Persuasion Profiling, Personalization, Bandits, and all that.

Dr. M.C.Kaptein
Assistant Professor
Statistics and Research Methods

February 19, 2016

Section 1

Background

Background

- MSc. Economic Psychology, Tilburg University
- PdEng. User System Interaction, University of Eindhoven
- Ph.D. Industrial Design, University of Eindhoven & Stanford University
- Post Doc. Marketing, Aalto School of Economics, Helsinki
- Assistant Professor Artificial Intelligence, Radboud University, Nijmegen
- Founder & Chief Scientist, PersuasionAPI & Science Rockstars, Amsterdam / Barneveld. Acquired by Webpower bv.

Current appointments

- Assistant Professor, Statistics and Research Methods, University of Tilburg
- Speaker for The Next Speaker
- Author: "Persuasion Profiling" (in Dutch "Digitale Verleiding")

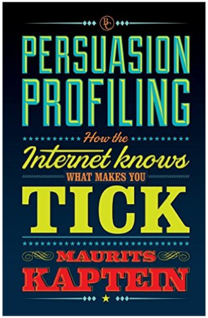

New book!

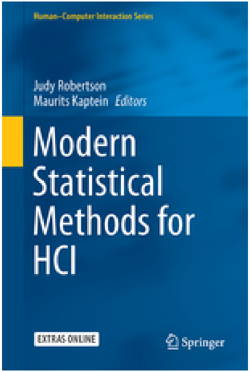
Human–Computer Interaction Series
Judy Robertson
Maurits Kaptein Editors
Modern Statistical Methods for HCI
EXTRAS ONLINE
Springer

Section 2

Persuasion Profiling

Persuasion in E-Commerce

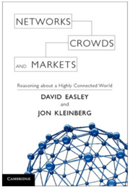

Networks, Crowds & Markets

Easley & Kleinberg

Recommended price: $ 14.99

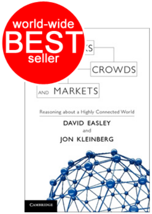

Networks, Crowds & Markets

Easley & Kleinberg

Recommended price: $ 14.99

Main Effects of Persuasion

- Average effect of the little "button": Willingness to pay increase by $> 30\%$
- Similar effects in different studies: Probability of purchase increased by 5 to 25%

Distinct Persuasion Strategies

- Scarcity
- Authority
- Social proof
- Liking
- Reciprocity
- Commitment
- ...

Estimating individual level effects of persuasion[1]

- Measure repeated responses to persuasive messages
- Use hierarchical models to estimate individual level effects
- Unique opportunity to estimate individual level effects

[1] Kaptein & Eckles (2012). Heterogeneity in the effects of online persuasion. *Journal of Interactive Marketing*

Results

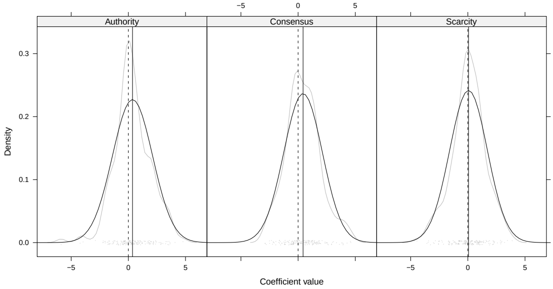

Authority
Consensus
Scarcity
Density
0.0
0.1
0.2
0.3
−5
0
5
Coefficient value

The persuasion profile[2]

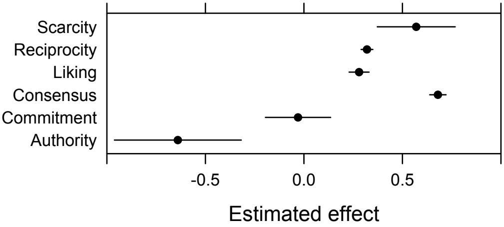

[2] Kaptein, Eckles, & Davis (2011). Envisioning Persuasion Profiles. *ACM Interactions*

Using persuasion profiles for email compliance[3]

- Susceptibility estimated based on behavioural response
- Selection of strategies Adaptive, Original, Pre-tested, Random
  - More on adaptive later ...
- Large differences in success probability
- $N = 1129$

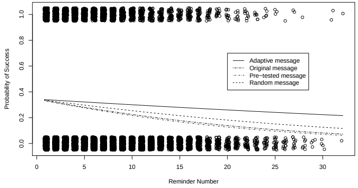

[3] Kaptein & van Halteren (2012).
Adaptive Persuasive Messaging to Increase Service Retention.
Journal of Personal and Ubiquitous Computing

Section 3

The Multi-Armed Bandit Problem

Slot machines

Formal presentation

- Each each timpoint,
- we select and action,
- and observe a response.
- We wish to "optimize" the response using a selection *Policy*

- For $t = 1, \ldots, t = T$
- Select and action a_t out of $\mathcal{A}_t$. Often actions $k = 1, \ldots, k = K$.
- Observe reward r_t (generated by some unknown distribution $F_k(r|\theta_k)$)

- Play according to some policy $\Pi : \{a_1, \ldots, a_{t-1}, r_1, \ldots, r_{t-1}\} \mapsto a_t$

Aim of a “good” policy

- Well, get as much reward as possible!
- Or, equivalently, minimize *Regret*
- Thus, maximize $r(t) = \sum_{t=1}^{T} r_t$
- Or, minimize: $R(t) = \sum_{t=1}^{T} (\Pi^*(t) - \Pi(t))$

Exploration-Exploitation tradeoff

- Should we take the action we think is best? Or should we learn more?
- Suppose observations $X_k \sim \text{Bern}(p_k)$
- Explore: $p_1 > p_2$? Play alternating arms to learn.
- Exploit: Play arm 1.

Very general trade-off: Exploring the outcomes of uncertain actions, versus choosing actions that one beliefs to be good.

Omnipresence of the tradeoff

Exploration vs. exploitation found in many places:

- Clinical trial: which medicine to subscribe?
- Online content selection: Which ad, news article, or product to show?
- Job choices: Try something new, vs. stick to what you have?
- Food choices: Try a new dish, stick to one you like
- Etc. etc.

Also known as *Earning vs. Learning*.

Extension: contexts

- A common generalization: We first see the state of the world (the *context*), then take and action, and then observe the response.

- For $t = 1, \ldots, t = T$
- **Observe the world,** $x_t \in \mathcal{X}_t$
- Select and action a_t out of $\mathcal{A}_t$. Often actions $k = 1, \ldots, k = K$.
- Observe reward r_t (generated by some unknown distribution $F_k(r|\theta_k)$)

- Play according to some policy $\Pi : \{x_1, \ldots, x_{t-1}, a_1, \ldots, a_{t-1}, r_1, \ldots, r_{t-1}\} \mapsto a_t$

What does this have to do with Persuasion Profiling (or Personalized communication)?

- We communicate in sequence ($=t$)
- We observe the state of the world, state of the receiver, etc. ($=$ context)
- We select and message ($=$ action)
- We observe the effect of the message ($=r_t$)
- We need a *way to select messages* to maximize the effect.

cMAB provides a formalization for personalization attempts!

Our standard solution: the "experiment"

What is the regret of a simple experiment for choosing between two messages?

- Simple choices: Two actions with p_1 and p_2 for succes
- Obviously: p_1, p_2 unknown at start
- N recipient in trial, P total recipients in population
- Always $Pr(Wrong|experiment) > 0$

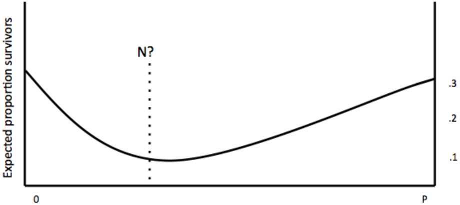

The experiment: Regret

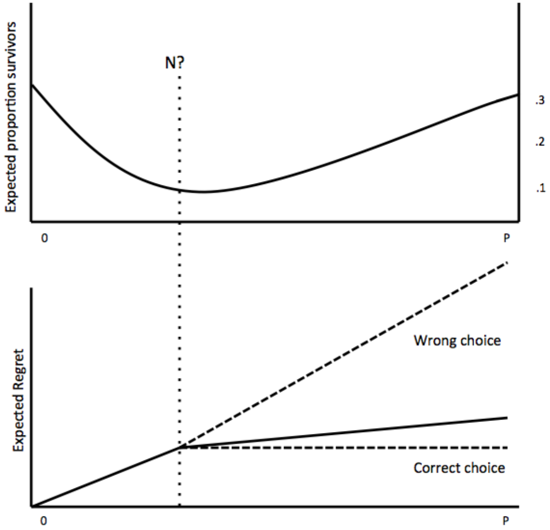

Optimal policies: Regret

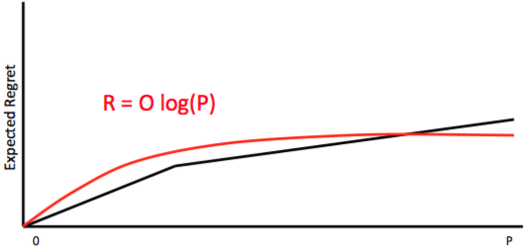

Thompson sampling: An alternative Policy

- We can select an action according to its probability of being optimal.
- If we believe action 1 is better, we play it more often.
- Compute or sample from $\Pr(\theta|\mathcal{D})$.

$$\int 1 \left[\mathbb{E}(r|a,\theta) = \max_{a'} \mathbb{E}(r|a',\theta) \right] \Pr(\theta|\mathcal{D}) d\theta \tag{1}$$

  where 1 is the indicator function.
- Thompson sampling is an *asymptotically optimal* strategy.

cMAB and Personalization

- cMAB framework used for advertising personalization by web companies (e.g., Yahoo! personalized news selection)
- We used a cMAB framework for personalizing persuasive messages: Persuasion profiling
  - Context: Identifier of the person
  - Actions: The persuasive message (Authority, Social proof, etc.)
  - Reward: Click on the product
  - Goal: Maximizes the number of clicks

Section 4

StreamingBandit: Software

Streaming Bandit [4]

If we have a consistent formalization, we can setup a general solution to examine different policies.

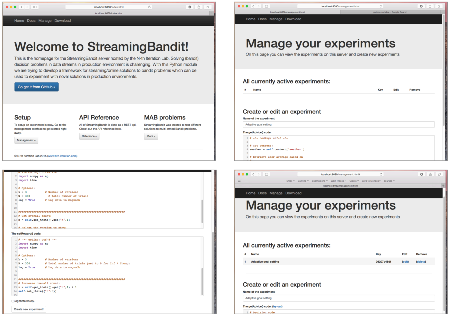

[4] Kaptein, M.C. & Kruijswijk, J. (2015).
Available on Github. `https://github.com/MKaptein/streamingbandit`

Design choice: learning and choosing

We identify two steps:

1. The *summary* step: In each summary step we update our model given the observed result.
2. The *decision* step: In the decision step, we select an action given the model.

Jointly, this implements a Policy.

Streaming Bandit in Practice:

Personalizing interest rates for small size consumer loans.

- Observe features of a customer requesting a loan
- Select an interest rate
- Observe acceptance of loan
- Objective: Choose interest such as to maximize profit

Streaming Bandit for Persuasion Profiling:

Personalizing persuasive messages

- Observe a user coming to a webpage
- Select a persuasive “label” to place over a product
- Observe click behavior

Note that we use hierarchical models to pool data over customers

Section 5

Future Vision

Personalized feedback and treatment selection

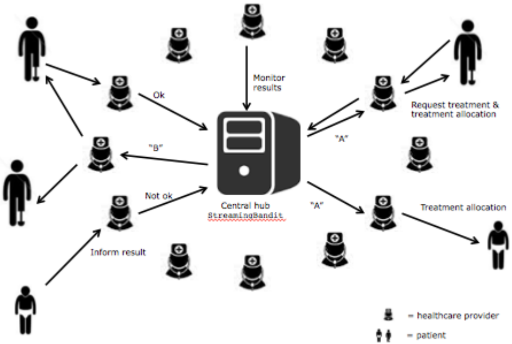

Discussion points

- Feasibility of Persuasion Profiling
- Usefulness of formal presentation of Personalization as Bandit problem
- Possibility to develop a general framework (software package) for personalization

Questions?

Maurits Kaptein
Archipelstraat 13
6524LK, Nijmegen
0621262211
maurits@mauritskaptein.com