Using Behavioral Economics to Rank Newsfeed

Facebook Data Science

Mike Bailey with Alex Peysakhovich and Ta Chiraphadhanakul

Motivation

Search for people, places and things

Janelle | Home

Janelle Vreeland
Edit Profile

Update Status | Add Photos/Video

What's on your mind?

Viewing most recent stories · Back to top

Silent Film Fans
1 min ·

Happy Birthday to legendary animator Ub Iwerks (1901–1971). Seen here (on the right) with his boss, Walt Disney, who called him "The greatest animator in the world".

Like · Comment · Share · 1

Eric Idle
1 min · Twitter ·

Well we have chosen the new Python single, and I think you are going to be surprised and I hope delighted by who is on it. Out in April.

Like · Comment · Share · @EricIdle on Twitter

YOUR ADS | Create Ad
Provide Feedback | Today | Monthly

24	12	8
Ads	Campaigns	Page Likes

Promote Your Post

This post "The Florentine is rolling out ..." is getting more engagement than 75% of your recent posts. Get more likes, comments and shares by promoting your post.

Promote

5 upcoming birthdays

Chi-Fi 0 on March 29

TRENDING | Learn More

Blue Line: CTA Blue Line derailment at O'Hare injures at least 30

Mila Kunis: Mila Kunis Pregnant, Expecting First Child With Fiance Ashton Kutcher

Mark Cuban: Cuban: NFL '10 Years from Implosion'

See More

MUSIC PAGES YOU MAY LIKE | See All

Bobby Rydell
105,235 people like him.
Like

SPONSORED | Create Ad

Learn about ASU Online
asuonline.asu.edu

Ranked Top Tier University - U.S. News & World Report. 100% online. Apply now.

Try Fandor for 2 Weeks

News Feed
the right content to the right
people at the right time

What is the right content?

Who are the right people?

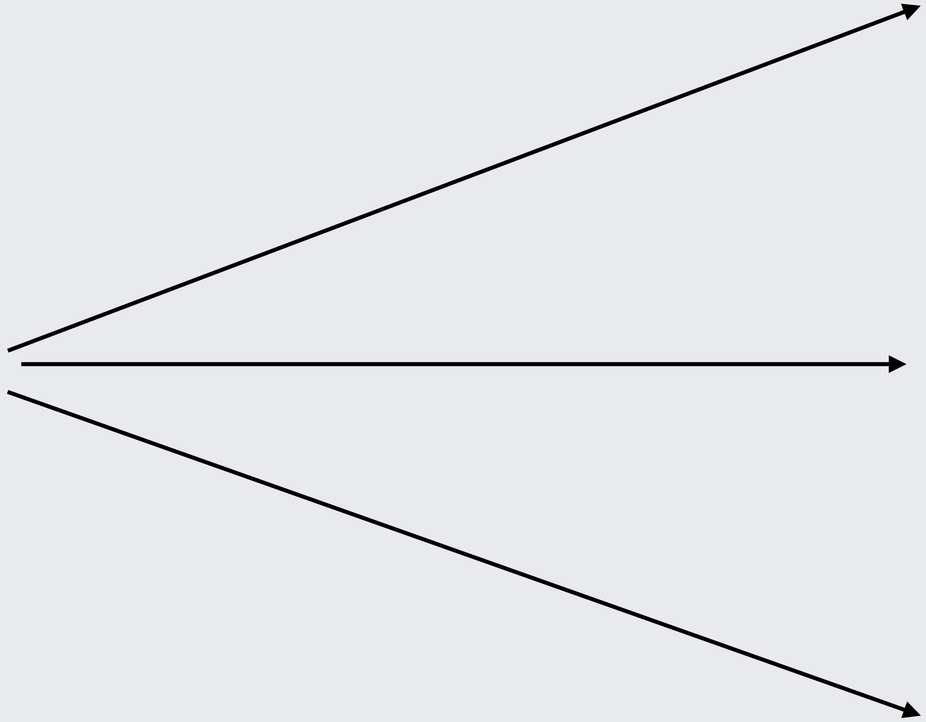

When is the right time?

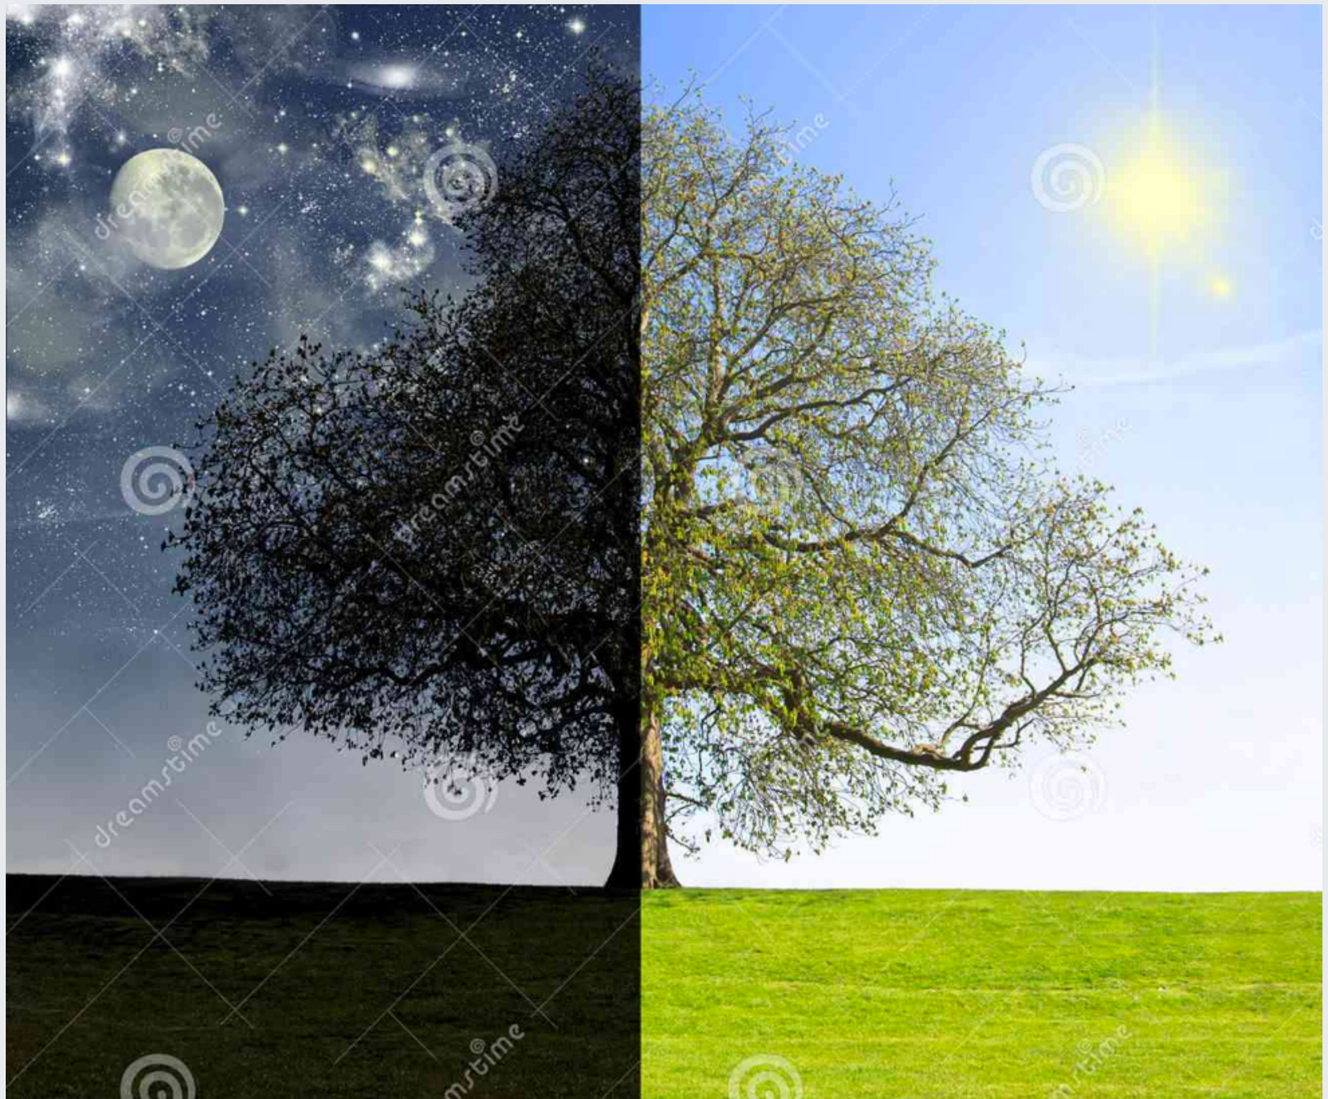

1. Big Data
2. Machine Learning
3. …
4. Profit

Optimize for likes

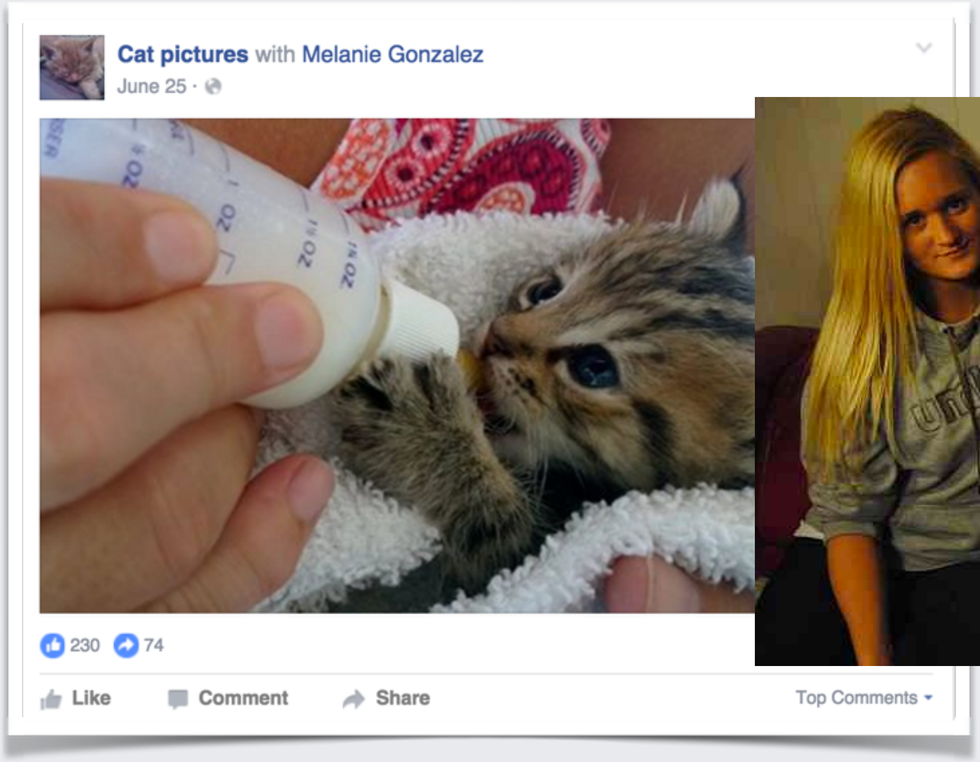

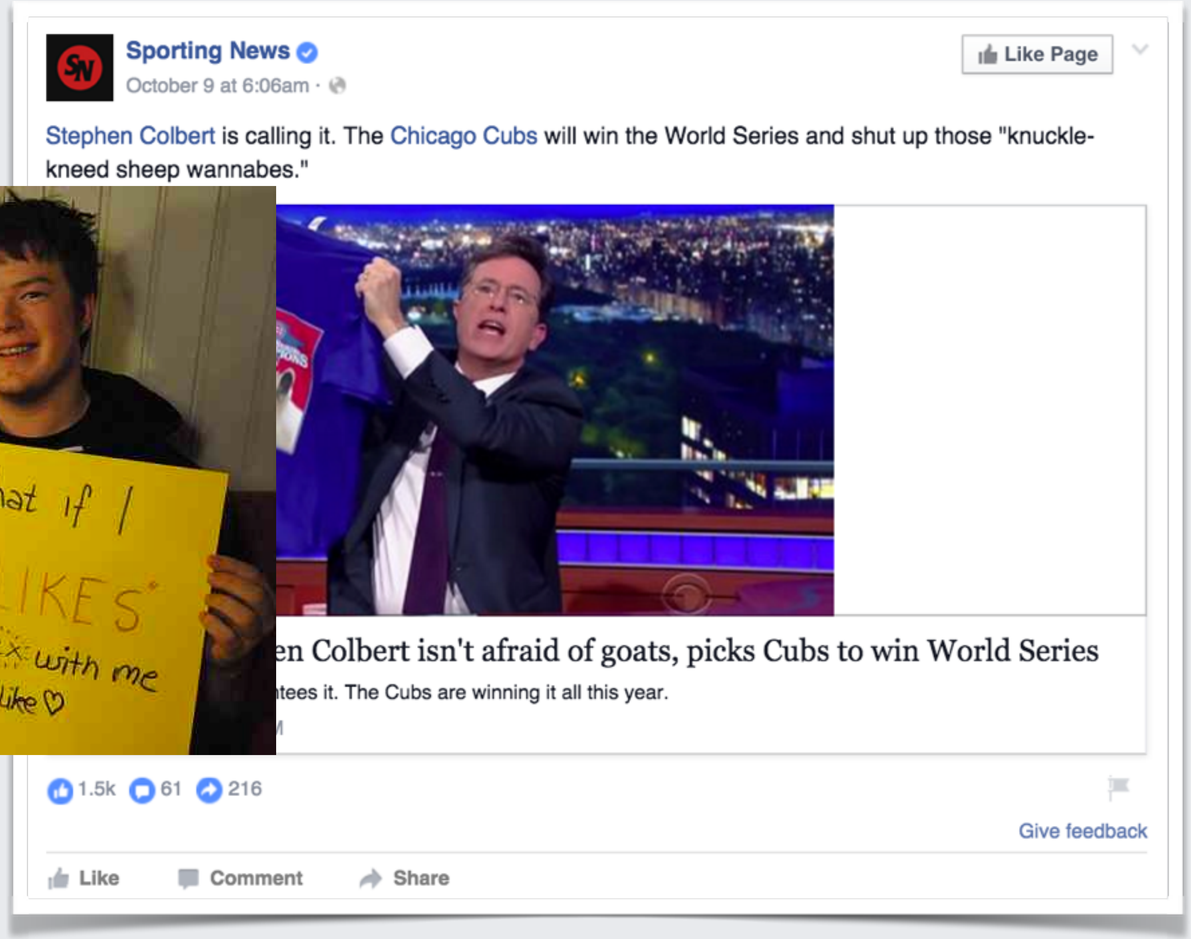

Optimize for shares

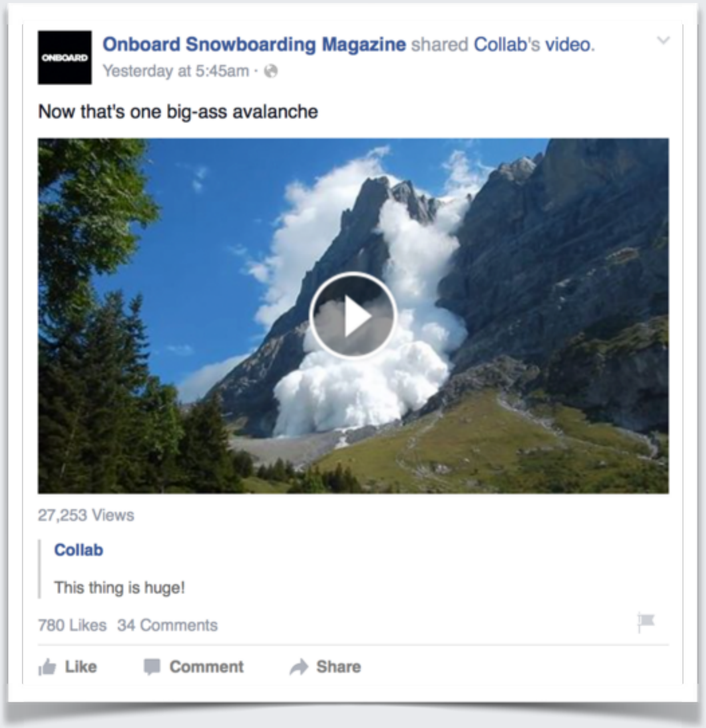

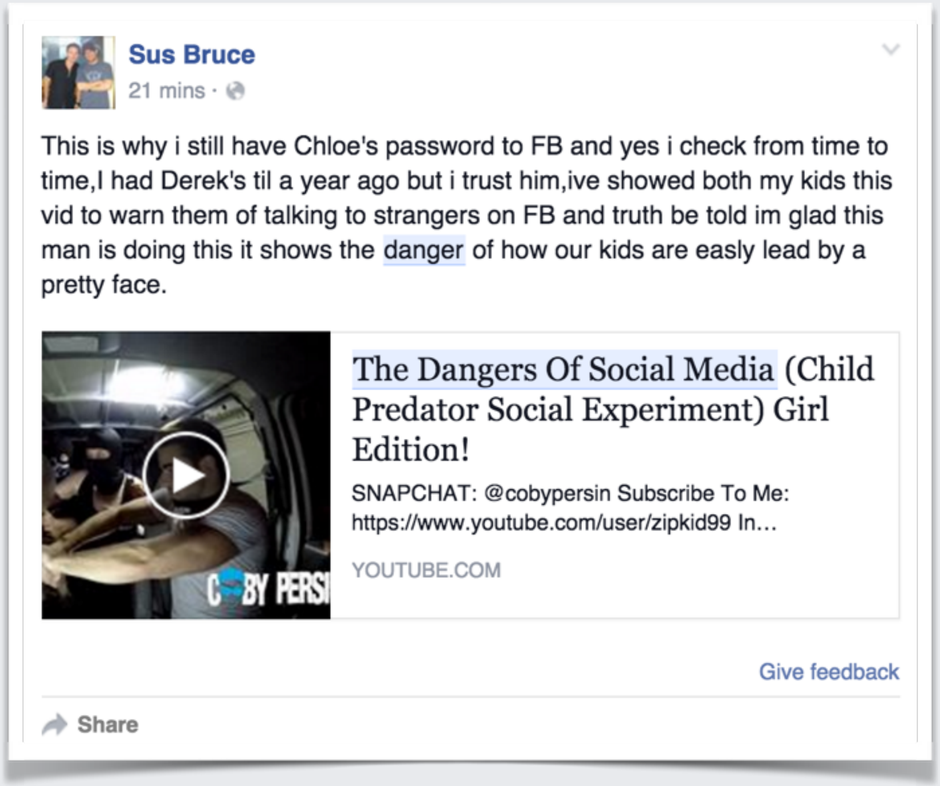

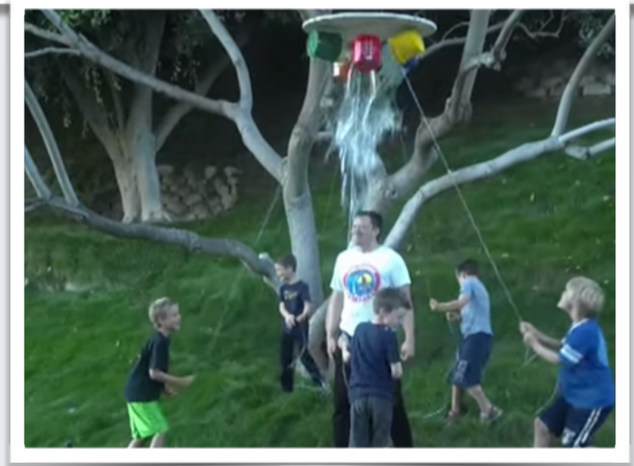

Optimize for comments

National Rifle Association
Yesterday at 12:47pm ·

Verizon Fios has dropped Outdoor Channel and Sportsman Channel - television networks geared to tell compelling outdoor lifestyle stories and strong advocates for all areas of the shooting lifestyle. What gives?

Silencing Your Access to NRA Programs
Act Now and Demand Access to the Best in Outdoor Television
NRAILA.ORG | BY NRA-ILA

3.7k 1.4k

Like Comment Share Most Relevant

Planned Parenthood
7 hrs ·

Powerful.

The First Legal Abortion Providers Tell Their Stories
Seven physicians on the pre- and post-Roe world.
P.PPFA.ORG | BY ALEX RONAN

2.2k 610

Like Comment Share Most Relevant

Donald J. Trump
4 hrs · Edited ·

I will be hosting Saturday Night Live on November 7th. Looking forward to it!

68k 2 1 1 5.3k

Like Comment Share Most Relevant

Write a comment...

Mark Smith TRUMP 2016!!!!
Like · Reply · 1,506 · 4 hrs
122 Replies · 4 mins

Brenda G Short You are the ONLY reason I would watch snl.
Like · Reply · 1,091 · 4 hrs
62 Replies · 12 mins

View more comments 2 of 5,142

Optimize for ??

likes + 10*comments + 5*shares?

Ranking

- **Long term objectives are hard**
  - We could optimize for content that keeps people engaged 30 days later, in practice this is really difficult to do.
- **Weight tuning**
  - Users are different and ranking should reflect that. Our ranking should understand heterogeneity.
  - Not only do we need to worry about a misspecified objective function, each user might be operating under a different objective function.

How Not to Drown in Numbers

If you're trying to build a self-driving car or detect whether a picture has a cat in it, big data is amazing. But here's a secret: If you're trying to make important decisions about your health, wealth or happiness, big data is not enough.

The problem is this: The things we can measure are never exactly what we care about. Just trying to get a single, easy-to-measure number higher and higher (or lower and lower) doesn't actually help us make the right choice. For this reason, the key question isn't "What did I measure?" but "What did I miss?"

FACEBOOK'S HUMAN-POWERED ASSISTANT MAY JUST SUPERCHARGE AI

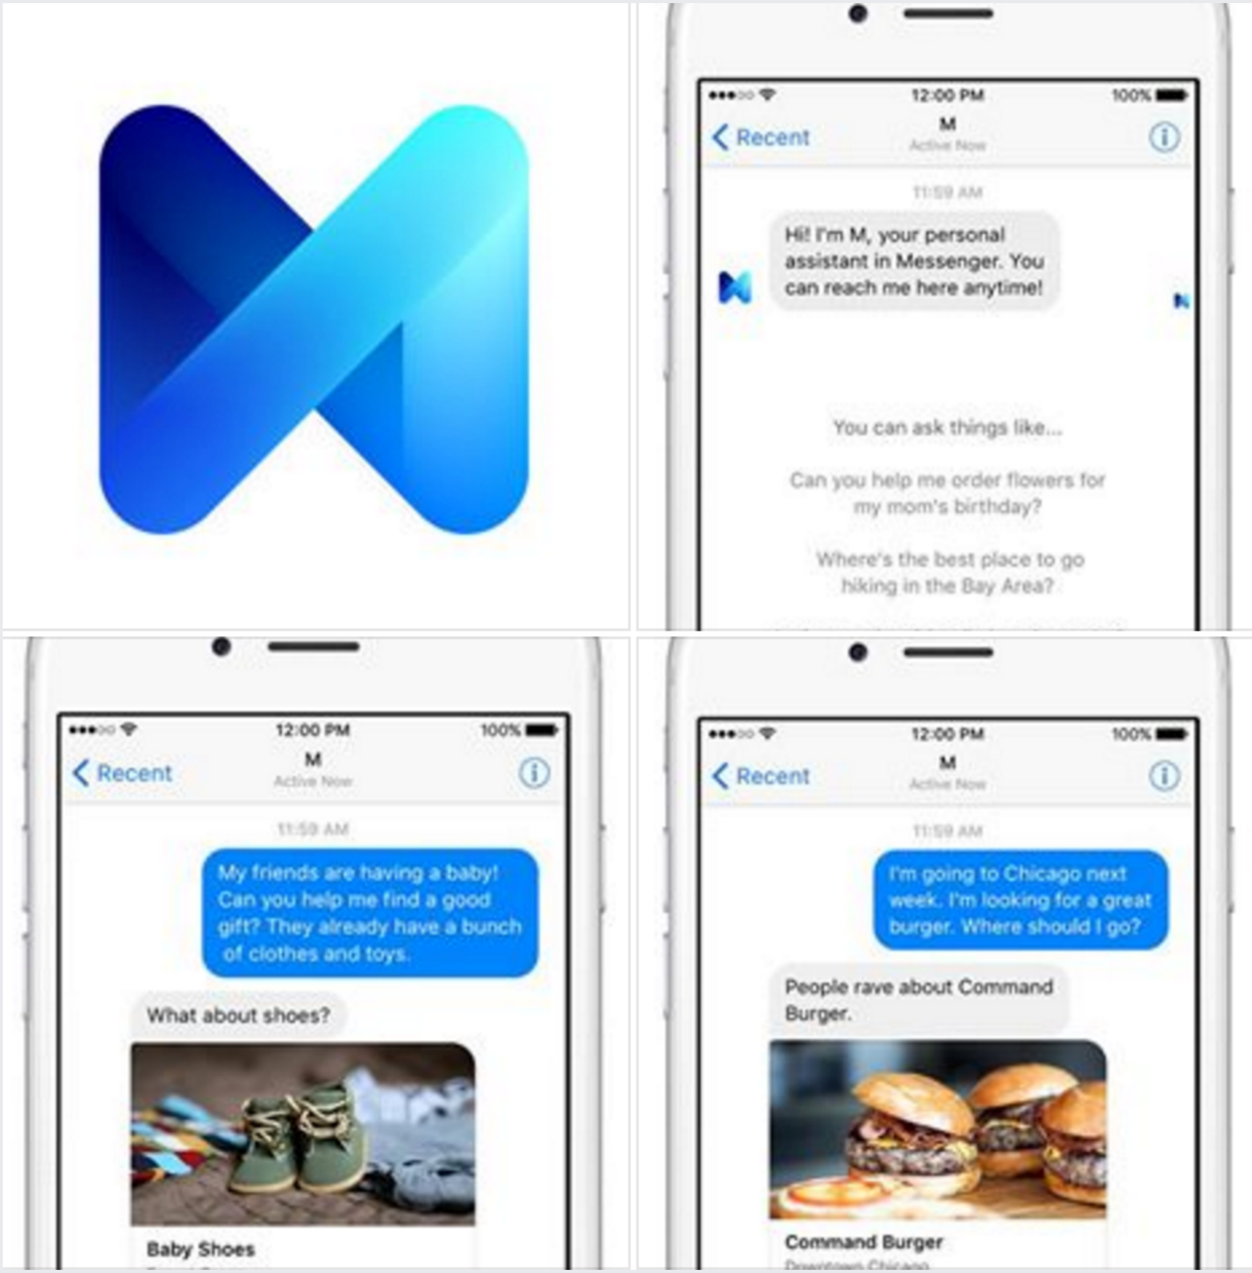

Measurement

- Not a problem with Machine Learning - we’re just measuring the wrong things.
- Ask users what the right content is then use that to inform our algorithms.
- Easier said than done.

Facebook Survey

2 of 11

How much do you agree with this statement?
This ad is high-quality.

Strongly disagree	Disagree	Neither agree nor disagree	Agree	Strongly agree

Women Get It Free · Website
Sponsored ·

Like Page

I couldn't believe what happened after signing up for Get It Free!

FREE Samples in the Mail

GET IT FREE – CLICK HERE

Like · Comment · Share · 2,771 635 448

Search Facebook

Home

News Feed

Messages 20+

Events 1

Update Status | Add Photos/Video | Create Photo Album

What's on your mind?

Backchannel
8 hrs ·

How 30 Random People in Knoxville May Change Your Facebook News Feed

In its never-ending quest to make its flagship feature more relevant, Zuckerberg and company are augmenting algorithms with direct questions.

29 **33rd birthday at ele...** Happening Now

TRENDING

- **Mission District, San Francisco**: 1 dead, several injured in 4-alarm fire in mixed-use building
- **Malaysia Airlines Flight 370**: Malaysia ends search for survivors, declares plane's loss an accident
- **Amal Clooney**: Lawyer points to robes, answers 'Ede & Ravenscroft' when asked about her court attire
- **Comcast**: Media company apologizes after customer's name appears as profanity on monthly bill
- **Yosemite National Park**: Rare Sierra Nevada red fox spotted in national park for 1st time in 100 years
- **Mongolia**: 200-year-old mummified remains of Buddhist monk found sitting in lotus position, report says
- **Gustine, Texas**: Students forced to lower pants for inspection after feces found on gym floor
- **UC Davis Medical Center**: California hospital confirms admission of suspected Ebola patient

1. Big Data
2. Machine Learning
3. Small Data
4. Profit

Learning Preferences

Behavioral economics to the rescue

- Economists have long studied how to estimate and model user preferences.
- Goes back to Bernoulli, Bentham trying to understand declining marginal value.
- Given a choice set, X, preferences are “well defined” if a preference ranking exists over those choices (e.g. no cycles, no incomparables).

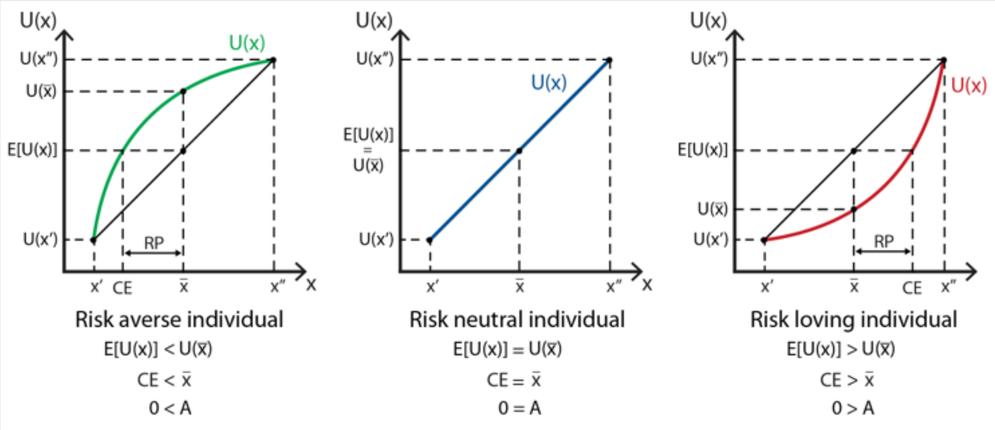

Psychophysics

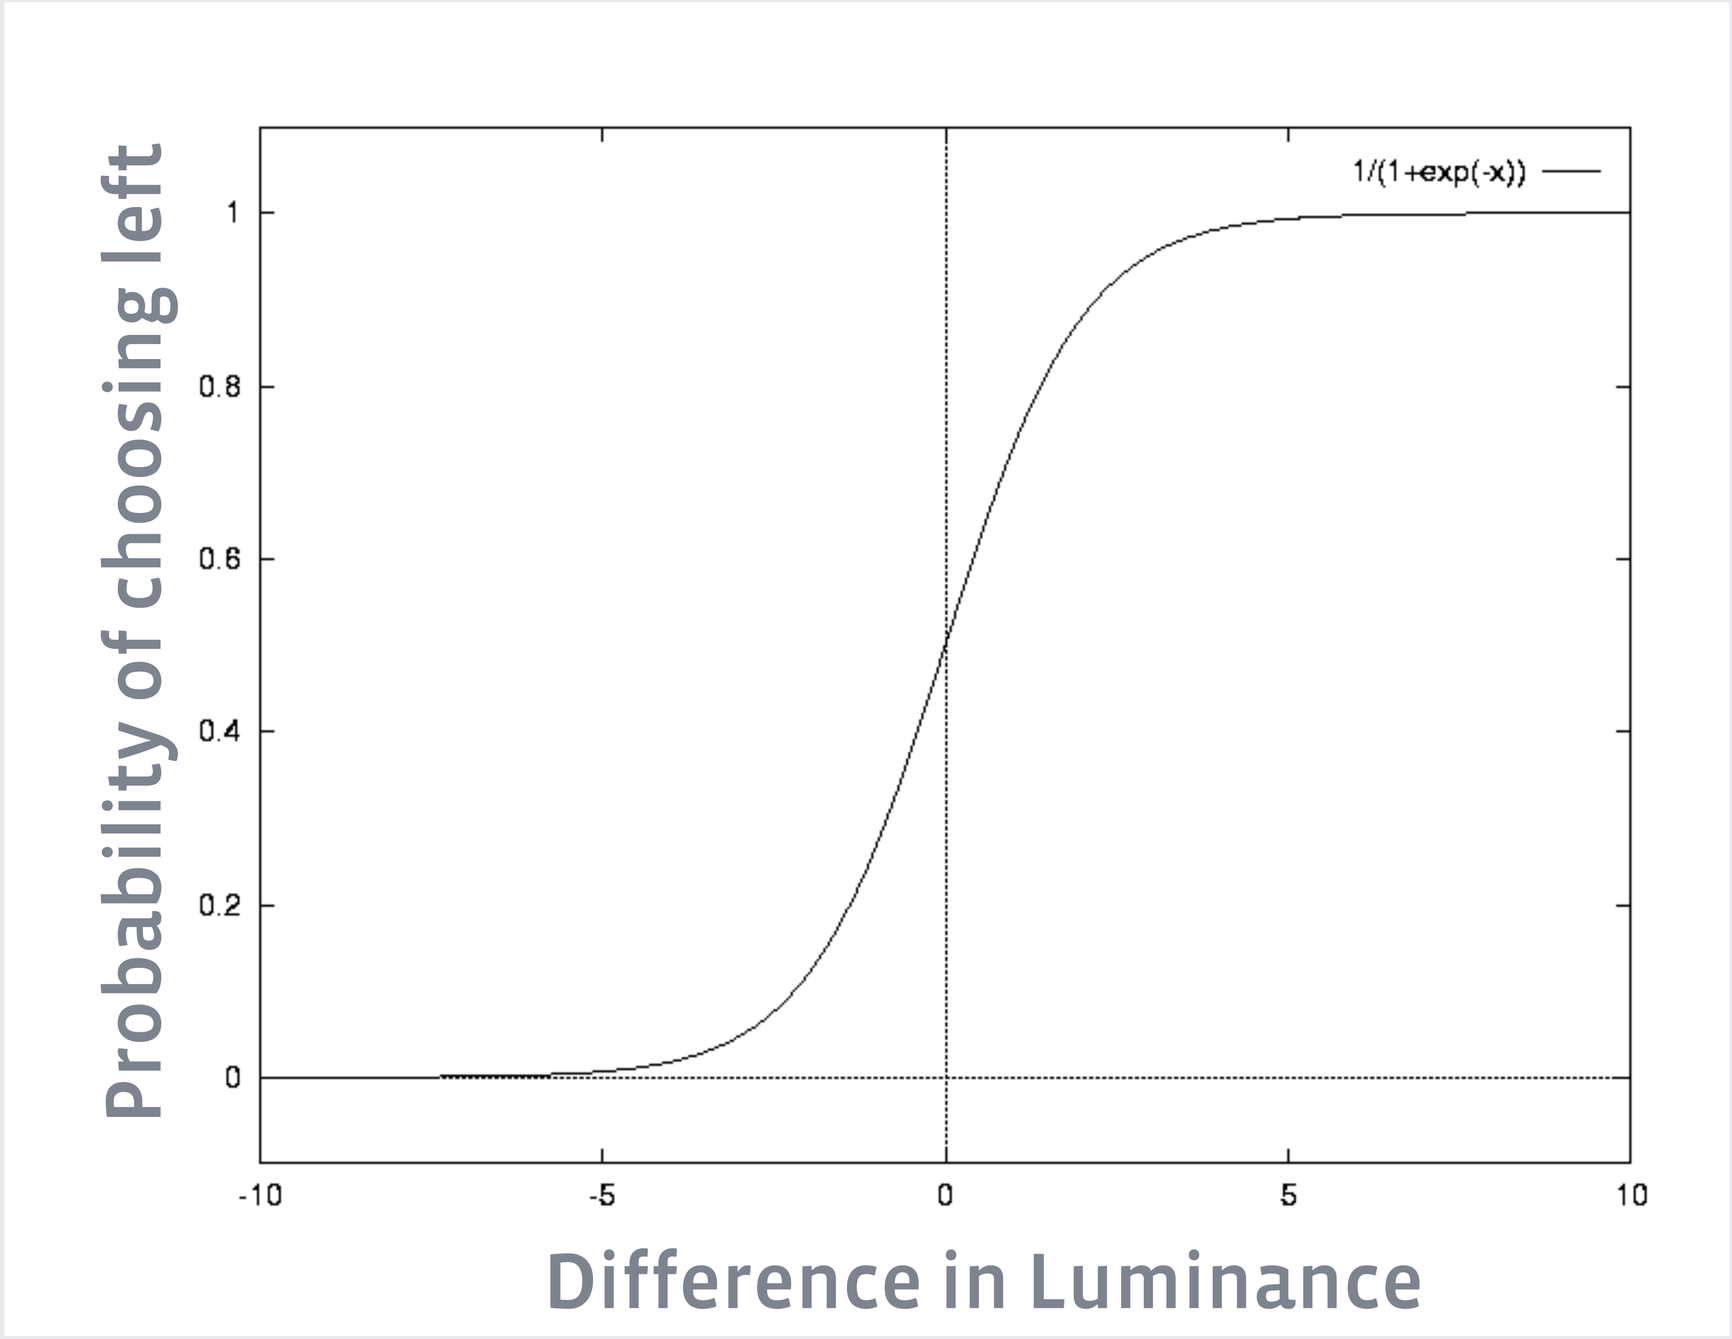
Probability of choosing left
1/(1+exp(-x))
1
0.8
0.6
0.4
0.2
0
-10
-5
0
5
10
Difference in Luminance

Random Utility Model

- Assume everyone has the same underlying utility function: $V(x_i) = \sum_k \beta^k x_i^k$
  - e.g., 5 x P(like) + 20 x P(comment) + 40 x P(share) + ...
- Each story *i* is perceived with noise: $U_i = V(x_i) + \epsilon_{ni}$
- User picks the story with higher utility

$$\mathbb{P}(\text{choose left}) = \mathbb{P}(U_l > U_r) = \mathbb{P}(V(x_l) - V(x_r) > \epsilon_{nr} - \epsilon_{nl})$$

- Estimate the weights using logistic regression

Random Utility Model

- Estimated weights are dimensionless (up to affine transformations)
- Coefficient ratios have well defined meaning (marginal rates of substitution)
- For ranking, we only care about the relative weights

Likert scale

- **Absolute Scale:** a “5” beats all “4”s:
  - For n comparisons you can build (n^2 - n)/2 pairwise comparisons.
  - Requires users to apply same scale across all comparisons.
- **Sensitive to framing**:
  - Question and labels matter!
- **Need lots of guidance to give consistent answers.**

Pairwise Comparisons

- **Relative Scale:**
  - Can’t determine strength of preference.
  - More deeply linked with choice task.
- **Task can be difficult:**
  - A “skip” option is often used.
- **Possibly intransitive**:
  - A > B > C > A

Improve Your News Feed

Select the story you would most want to see in your News Feed

2 of 7

Can't decide

Ta Virot Chiraphadhanakul
January 25 · Instagram ·

Like · Comment · Share

Write a comment...

Alex Peysakhovich
49 mins ·

Friendship can overcome any distance.

Friendship Around the World
True stories and stats about friendship.
WWW.FACEBOOKSTORIES.COM/FRIENDSHIP

Like · Comment · Share

22 people like this.

Write a comment...

We'd like to improve Facebook for your friends
We'd like to hear your feedback! Please answer a few questions about your News Feed.

Give Feedback

TRENDING

- **Surf City, North Carolina**: Police Chief Retires After Coming Under Scrutiny for 'Black Lives Matter' Post
- **#DF15**: Salesforce Holds Annual Dreamforce Conference in San Francisco Through Sept. 18
- **Disney Princesses**: Artist Shares Drawings of Disney Characters as Pin-Up Models on Instagram

See More

Random Utility

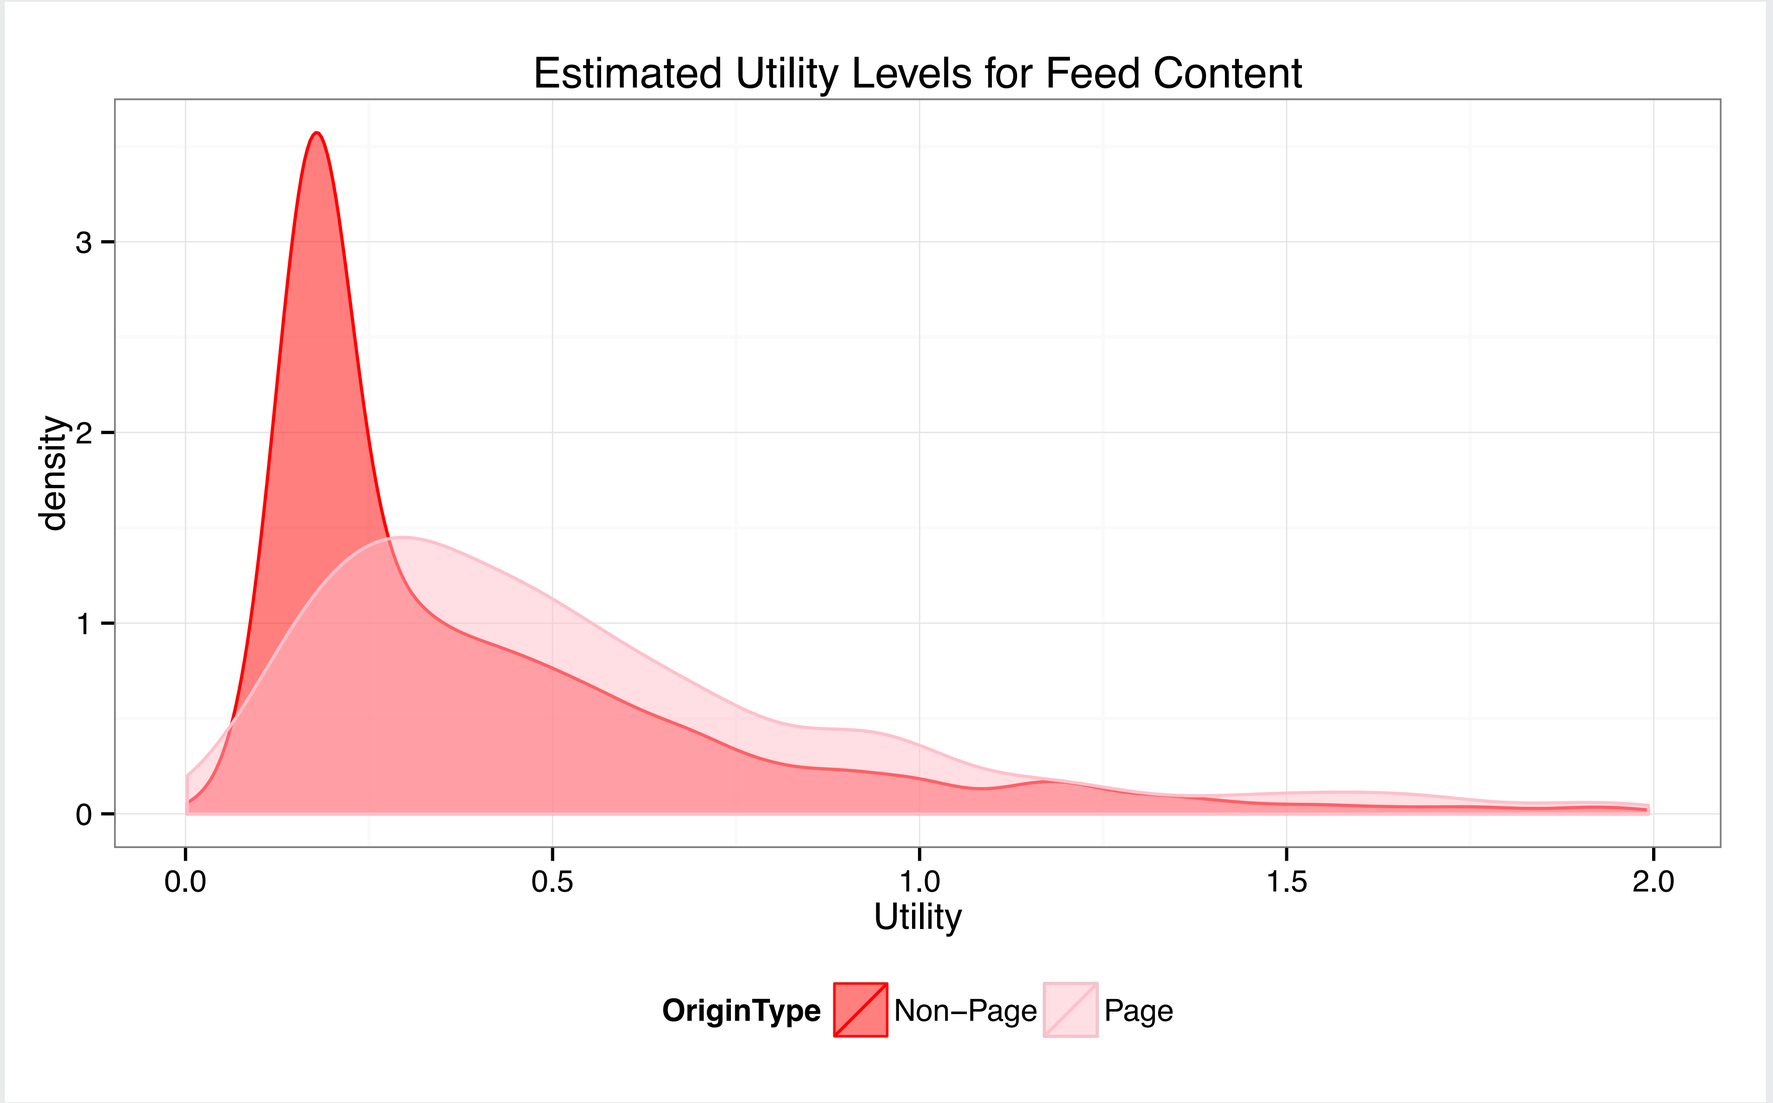
Estimated Utility Levels for Feed Content
density
3
2
1
0
0.0
0.5
1.0
1.5
2.0
Utility
OriginType
Non−Page
Page

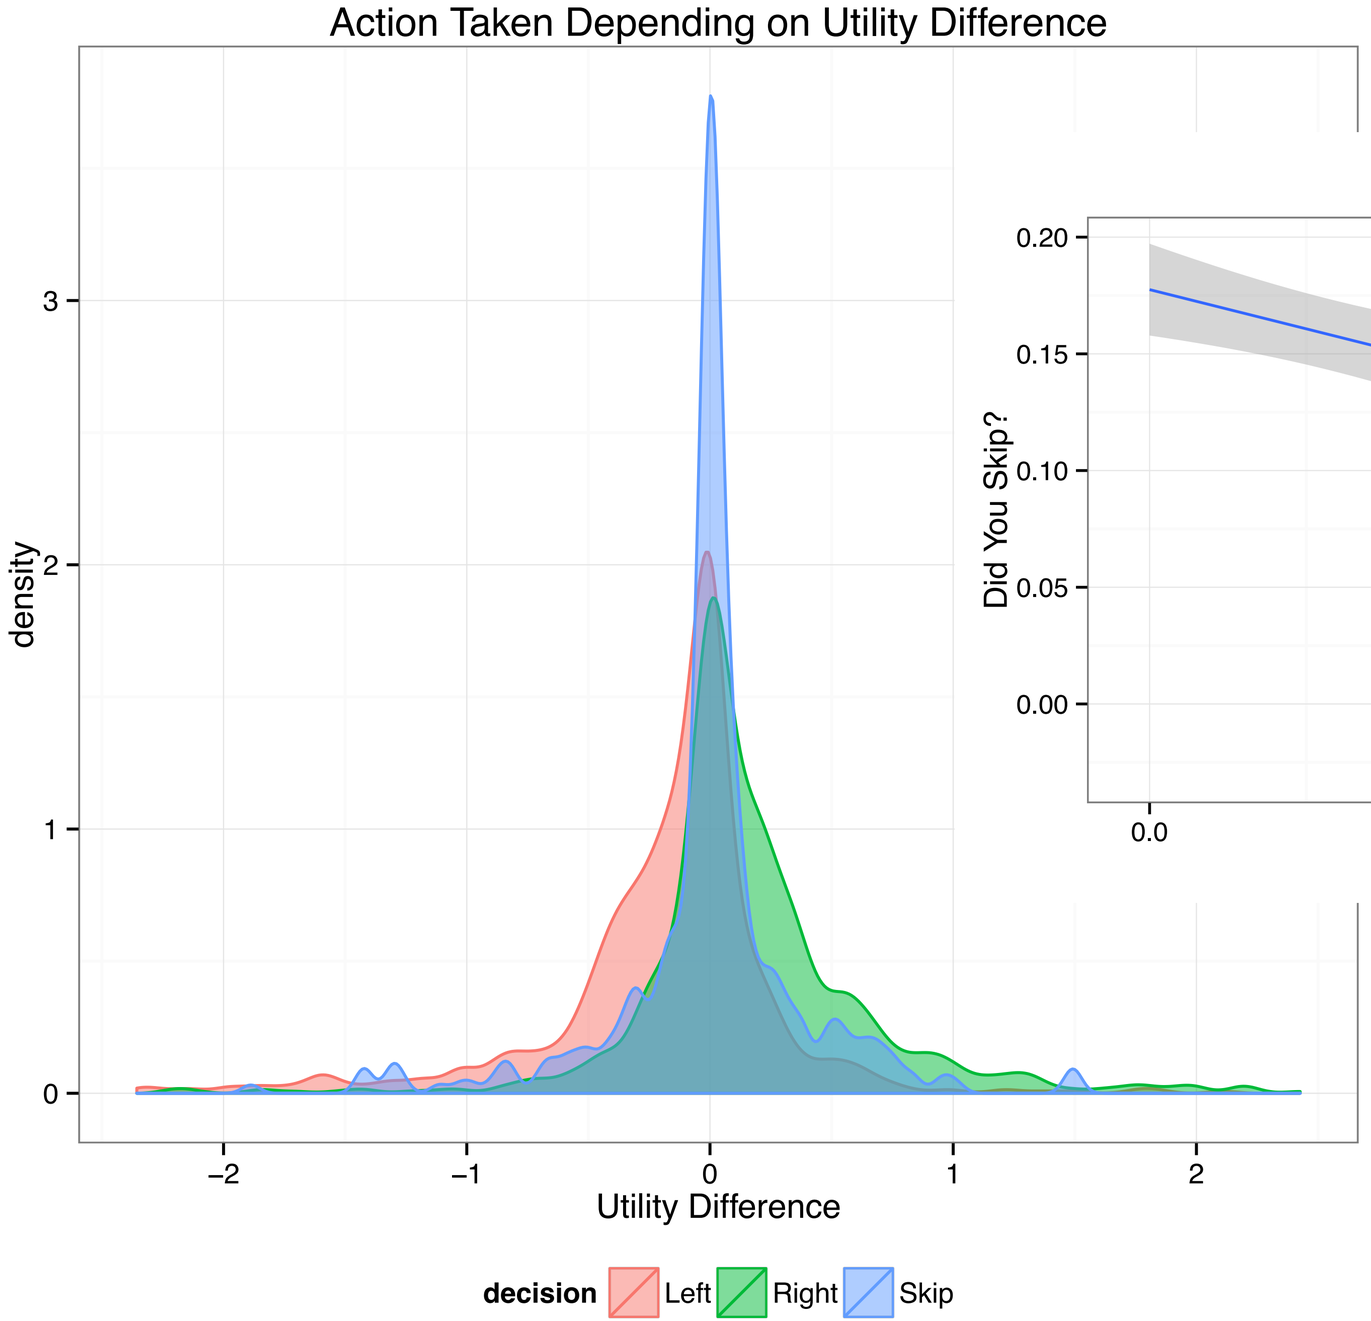
Action Taken Depending on Utility Difference
density
0
1
2
3
−2
−1
0
1
2
Utility Difference
decision
Left
Right
Skip

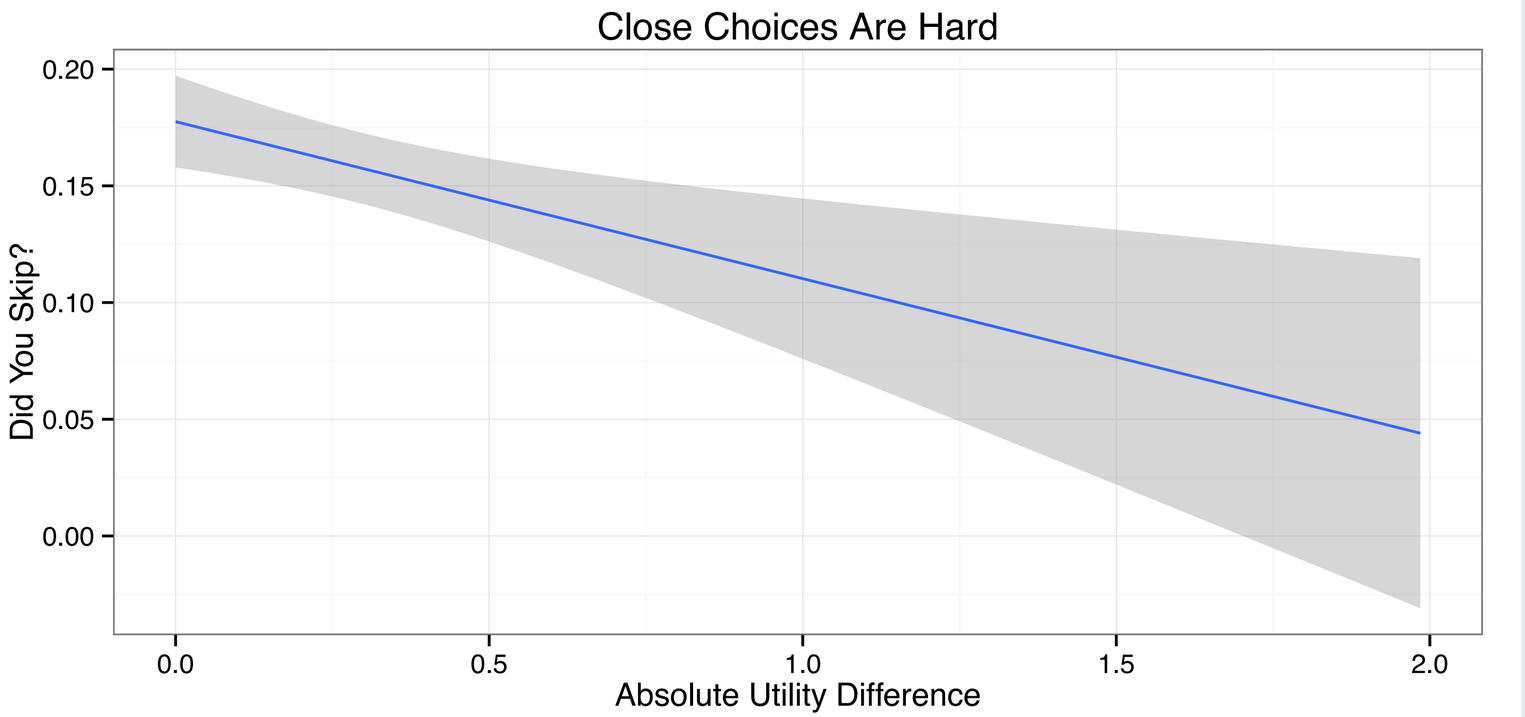
Close Choices Are Hard
Did You Skip?
0.00
0.05
0.10
0.15
0.20
0.0
0.5
1.0
1.5
2.0
Absolute Utility Difference

Scale vs. Pairwise

(this section not presented during class)

Findings:

- Likert scale is more efficient in eliciting user preferences under ideal conditions, however in most realistic conditions pairwise comparisons are preferred because of:
    - **differential item functionality** (i.e., the fact that different scale points may mean different things to different people) or
    - **low quality inputs** (e.g., lack of attention or understanding by survey participants or noisily measured input features)
- In practice: PC has been more effective for our objectives - task is easier and data is cleaner.

Approach:

- Generate n users, f features, and degree of heterogeneity in preferences (designated b).
- For each user, simulate 10 survey choices from a likert scale task and a PC task.
- We then determine how well we can uncover user preferences:
  - How much error is there in estimated marginal rates of substitution?
  - How well does our model recover the true ranking function? (transposition score)
- We then introduce DIF, rater noise and measure how well each holds up.

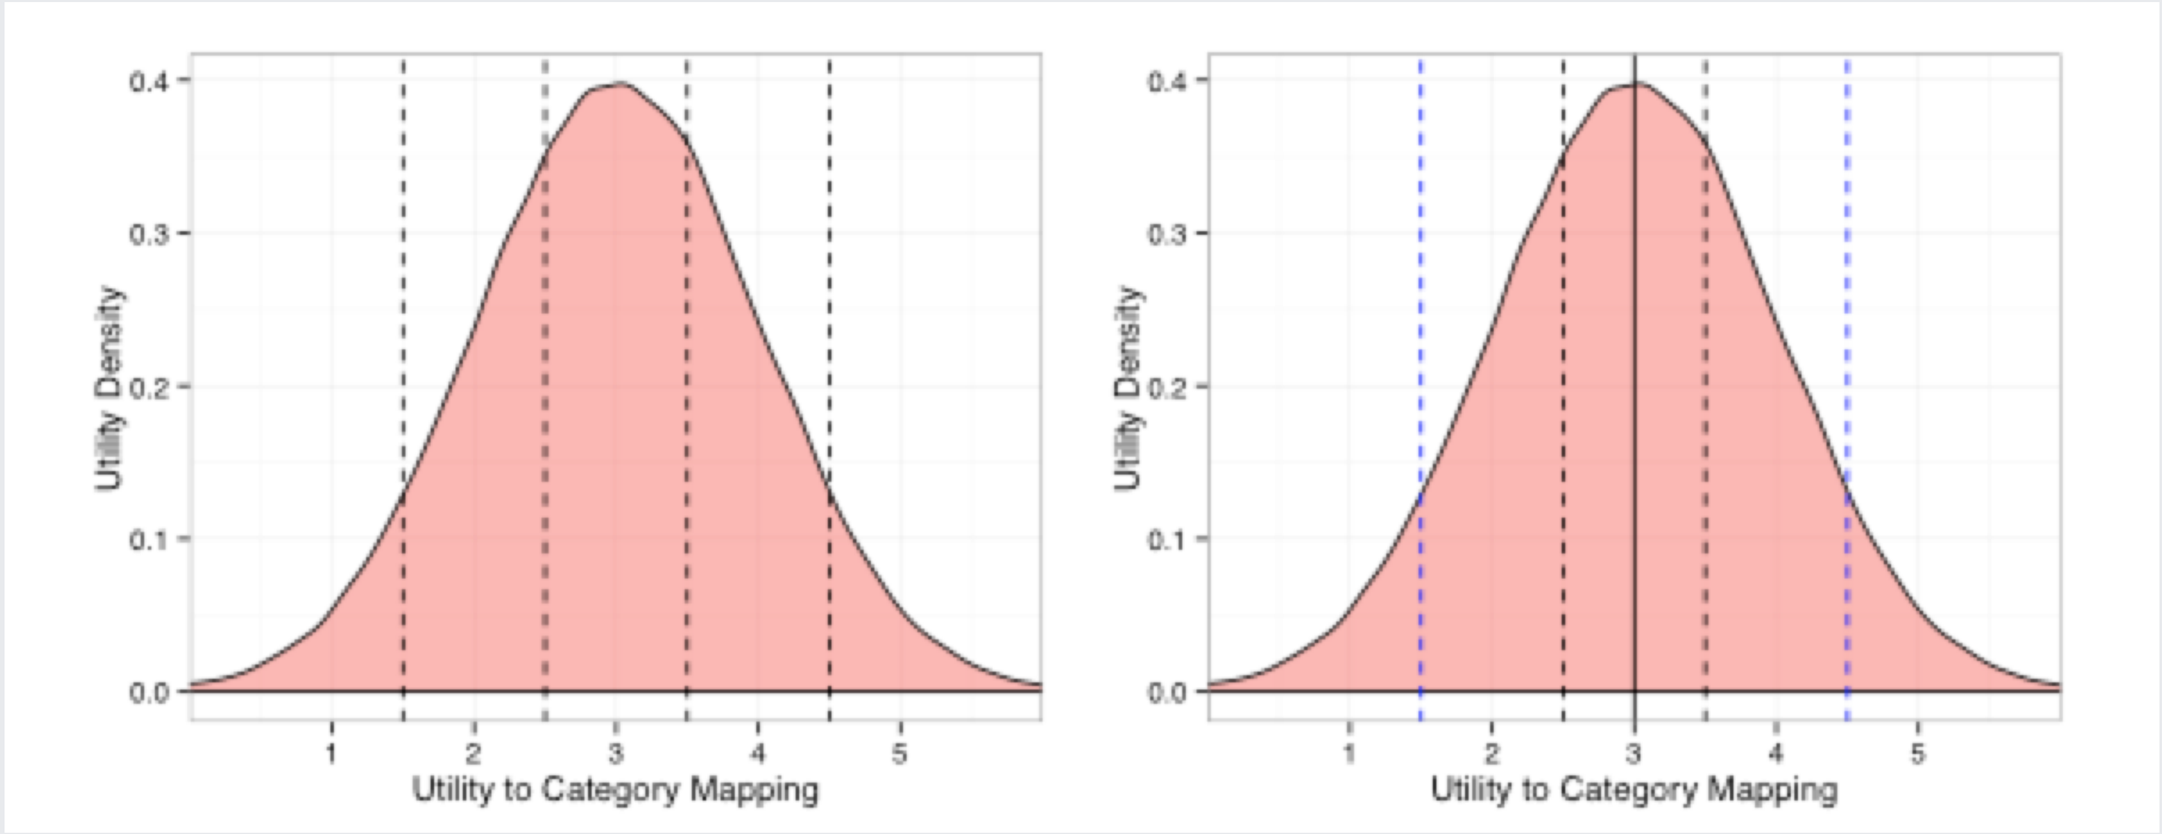
0.4
0.3
0.2
0.1
0.0
Utility Density
1
2
3
4
5
Utility to Category Mapping
0.4
0.3
0.2
0.1
0.0
Utility Density
1
2
3
4
5
Utility to Category Mapping

Likert more efficient

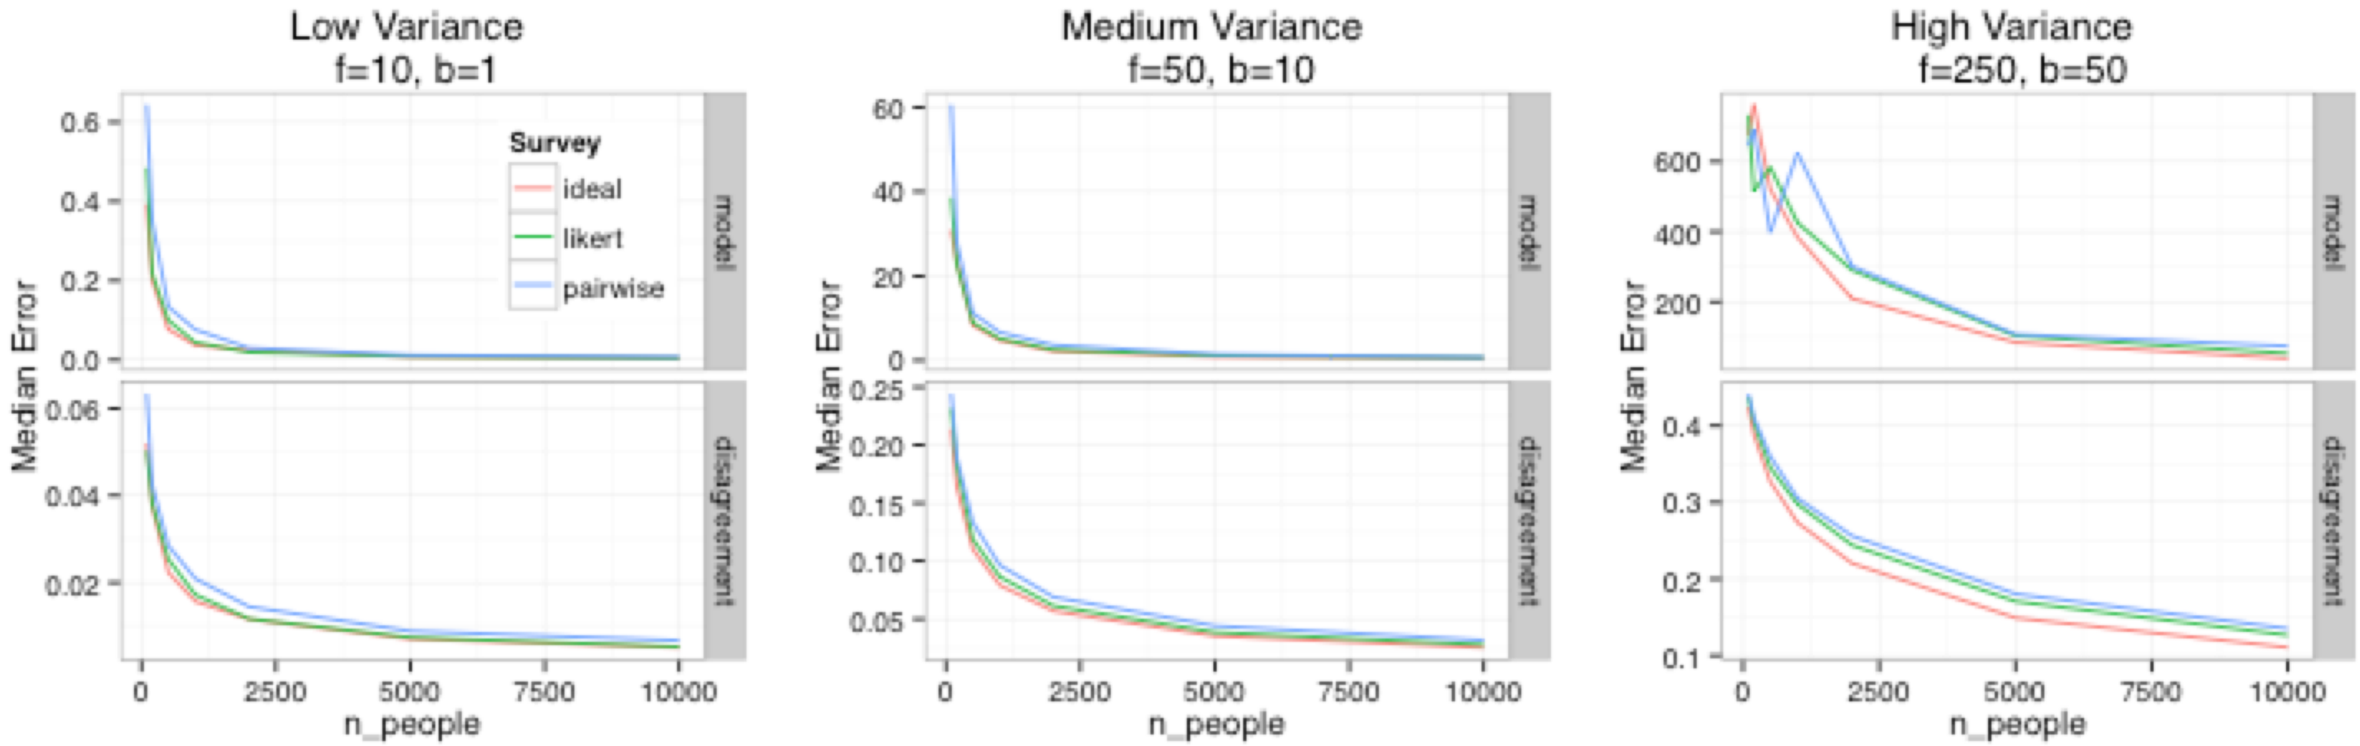

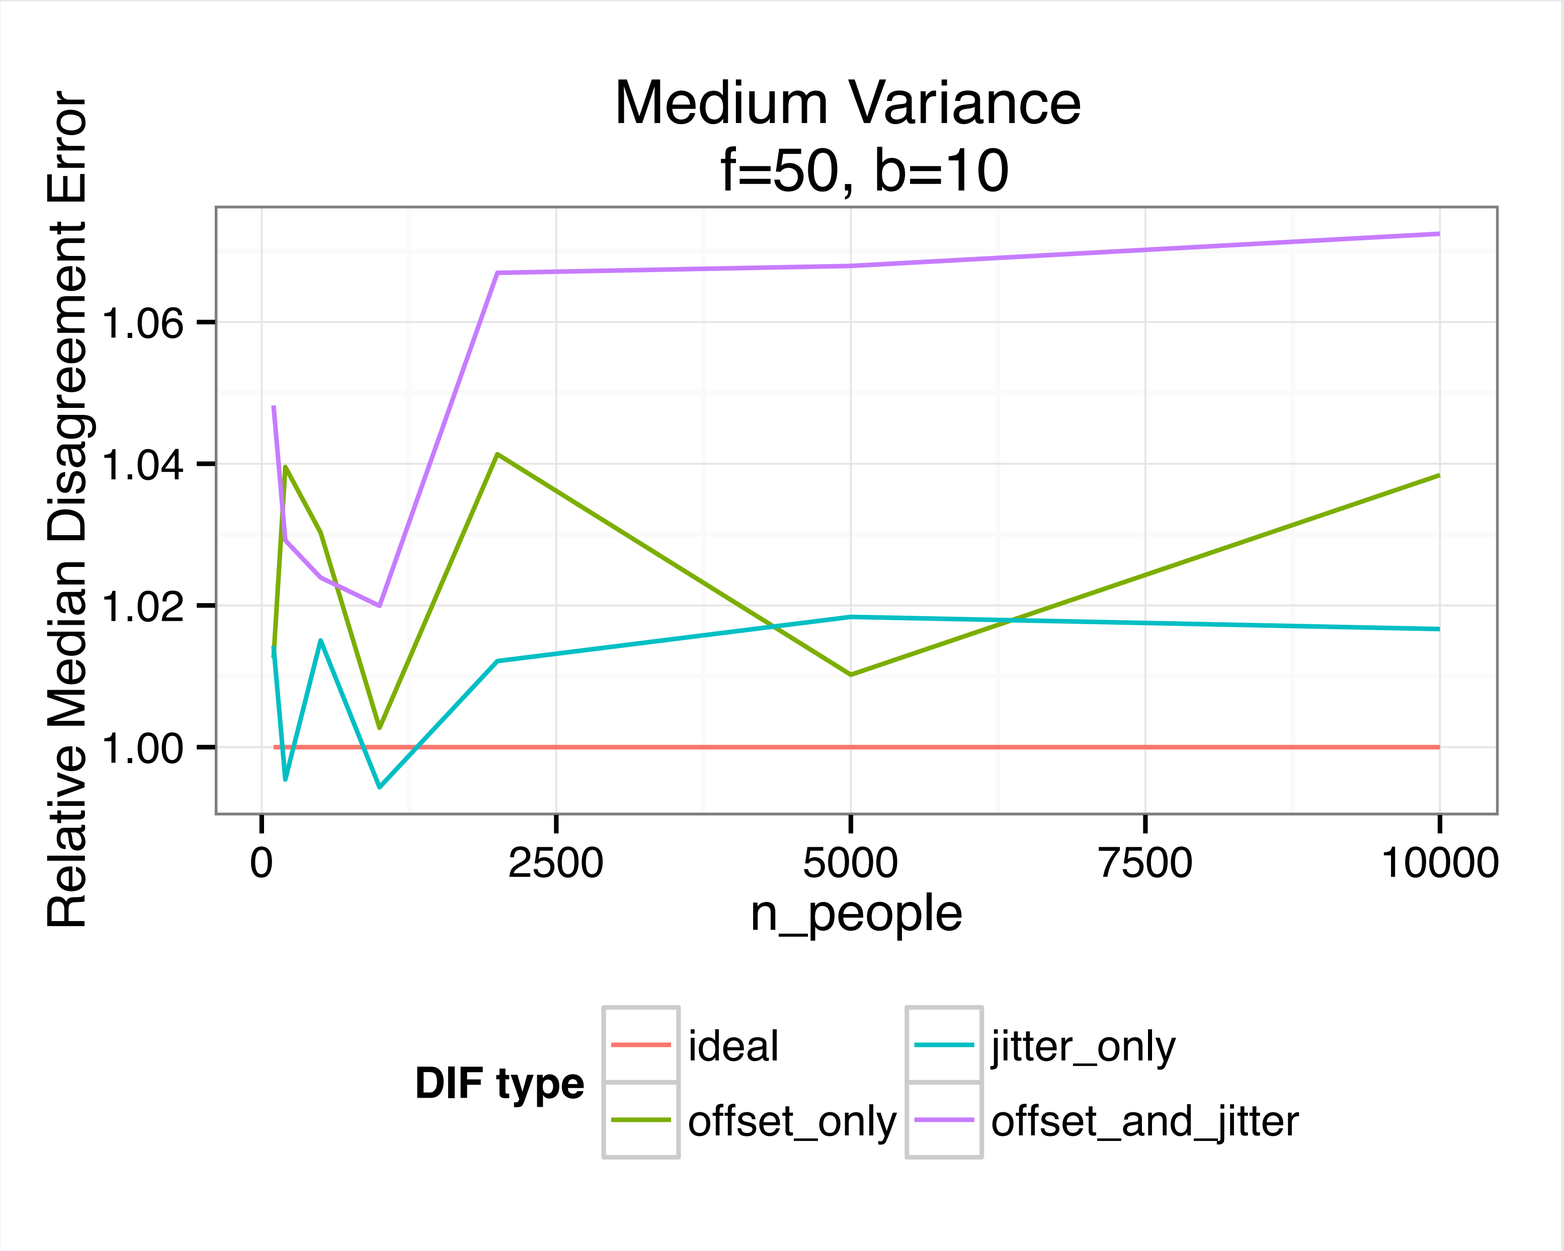

Adding DIF decreases the accuracy of likest-elicited ranking functions

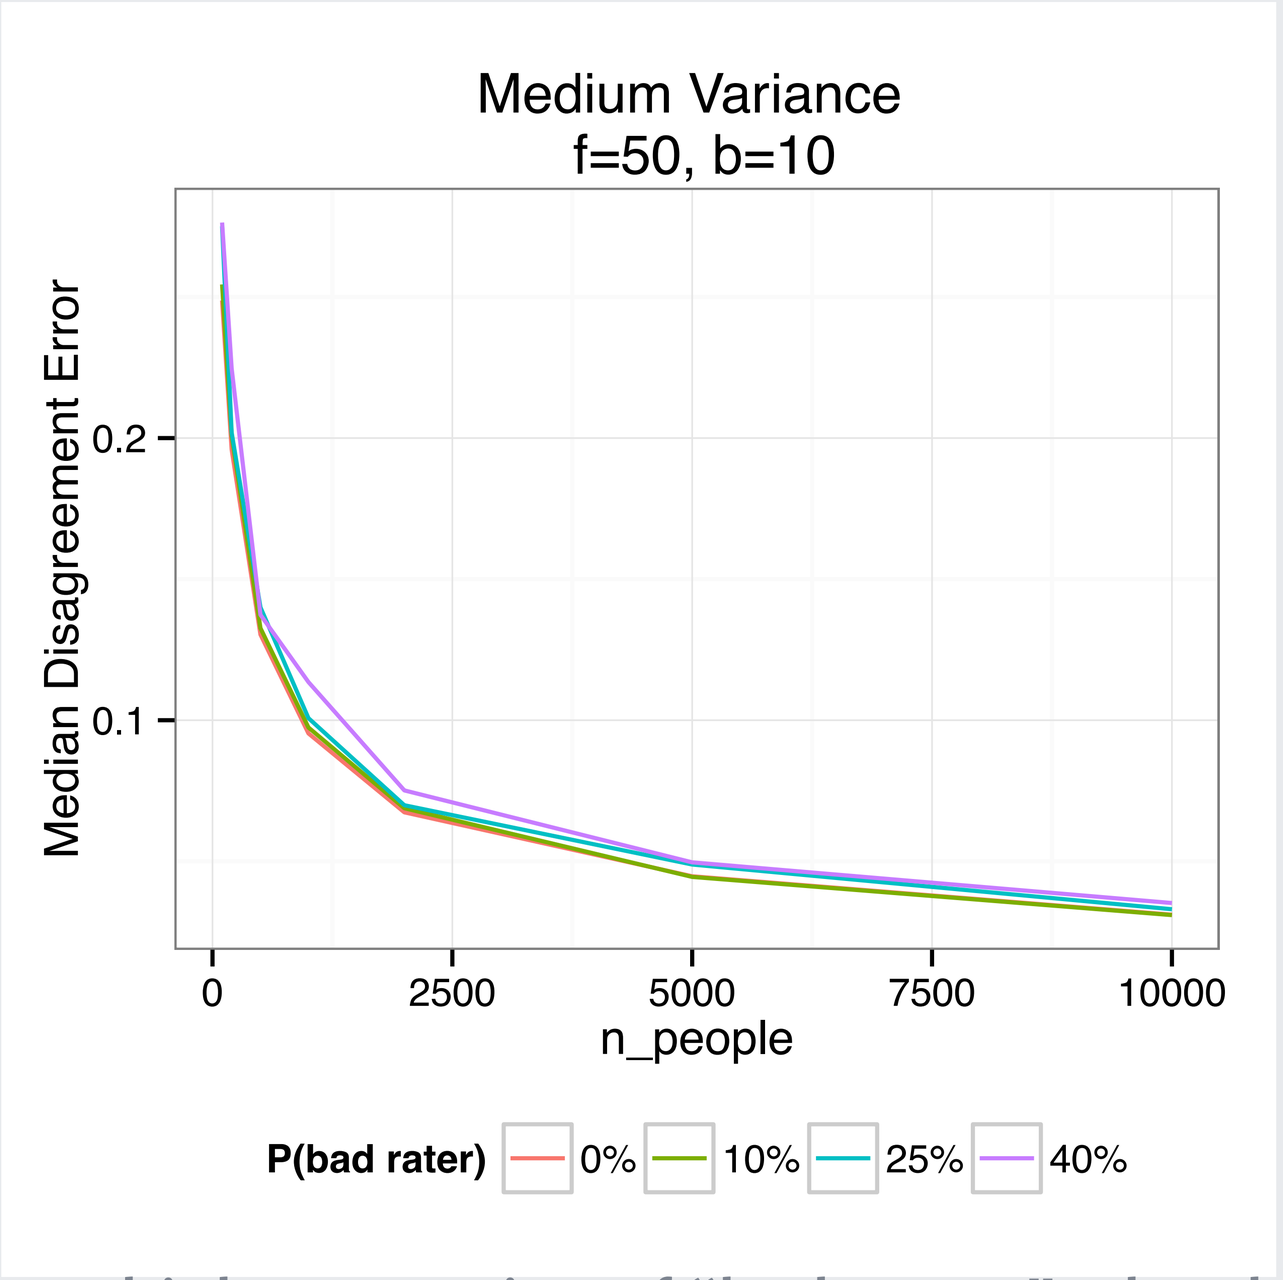

PC is robust to a high proportion of “bad raters” who choose randomly

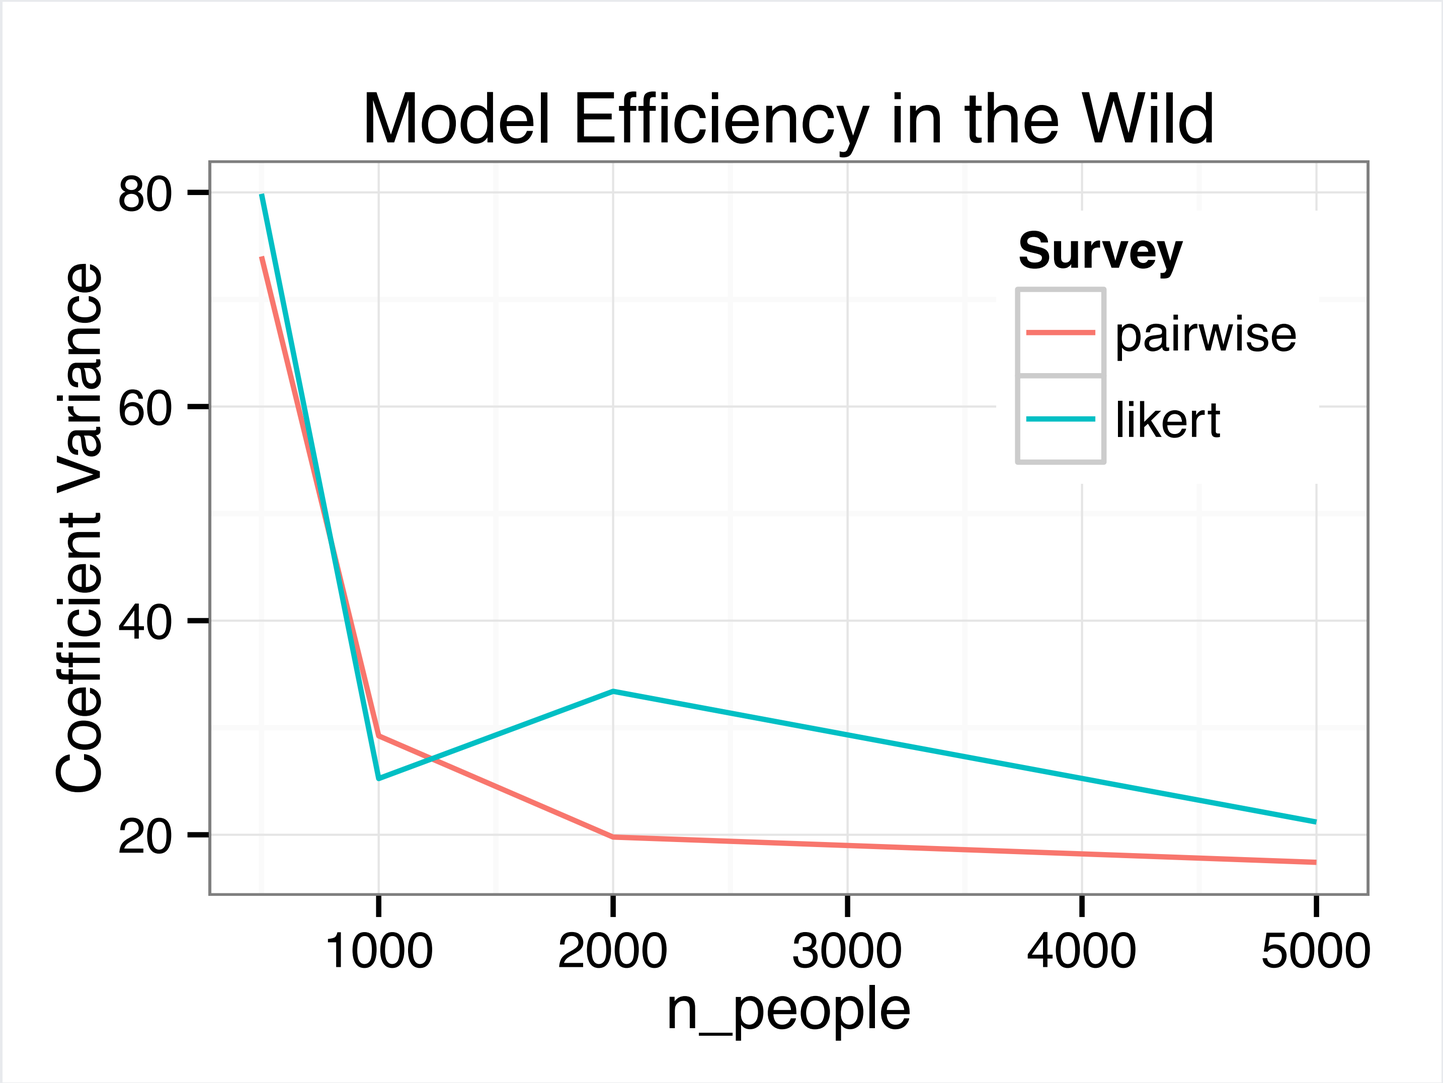
Model Efficiency in the Wild
Coefficient Variance
80
60
40
20
1000
2000
3000
4000
5000
n_people
Survey
pairwise
likert

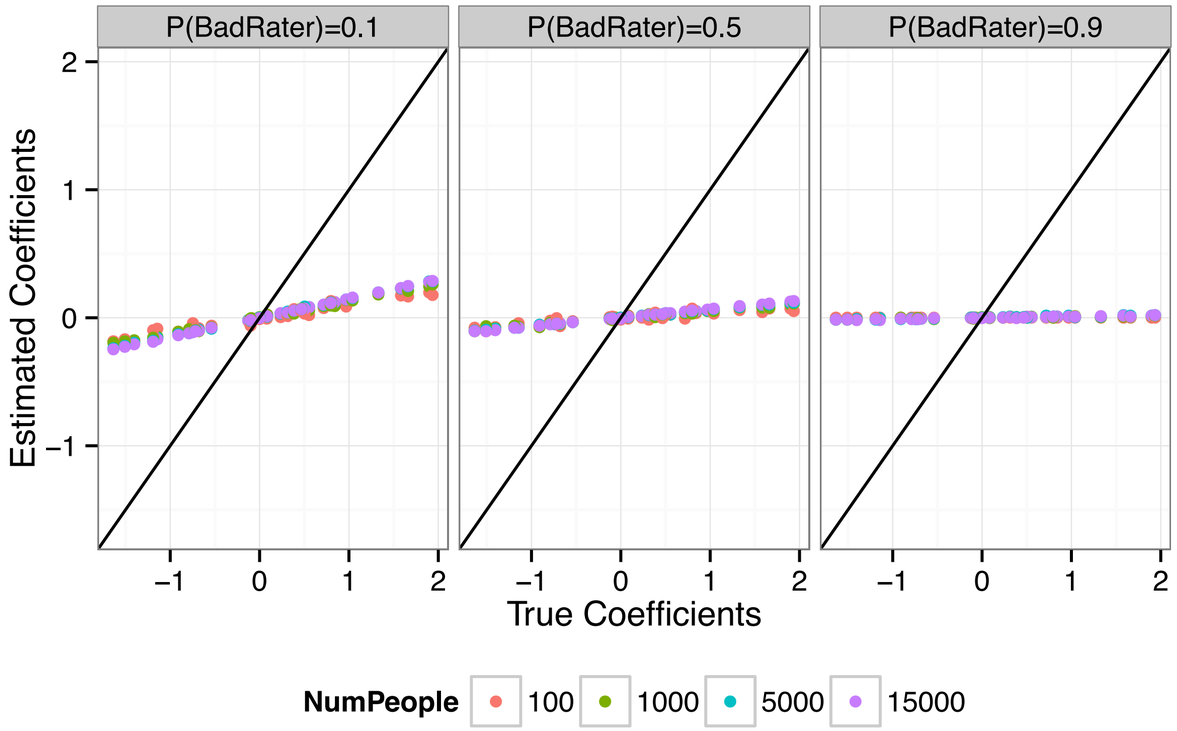
P(BadRater)=0.1
P(BadRater)=0.5
P(BadRater)=0.9
Estimated Coefficients
True Coefficients
NumPeople
100
1000
5000
15000

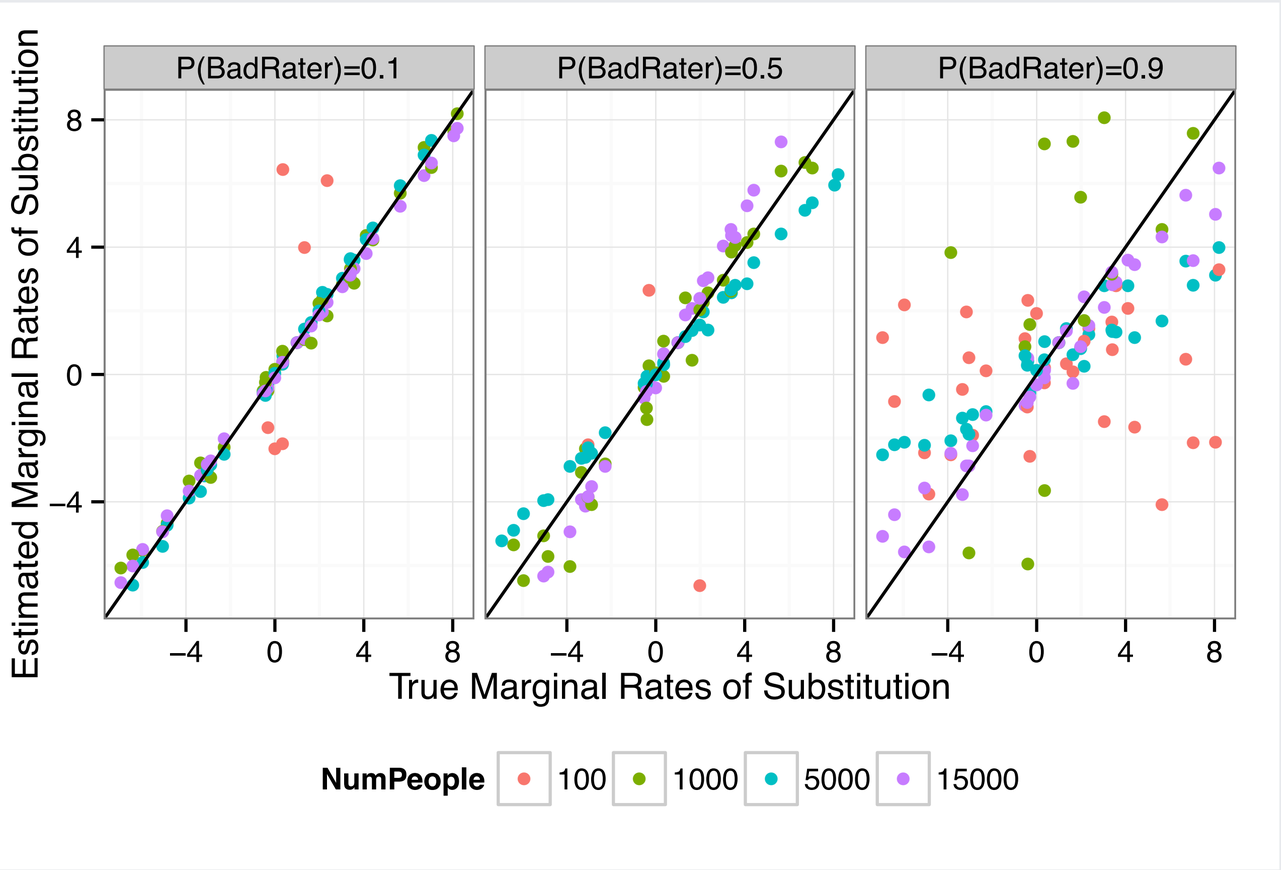
P(BadRater)=0.1
P(BadRater)=0.5
P(BadRater)=0.9
Estimated Marginal Rates of Substitution
True Marginal Rates of Substitution
NumPeople
100
1000
5000
15000

facebook