Recommender systems: biased advice and its nature

Emilio Calvano - U. Bologna & CSEF
Bruno Jullien - Toulouse School of Economics

IDEE Meeting
Bologna - October 2016
In progress!!

Recommender Systems defined

- Recommender Systems are software tools and techniques providing suggestions for items to be of use to a consumer.

The recommender problem

Estimating a utility function to predict how a user will like an item based on, past behavior, relation to other users, Item similarity, Context...

Xavier Amatriain
(Engineer - director of algorithm engineering at Netflix)

On Amazon...

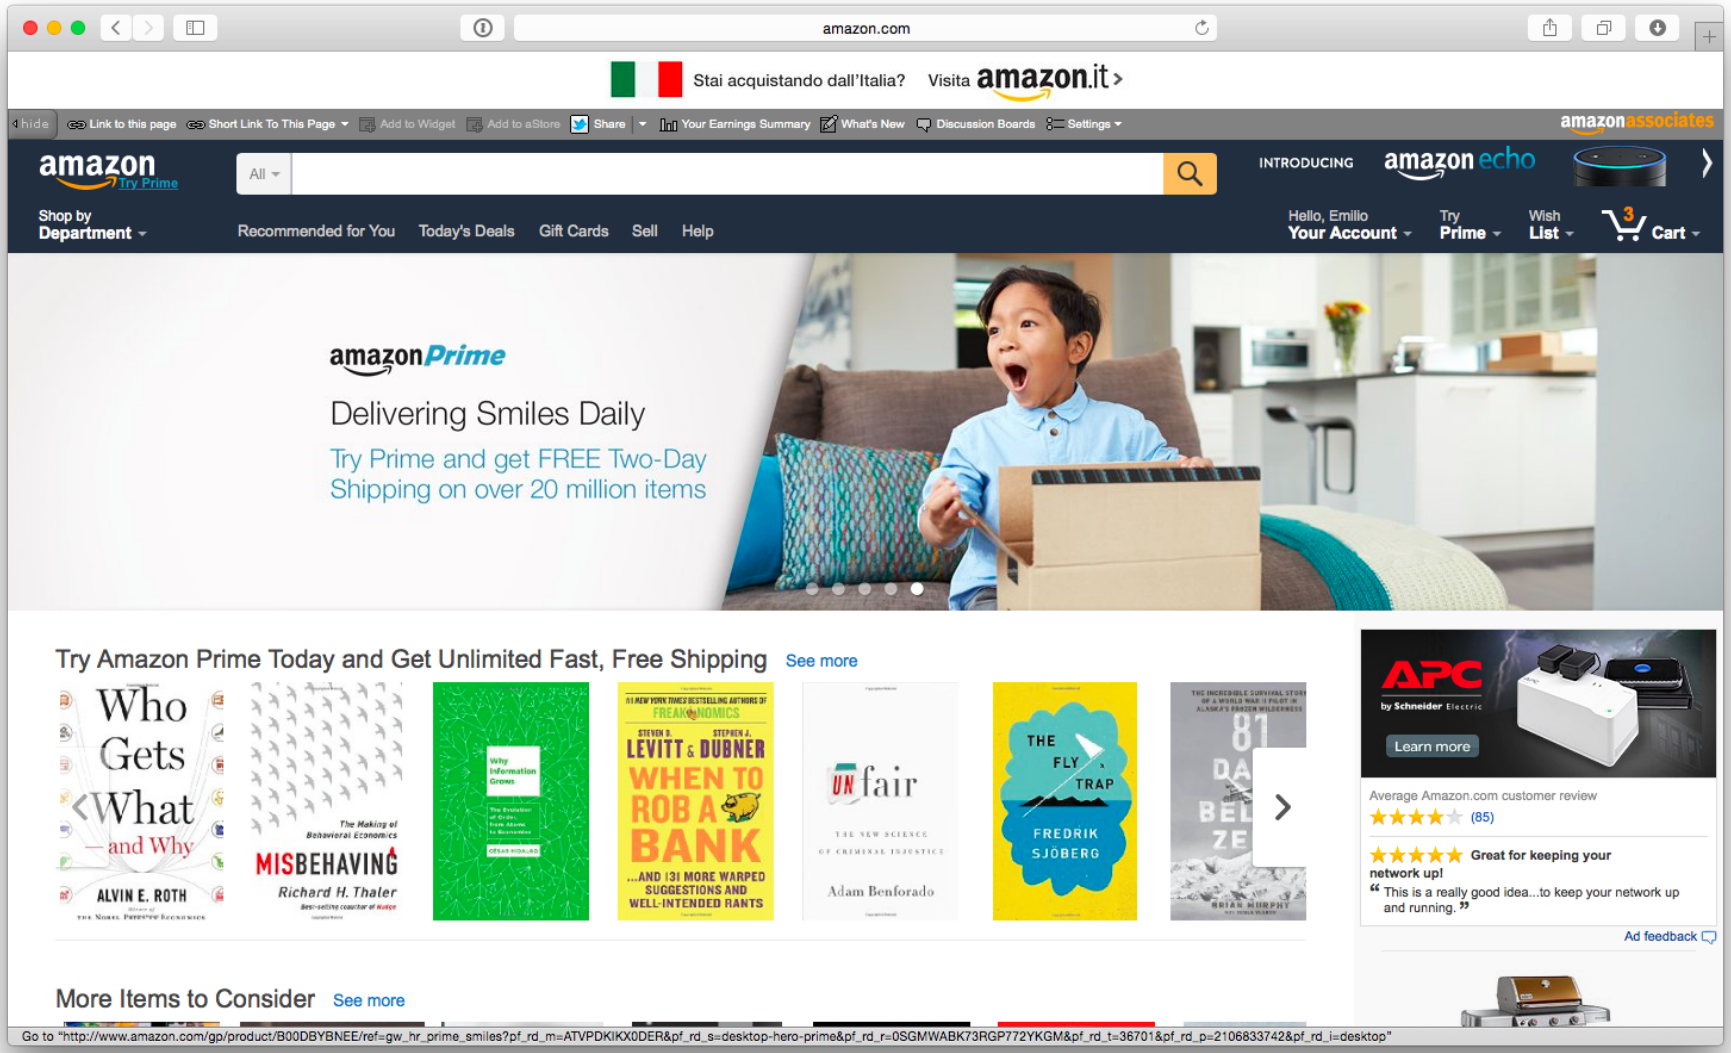
amazon.com
Stai acquistando dall'Italia? Visita amazon.it
amazon Try Prime
Shop by Department
Recommended for You
Today's Deals
Gift Cards
Sell
Help
INTRODUCING amazon echo
Hello, Emilio Your Account
Try Prime
Wish List
Cart
amazonPrime
Delivering Smiles Daily
Try Prime and get FREE Two-Day Shipping on over 20 million items
Try Amazon Prime Today and Get Unlimited Fast, Free Shipping
See more
Who Gets What — and Why
ALVIN E. ROTH
MISBEHAVING
Richard H. Thaler
LEVITT & DUBNER
WHEN TO ROB A BANK
...AND 131 MORE WARPED SUGGESTIONS AND WELL-INTENDED RANTS
unfair
Adam Benforado
THE FLY TRAP
FREDRIK SJÖBERG
APC by Schneider Electric
Learn more
Average Amazon.com customer review
(85)
Great for keeping your network up!
"This is a really good idea...to keep your network up and running."
Ad feedback
More Items to Consider
See more
Go to "http://www.amazon.com/gp/product/B00DBYBNEE/ref=gw_hr_prime_smiles?pf_rd_m=ATVPDKIKX0DER&pf_rd_s=desktop-hero-prime&pf_rd_r=0SGMWABK73RGP772YKGM&pf_rd_t=36701&pf_rd_p=2106833742&pf_rd_i=desktop"

On spotify...

Spotify®

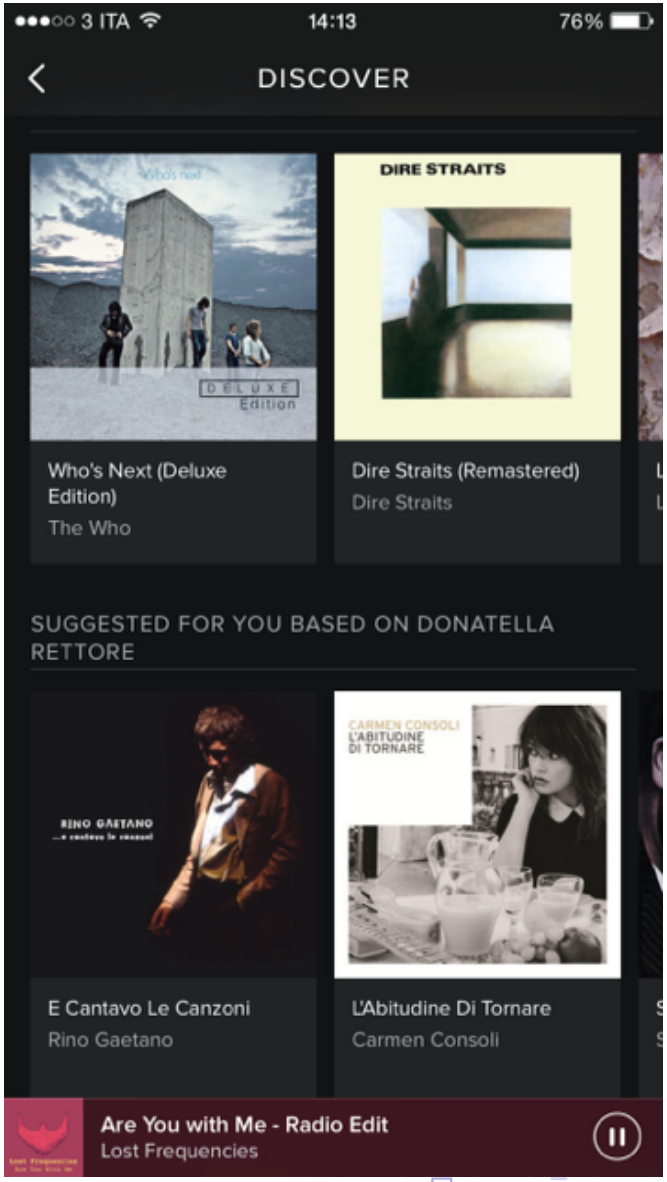
3 ITA
14:13
76%
DISCOVER
DIRE STRAITS
Who's Next (Deluxe Edition)
The Who
Dire Straits (Remastered)
Dire Straits
SUGGESTED FOR YOU BASED ON DONATELLA RETTORE
RINO GAETANO
CARMEN CONSOLI
L'ABITUDINE DI TORNARE
E Cantavo Le Canzoni
Rino Gaetano
L'Abitudine Di Tornare
Carmen Consoli
Are You with Me - Radio Edit
Lost Frequencies

On Netflix...

- Hugely popular DVD rental company (now mostly streaming).
- 50M subscribers.
- 7B hours/quarter.
- 90 minutes a day (Avg)
- 150 choices (clicks) a day.

- Hugely powerful and popular tools in electronic commerce.
- Address 'long tail' / 'abundance of choice' problem
- Allow keep users engaged: inducing repeated purchases / renewed subscriptions
- (Anecdotal/Observational) Evidence that they shape consumption choices.
- **Key challenge: build consumer trust in the system**
- When to *not* recommend?

Questions

- Do firms systematically distort recommendations?
- If yes, in which direction?
- Which factors are most conducive to distortion?

Questions

- Do firms systematically distort recommendations?
- If yes, in which direction?
- Which factors are most conducive to distortion?

- 1 recommender, 1 consumer, set of items
- Recommender has no preferences over consumption choices

Questions

- Do firms systematically distort recommendations?
- If yes, in which direction?
- Which factors are most conducive to distortion?

- 1 recommender, 1 consumer, set of items
- Recommender has no preferences over consumption choices

1. Theory: distorted recommendations due to reputational incentives
2. Lab evidence

Model

- 1 Consumer w/ known reservation utility m.
- 1 risky prospect, say a movie. Mean $\mu = m$.
- 1 Recommender with private signal on C's utility from movie
- R should -recommend- or -not- the movie
- C always follows recommendation
- (Alternative interpretation 'niche' vs 'mass' movie)

Model

- 1 Consumer w/ known reservation utility m.
- 1 risky prospect, say a movie. Mean $\mu = m$.
- 1 Recommender with private signal on C's utility from movie
- R should -recommend- or -not- the movie
- C always follows recommendation
- (Alternative interpretation 'niche' vs 'mass' movie)

- R can be either high or low accuracy, Pr(high)$= p_0$
- Accuracy is not directly observable
- C infers accuracy from experience $u \rightarrow p_1$

Model

- 1 Consumer w/ known reservation utility m.
- 1 risky prospect, say a movie. Mean $\mu = m$.
- 1 Recommender with private signal on C's utility from movie
- R should -recommend- or -not- the movie
- C always follows recommendation
- (Alternative interpretation 'niche' vs 'mass' movie)

- R can be either high or low accuracy, Pr(high)$= p_0$
- Accuracy is not directly observable
- C infers accuracy from experience $u \rightarrow p_1$

- R's payoff $\pi(p_1)$ increasing in p_1

Model

- 1 Consumer w/ known reservation utility m.
- 1 risky prospect, say a movie. Mean $\mu = m$.
- 1 Recommender with private signal on C's utility from movie
- R should -recommend- or -not- the movie
- C always follows recommendation
- (Alternative interpretation 'niche' vs 'mass' movie)

- R can be either high or low accuracy, Pr(high)$= p_0$
- Accuracy is not directly observable
- C infers accuracy from experience $u \to p_1$

- R's payoff $\pi(p_1)$ increasing in p_1

- Low type wants to persuade C that he is high. Focus on his incentive and fix the strategy of the high type.

Can 'hedonic' recommendations be an EQM?

- Suppose R observes a binary signal: $E[u|info] > m$ or $E[u|info] < m$
- Hedonic = do 'right thing:' -recommend- iff $E[u|info] > m$

Can 'hedonic' recommendations be an EQM?

- Suppose R observes a binary signal: $E[u|info] > m$ or $E[u|info] < m$
- Hedonic = do 'right thing:' -recommend- iff $E[u|info] > m$

Key Obs - Movies are informative experiments

High R more likely to 'deliver' $u > m$ conditional on -recommending- $\Rightarrow$ from C's perspective 'Movie' is an informative experiment about type

$$\begin{cases} p_1 > p_0 & \text{if movie is good } (u > m) \\ p_1 < p_0 & \text{if movie is bad } (u < m) \end{cases}$$

Movies as 'risky bets'

From R's perspective rec. a movie is a 'risky reputational bet'
Not recommending preserves the status quo: $p_1 = p_0$

Strategic incentive to be cautious

- Bayesian updating implies that the expected posterior moves in the direction of the **true** state:

$$E[p_1|\text{low}] < p_0$$

$\Rightarrow$ **R never recommends to watch**

Strategic incentive to be cautious

- Bayesian updating implies that the expected posterior moves in the direction of the **true** state:

$$E[p_1|\text{low}] < p_0$$

$\Rightarrow$ **R never recommends to watch**

Proposition

R has a strategic incentive to be cautious (SIC) if $\pi(p_1)$ is concave or not too convex.

E.g. $\pi'' < 0$ diminishing returns from reputation.

Application: favorable selection

- Repeated interaction $t = 1, 2$
- Mass 1 of C, index i.
- Subscription price f so period profit: $f * d_t(p_{t-1})$

$$d_t = \Pr\{v(p_{t-1}) - f + \varepsilon_{it} \geq 0\}$$

- ε_{it} uniform i.i.d. across i.
- Taste persistence $\varepsilon_{i1} = \varepsilon_{i2}$

In period 2 you have a favorable selection of consumers from period 1

- If $p_1 \geq p_0$ all returning customers subscribe: $d_1 = 1$
- If $p_1 < p_0$ some returning customers subscribe

Proposition

$\pi(p_1) = f * d_2(p_1)$ is linear and increasing for $p_1 < p_0$ and flat for $p_1 \geq p_0$. So it is concave. Therefore there is a strategic incentive to be cautious in period 1.

Equilibrium

- Posterior is monotone in 'experience:' $p_1(u)$ increasing.

Proposition

If SIC then equilibrium is overly cautious. (full characterization). Eqm probability of -recommend- lower than the surplus maximizing one.

- R never plays -recommend- when he should not
- R sometimes plays -n/recommend- when it should.

- Distortion increases with prior reputation p_0

Mitigating factors

Mitigating factor 1 - Endogenous accuracy

- Idea: more data $\rightarrow$ better recommendations
- Accuracy through traffic: more users $\rightarrow$ more data
  - No effect on recommendation policy
- Accuracy through experimentation: more -recommend- $\rightarrow$ more data
  - Equilibrium is **not** necessarily cautious

Mitigating factor 2 - Naive consumers

- Idea: demand is monotone increasing in both p_1 and u.
  - Equilibrium is **not** necessarily cautious

Lab experiment

1. Do recommendations affect choices?
2. Do subjects learn from experience?
3. Do they cancel their subscription if they loose confidence in R's ability?
4. Do distorted recommendations enhance profits?

...

- (!!!) We run the pilot this Thursday in Bologna
- 63 subjects so far.

- Non strategic setting
- Maximum 80 periods – 2 stages per period
- Stage 2 choice assisted by a recommendation

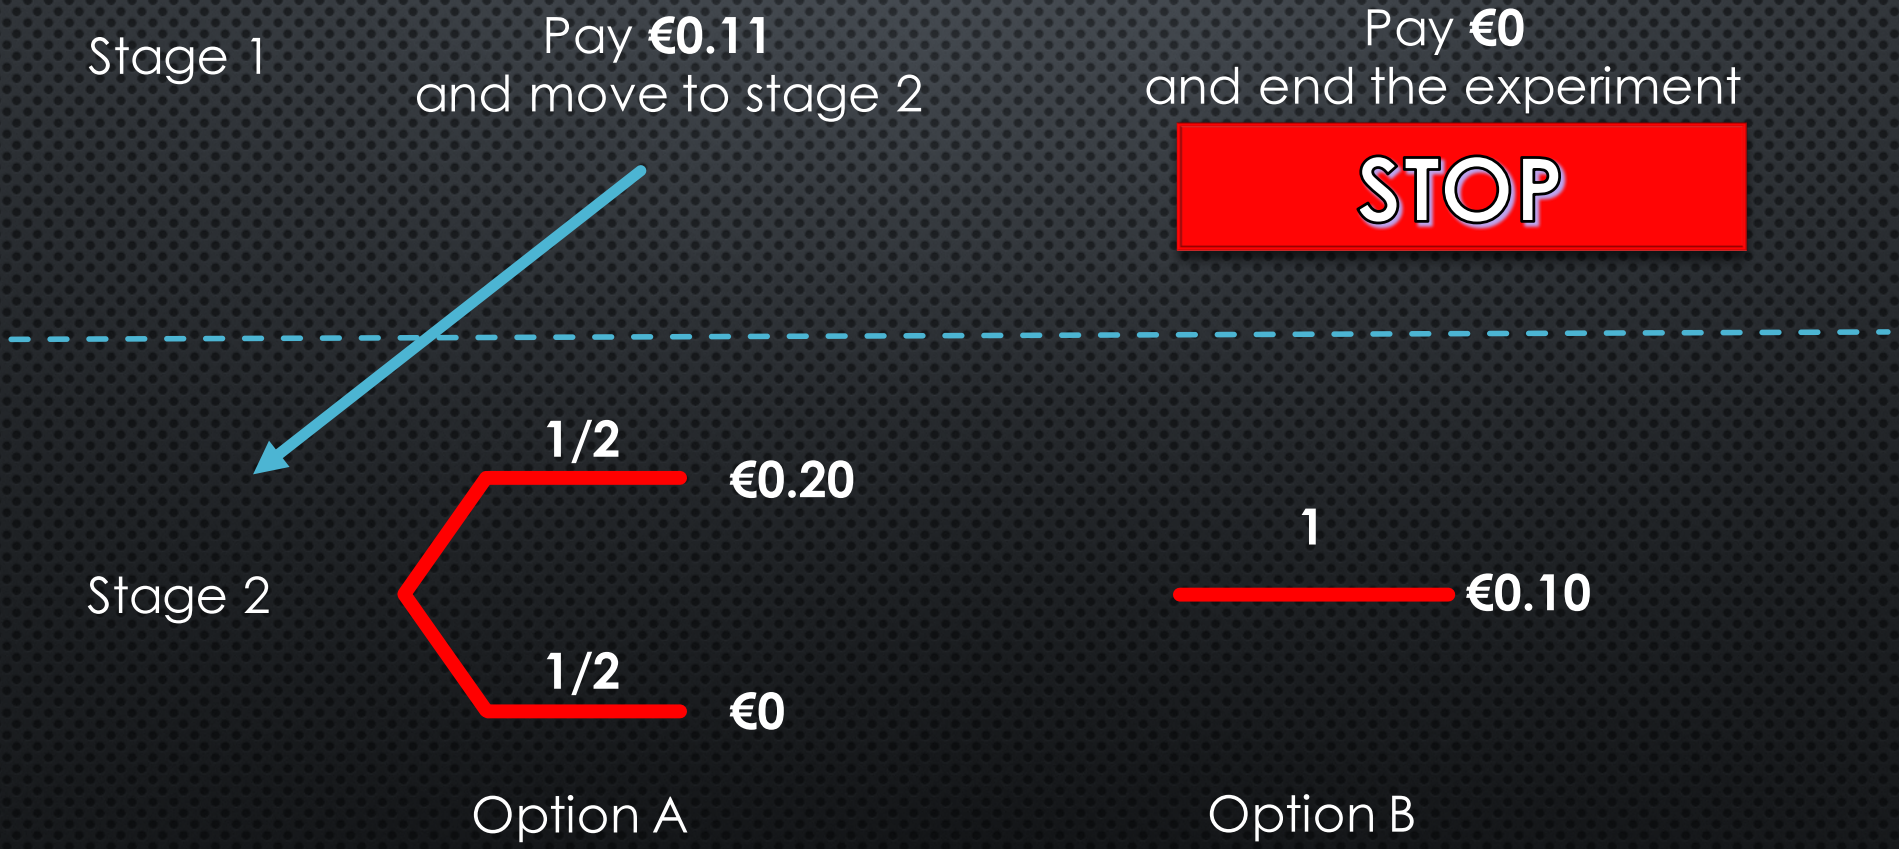

Stage 1

Round primo ciclo	Somma dei punti accumulati primo ciclo	Round secondo ciclo	Somma dei punti accumulati nel secondo ciclo	Punti iniziali
1	-5	0	0	2500
2	40	0	0	2500
3	35	0	0	2500
4	30	0	0	2500
5	80	0	0	2500
0	0	1	75	2500

Scegli tra le due opzioni:

Continua e paga 55 punti

Stop

Stage 2

Scegli tra le due opzioni:

Apri la busta

Non aprire la busta

Suggerimento: Apri la busta

50 punti

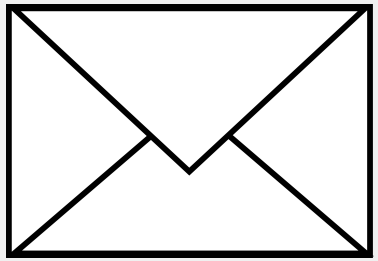

Valore estratto dalla busta: 100 punti

Continua

Round: 2

Ciclo: A

RECOMMENDATIONS

- GENERATE INFORMATIVE SIGNALS
- 2 TREATMENTS BY ACCURACY
  - VERY ACCURATE: PROB. RECOMMENDATION ERROR 9/10).
  - ACCURATE: PROB. RECOMMENDATION ERROR 1/4).
- 2 TREATMENTS BY STRATEGY
  - HEDONIC TREATMENT: R MAXIMIZES EXPECTED PAYOFF
  - STRATEGIC TREATMENT: R RECOMMENDS –OPEN– LESS FREQUENTLY
    - CAUTIOUS: PROB. RECOMMENDING NICHE 2/3 LOWER
    - VERY CAUTIOUS: PROB. RECOMMENDING NICHE 1/2 LOWER
- WHAT DO SUBJECTS KNOW? NOT MUCH….

"RECOMMENDATIONS ARE AUTOMATICALLY GENERATED BY THE COMPUTER EMPLOYING INNOVATIVE MATHEMATICAL ALGORITHMS. THESE ALGORITHMS CAN, WITH SOME PRECISION, PREDICT THE CONTENT OF THE ENVELOPE CORRESPONDING TO EACH ROUND. […] RECOMMENDATIONS DO NOT DEPEND IN ANY WAY ON THE CHOICES MADE DURING THIS STUDY. "

HOW DOES THE EXPERIMENT WORK?

- CONDITIONAL ON PARTICIPATING AND OPENING THE ENVELOPE, SUBJECTS LEARN ABOUT THE ACCURACY OF THE RECOMMENDER
  - IF REC=OPEN, CHOICE=OPEN AND ENVELOPE=**€0.20** THEN GOOD NEWS
  - IF REC=N/OPEN, CHOICE=OPEN AND ENVELOPE=**€0.20** THEN BAD NEWS
  - IF REC=N/OPEN, CHOICE=OPEN AND ENVELOPE=**€0** THEN GOOD NEWS
  - IF REC=OPEN, CHOICE=OPEN AND ENVELOPE=**€0** THEN BAD NEWS

- WE USE AS CONTROL FOR INFORMATION THIS CUMULATIVE MEASURE:

$$X_{it} = \#\text{bad news}_t \,/\, \#\text{news}_t$$

- Do individuals Follow recommendations?
- Yes - 88 % of all choices are recommended

% choices taken against advice

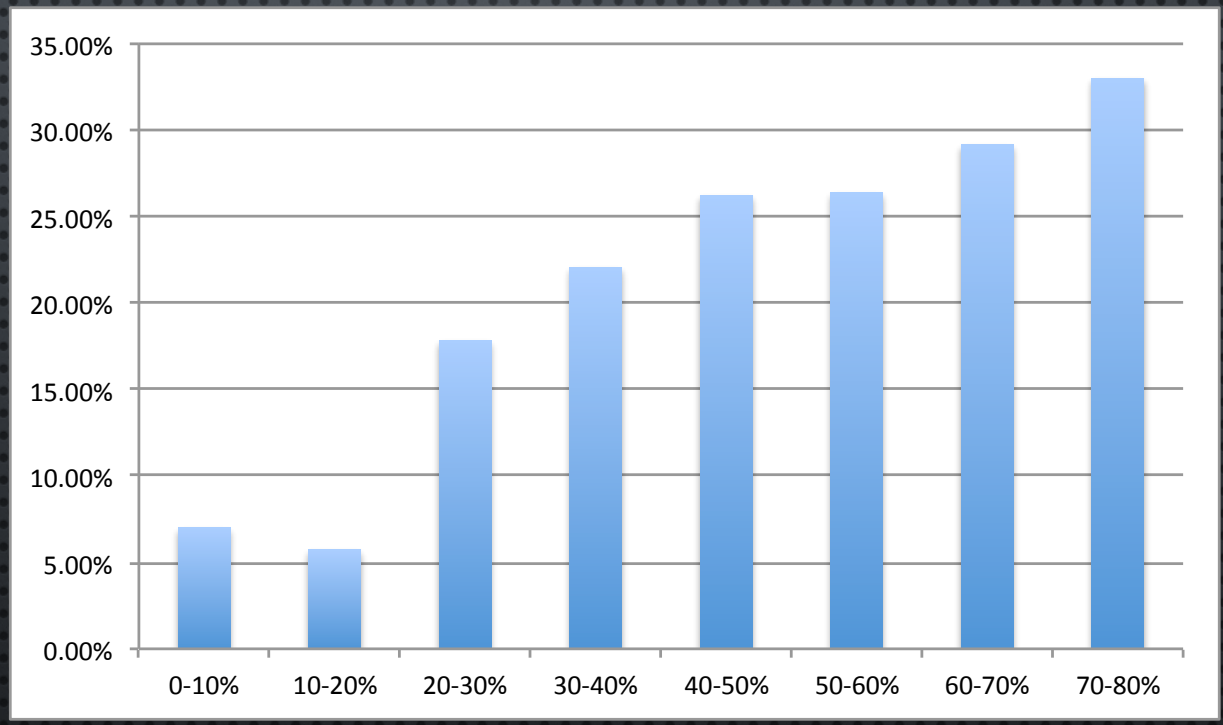

X_{it} = #bad $news_t$ / #$news_t$

X_{it}

Data on 63 subjects -2 treatments per subject (Pilot experiment)

- DO INDIVIDUALS LEARN ABOUT ACCURACY?

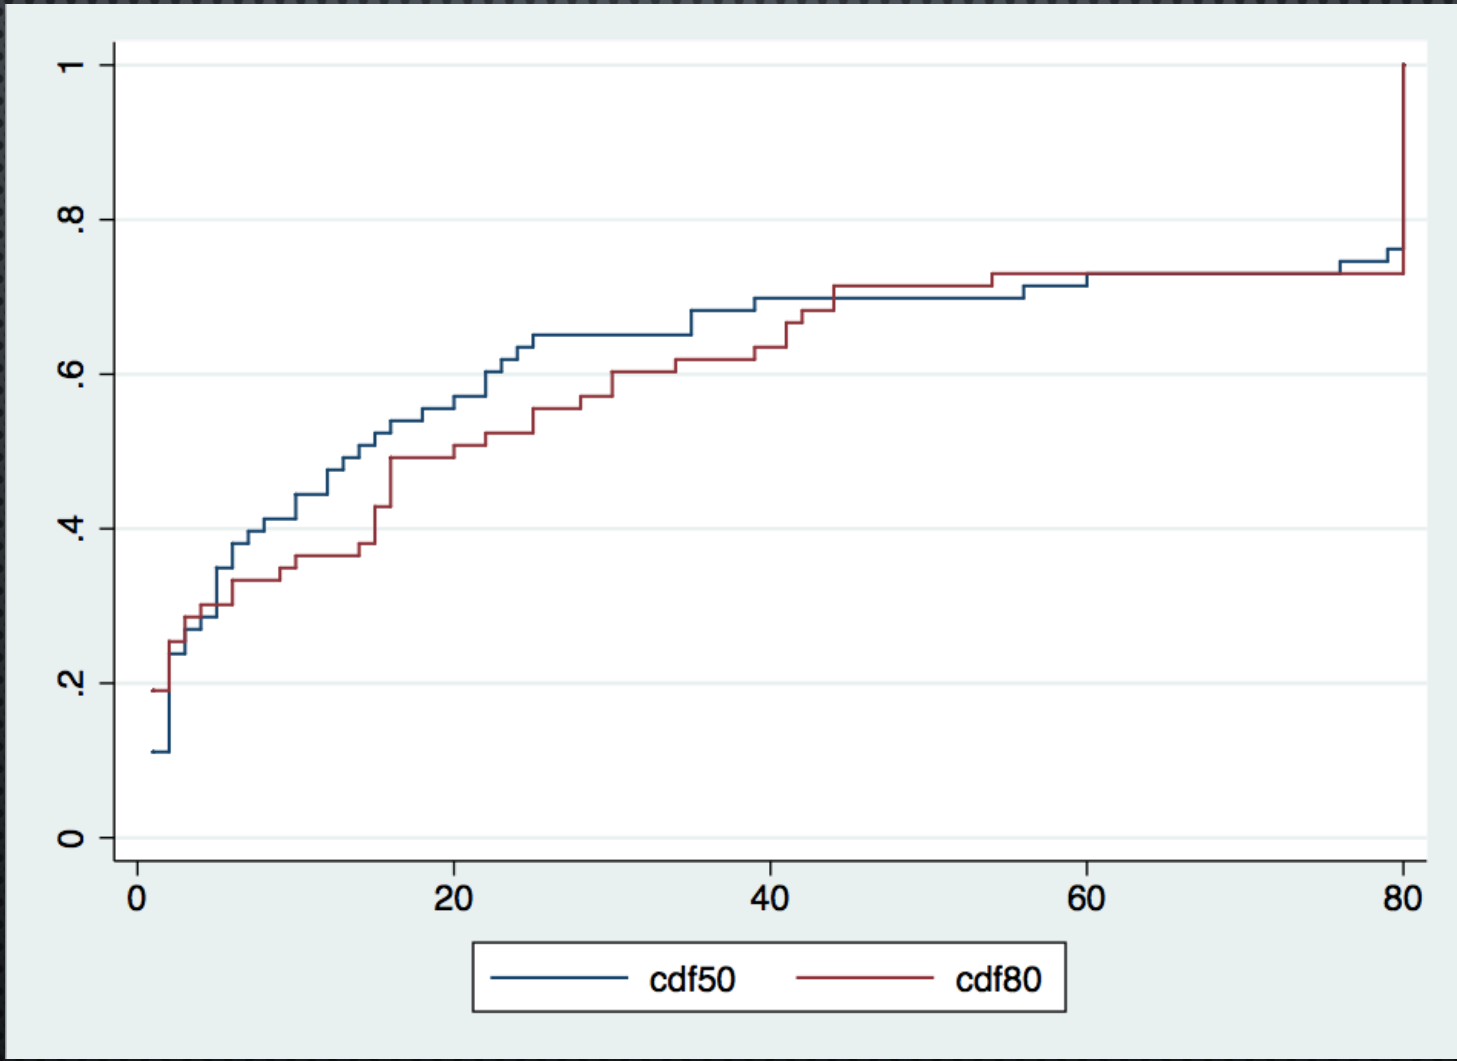

- ARE THEY MORE LIKELY TO QUIT AFTER OBSERVING MORE ERRORS?

$$\frac{\sum_i x_{iT}}{\sum_i \sum_t x_{it}}$$

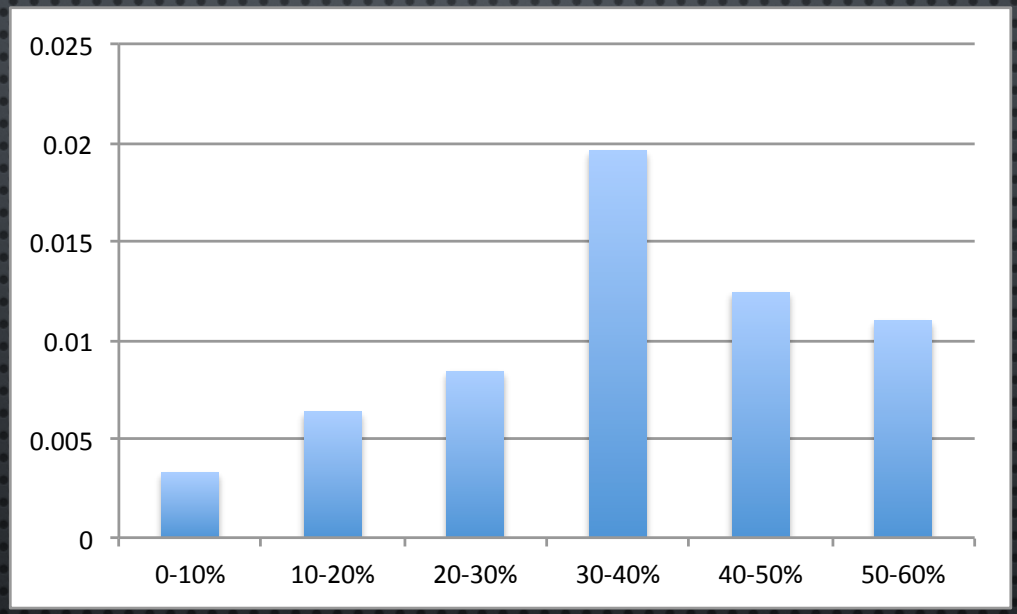

X_{it} = #bad news$_t$ / #news$_t$

X_{it}

Data on 63 subjects -2 treatments per subject
(Pilot experiment)