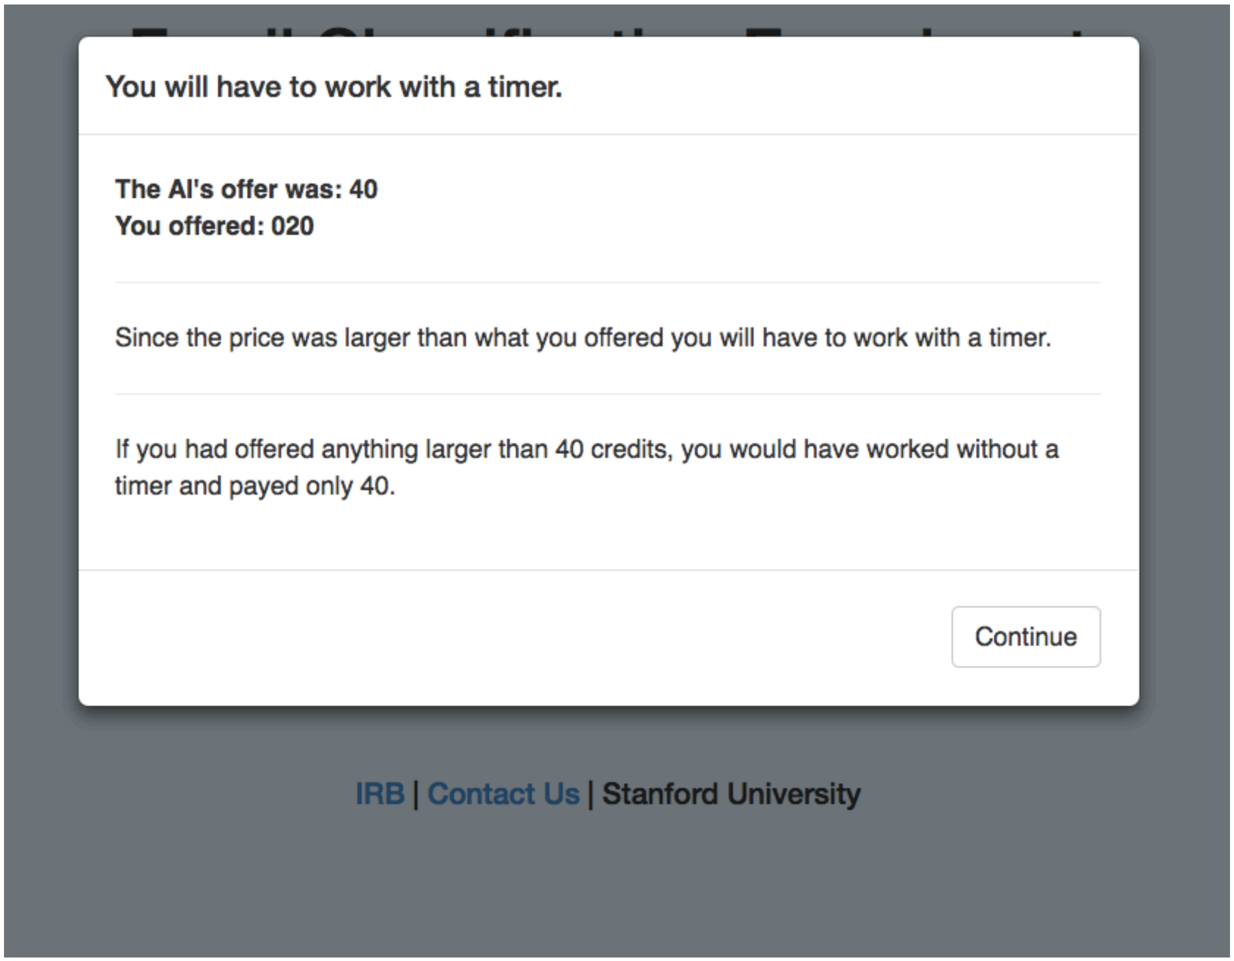
You will have to work with a timer.
The AI's offer was: 40
You offered: 020
Since the price was larger than what you offered you will have to work with a timer.
If you had offered anything larger than 40 credits, you would have worked without a timer and payed only 40.
Continue
IRB | Contact Us | Stanford University

PBDM

Email Classification Experiment

Click Here For Instructions

Phase 4: Email Classification

Total Credits: 42

Stage: 1 / 20 , Phase Credits: 0

Time left:

1

```
Subject: fw : first delivery - rodessa operating co .
daren ,
please read this memo and the one dated 2 / 21 / 01 from vance . he told me all new p
roduction should be entered at ifhsc . the deal in question is 634075 . do you want i
t changed to gas daily ? please advise . thanks .
bob
- - - - - - - - - - - - - - - - - - - - - - forwarded by robert cotten / hou / ect on
03 / 21 / 2001 11 : 02 am - - - - - - - - - - - - - - - - - - - - - - - - - - -
from : vance l taylor / enron @ enronxgate on 03 / 19 / 2001 02 : 24 pm
to : tom acton / corp / enron @ enron , robert cotten / hou / ect @ ect
cc : julie meyers / hou / ect @ ect , lisa hesse / hou / ect @ ect , donald p reinhar
dt / enron @ enronxgate , susan smith / enron @ enronxgate , melissa graves / enron @
enronxgate , cynthia hakemack / hou / ect @ ect
subject : fw : first delivery - rodessa operating co .
tom / bob ,
the following production is now on - line and a ticket should be created and entered
into sitara based on the following :
counterparty meter volumes price period global no .
hesco gathering co . , llc 9876 85 mmbtu / d 100 % gd less $ 0 . 17 3 / 13 - 3 / 31 9
6057368
fyi , susan will create and submit a committed reserves firm ticket for the remaining
term of the deal beginning with the month of april . additionally , this is a produce
r svcs . deal and should be tracked in the im wellhead portfolio . . . attached to th
```

DEMAND ESTIMATION

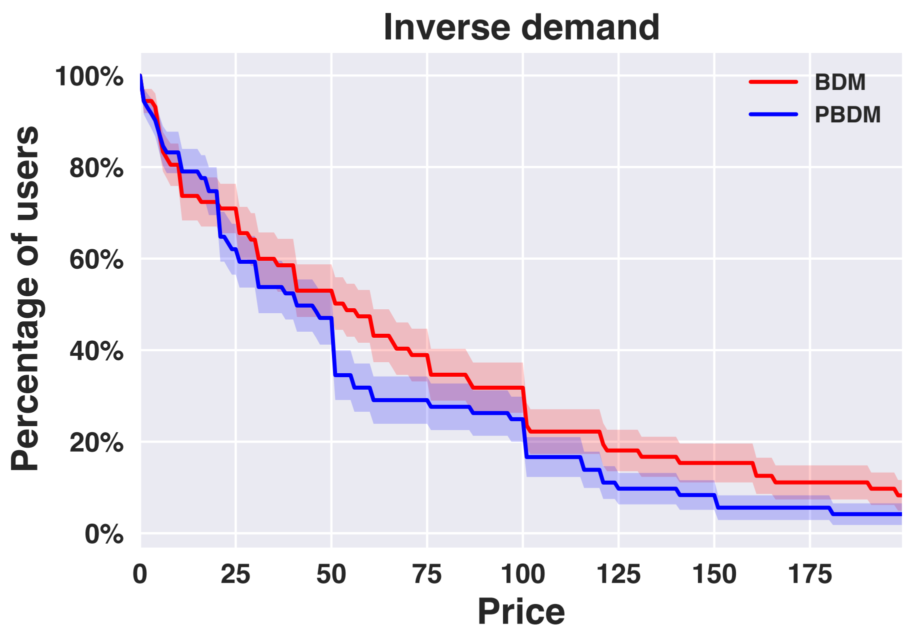

ESTIMATING CAUSAL EFFECTS

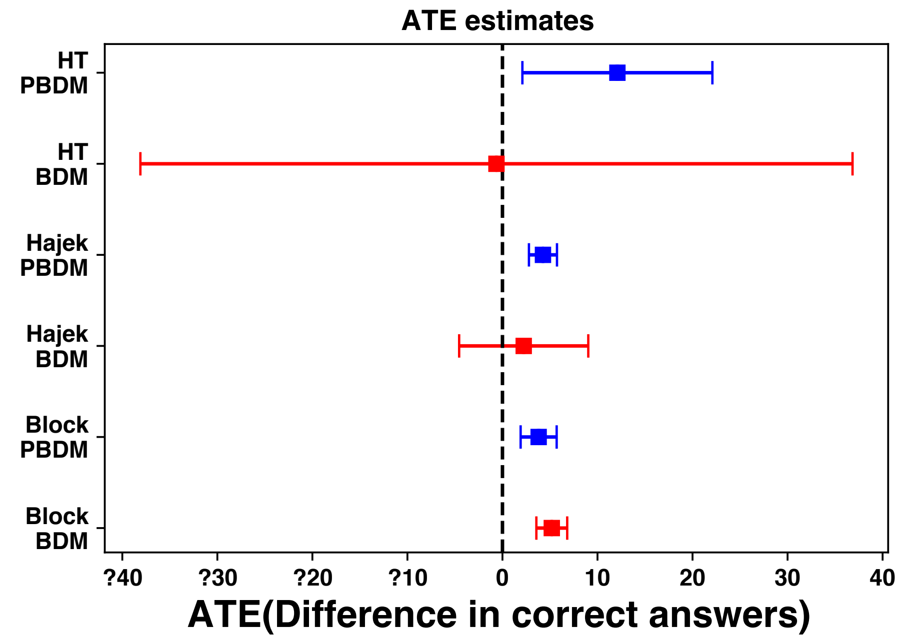

	BDM	PBDM
Percentage treated	0.31	0.52
Hajek ATE	2.23	4.26
Standard Error	3.72	0.96
Average Budget Regret	65	45

VARIANCE OF ESTIMATORS

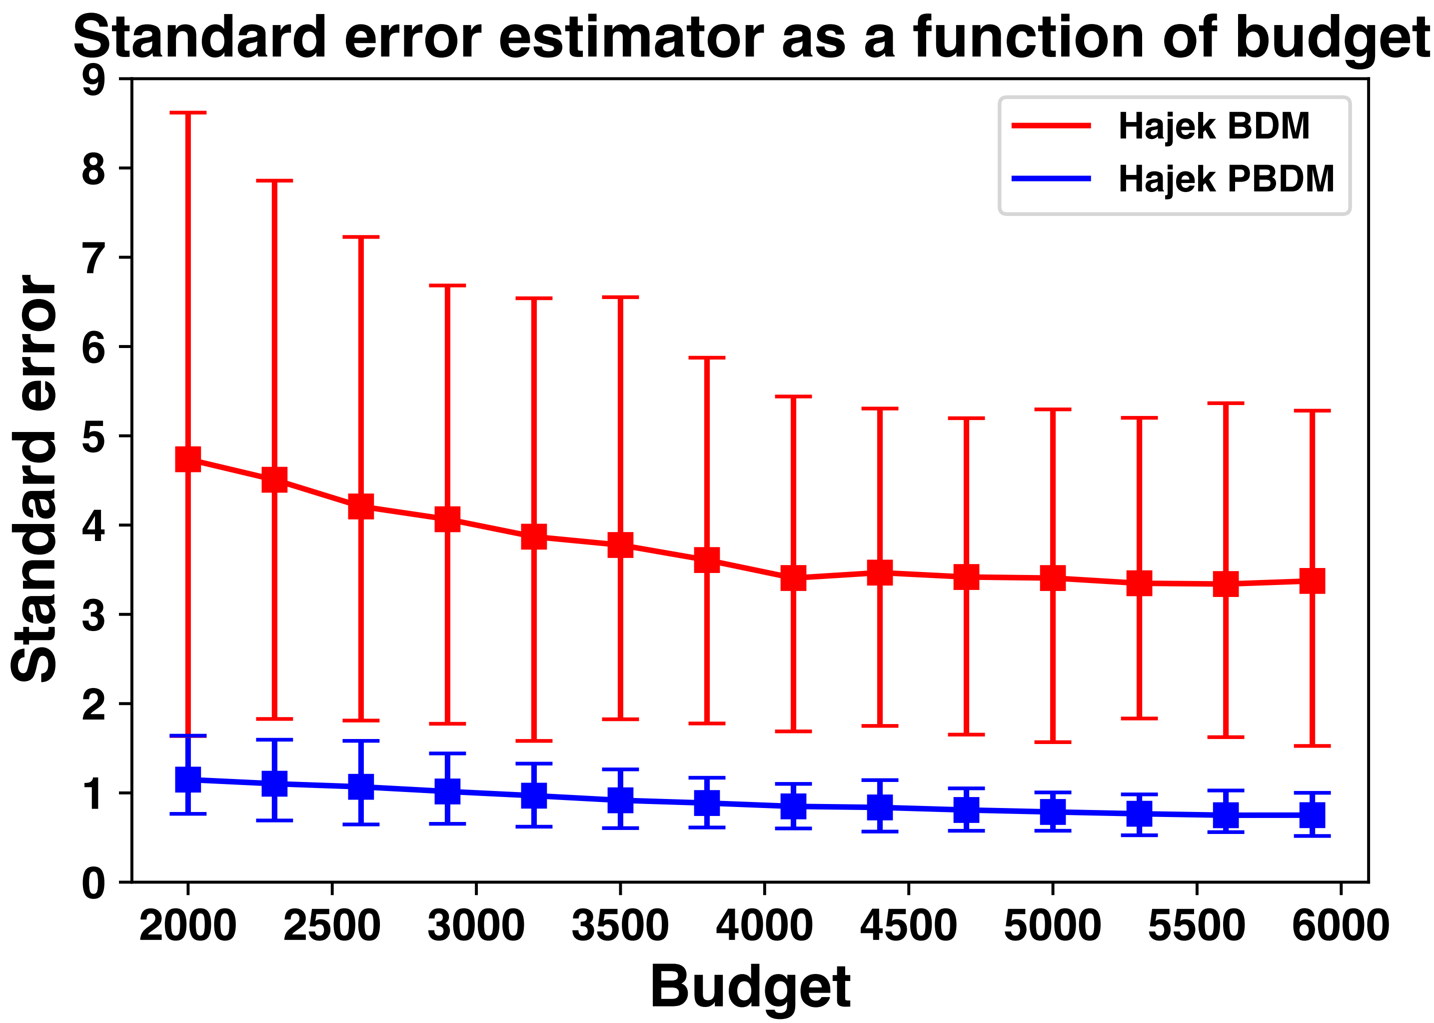

PBDM SMOOTHNESS

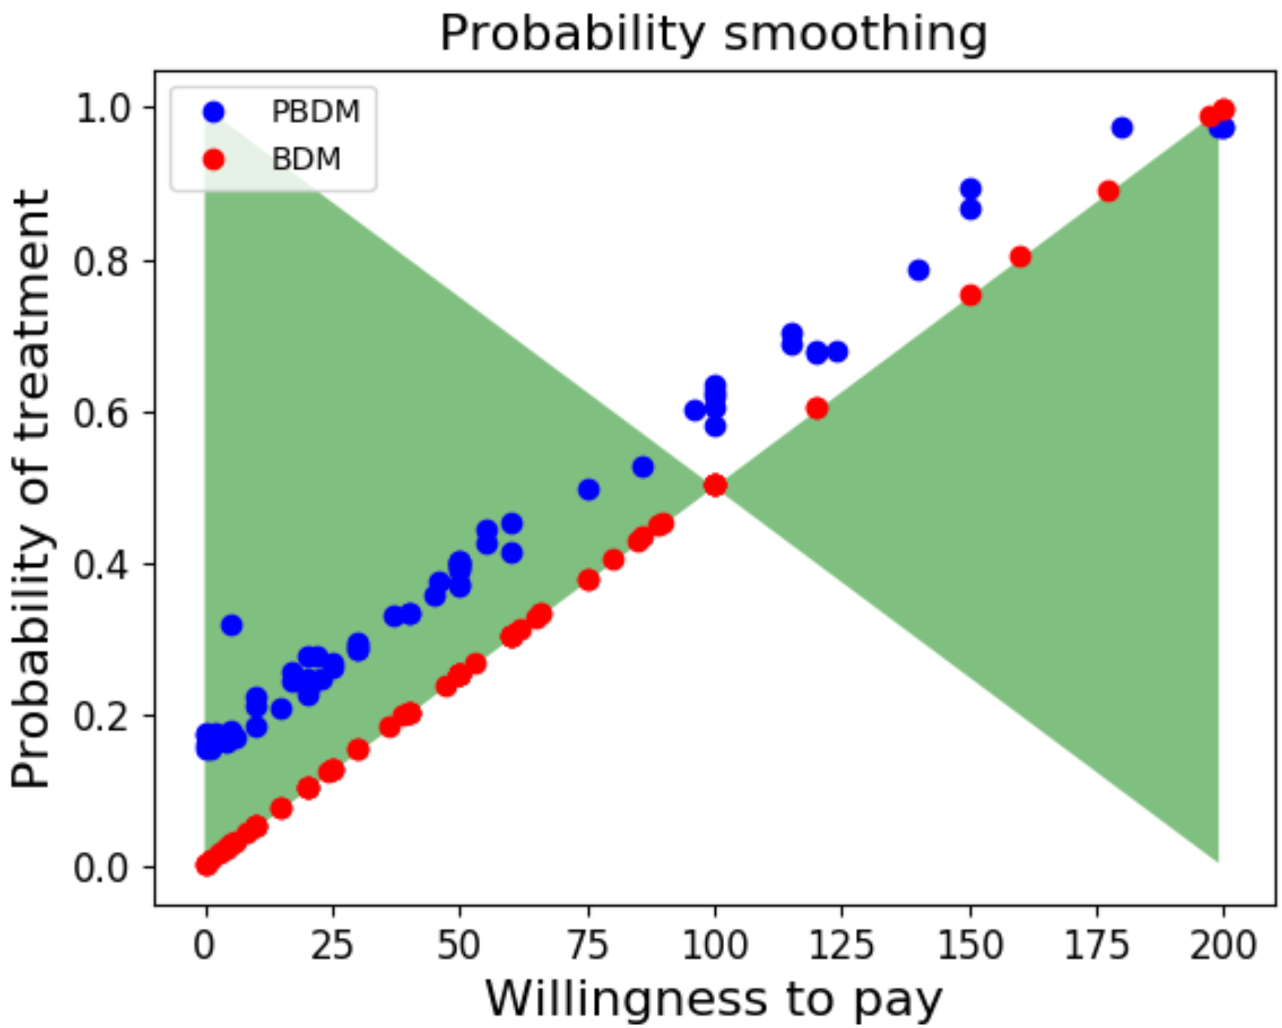

SUMMARY

- Presented a way to introduce personalization using machine learning to experiments without losing the causal interpretation.
- Showed that personalization can reduce the cost of unnecessary subsidies in this kind of experiments.
- Evaluated our methods on a Mechanical Turk experiment and found that even though for the small sample size we were not able to find big differences in estimation preciseness, the amount of subsidy given to users was cut to half for our algorithm.

FUTURE WORK

- Currently working on the estimation of heterogeneous treatment effects.
- Optimal strategy when balancing the treatment conditional on willingness to pay and the treatment conditional on price paid.
- Expanding this work to other types of mechanisms and notions of incentive compatibility.
- Looking for applications where we can predict willingness to pay from observed characteristics.

THANK YOU

WERE USERS UNDERSTANDING THE MECHANISM?

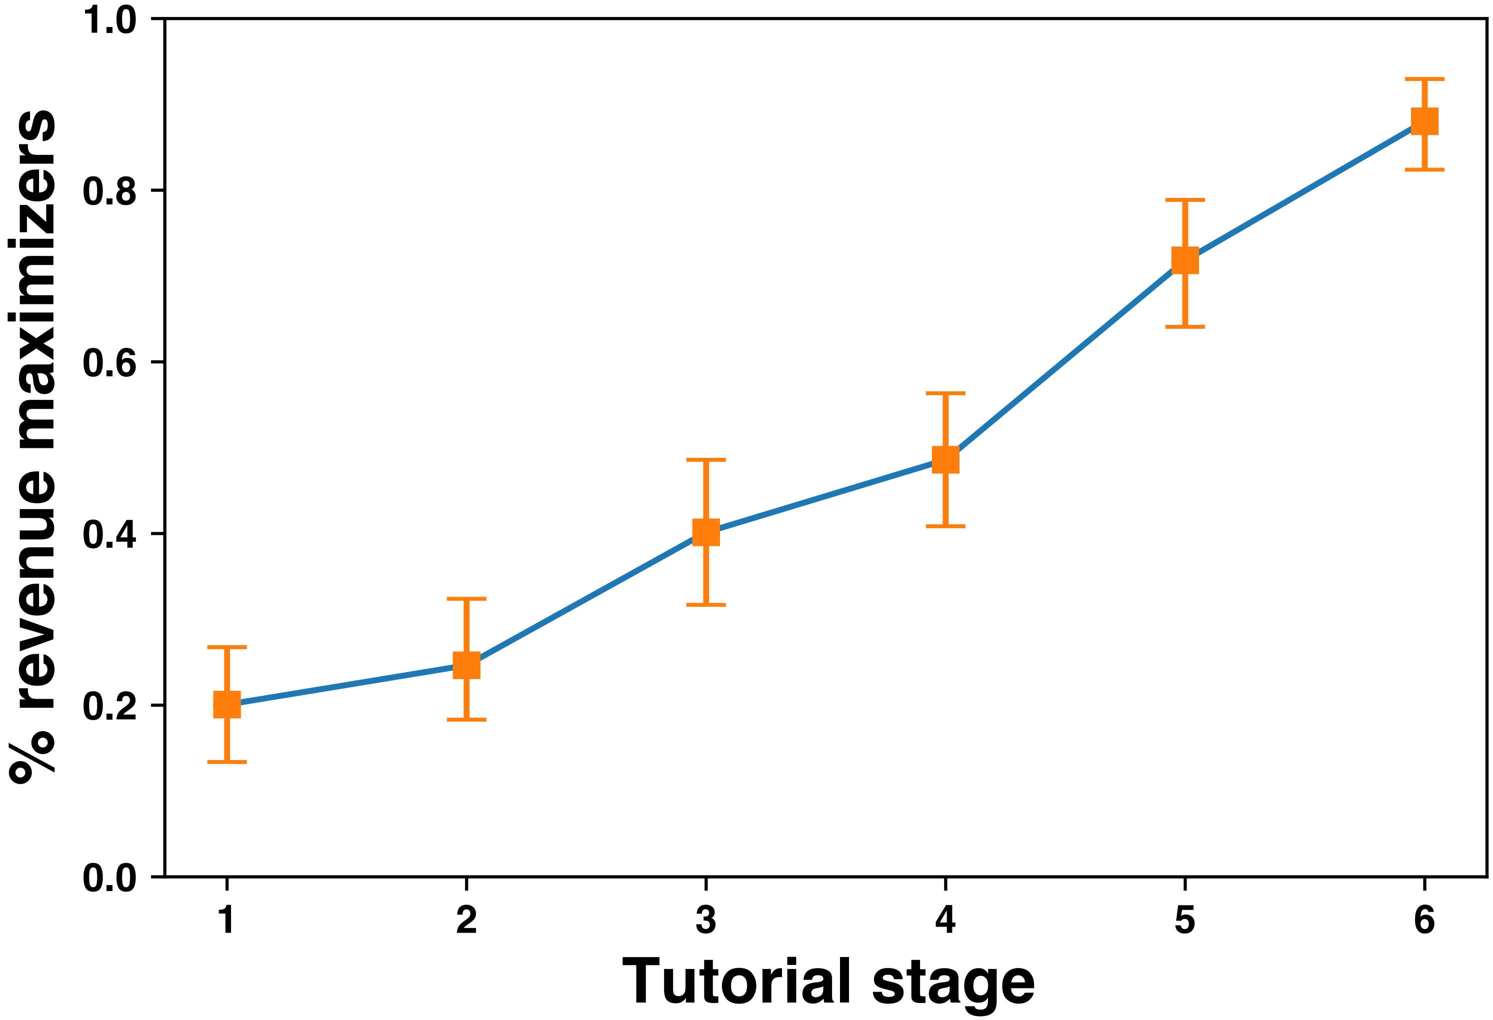

WAS THE ALGORITHM LEARNING?

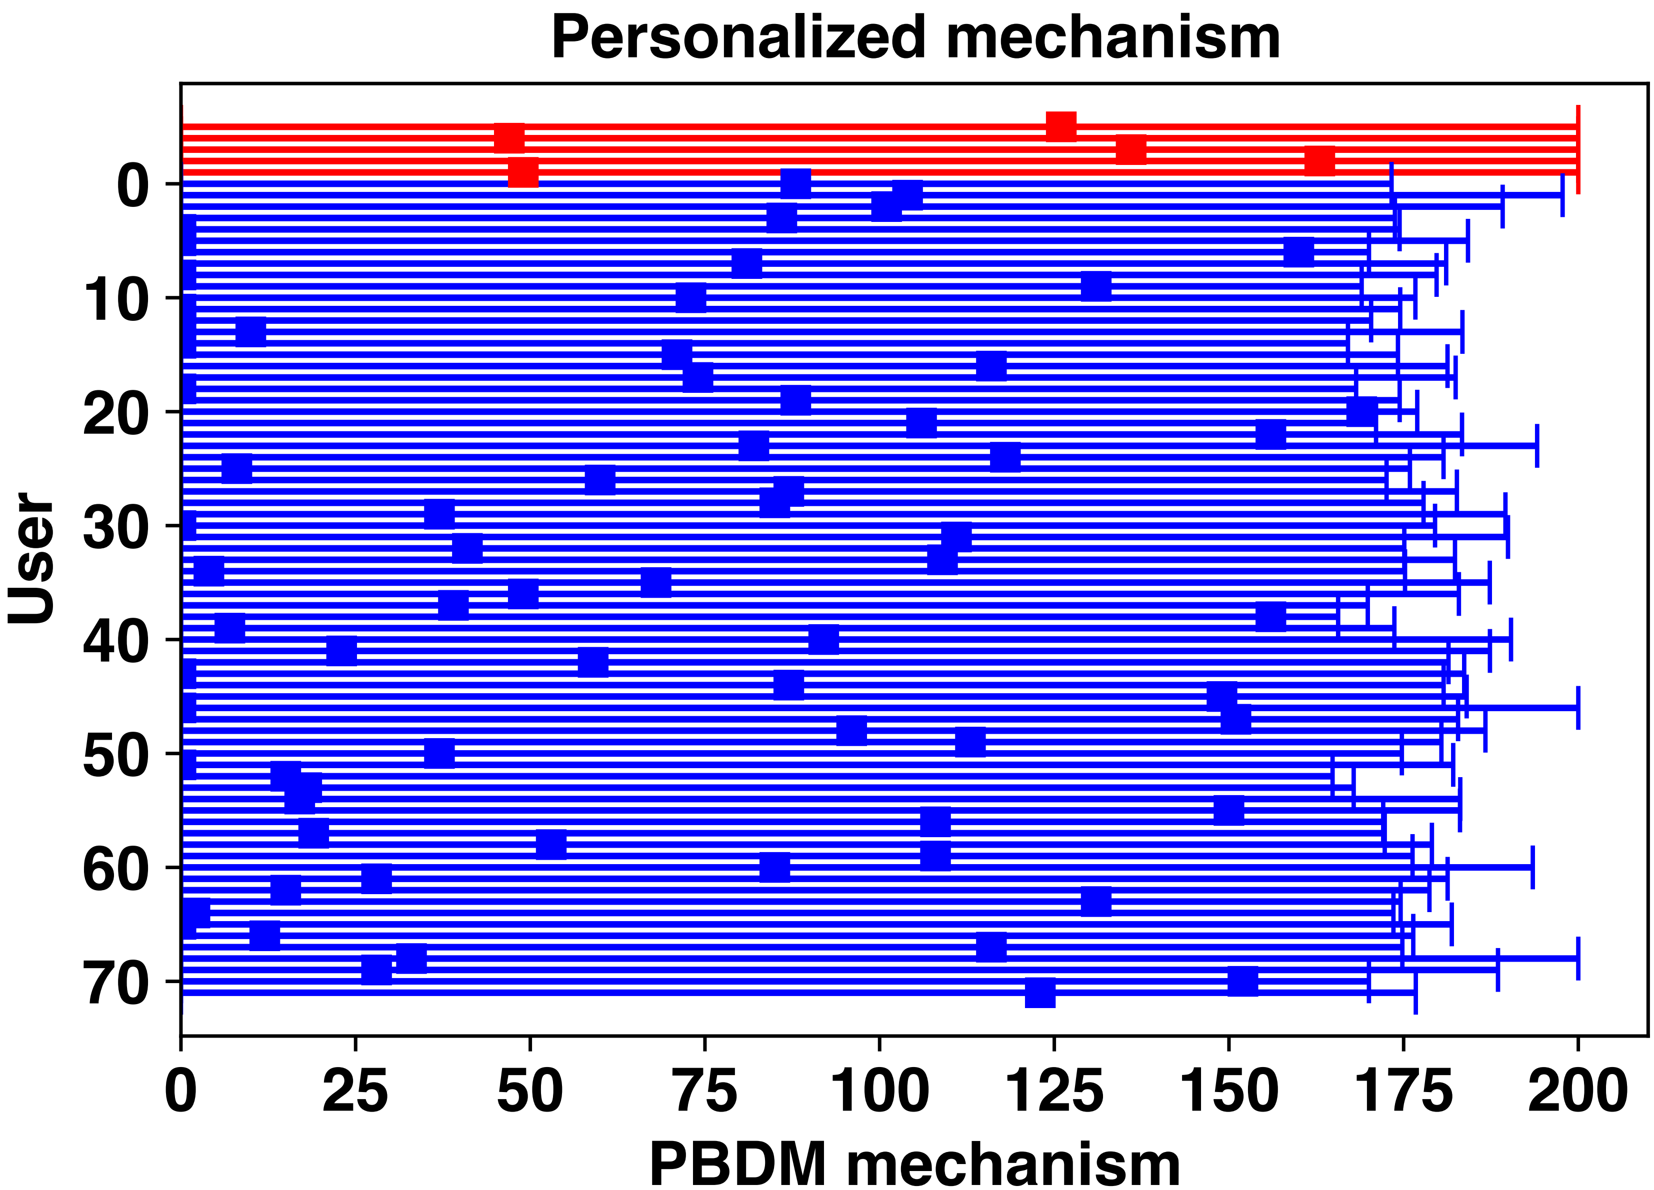

ESTIMATING THE PROBABILITY OF TREATMENT

- We can estimate the probability of assignment by sampling over all possible permutations and simply averaging how many times a given user would have been treated.

Propensity Scores

$$p_i(\mathbf{X}, \mathbf{W}, \mathbf{Y}(0), \mathbf{Y}(1), \pi) = \mathcal{I}(W_i < \hat{w}(x_i; \{x_j\}_{j:\pi(j)<\pi(i)}),$$

$$q(X_i, W_i, Y(0), Y(1)) = \sum_{\pi \in S_n} p_i(\mathbf{X}, \mathbf{W}, \mathbf{Y}(0), \mathbf{Y}(1), \pi)\Pr(\pi).$$

PROPENSITY SCORES AND ARRIVAL ORDER RANDOMNESS

- Estimate probability of assignment at a given arrival position.
- Then, assume random arrival order and take average over order

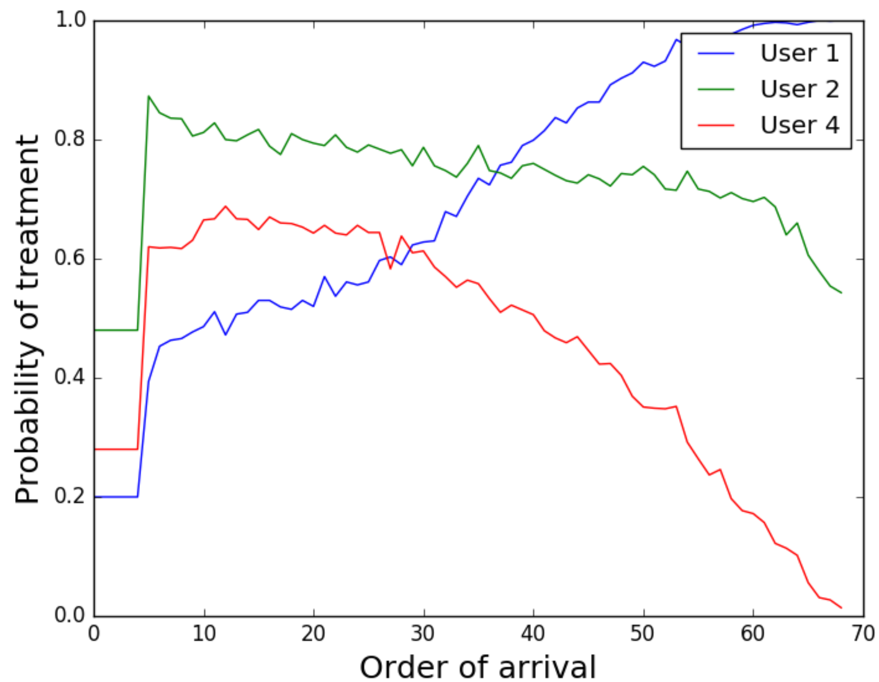

Treatment probability:	0.69	0.72	0.45
Willingness to pay:	50	120	70