Powering Causal Inference with EconML

Eleanor W. Dillon

Microsoft Research, ALICE

Microsoft

Current Team:

Keith Battocchi

Eleanor Dillon

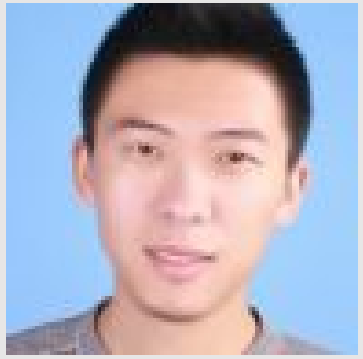

Kevin Gao

Darya Moldavskaya

Fabio Vera

Distinguished Alumni:

Maggie Hei

Greg Lewis

Miruna Oprescu

Pedro Sant'Anna

Vasilis Syrgkanis

EconML is a Python package that applies the power of machine learning techniques to estimate individualized causal responses from observational or experimental data.

Flexible model forms avoid strong assumptions and can estimate personalized responses to treatment

Unified API brings together all the latest advances in causal machine learning and econometrics

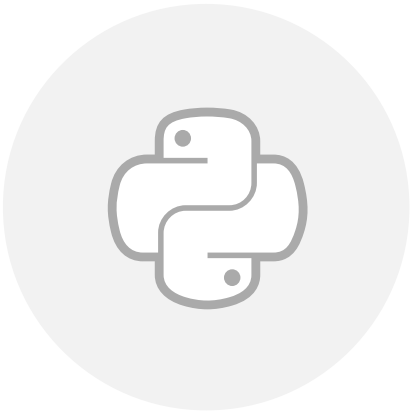

Familiar Interface built on standard Python packages make causal analysis quicker and easier for a broad set of users

Motivation and History

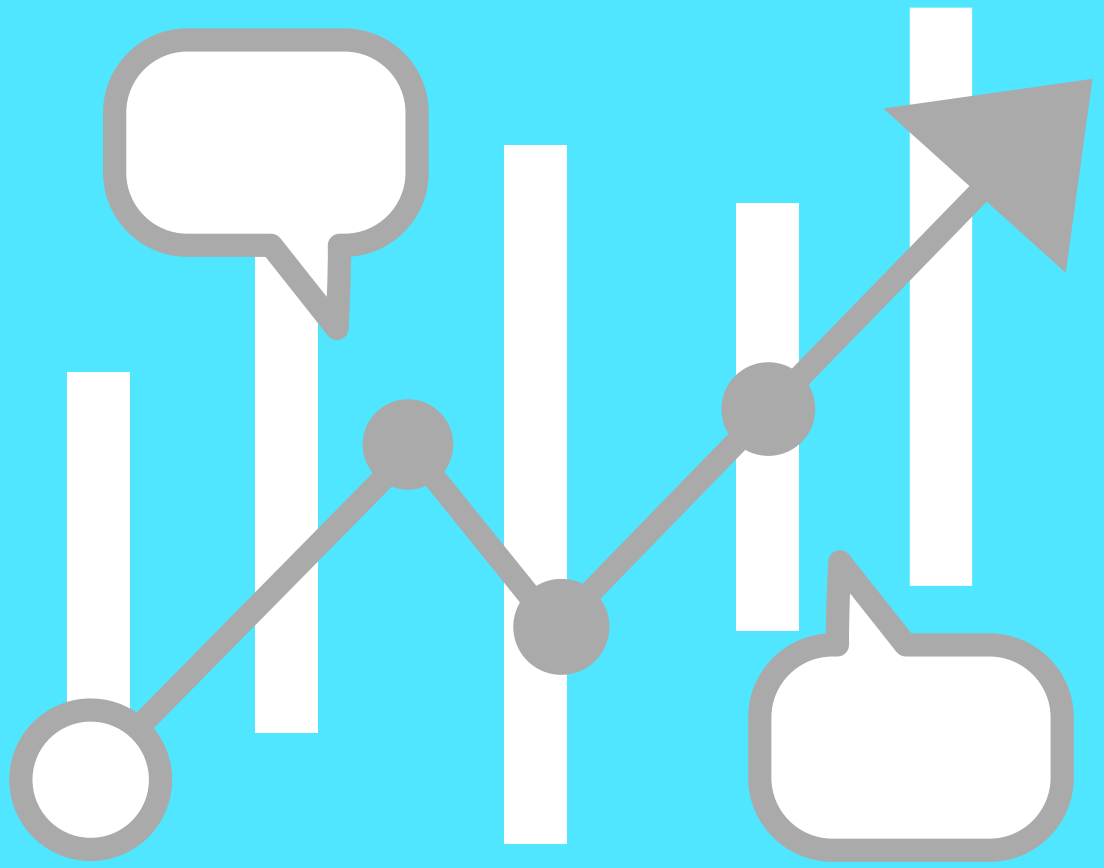

- Original mission statement: automate causal inference

- First version March 2019
- Now averaging nearly 200k downloads/month
- Many applied use cases internally and externally

Limits / Barriers / Choices Made

- What language?
  - Python reaches data scientists and users in businesses. Less popular for academic economists
- Big data
  - Recent additions: federation package to split any model with a linear final stage across multiple machines. Wrapper to run on Spark.
- Completeness vs. coherent API
  - Focus on heterogeneous treatment effects under unconfoundedness or instrumental variables assumptions
  - Toying with a panel data module
  - Should really have matching estimators

Automating Causal Inference?

Frame question

User/ decision-maker

Tools could help

Assemble data

Biggest hurdle to automation

Automate for specific cases?

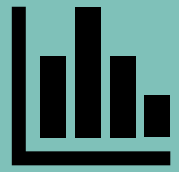

Select and estimate causal model

EconML automates

Sensitivity checks

Automate production

Tools can help decision

Interpret results

Tools can help

Early customer-facing application: Responsible AI Toolbox

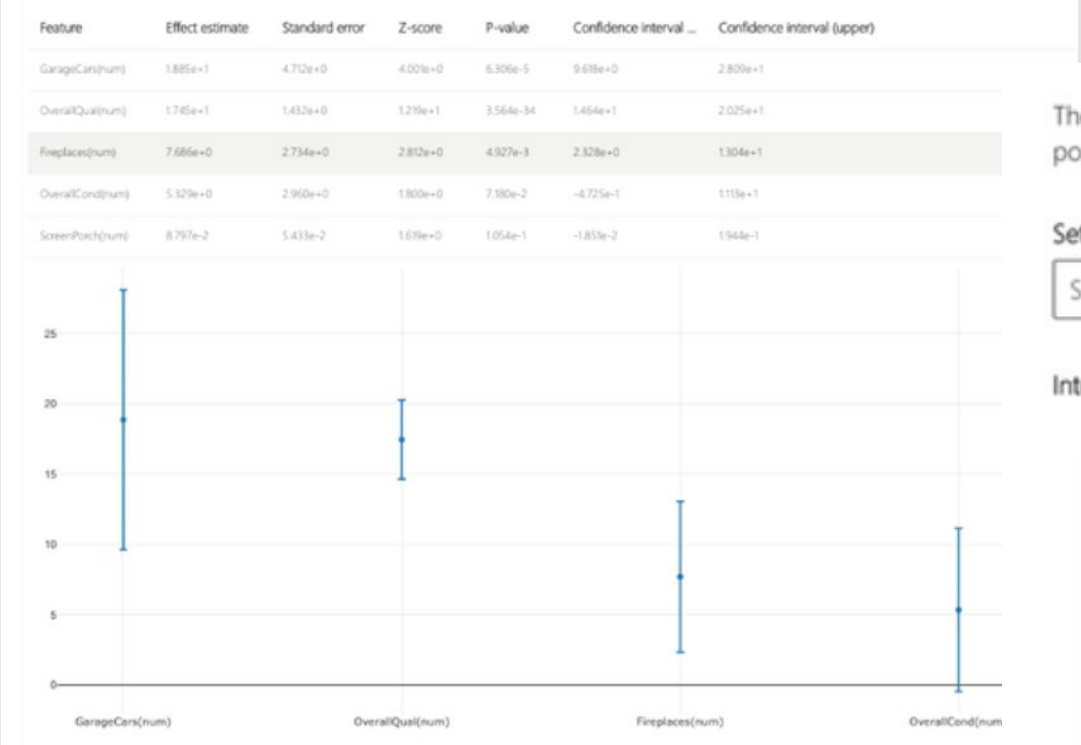

Feature	Effect estimate	Standard error	Z-score	P-value	Confidence interval ...	Confidence interval (upper)
GarageCars(num)	1.885e+1	4.712e+0	4.001e+0	6.306e-5	9.618e+0	2.809e+1
OverallQual(num)	1.745e+1	1.432e+0	1.219e+1	3.564e-34	1.464e+1	2.025e+1
Fireplaces(num)	7.686e+0	2.734e+0	2.812e+0	4.927e-3	2.328e+0	1.304e+1
OverallCond(num)	5.329e+0	2.960e+0	1.800e+0	7.180e-2	-4.725e-1	1.113e+1
ScreenPorch(num)	8.797e-2	5.433e-2	1.619e+0	1.054e-1	-1.851e-2	1.944e-1

Aggregate causal effects

These tools help build policies for future interventions. You can identify what parts of your sample experience the largest response populations should be targeted for particular interventions.

Set treatment feature

ScreenPorch

Interpretable recommended global treatment policy for sample size (n) = 730

OverallCond <= 6.5	EnclosedPorch <= 15	EnclosedPorch > 15
	n = 508 Recommended treatment = increase	n = 73 Recommended treatment = decrease
OverallCond > 6.5	OpenPorchSF <= 151	OpenPorchSF > 151
	n = 138 Recommended treatment = decrease	n = 11 Recommended treatment = increase

Treatment policy rules

Links

GitHub - py-why/EconML

Microsoft Responsible AI Toolbox - Microsoft Responsible AI

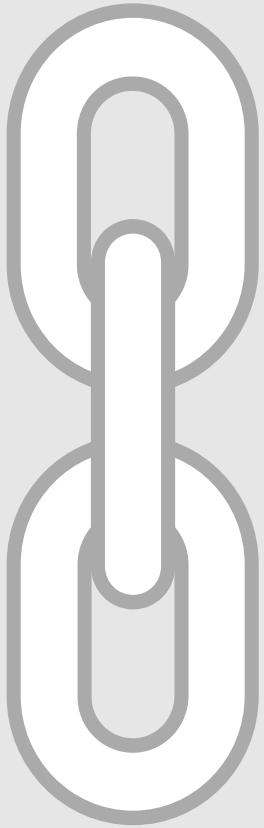

Case Study: TripAdvisor Membership

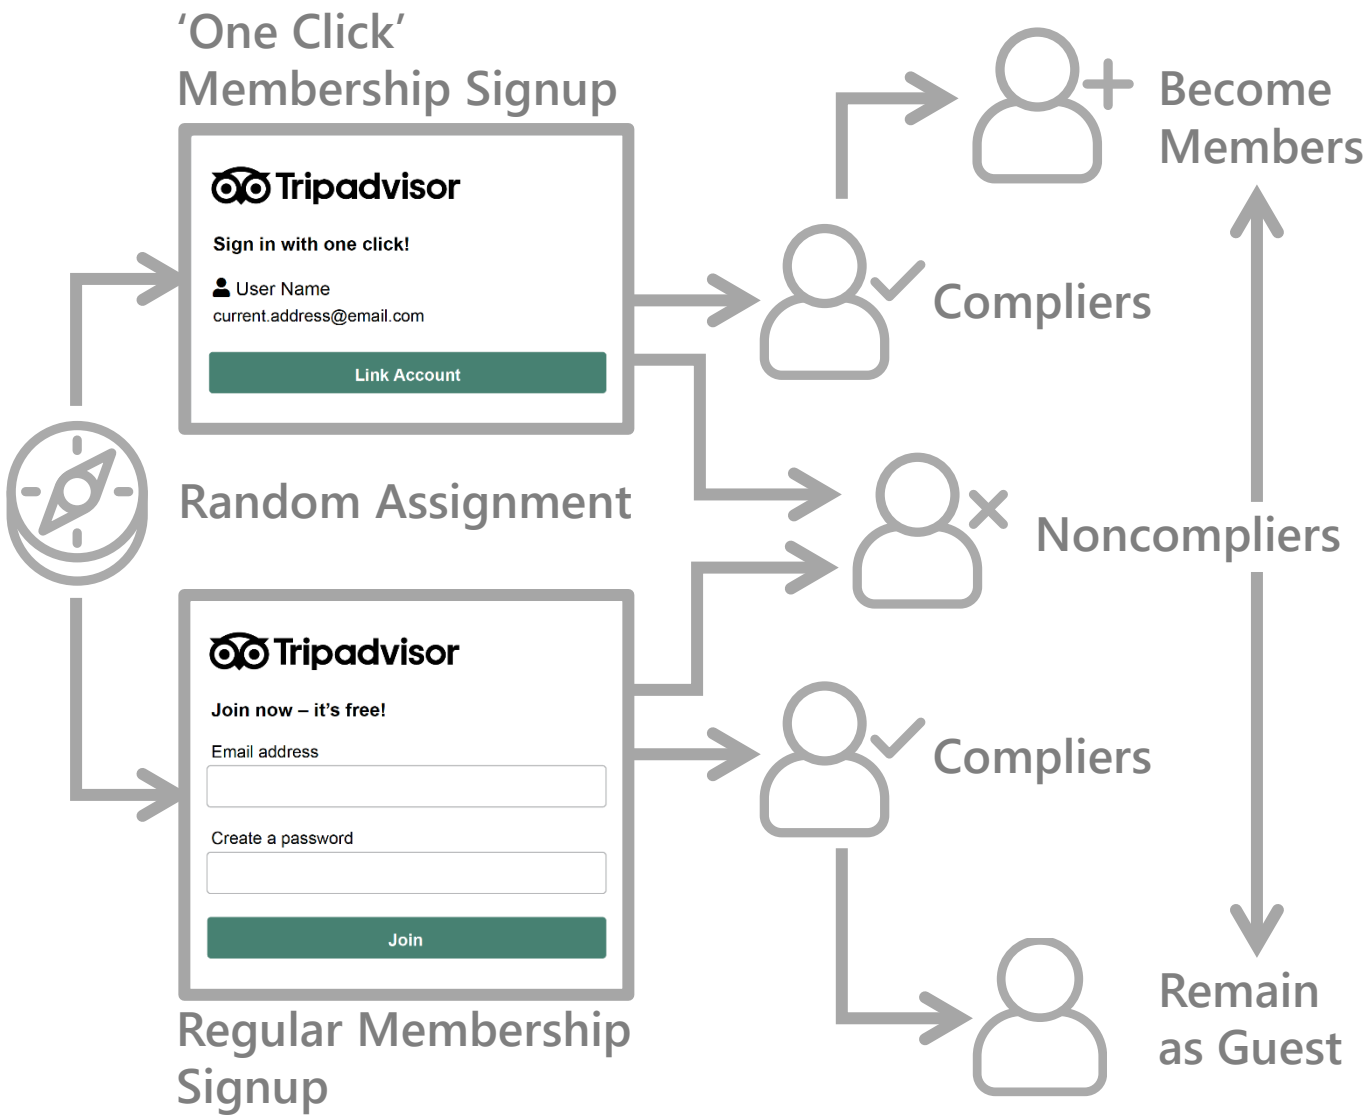

Question: does **membership** drive **engagement and bookings** on their platform?

TripAdvisor randomly offered some users an easier path to membership, but **couldn't force** users to become members

Insights from ALICE-TripAdvisor Collaboration

TripAdvisor learned the average treatment effect of membership

and which users are most affected by membership