Topics from experimentation in the tech industry

Steve Howard

November 16, 2023

Experiments come in many flavors...

UI / UX experiments

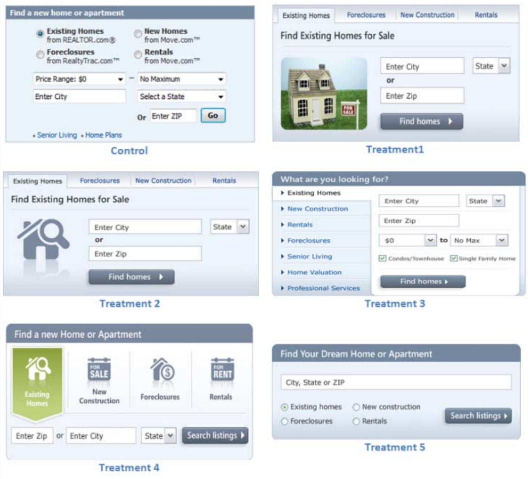

Kohavi et al. (2012)

Ridesharing dispatch

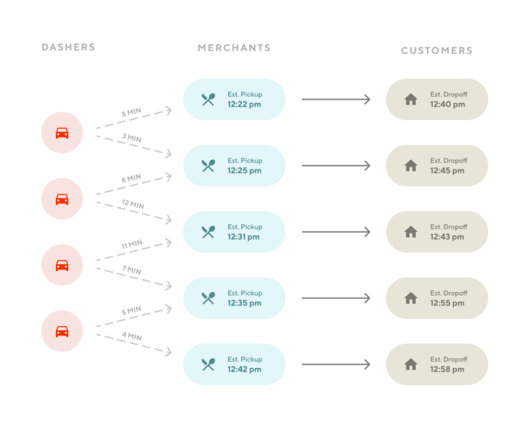

doordash.engineering/2018/02/13/switchback-tests-and-randomized-experimentation-under-network-effects-at-doordash/

Pricing experiments

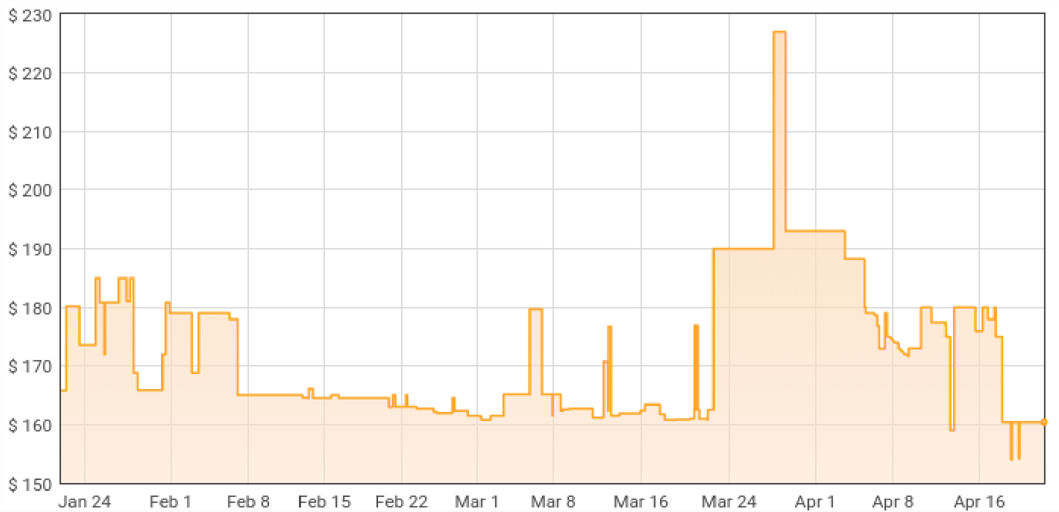

From keepa.com

Ranking experiments

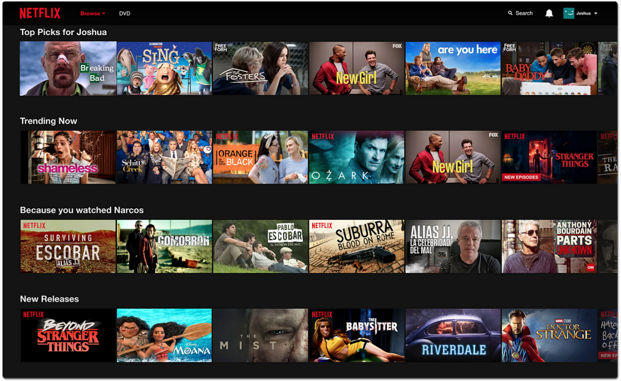

netflixtechblog.com/interleaving-in-online-experiments-at-netflix-a04ee392ec55

Ads blindness

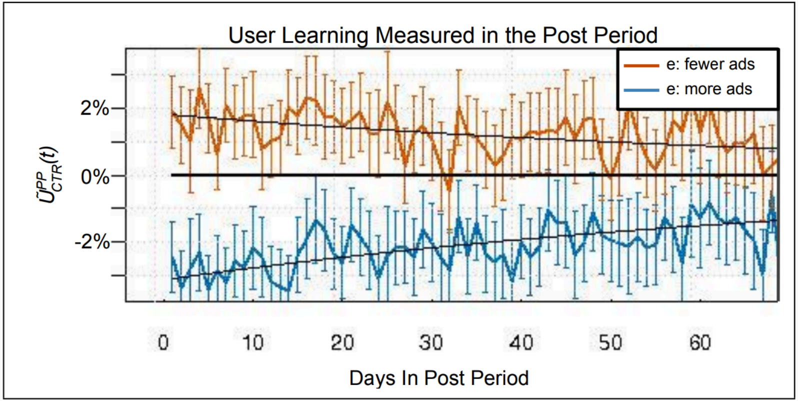

Hohnhold et al. (2015)

1. A brief tour
    - Hash randomization and ramping
    - Novelty effects, long-term effects, surrogates
    - Variance reduction
    - Interference
    - Sequential testing
    - ...
2. Augmented inverse propensity weighting: a really useful technique

Part I: a brief tour through some aspects of experimentation in tech

Hash randomization and ramping

Hash randomization

- Deterministic map from ID to [0, 1]
- Allows consistent ramping
- No need to store assignments

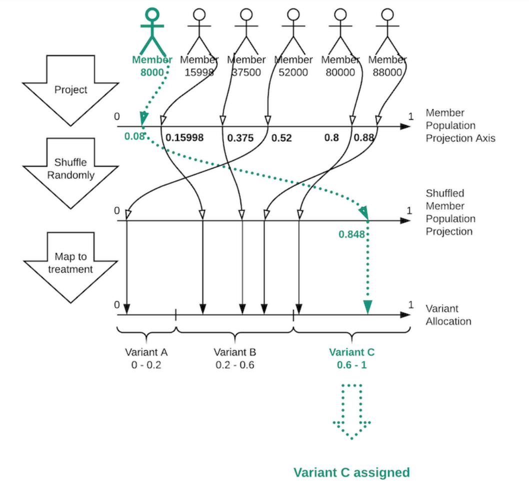

engineering.linkedin.com/blog/2020/a-b-testing-variant-assignment

Ramping / staggered entry

- Week 1: 5% canary
- Weeks 2-3: 50% for max power
- Weeks 4-7: 95% holdout

(Just an example)

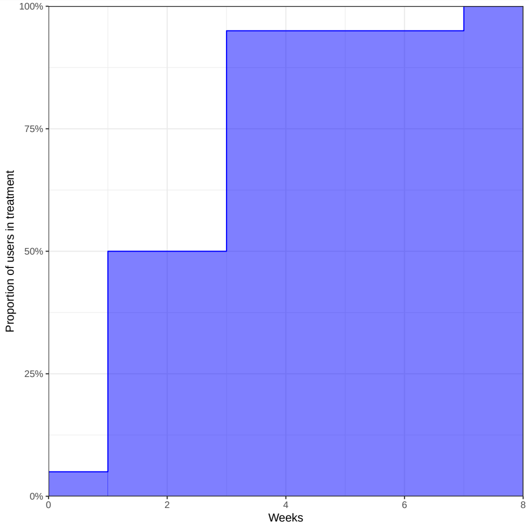

Regression detection

Catch regressions as early as possible

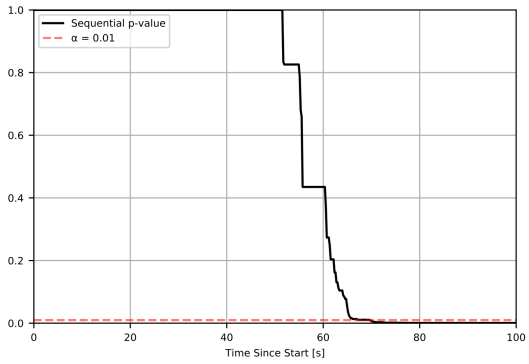

Lindon et al. (2022)

Novelty effects, long-term effects, surrogates

Novelty effects

- Users react to changes
- Or, power users enroll sooner
- Major source of false positives

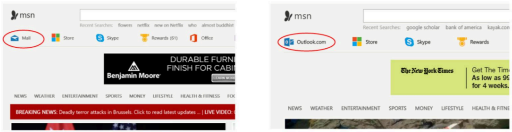

Figure 1: Screenshots of treatment with the Mail app button (left), and control with the Outlook.com button (right) in the msn.com experiment.

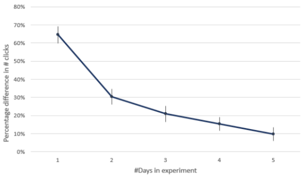

Figure 2: The percentage difference in number of clicks on the Mail/Outlook.com button, each day, between treatment and control in the msn.com experiment.

Sadeghi et al. (2021)

Learning effects

- Can take months to see long-term effects
- Can be hard to estimate

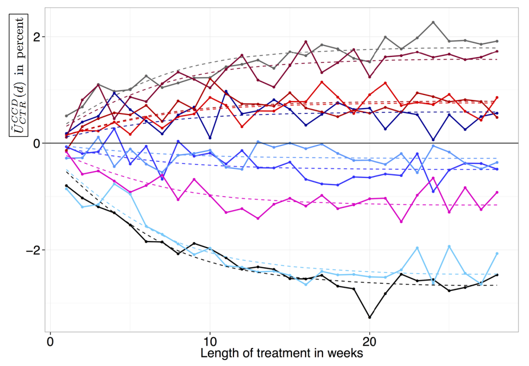

Hohnhold et al. (2015)

Surrogates: overview

- Surrogates are a popular approach to estimate long-term effects
- Basic idea
  - Regress long-term “true north” metric on short-term “surrogate” metrics
  - Use the resulting combination of surrogate metrics as an experiment outcome metric
- Example from Athey et al. (2019):
  - Treatment: job training program in California in the 1980s
  - Long-term metric: employment rate over nine years
  - Surrogate metrics: employment over each of first six quarters

Surrogates: example data

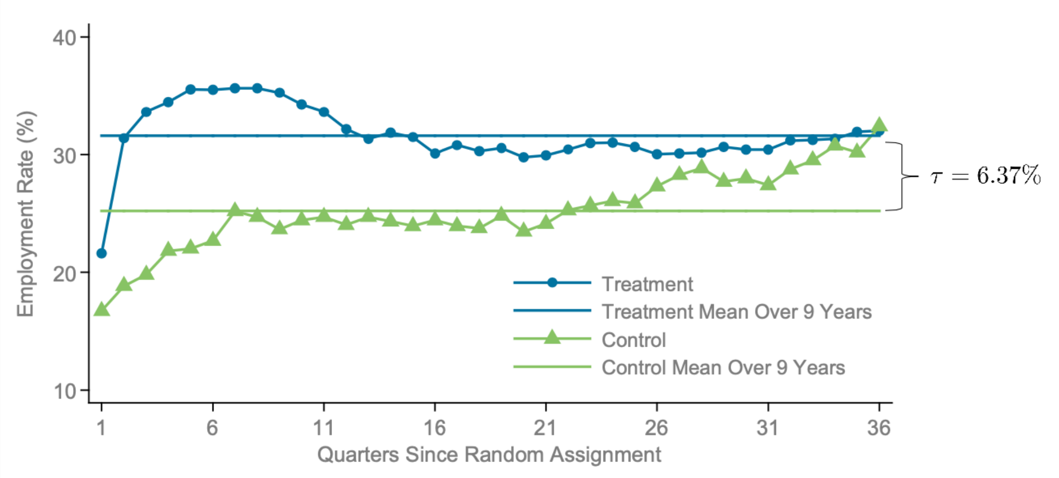

Athey et al. (2019)

Surrogate estimation

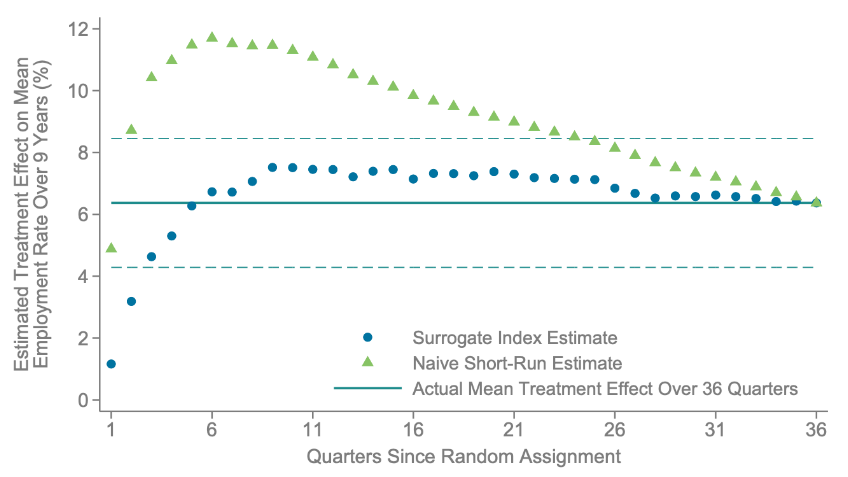

Athey et al. (2019)

Variance reduction

Covariate adjustment for variance reduction

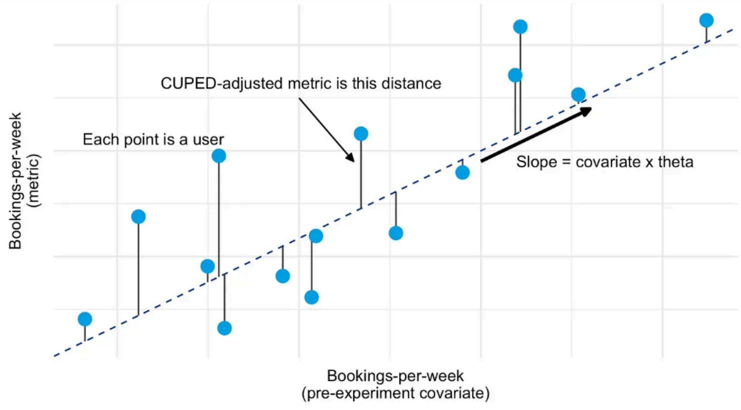

booking.ai/

how-booking-com-increases-the-power-of-online-experiments-with-cuped-995d186fff1d

Interleaving for ranking experiments

- Increase sample size
- Avoid chance imbalance from "power users"
- Main challenge is position bias

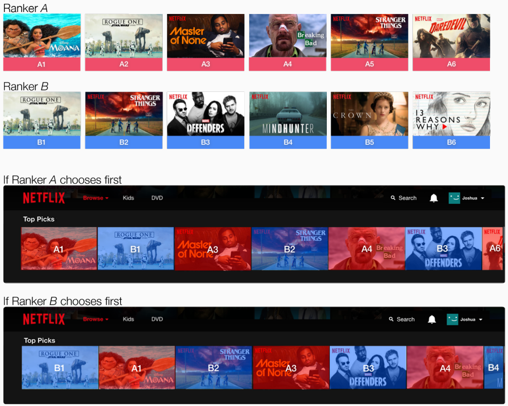

netflixtechblog.com/interleaving-in-online-experiments-at-netflix-a04ee392ec55

Interference

Interference

- Ideal experiment:
    1. (A) Treat everyone
    2. (B) Go back in time and treat no one
    3. (C) Compare
- Randomization ensures treated and control groups resemble the full population
- SUTVA ensures treated and control group *outcomes* are representative of the above counterfactuals (A) and (B)
- **Interference**: treatment for one unit affects outcome for another (violation of SUTVA)
- Outcomes in treated (control) group may not be representative of outcomes if we were to treat everyone (no one)

Examples of interference

- Social network sharing
- Ridesharing marketplaces
- Product marketplaces
- Ad auctions
- Producer-side ranker experiments

Coarse randomization

Randomization unit	Bias axis	Variance axis
User sessions	↑	↓
Users		
Fine spatial units (geohash)		
Time interval (hour)		
Coarse spatial units (city)		

eng.lyft.com/experimentation-in-a-ridesharing-marketplace-b39db027a66e

Switchback design

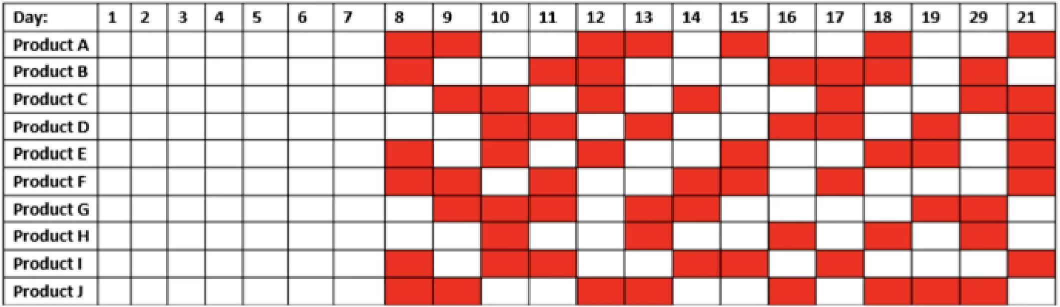

Day:	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	29	21
Product A																					
Product B																					
Product C																					
Product D																					
Product E																					
Product F																					
Product G																					
Product H																					
Product I																					
Product J																					

www.amazon.science/blog/the-science-of-price-experiments-in-the-amazon-store

Multiple randomization

- Randomize individual interactions
- Can choose design to target specific spillovers
- In this example:
    - Blue - Red = spillover effect within sellers
    - Green - Red = spillover effect within buyers
    - Black - Red = informative about total effect

j (Sellers)

i (Buyers)	1	2	3	4	5	6	7	8
1	C	C	C	C	C	C	C	C
2	C	C	C	C	C	C	C	C
3	C	C	C	C	T	T	T	T
4	C	C	C	C	T	T	T	T
5	C	C	C	C	T	T	T	T

Bajari et al. (2021)

Graph clustering

- Assume interference is limited by a known graph
- Use isolated clusters as randomization units
- There are many varieties; illustration shows “ego clustering”

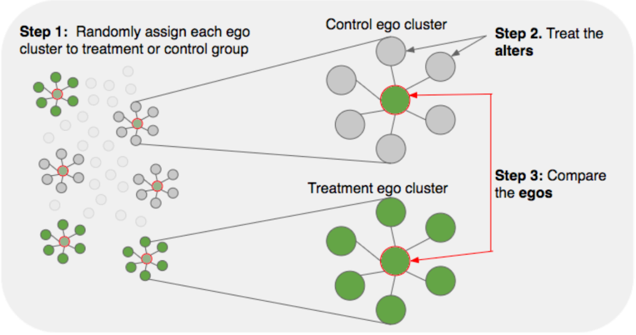

Saint-Jacques et al. (2019)

Budget split

Avoids interference due to scarce resources, e.g., in ad marketplaces

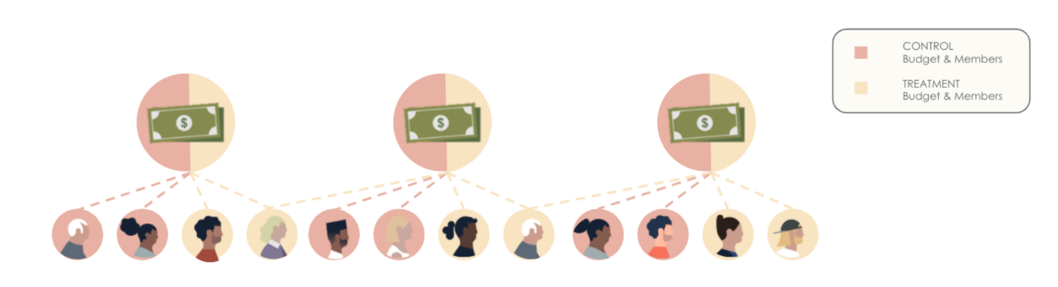

`engineering.linkedin.com/blog/2021/budget-split-testing`

Sequential testing

Sequential monitoring of experiment results is problematic.

Click on the button

The percentage of visitors who clicked on a tracked element.

Variation #3 is beating Original by +58.0%.

VARIATIONS	VISITORS	CONVERSIONS	CONVERSION RATE	IMPROVEMENT	CHANCE TO BEAT BASELINE
Variation #3	970	32	3.3% (±1.12%)	+58.0%	95.2%
Original BASELINE	1,006	21	2.1% (±0.88%)	---	---
Variation #1	999	11	1.1% (±0.65%)	-47.3%	3.9%
Variation #2	1,027	11	1.1% (±0.63%)	-48.7%	3.3%

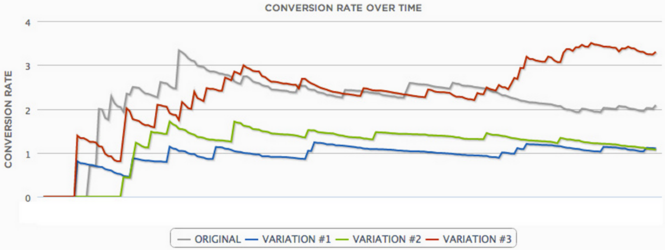

An old screenshot from optimizely.com

One path of p-values from a fair coin

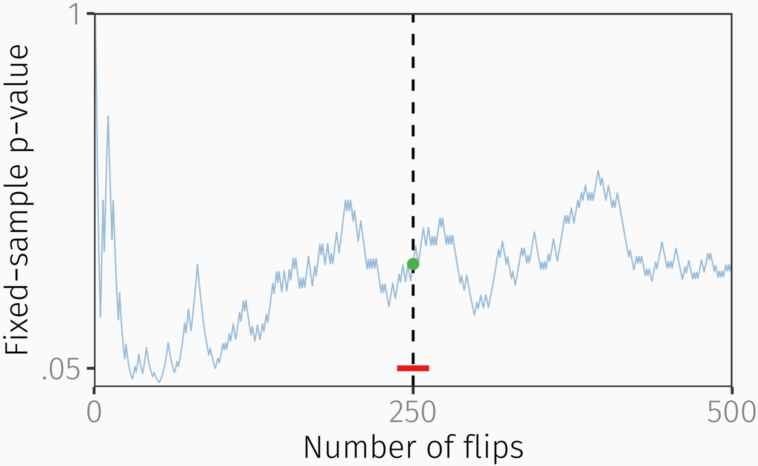

Let's look at many such paths...

With no bias, we only rarely conclude the coin is biased.

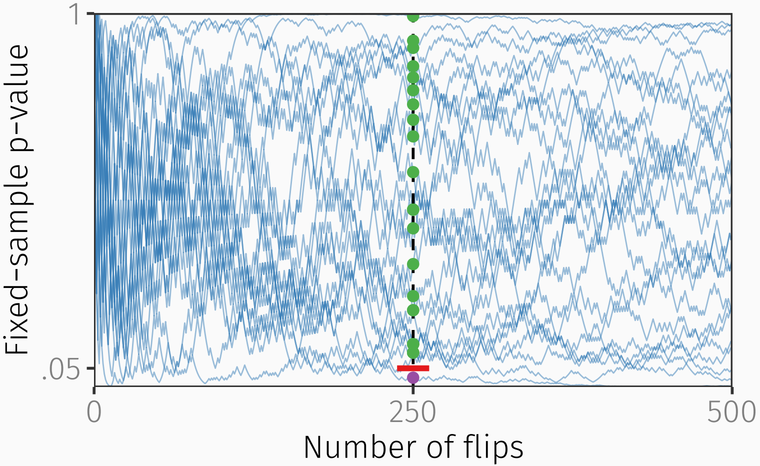

Just one out of 25 p-values is below 0.05.

Continuous monitoring of fixed-sample p-values inflates the false positive rate.

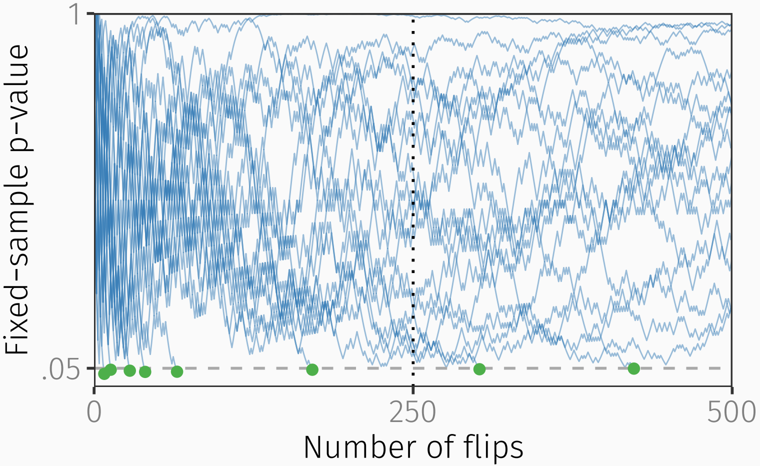

Here, with a fair coin, eight out of 25 paths reach significance.

The false positive rate grows arbitrarily large with enough flips.

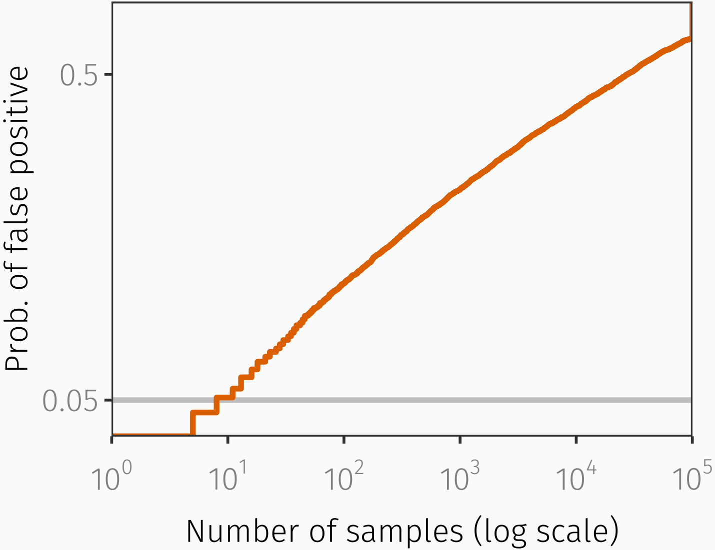

Solutions for sequential monitoring

Solution is basically to make p-value larger / confidence intervals wider.

Solutions for sequential monitoring

Solution is basically to make p-value larger / confidence intervals wider.

Two main appraoches:

- **Confidence sequences / always-valid inference**: allow peeking as frequently as desired (**more flexible**)
- **Group sequential tests**: allow peeking on a predetermined schedule (e.g., once a week) (**more powerful**)

So many more topics...

- Relative lift and ratio metrics
- Quantile treatment effects
- Heterogeneous treatment effects
- Offline evaluation and policy optimization
- Multiple testing
- Sequential allocation and optimization
- A/A tests and sample ratio mismatch
- Data engineering

Questions?

Part II: the Zen of Augmented Inverse Propensity Weighting

Some common analyses of randomized experiments

- Estimate average treatment effect / absolute lift
- Estimate relative lift
- Estimate a model for heterogeneous treatment effects
- Estimate the value of a personalized policy (and perhaps optimize such a policy)
- Estimate quantile treatment effects

We want to do all these things *efficiently*, making use of covariates to reducing variance.

Some common analyses of randomized experiments

- Estimate average treatment effect / absolute lift
- Estimate relative lift
- Estimate a model for heterogeneous treatment effects
- Estimate the value of a personalized policy (and perhaps optimize such a policy)
- Estimate quantile treatment effects

We want to do all these things *efficiently*, making use of covariates to reducing variance.

Augmented inverse propensity weighting (AIPW) provides a unified approach to solving all of these.

Setup: two-armed experiment with potential outcomes

- We run an experiment with two arms, treatment and control.
- Each units *(e.g., user)* is represented by a quadruple $(X_i, W_i, Y_i(0), Y_i(1))$
  - X_i: covariates measured before treatment *(e.g., pre-experiment weekly sessions)*
  - $W_i \in 0, 1$: treatment indicator *(e.g., send new notifications)*
  - $Y_i(0), Y_i(1)$: potential outcomes under control and treatment, respectively *(e.g., weekly sessions if assigned to control or treatment)*
- We observe either $Y_i(0)$ or $Y_i(1)$, not both: our *observed* sample consists of $(X_i, W_i, Y_i(W_i))$
  - We make the usual SUTVA / no-interference assumption.
- We assign W_i by i.i.d. coin flips with $P(W_i = 1) = p$.

Example potential outcomes table

Full table (unobserved):

	X_i	W_i	$Y_i(0)$	$Y_i(1)$
Unit 1	6	1	5	8
Unit 2	2	0	2	0
Unit 3	3	1	4	4
⋮	⋮	⋮	⋮	⋮

Observed table:

	X_i	W_i	$Y_i(0)$	$Y_i(1)$
Unit 1	6	1	?	8
Unit 2	2	0	2	?
Unit 3	3	1	?	4
⋮	⋮	⋮	⋮	⋮

The estimands and some (inefficient) estimators

Estimand	Definition	Naive estimator
ATE / Absolute lift	$\mathbb{E}\left[Y(1) - Y(0)\right]$	$\mathbb{E}_n\left[Y_i \mid W_i = 1\right] - \mathbb{E}_n\left[Y_i \mid W_i = 0\right]$

The estimands and some (inefficient) estimators

Estimand	Definition	Naive estimator
ATE / Absolute lift	$\mathbb{E}[Y(1) - Y(0)]$	$\mathbb{E}_n[Y_i \mid W_i = 1] - \mathbb{E}_n[Y_i \mid W_i = 0]$
Relative lift	$\frac{\mathbb{E}[Y(1)]}{\mathbb{E}[Y(0)]}$	$\frac{\mathbb{E}_n[Y_i \mid W_i = 1]}{\mathbb{E}_n[Y_i \mid W_i = 0]}$

The estimands and some (inefficient) estimators

Estimand	Definition	Naive estimator
ATE / Absolute lift	$\mathbb{E}\left[Y(1) - Y(0)\right]$	$\mathbb{E}_n\left[Y_i \mid W_i = 1\right] - \mathbb{E}_n\left[Y_i \mid W_i = 0\right]$
Relative lift	$\frac{\mathbb{E}[Y(1)]}{\mathbb{E}[Y(0)]}$	$\frac{\mathbb{E}_n[Y_i \mid W_i=1]}{\mathbb{E}_n[Y_i \mid W_i=0]}$
Policy value	$V(\pi) = \mathbb{E}\left[Y(\pi(X))\right]$	$\widehat{V}(\pi) = \mathbb{E}_n\left[\frac{1_{W_i=\pi(X_i)} Y_i}{\mathbb{P}(W_i=\pi(X_i))}\right]$

The estimands and some (inefficient) estimators

Estimand	Definition	Naive estimator
ATE / Absolute lift	$\mathbb{E}\left[Y(1) - Y(0)\right]$	$\mathbb{E}_n\left[Y_i \mid W_i = 1\right] - \mathbb{E}_n\left[Y_i \mid W_i = 0\right]$
Relative lift	$\frac{\mathbb{E}[Y(1)]}{\mathbb{E}[Y(0)]}$	$\frac{\mathbb{E}_n[Y_i \mid W_i=1]}{\mathbb{E}_n[Y_i \mid W_i=0]}$
Policy value	$V(\pi) = \mathbb{E}\left[Y(\pi(X))\right]$	$\widehat{V}(\pi) = \mathbb{E}_n\left[\frac{1_{W_i=\pi(X_i)} Y_i}{\mathbb{P}(W_i=\pi(X_i))}\right]$
CATE(x)	$\mathbb{E}\left[Y(1) - Y(0) \mid X = x\right]$	$Y_i(1) \sim \mu_1(X_i) \mid W_i = 1$ $Y_i(0) \sim \mu_0(X_i) \mid W_i = 0$ $\widehat{\text{CATE}}(x) = \hat{\mu}_1(x) - \hat{\mu}_0(x)$

The AIPW "imputation principle"

1. Regress $Y_i(1)$ on X_i among units with $W_i = 1$ to get $\hat{\mu}_1(x)$.
2. Regress $Y_i(0)$ on X_i among units with $W_i = 0$ to get $\hat{\mu}_0(x)$.
3. For each unit i compute "pseudo-outcomes"

$$\Gamma_i(0) = \hat{\mu}_0(X_i) + \frac{1 - W_i}{1 - p}\left(Y_i - \hat{\mu}_0(X_i)\right)$$
$$\Gamma_i(1) = \hat{\mu}_1(X_i) + \frac{W_i}{p}\left(Y_i - \hat{\mu}_1(X_i)\right)$$

(In practice, need to use cross-fitting. Also a good idea to use empirical propensity.)

The AIPW "imputation principle"

1. Regress $Y_i(1)$ on X_i among units with $W_i = 1$ to get $\hat{\mu}_1(x)$.
2. Regress $Y_i(0)$ on X_i among units with $W_i = 0$ to get $\hat{\mu}_0(x)$.
3. For each unit i compute "pseudo-outcomes"

$$\Gamma_i(0) = \hat{\mu}_0(X_i) + \frac{1 - W_i}{1 - p}\left(Y_i - \hat{\mu}_0(X_i)\right)$$

$$\Gamma_i(1) = \hat{\mu}_1(X_i) + \frac{W_i}{p}\left(Y_i - \hat{\mu}_1(X_i)\right)$$

(In practice, need to use cross-fitting. Also a good idea to use empirical propensity.)

Then treat $\Gamma_i(0), \Gamma_i(1)$ as if they were observed potential outcomes.

Why does this make sense?

Suppose only W_i is random (not $X_i, Y_i(0), Y_i(1)$), and treat $\hat{\mu}_0, \hat{\mu}_1$ as fixed. Then

$$\begin{aligned}
\mathbb{E}\Gamma_i(1) &= \mathbb{E}\left[\hat{\mu}_1(X_i) + \frac{W_i}{p}(Y_i - \hat{\mu}_1(X_i))\right] && (1)\\
&= \hat{\mu}_1(X_i) + \mathbb{E}\left[\frac{W_i}{p}(Y_i(1) - \hat{\mu}_1(X_i))\right] && (2)\\
&= \hat{\mu}_1(X_i) + \frac{\mathbb{E}W_i}{p}(Y_i(1) - \hat{\mu}_1(X_i)) && (3)\\
&= Y_i(1). && (4)
\end{aligned}$$

Why does this make sense?

Suppose only W_i is random (not $X_i, Y_i(0), Y_i(1)$), and treat $\hat{\mu}_0, \hat{\mu}_1$ as fixed. Then

$$\begin{aligned}
\mathbb{E}\Gamma_i(1) &= \mathbb{E}\left[\hat{\mu}_1(X_i) + \frac{W_i}{p}(Y_i - \hat{\mu}_1(X_i))\right] && (1)\\
&= \hat{\mu}_1(X_i) + \mathbb{E}\left[\frac{W_i}{p}(Y_i(1) - \hat{\mu}_1(X_i))\right] && (2)\\
&= \hat{\mu}_1(X_i) + \frac{\mathbb{E}W_i}{p}(Y_i(1) - \hat{\mu}_1(X_i)) && (3)\\
&= Y_i(1). && (4)
\end{aligned}$$

Meanwhile,

$$\operatorname{Var}\Gamma_i(1) = \frac{1-p}{p}(Y_i(1) - \hat{\mu}_1(X_i))^2. \qquad (5)$$

Imputed potential outcomes table

Full table (unobserved):

	X_i	W_i	$Y_i(0)$	$Y_i(1)$
Unit 1	6	1	5	8
Unit 2	2	0	2	0
Unit 3	3	1	4	4
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

Imputed table:

	X_i	W_i	$\Gamma_i(0)$	$\Gamma_i(1)$
Unit 1	6	1	6	10
Unit 2	2	0	2	2
Unit 3	3	1	3	5
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$

Efficient AIPW estimators

Estimand	Definition	AIPW estimator
ATE / Absolute lift	$\mathbb{E}\left[Y(1) - Y(0)\right]$	$\mathbb{E}_n\left[\Gamma_i(1) - \Gamma_i(0)\right]$

Efficient AIPW estimators

Estimand	Definition	AIPW estimator
ATE / Absolute lift	$\mathbb{E}\left[Y(1) - Y(0)\right]$	$\mathbb{E}_n\left[\Gamma_i(1) - \Gamma_i(0)\right]$
Relative lift	$\frac{\mathbb{E}[Y(1)]}{\mathbb{E}[Y(0)]}$	$\frac{\mathbb{E}_n[\Gamma_i(1)]}{\mathbb{E}_n[\Gamma_i(0)]}$

Efficient AIPW estimators

Estimand	Definition	AIPW estimator
ATE / Absolute lift	$\mathbb{E}\left[Y(1) - Y(0)\right]$	$\mathbb{E}_n\left[\Gamma_i(1) - \Gamma_i(0)\right]$
Relative lift	$\frac{\mathbb{E}[Y(1)]}{\mathbb{E}[Y(0)]}$	$\frac{\mathbb{E}_n[\Gamma_i(1)]}{\mathbb{E}_n[\Gamma_i(0)]}$
Policy value	$V(\pi) = \mathbb{E}\left[Y(\pi(X))\right]$	$\widehat{V}(\pi) = \mathbb{E}_n\left[\Gamma_i(\pi(X_i))\right]$

Efficient AIPW estimators

Estimand	Definition	AIPW estimator
ATE / Absolute lift	$\mathbb{E}[Y(1) - Y(0)]$	$\mathbb{E}_n[\Gamma_i(1) - \Gamma_i(0)]$
Relative lift	$\frac{\mathbb{E}[Y(1)]}{\mathbb{E}[Y(0)]}$	$\frac{\mathbb{E}_n[\Gamma_i(1)]}{\mathbb{E}_n[\Gamma_i(0)]}$
Policy value	$V(\pi) = \mathbb{E}[Y(\pi(X))]$	$\widehat{V}(\pi) = \mathbb{E}_n[\Gamma_i(\pi(X_i))]$
CATE(x)	$\mathbb{E}[Y(1) - Y(0) \mid X = x]$	Regress $\Gamma_i(1) - \Gamma_i(0)$ on X_i

Attractive properties of AIPW

- Unbiased / consistent no matter what outcome model you use

Attractive properties of AIPW

- Unbiased / consistent no matter what outcome model you use
- Lowest possible variance if the outcome model is good

Attractive properties of AIPW

- Unbiased / consistent no matter what outcome model you use
- Lowest possible variance if the outcome model is good
- Can estimate variance ignoring randomness of the outcome model. For example:

$$\widehat{\text{ATE}} = \frac{1}{n} \sum_{i=1}^{n} [\Gamma_i(1) - \Gamma_i(0)]$$

$$\widehat{\text{Var}}(\widehat{\text{ATE}}) = \frac{1}{n^2} \sum_{i=1}^{n} \left[\Gamma_i(1) - \Gamma_i(0) - \widehat{\text{ATE}}\right]^2.$$

(The story is more complicated for CATE estimation.)

Attractive properties of AIPW

- Unbiased / consistent no matter what outcome model you use
- Lowest possible variance if the outcome model is good
- Can estimate variance ignoring randomness of the outcome model. For example:

$$\widehat{\text{ATE}} = \frac{1}{n}\sum_{i=1}^{n}\left[\Gamma_i(1) - \Gamma_i(0)\right]$$

$$\widehat{\text{Var}}(\widehat{\text{ATE}}) = \frac{1}{n^2}\sum_{i=1}^{n}\left[\Gamma_i(1) - \Gamma_i(0) - \widehat{\text{ATE}}\right]^2.$$

(The story is more complicated for CATE estimation.)

Generalizes to non-constant propensity score: replace p with $e(x) = P\left(W_i = 1 \mid X_i = x\right)$.

AIPW for policy optimization

Recall AIPW estimate of the value of a personalized policy π:

$$\widehat{V}(\pi) = \mathbb{E}_n \left[\Gamma_i(\pi(X_i)) \right].$$

Can treat this as an "empirical welfare" objective and optimize over π.

Reduces to cost-sensitive multiclass classification.

In the binary treatment case, this is weighted binary classification:

$$\begin{aligned} \text{Label:} \quad & 1_{\Gamma_i(1) > \Gamma_i(0)} \\ \text{Weight:} \quad & |\Gamma_i(1) - \Gamma_i(0)| \end{aligned}$$

Further reading on AIPW

- Robins et al. (1994) is generally cited as the origin of AIPW
- Chernozhukov et al. (2016) did a lot to popularize the use of cross-fitting together with AIPW-style estimators
- Jin and Ba (2021) focus on the variance reduction in randomized experiments, ”ratio metrics” such as click-through rates
- Dudik et al. (2011) is a great source on policy evaluation with AIPW
- Kennedy (2023) analyzes the ”DR-learner” for CATE estimation
- Angelopoulos et al. (2023) discusses variance reduction for M-estimators

- Angelopoulos, Anastasios N., Stephen Bates, Clara Fannjiang, Michael I. Jordan, and Tijana Zrnic (2023). *Prediction-Powered Inference*. URL: http://arxiv.org/abs/2301.09633. preprint.
- Athey, Susan, Raj Chetty, Guido W. Imbens, and Hyunseung Kang (2019). *The Surrogate Index: Combining Short-Term Proxies to Estimate Long-Term Treatment Effects More Rapidly and Precisely*. URL: https://www.nber.org/papers/w26463. preprint.
- Bajari, Patrick, Brian Burdick, Guido W. Imbens, Lorenzo Masoero, James McQueen, Thomas Richardson, and Ido M. Rosen (2021). *Multiple Randomization Designs*. URL: http://arxiv.org/abs/2112.13495. preprint.
- Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins (2016). "Double/Debiased Machine Learning for Treatment and Causal Parameters".
- Dudik, Miroslav, John Langford, and Lihong Li (2011). *Doubly Robust Policy Evaluation and Learning*. URL: http://arxiv.org/abs/1103.4601. preprint.
- Hohnhold, Henning, Deirdre O'Brien, and Diane Tang (2015). "Focusing on the Long-term: It's Good for Users and Business". In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '15*. The 21th ACM SIGKDD International Conference. Sydney, NSW, Australia: ACM Press, pp. 1849–1858.
- Jin, Ying and Shan Ba (2021). "Towards Optimal Variance Reduction in Online Controlled Experiments".

Kennedy, Edward H. (2023). *Towards Optimal Doubly Robust Estimation of Heterogeneous Causal Effects*. URL: http://arxiv.org/abs/2004.14497. preprint.

Kohavi, Ron, Alex Deng, Brian Frasca, Roger Longbotham, Toby Walker, and Ya Xu (2012). "Trustworthy Online Controlled Experiments: Five Puzzling Outcomes Explained". In: *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, pp. 786–794.

Lindon, Michael, Chris Sanden, and Vaché Shirikian (2022). "Rapid Regression Detection in Software Deployments through Sequential Testing".

Robins, James M., Andrea Rotnitzky, and Lue Ping Zhao (1994). "Estimation of Regression Coefficients When Some Regressors Are Not Always Observed". *Journal of the American Statistical Association* 89 (427), pp. 846–866.

Sadeghi, Soheil, Somit Gupta, Stefan Gramatovici, Jiannan Lu, Hao Ai, and Ruhan Zhang (2021). *Novelty and Primacy: A Long-Term Estimator for Online Experiments*. URL: http://arxiv.org/abs/2102.12893. preprint.

Saint-Jacques, Guillaume, Maneesh Varshney, Jeremy Simpson, and Ya Xu (2019). *Using Ego-Clusters to Measure Network Effects at LinkedIn*. URL: http://arxiv.org/abs/1903.08755. preprint.

Questions?

`steve@stevehoward.org`

Bonus: quantiles

Suppose we want to estimate the π-quantile of treatment outcomes $Y(1)$.

Bonus: quantiles

Suppose we want to estimate the π-quantile of treatment outcomes $Y(1)$.

This is defined by the population estimating equation

$$\theta^\star = \operatorname*{solve}_\theta \mathbb{E}\left[1_{Y(1)\leq\theta} - \pi\right],$$

and ordinarily our (M-)estimator would be the sample quantile

$$\hat{\theta} = \operatorname*{solve}_\theta \frac{1}{n}\sum_{i=1}^n 1_{Y_i(1)\leq\theta} - \pi,$$

but $Y_i(1)$ is missing for units with $A_i = 0$. So **apply AIPW to the estimating equation** itself:

$$\hat{\theta}_{\text{AIPW}} = \operatorname*{solve}_\theta \frac{1}{n}\sum_{i=1}^n \left[1_{\hat{\mu}_1(X_i)\leq\theta} + \frac{A_i}{p}\left(1_{Y_i(1)\leq\theta} - 1_{\hat{\mu}_1(X_i)\leq\theta}\right)\right].$$