Blend me in:

Privacy-Preserving Input Generalization for Personalized Online Services

Alegria Baquero
U. of California, Irvine

Allan Schiffman
CommerceNet

Jeff Shrager
CommerceNet &
Stanford U.

AMS@COMMERCE.NET

If someone asks you a question they've got no business asking, you're under no obligation to tell them the truth.

- Leonard Schiffman

Thesis: Lie!

(Personalized) Online Services...

Over-collect data from subject's POV

Untrustworthy re motives & competence

Results aren't very personalized

They don't need/deserve the whole truth.

Outline

Blend Me In: The Idea

Background: Models & Analysis

Quantifying Privacy and Accuracy

Case Study: Clinical Trial Recommendation

Protoype Input Generalization UI

(Persuasive) Summary

The Idea

Goal: "Blend in" to the crowd

Motive: personalized services over-collect

Consider the future: "Forward Privacy"

Threat model: service as adversary

Therefore, generalize input:

Perturb data to gain privacy (e.g., lie)

...but so as to preserve output accuracy

Background & Motivation

Privacy Models: Public Databases

Database tradition

Post collection

Enemy: analysts

Generalization & suppression

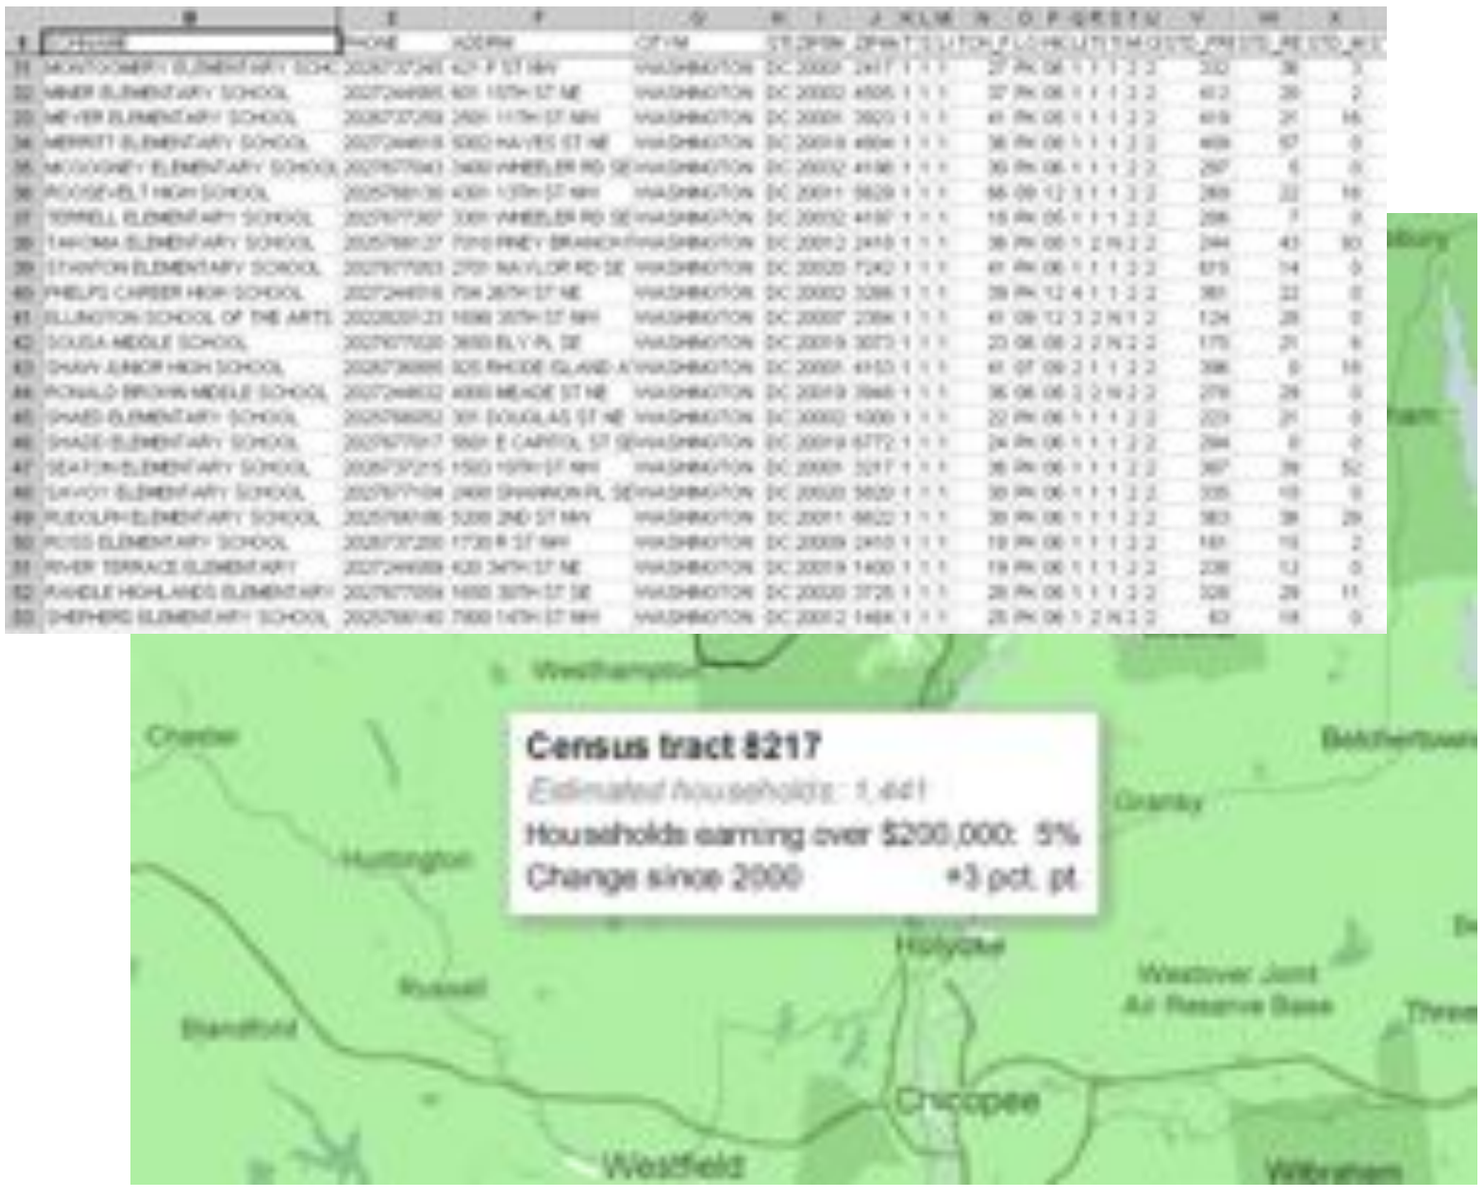

Privacy Models: Sensor Streams

Traffic analysis tradition

Enemy: stream subscribers

Masking & de-identification

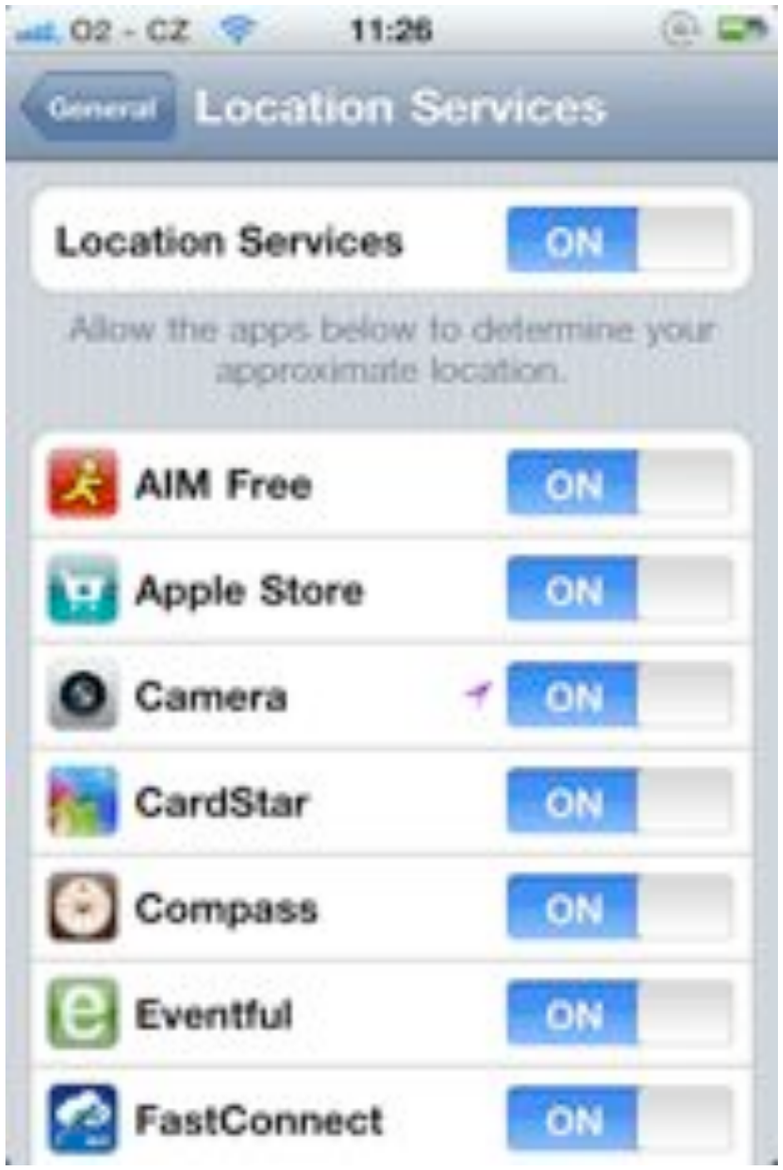

Privacy Models: Personalization

Consumer Protection Tradition

Enemy: 3rd party

Explicit Policies (UI & tools)

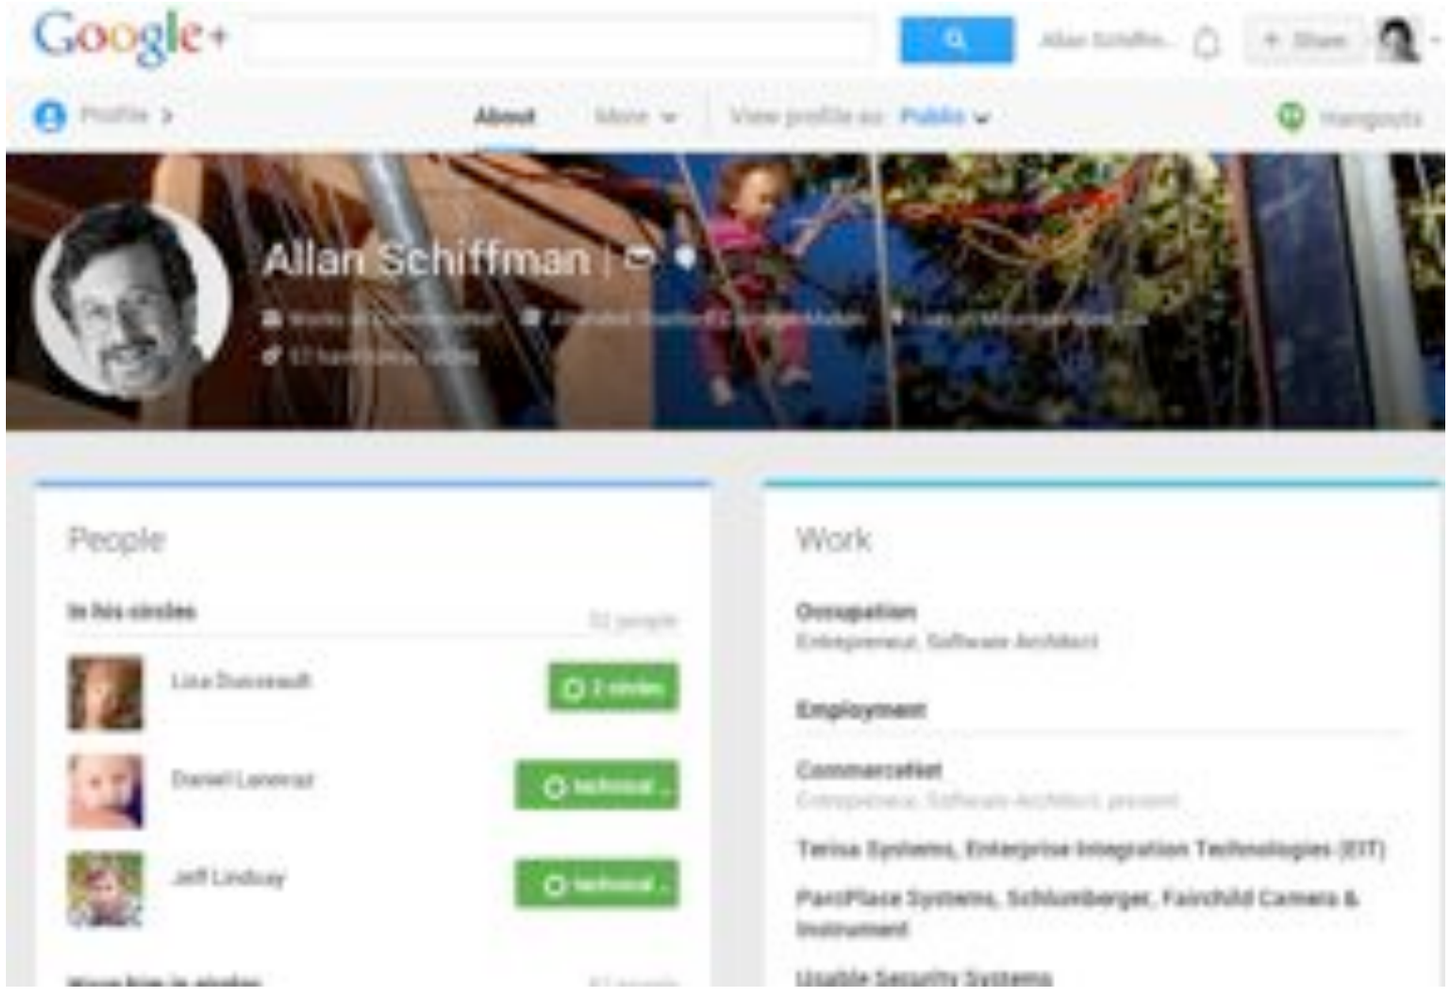

Personalized Services

Search &
Recommendation

Complaints:

- Misaligned incentives
- Bad Policies
- UI mismatch
- Security too hard
- Mgmt changes

What is revealed?

Direct Identifiers

Name
National ID
Home Address
Tel #
Email Address
& ...

(Implicit) Quasi-Identifiers

User Agent
Geocode
Cookies
IP Address
& ...

(Explicit) Quasi-Identifiers

Postal Code
Gender
Birthday
Current Query
& ...

Threat Analysis

Identification, re-identification

Association with sensitive attributes/groups

Heath data especially

The inference problem

Important: the service is the adversary

Forward Privacy: Definition

Privacy preserved under future conditions

By analogy with “Forward Secrecy”

Also see “work advantage” (re unknown adversary) – more is better

Forward Privacy: Desirable

The future is uncertain

Your concerns may change

Site's incentives may change

Adversary's advantages will increase

More background knowledge

Better technology

So, minimize the truth you tell

Input Generalization

Reduce precision of micro-data

Before input is provided ("ex-ante")

Typically, supply inaccurate value

Goal: choose best trade-off (min-max)

Minimize privacy loss, maximize utility

Necessarily knowledge-based

Quantifying Privacy and Accuracy

β **vs.** **s**

Privacy Gain:
“Blending Factor” ß

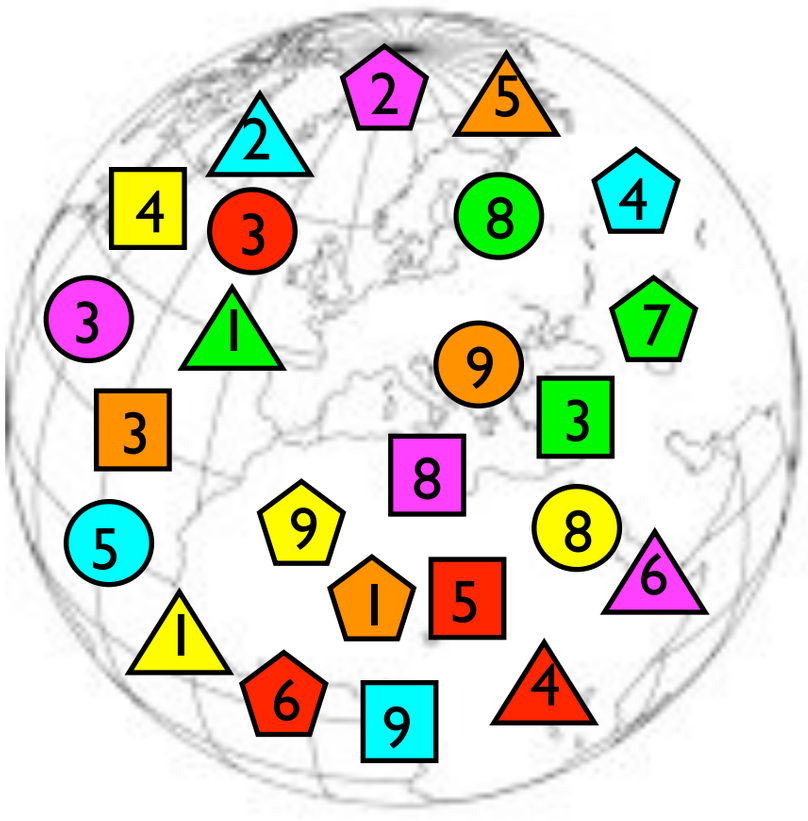

Privacy Gain:
“Blending Factor” ß

Privacy Gain: "Blending Factor" ß

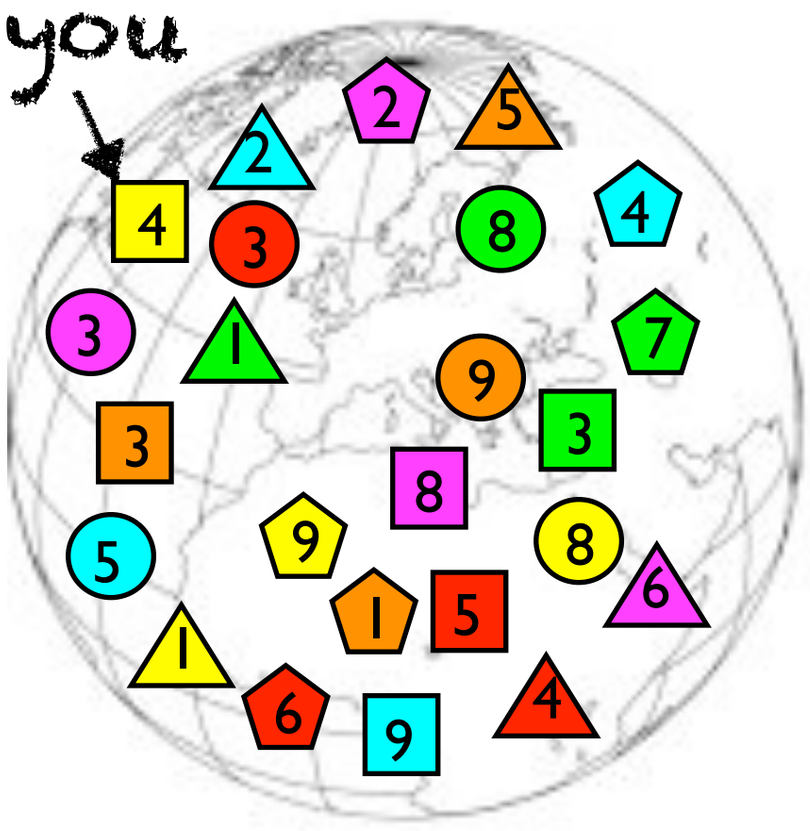

Goal: generalize your label (4) to a range (2-7)

Privacy Gain: "Blending Factor" ß

Goal: generalize your label (4) to a range (2-7)

$$\beta = \frac{\text{false pop. size}}{\text{true pop. size}}$$

Privacy Gain: "Blending Factor" ß

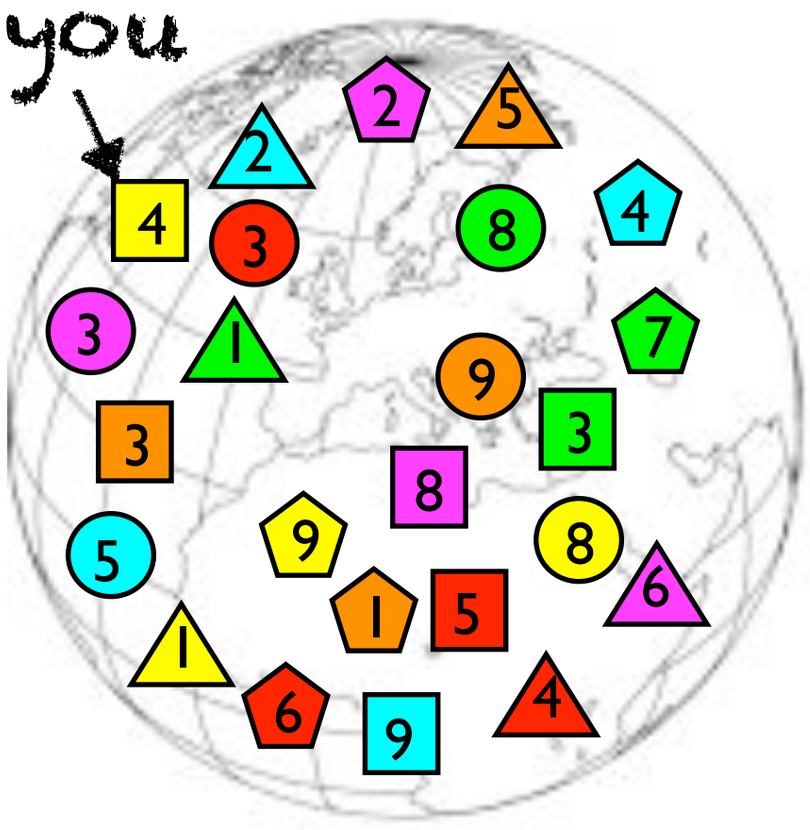

Goal: generalize your label (4) to a range (2-7)

$$\beta = \frac{\text{false pop. size}}{\text{true pop. size}} = \frac{|\{e: 2 < L(e) < 7\}|}{|\{e: L(e) = 4\}|}$$

Privacy Gain: "Blending Factor" ß

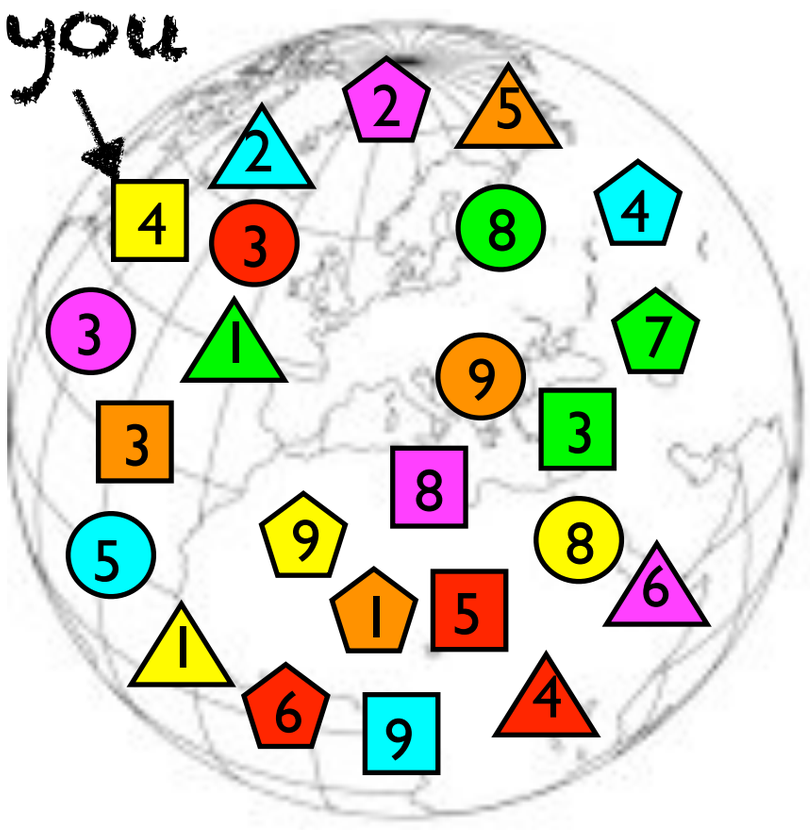

Goal: generalize your label (4) to a range (2-7)

$$\beta = \frac{\text{false pop. size}}{\text{true pop. size}} =$$

$$\frac{|\{e: 2 < L(e) < 7\}|}{|\{e: L(e) = 4\}|} =$$

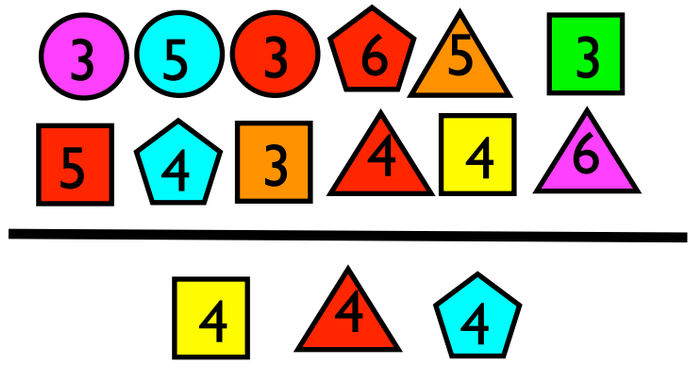

Privacy Gain: "Blending Factor" ß

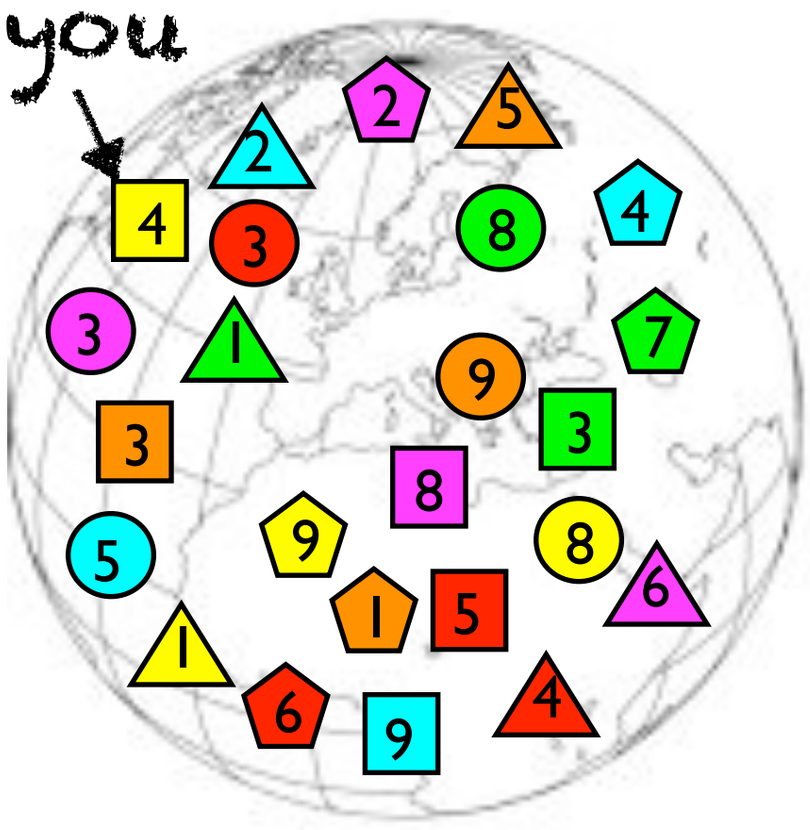

Goal: generalize your label (4) to a range (2-7)

$$\beta = \frac{\text{false pop. size}}{\text{true pop. size}} =$$

$$\frac{|\{e: 2 < L(e) < 7\}|}{|\{e: L(e) = 4\}|} =$$

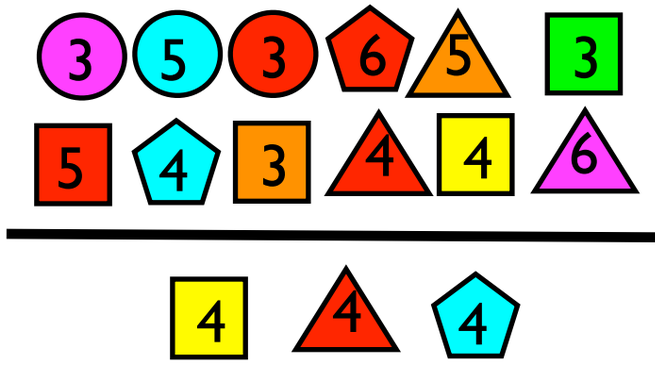

$$= 4$$

Output Accuracy (s)

"Truth"
Output

A	
B	
C	
D	
E	
F	
G	
H	

Output Accuracy (s)

"Truth" Output	*"Lies" Output*
A	G
B	A
C	H
D	B
E	F
F	D
G	Z
H	

Output Accuracy (s)

"Truth" Output	*"Lies" Output*
A	G
B	A
C	H
D	B
E	F
F	D
G	Z
H	

Purple: false negative (C&E)
Red: false positive (Z)

Output Accuracy (*s*)

"Truth" Output	*"Lies" Output*
A	G
B	A
C	H
D	B
E	F
F	D
G	Z
H	

$$s = \frac{|\ \text{correct}\ |}{|\ \text{correct + incorrect}\ |}$$

Purple: false negative (C&E)
Red: false positive (Z)

Output Accuracy (*s*)

"Truth" Output	*"Lies" Output*
A	G
B	A
C	H
D	B
E	F
F	D
G	Z
H	

$$s = \frac{|\ \text{correct}\ |}{|\ \text{correct + incorrect}\ |} = \frac{|\{G,A,H,B,F,D\}|}{|\{G,A,H,B,F,D\}| + |\{C,E\}| + |\{Z\}|}$$

Purple: false negative (C&E)
Red: false positive (Z)

Output Accuracy (*s*)

"Truth" Output	*"Lies" Output*
A	G
B	A
C	H
D	B
E	F
F	D
G	Z
H	

$$s = \frac{|\text{ correct }|}{|\text{ correct + incorrect }|} = \frac{|\{G,A,H,B,F,D\}|}{|\{G,A,H,B,F,D\}| + |\{C,E\}| + |\{Z\}|}$$

$$s = \frac{6}{9} = 67\%$$

Purple: false negative (C&E)
Red: false positive (Z)

Case Study: Personalized Medical Recommendations

About Clinical Trials

Medical progress & treatment

US:130k active

Full registry

Description & in/out criteria

Scientists search

Clinical Trial Recommendation

Search is hard for Doctors & Patients

Ask questions and recommend trials

Choose trial: useful & patient eligible.

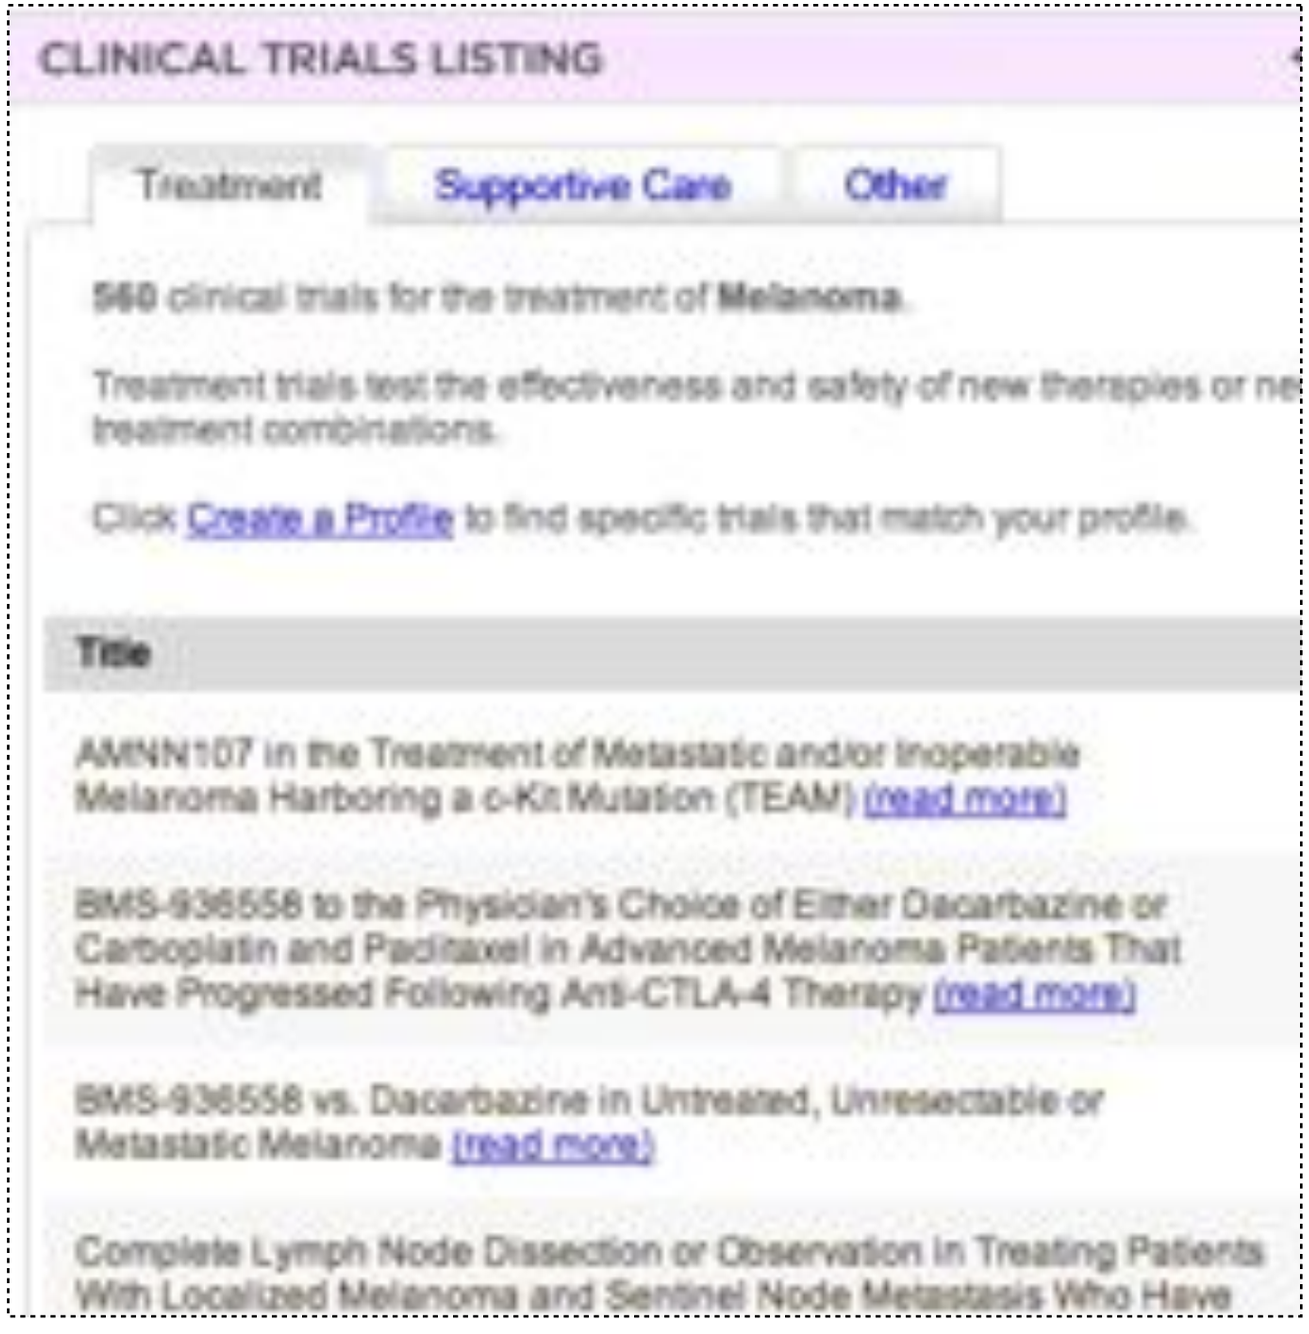

Service questionnaire

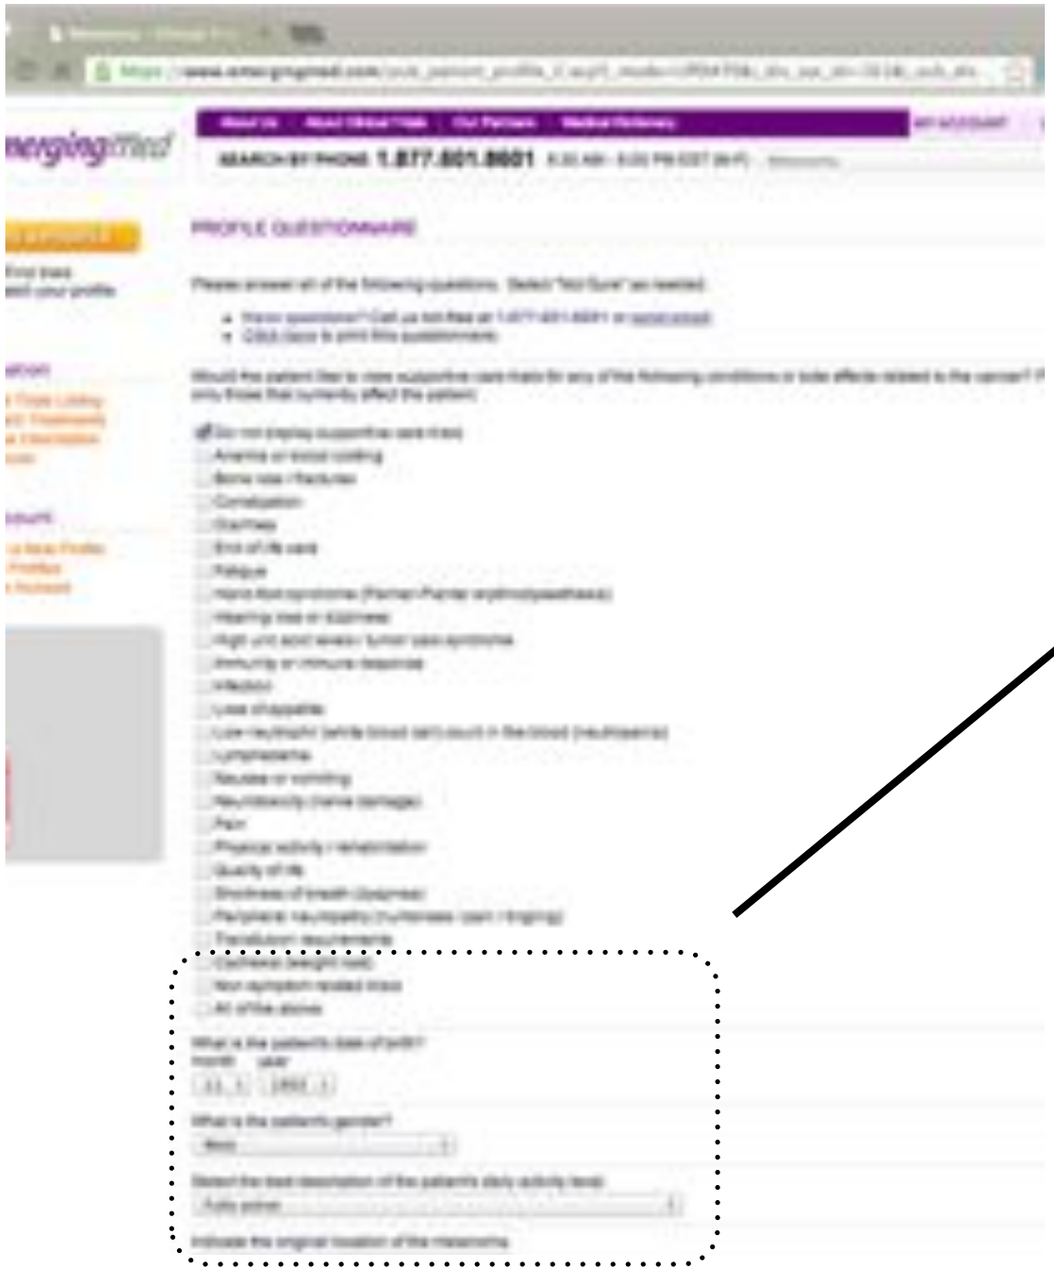

Study methodology

Create robust medical profiles for “patients”

Required special expertise

Collect service output for varying inputs

“True” profile

Input generalized profiles

Best-effort & randomized

What to generalize

Yes: Personal Profile

Zip Code

Date of Birth

Yes: treatment history

No: Diagnosis

Drug hierarchy

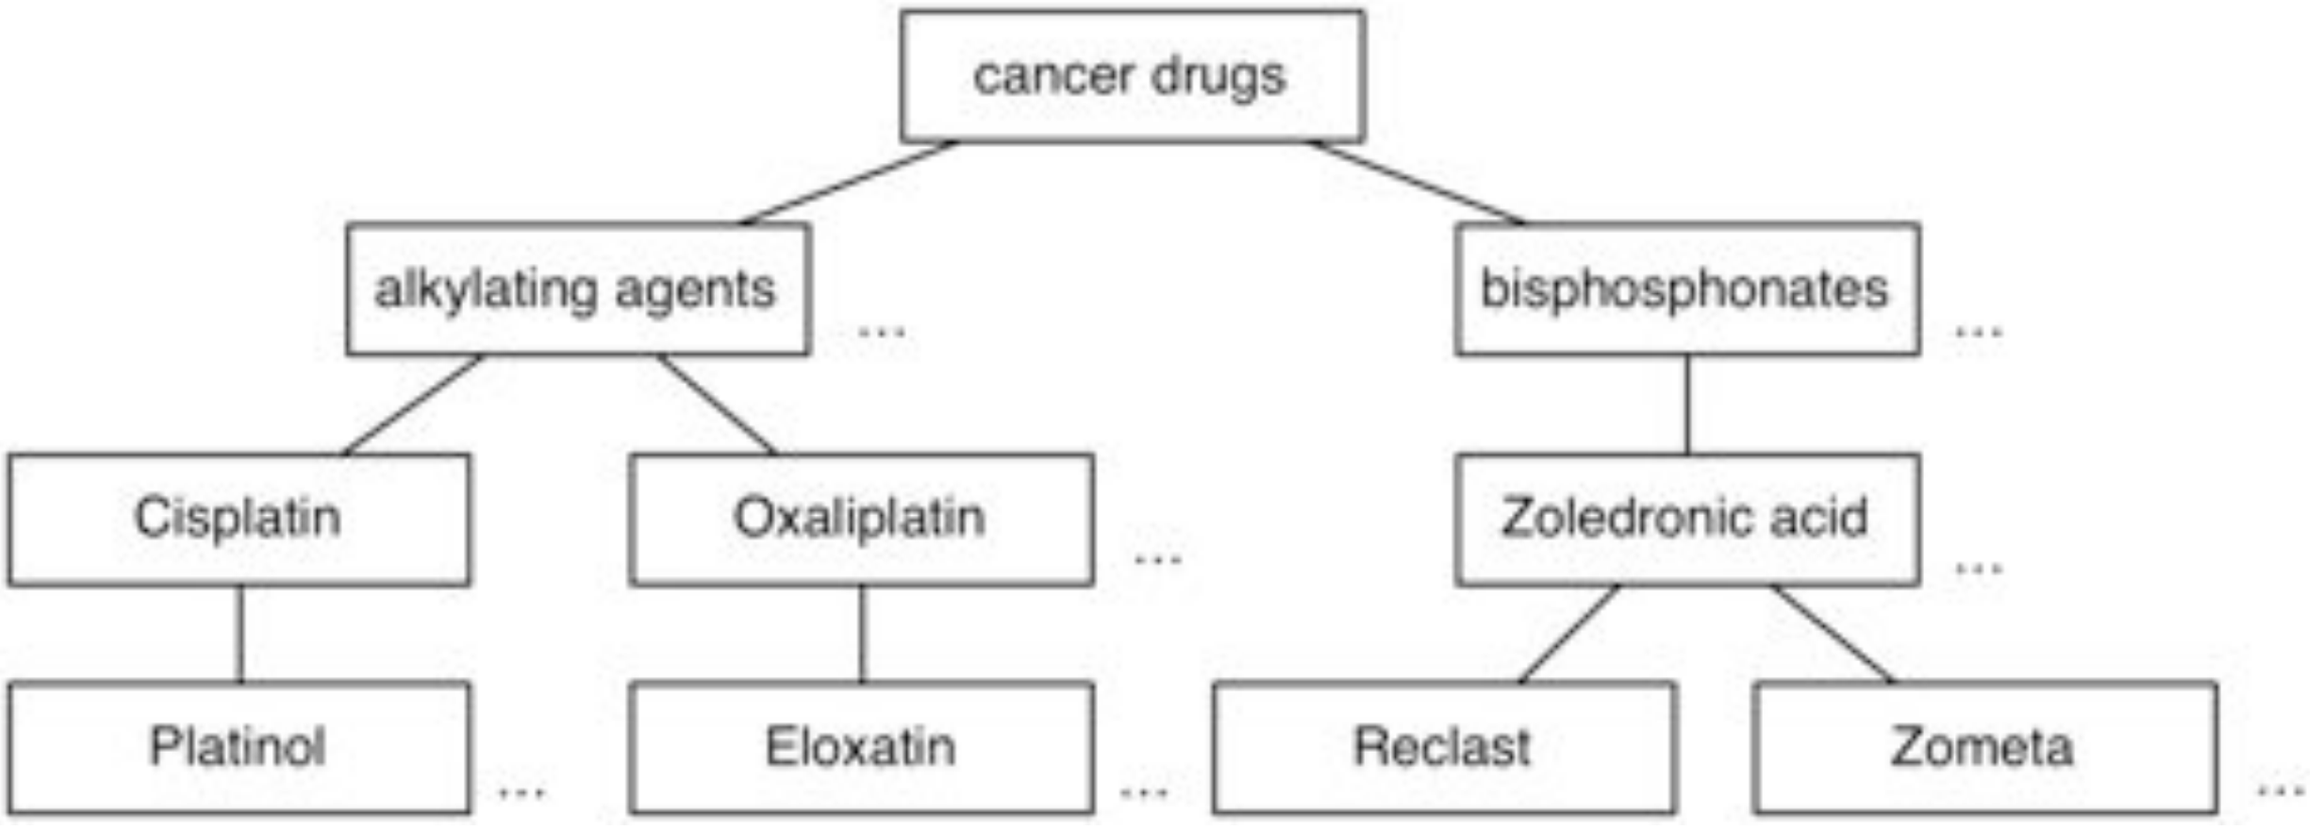

Results: high ß and *S*

		Original values				Generalized values							
		zip code	dob	prescriptions	T	zip code	dob	prescriptions	G	F_P	F_N	accuracy	β
melano.	Patient 1	94025	Jan 1955	Vemurafenib, Dacarbazine	20	95056	Jun 1951	Vemurafenib, Carmustine	20	0	0	100%	1.7×10^5
	Patient 2	10016	Jan 1965	Ipilimumab, Dacarbazine	3	11222	Aug 1960	Ipilimumab, Temozolomide	3	0	0	100%	1.7×10^5
	Patient 3	60601	Jan 1975		3	60202	Nov 1977		3	0	0	100%	9.6×10^3
colorectal	Patient 4	94025	Jan 1955	Bevacizumab, Erlotinib, Fluorouracil, Oxaliplatin	284	95129	Jul 1957	Bevacizumab, Gefitinib, Fluoracil, Cisplatin	283	2	3	98.3%	8.6×10^5
	Patient 5	10016	Jan 1955	Fluorouracil, Oxaliplatin	3	10024	Jan 1958	Capecitabine, Cisplatin	3	0	0	100%	1.4×10^6
	Patient 6	60601	Jan 1975		3	60621	Feb 1972		3	0	0	100%	9.6×10^3
lung	Patient 7	94025	Jan 1955		21	94544	May 1950		21	1	1	90.9%	9.6×10^3
	Patient 8	10016	Jan 1965	Carboplatin, Oxaliplatin	3	10105	Oct 1967	Lomustine, Cyclophosphamide	3	0	0	100%	3.1×10^6
	Patient 9	60601	Jan 1975	Cetuximab, Carboplatin, Oxaliplatin	250	60706	Sep 1980	Gefitinib, Cisplatin, Cyclophosphamide	248	2	4	97.6%	1.5×10^7

T is the true output cardinality, G the generalized output cardinality, F_P the number of false positives, and F_N the number of false negatives.

Blend me in:

The movie

Lt. Kaffee: I want the truth!
Col. Jessup: You can't handle the truth!

- A Few Good Men (1992)

Blend Me In Prototype

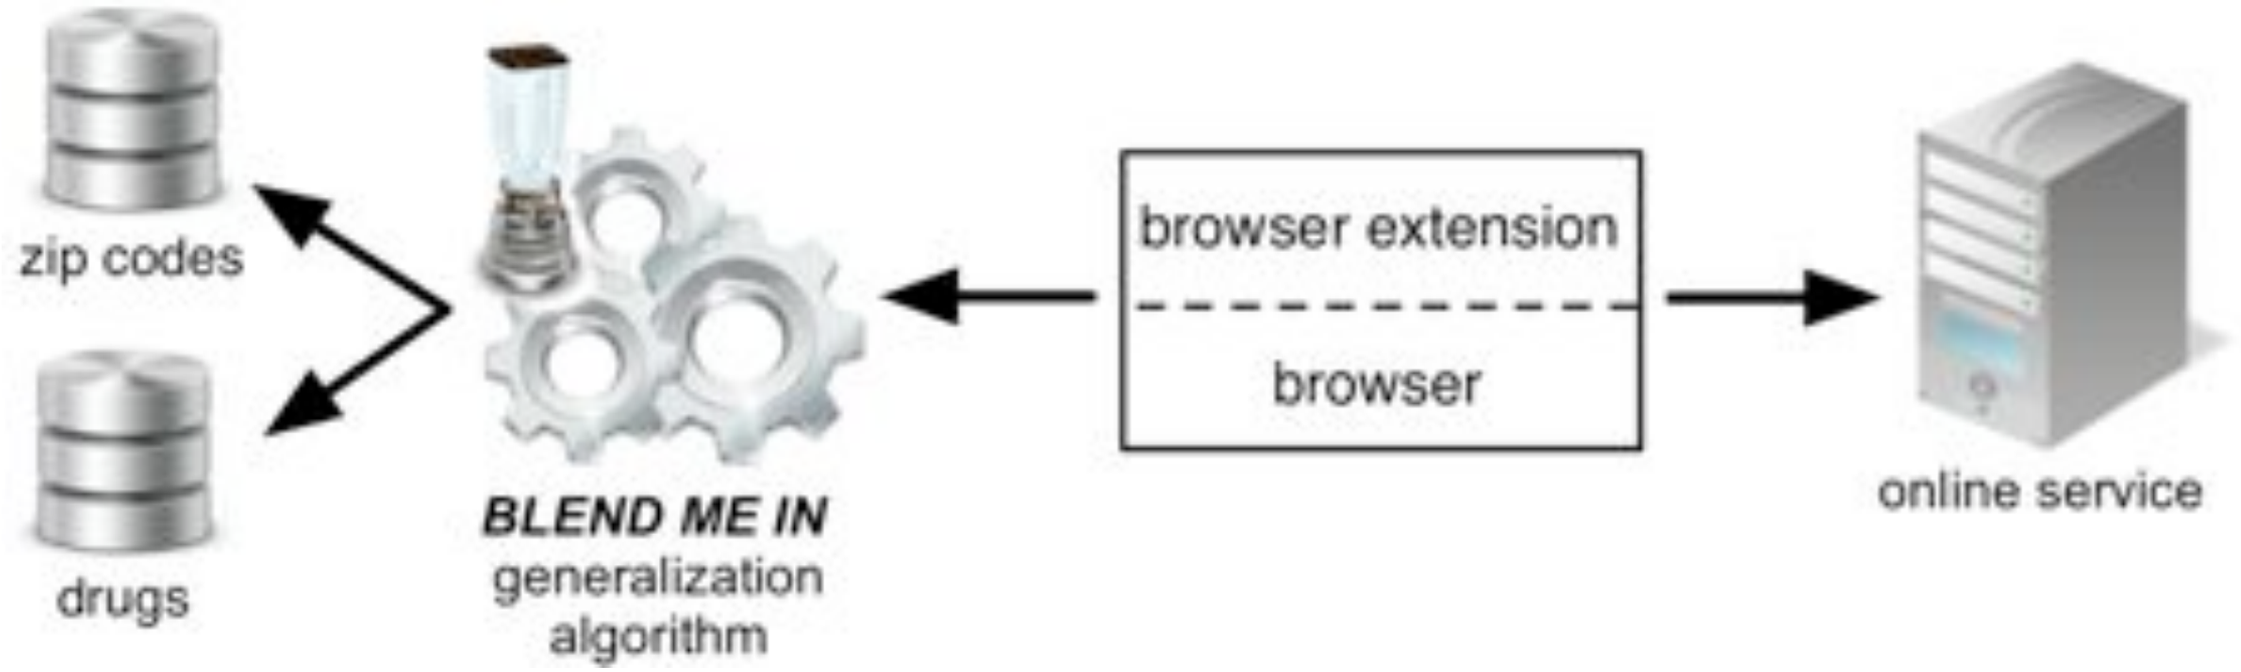

UI for zipcode

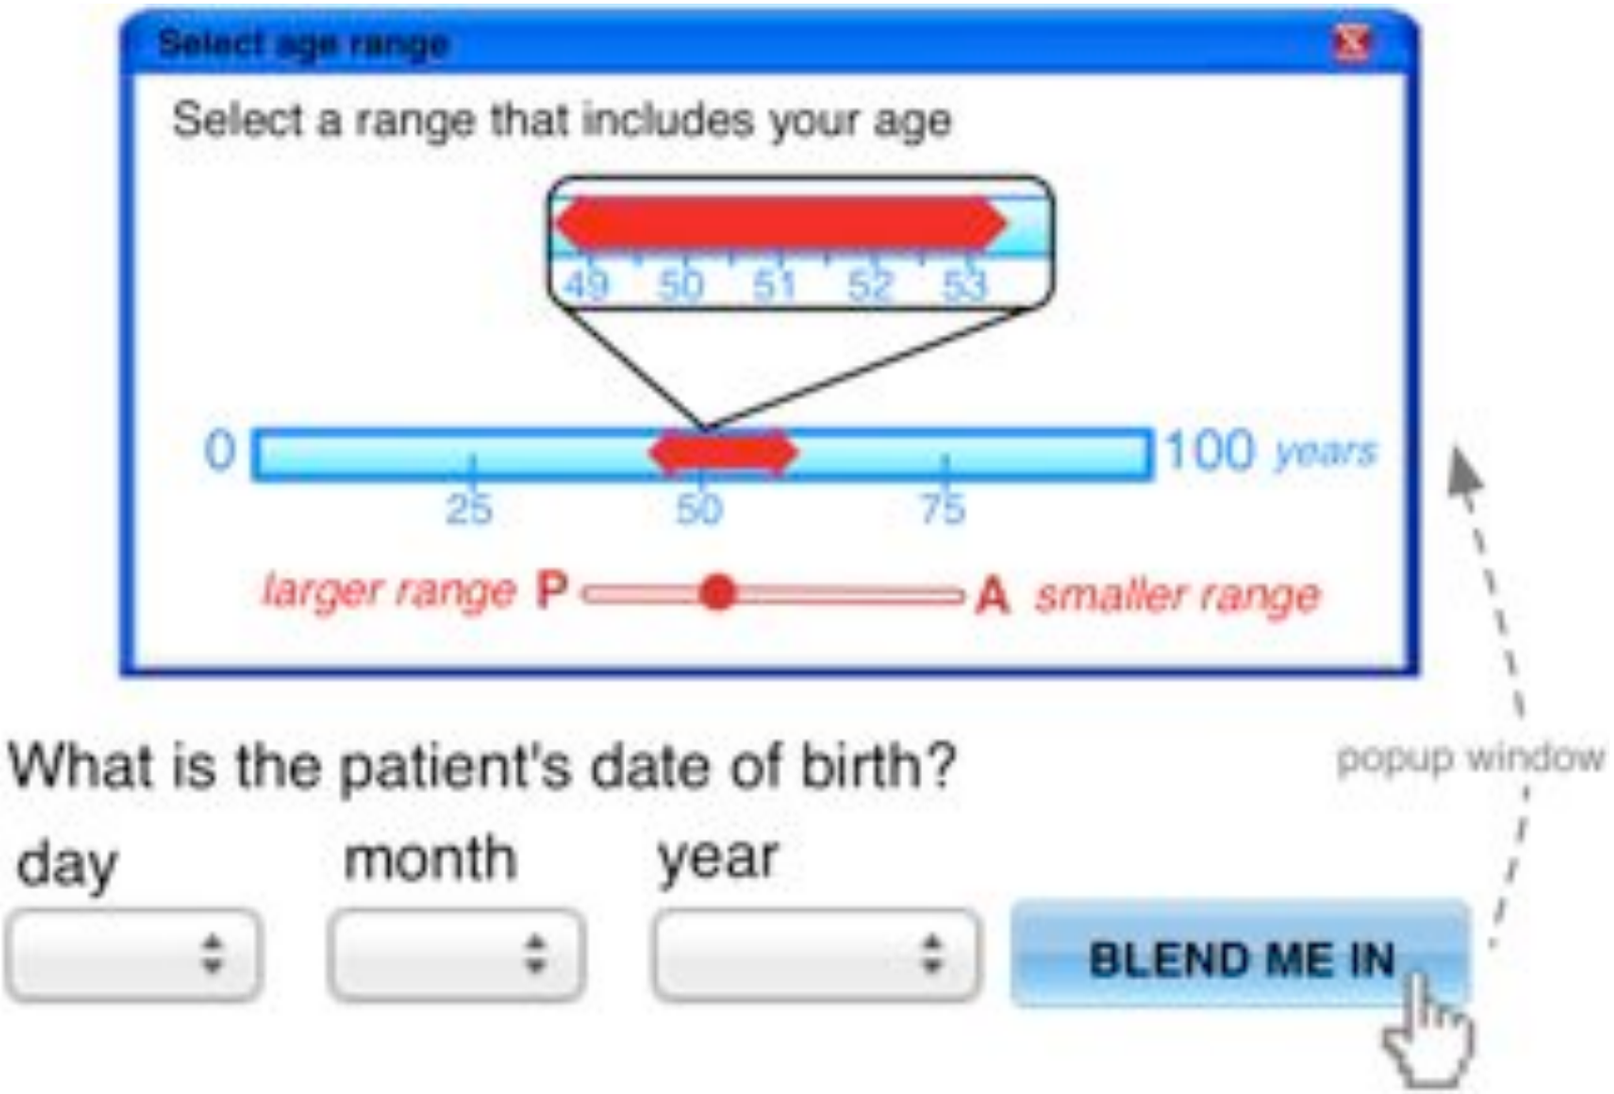

UI widget for date of birth

UI for Drug Treatment History

Is *Blend Me In* Scalable?

Collective intelligence script ecosystem

Library of UI widgets for enlightened sites

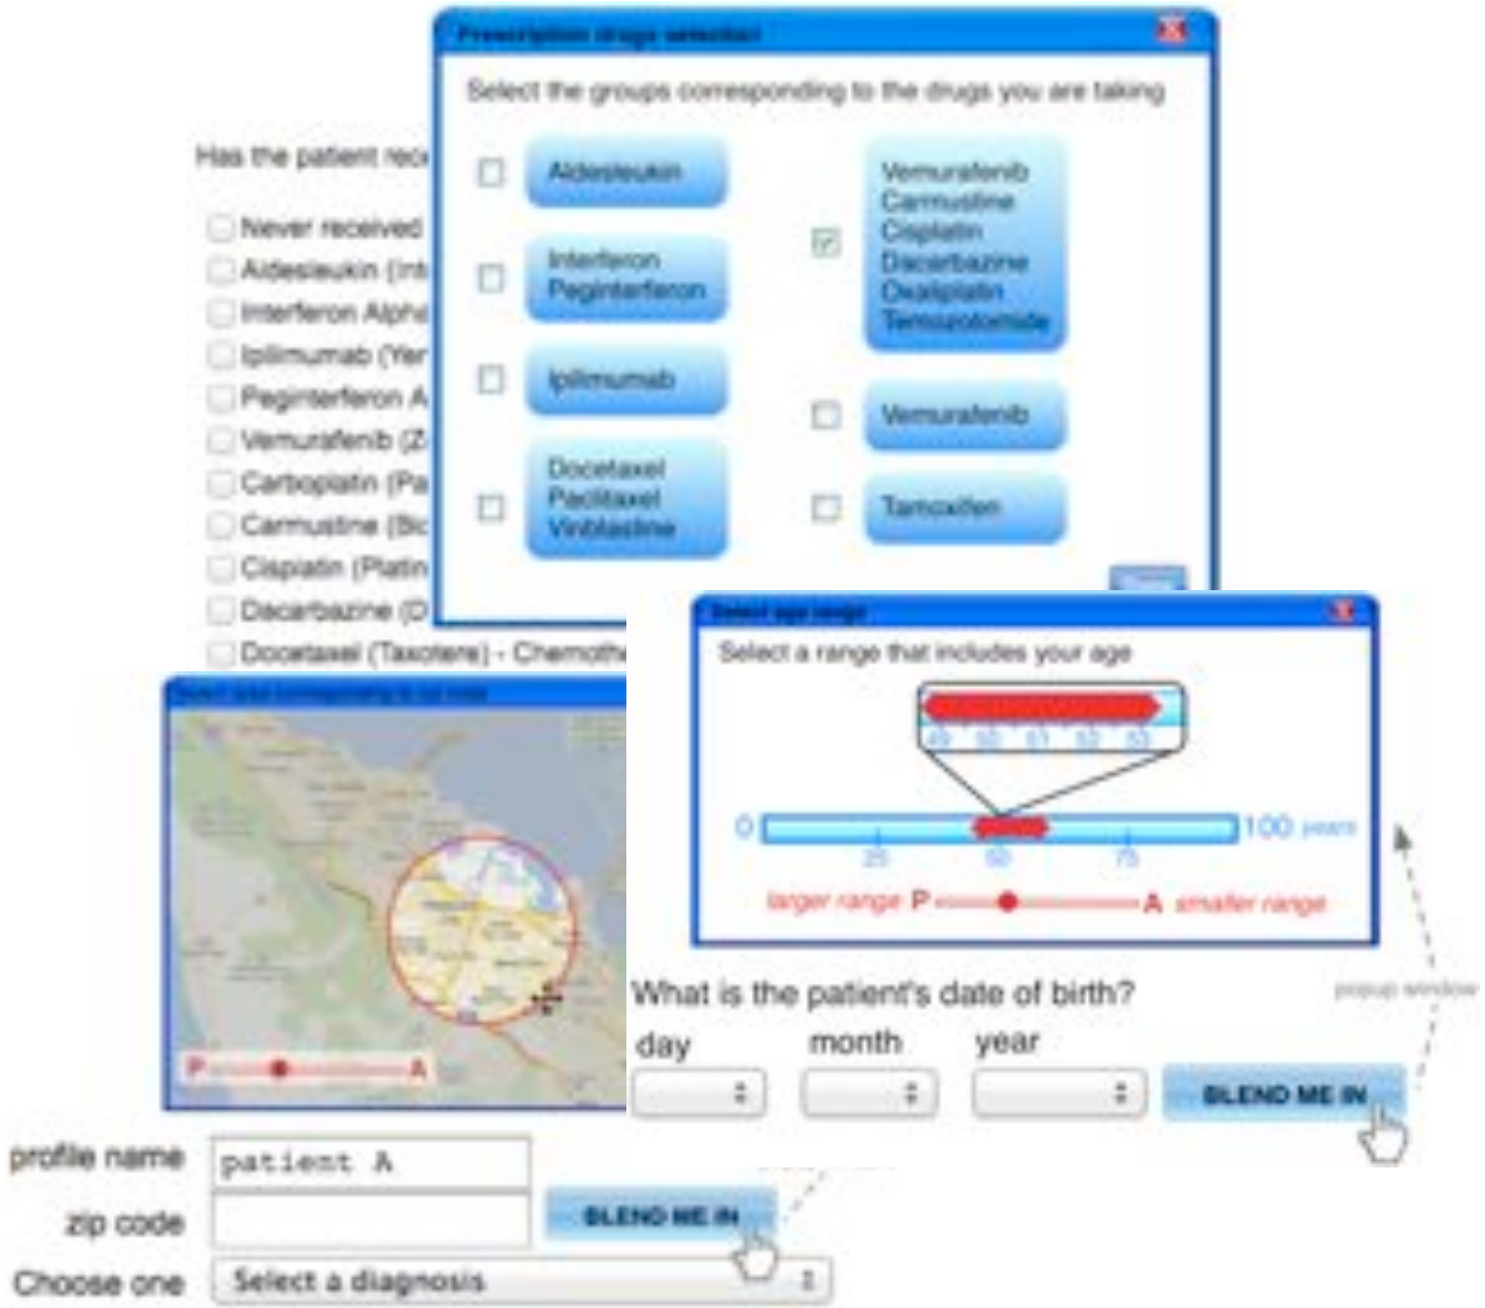

Philosophical Conclusion

Privacy vs. Utility Tradeoff

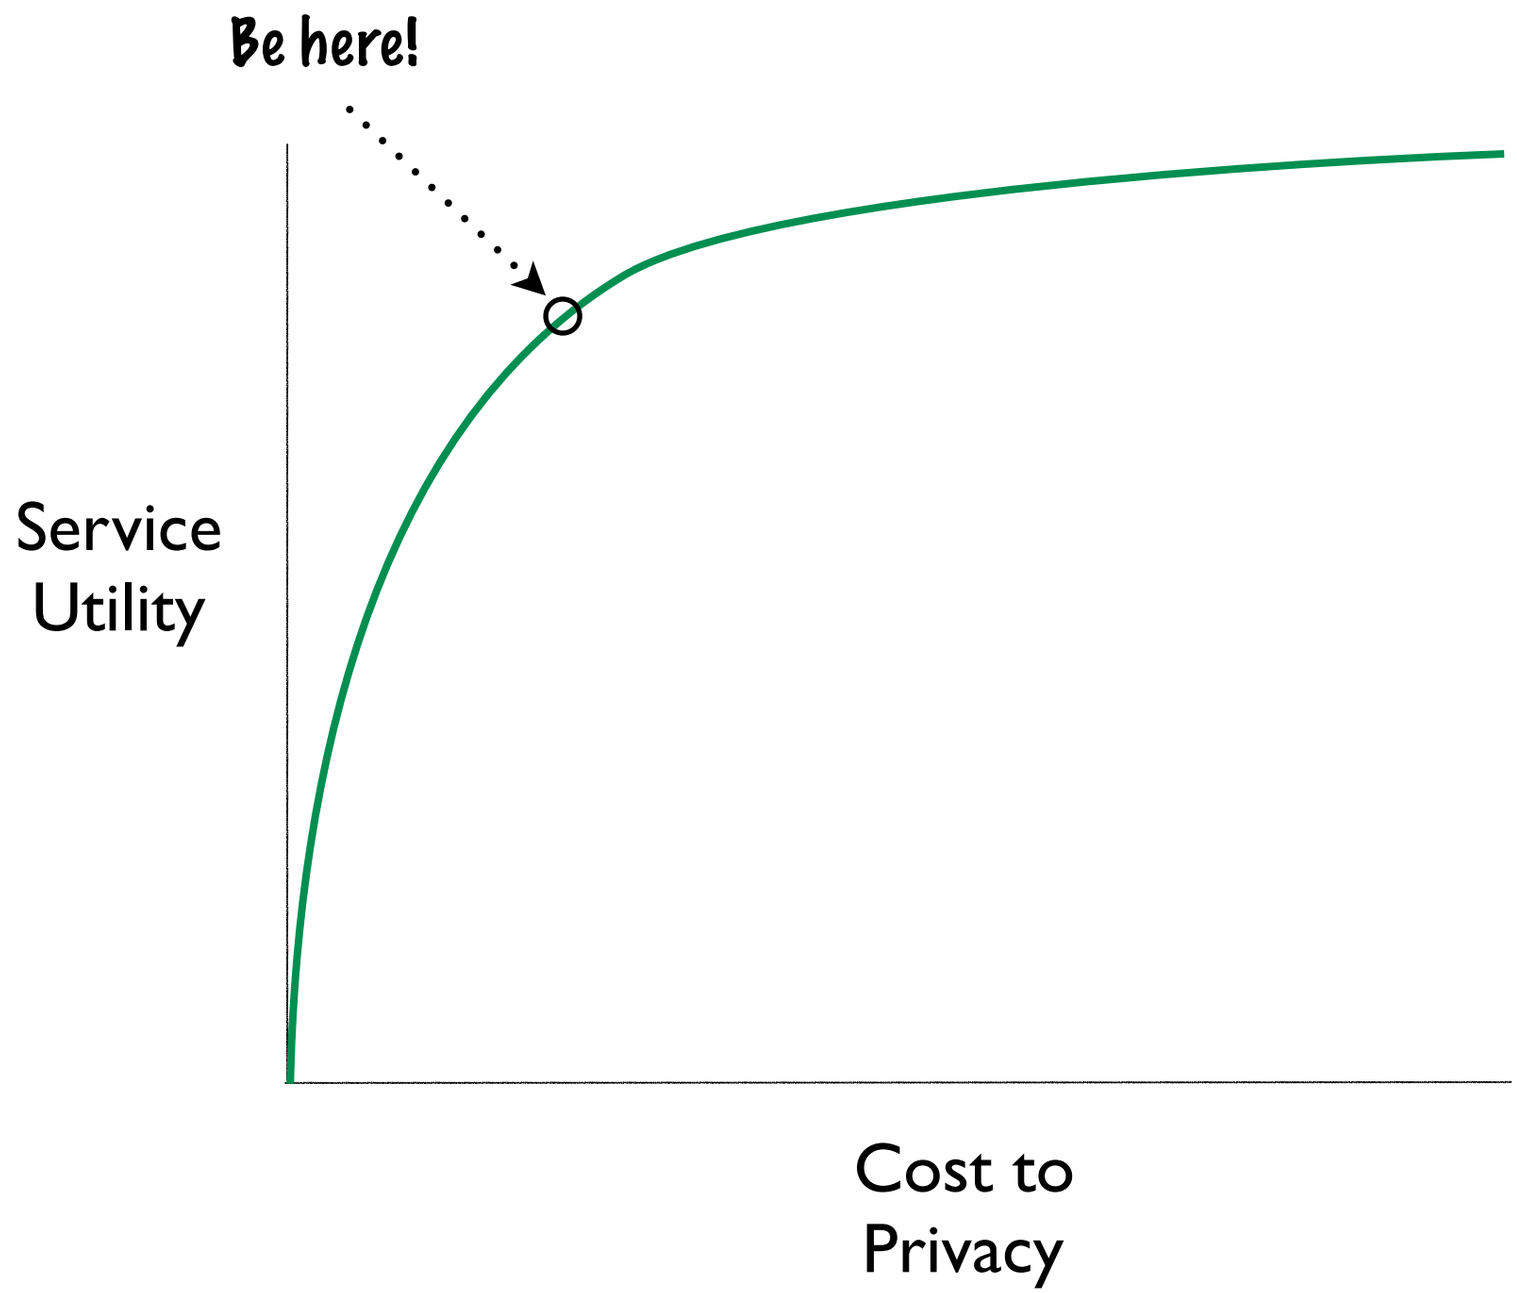

Privacy vs. Utility Tradeoff

(from Krause & Horvitz)

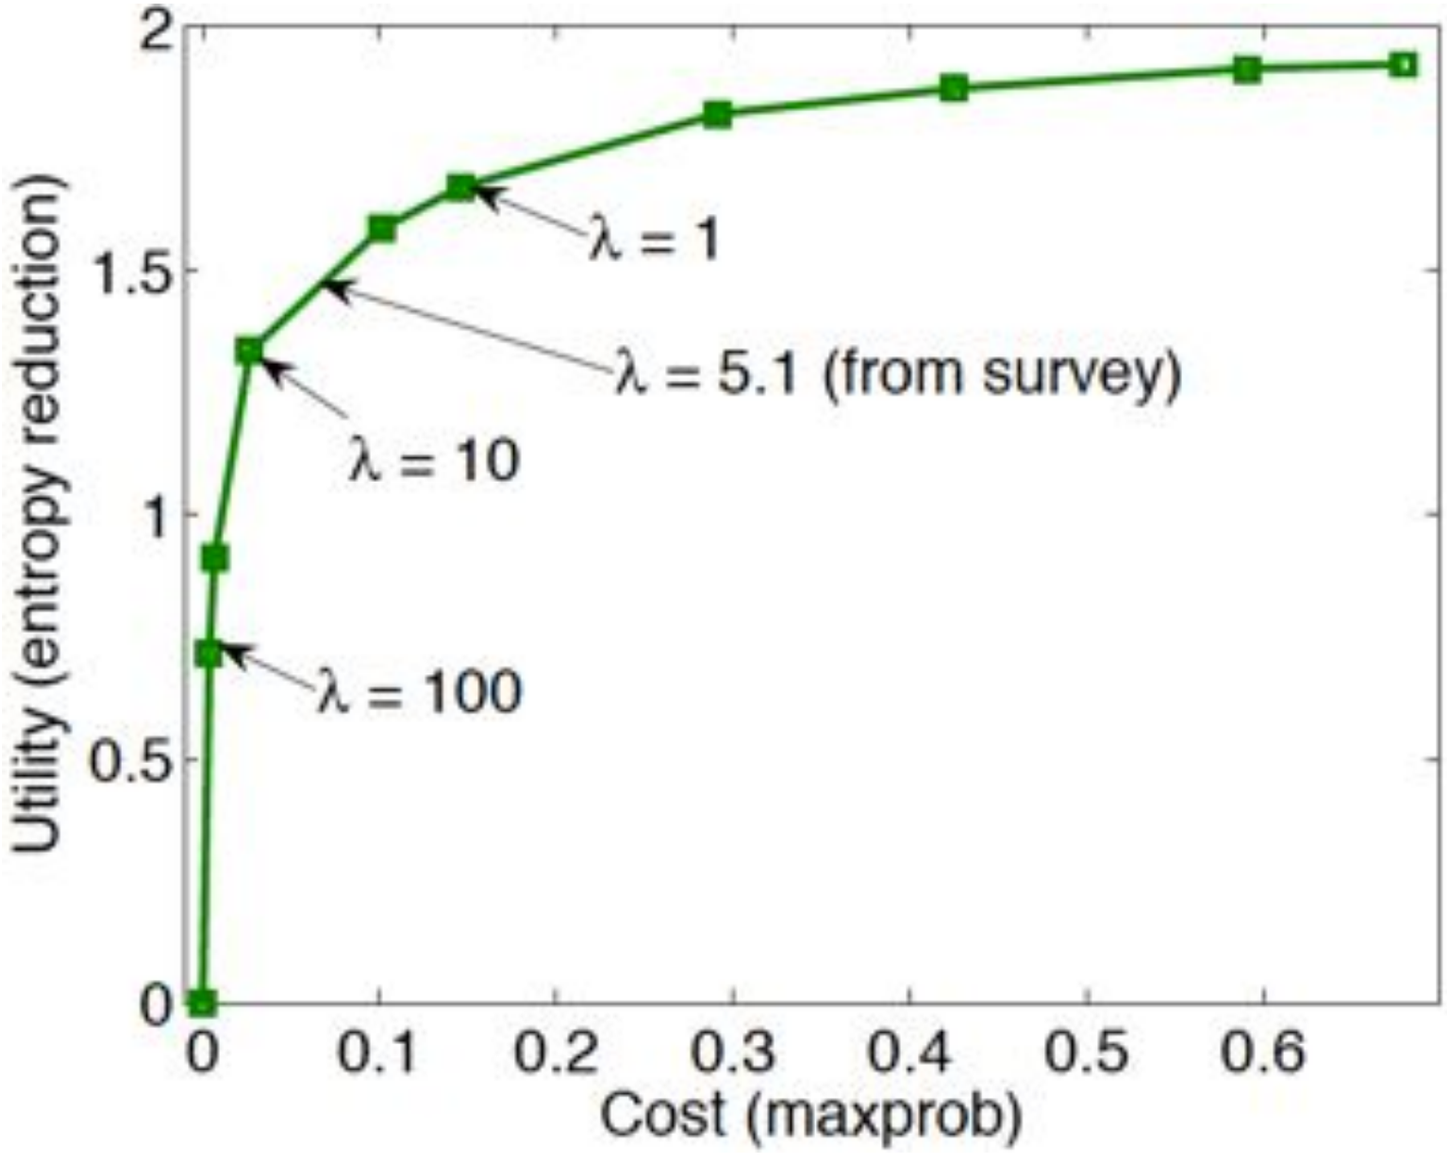

Blend Me In: The Idea

"Blend in" to the crowd

Motive: personalized services over-collect

"Forward Privacy" – consider the future

Threat model: service as adversary

Therefore, generalize input:

Perturb data to gain privacy (e.g., lie)

...but so as to preserve output accuracy

Truth is so precious that she should always be attended by a bodyguard of lies.

- Winston Churchill