Confidential Computing - Federate Private Data Analysis

Richard Nock

http://users.cecs.anu.edu.au/~rnock/

Australian National University

THE UNIVERSITY OF SYDNEY

CSIRO

www.data61.csiro.au

Confidential Computing project

Lead: Dr. Stephen Hardy

Engineering	Research	Business
Mr. Brian Thorne	Dr. Richard Nock	Mr. Warren Bradey
Dr. Mentari Djatmiko	Mr. Giorgio Patrini	Ms. Shelley Copsey
Dr. Guillaume Smith	Dr. Roksana Borelli	
Dr. Wilko Henecka	Dr. Arik Friedman	
Dr. Hamish Ivey-Law	Pr. Hugh Durrant-Whyte	
Dr Max Ott		

+ PhD students / interns: Raphaël Canyasse (Ecole Polytechnique), Alexis Le Dantec (Ecole Polytechnique), Giorgio Patrini (ANU)

Outline

Outline

Confidential Computing
/
N1 Analytics

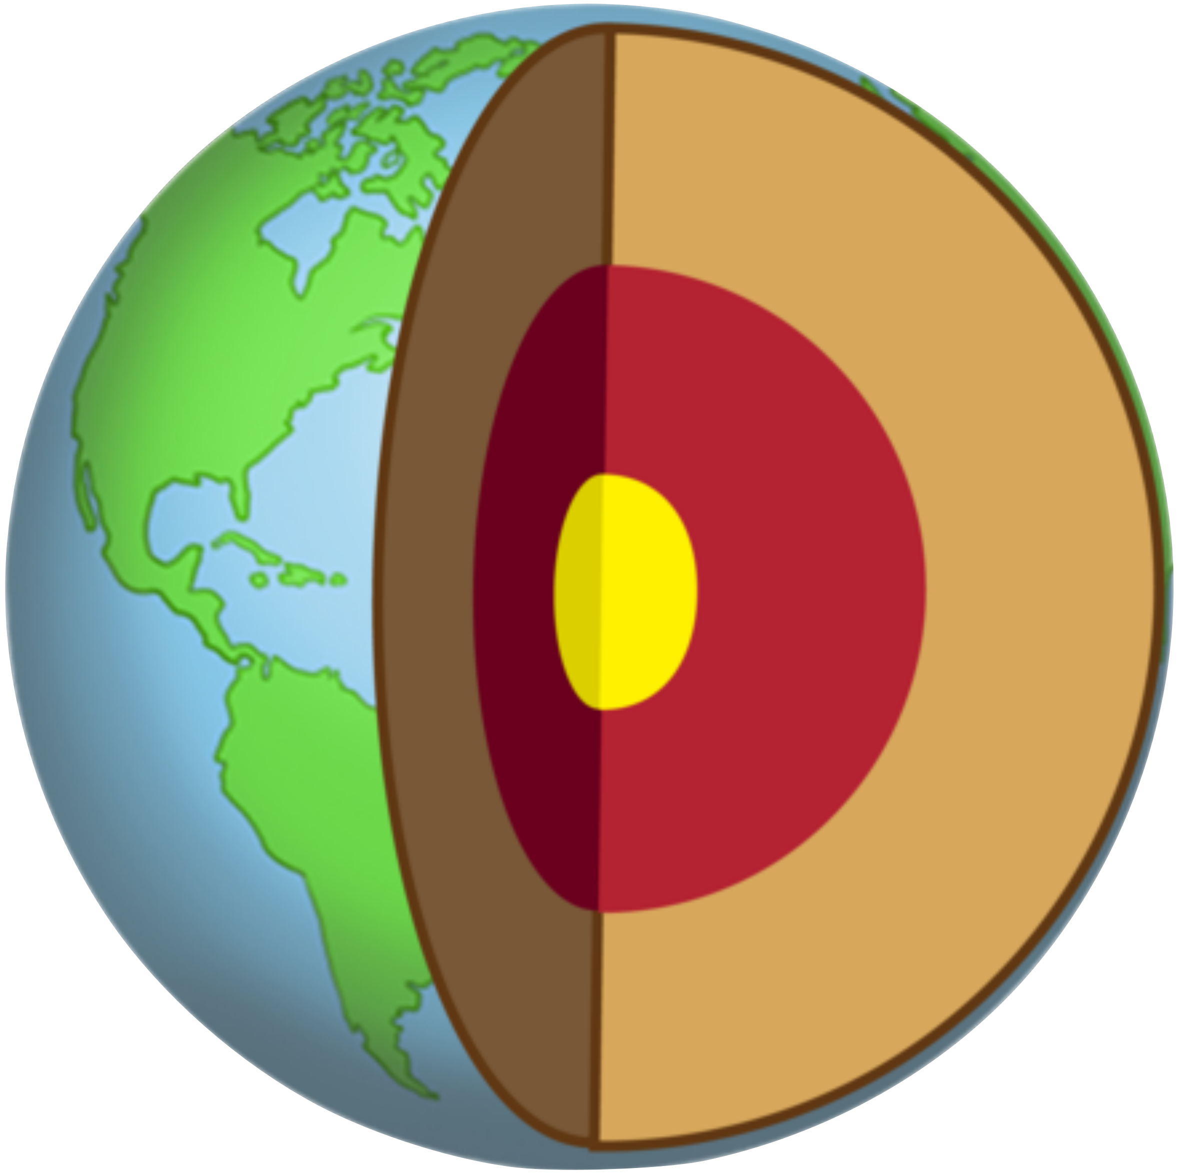

Outline

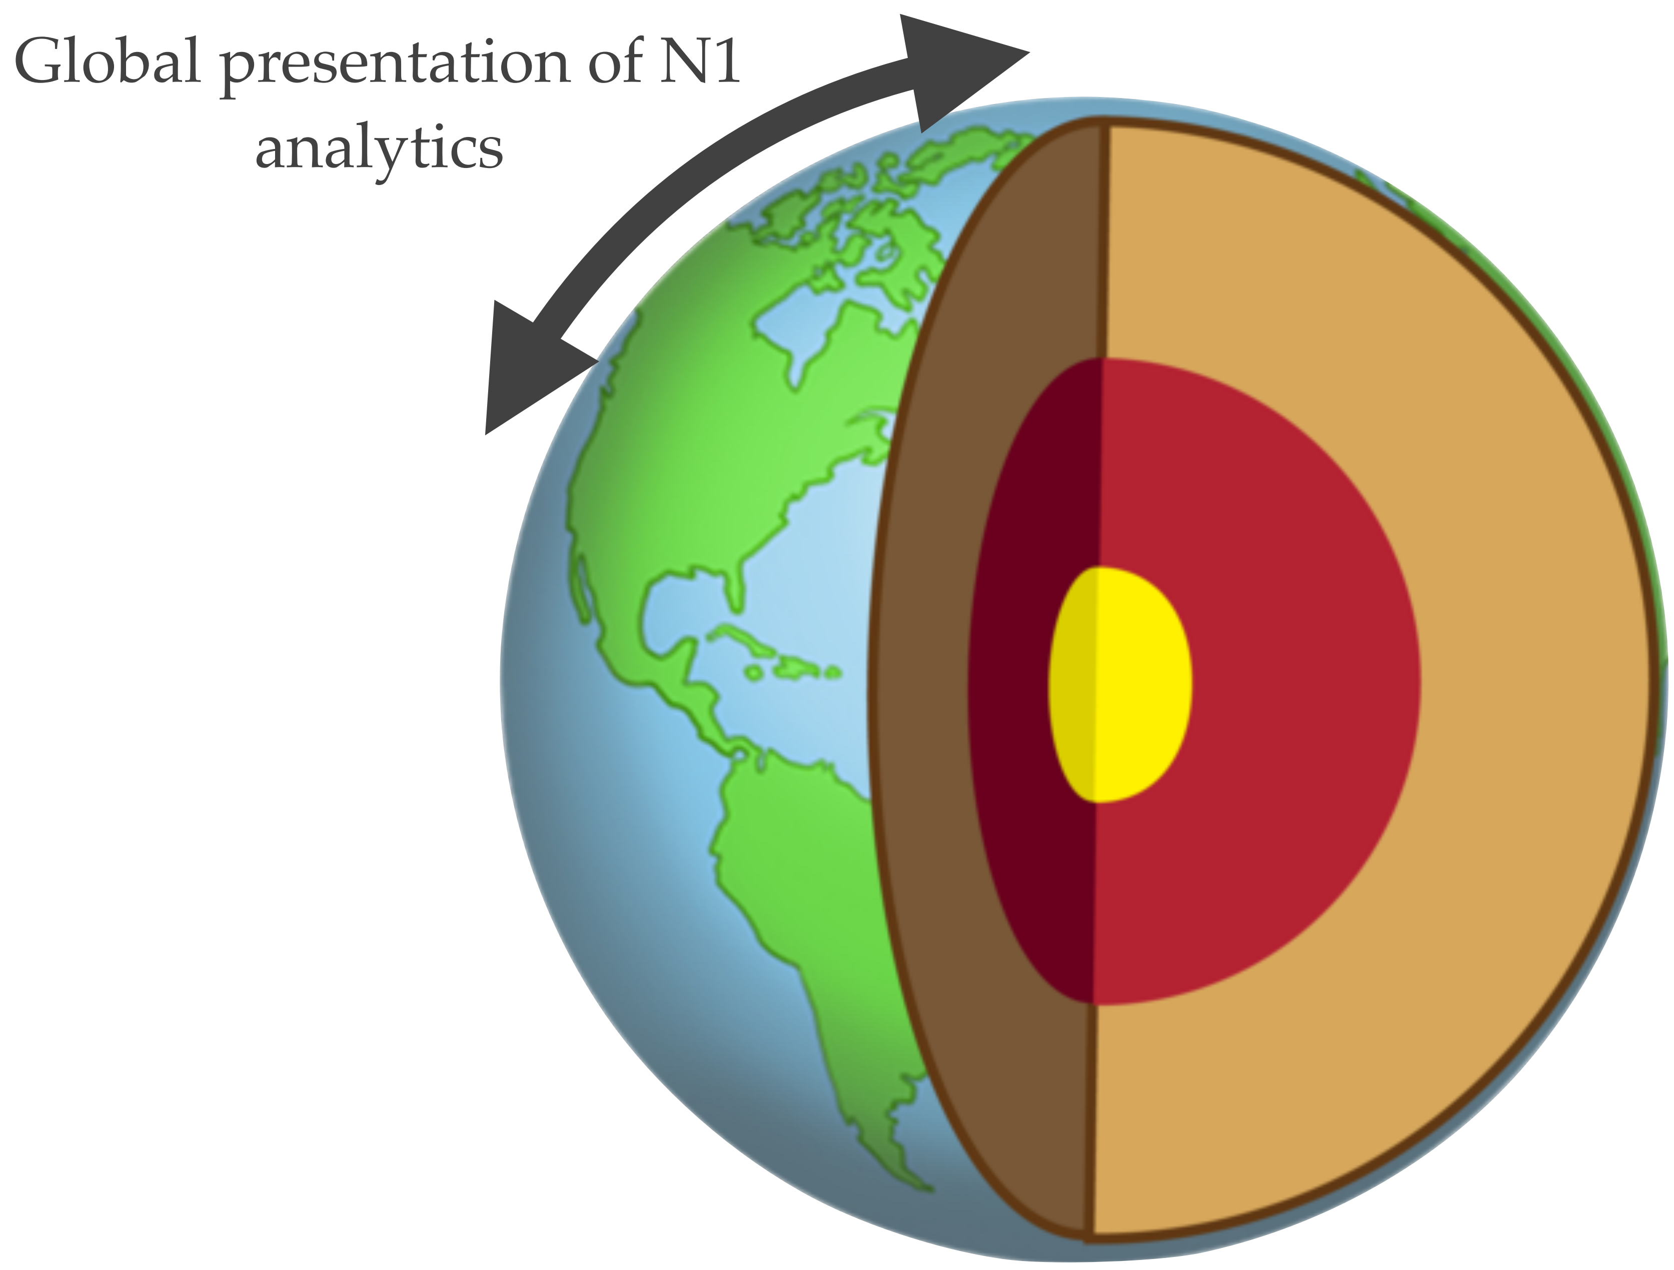

Outline

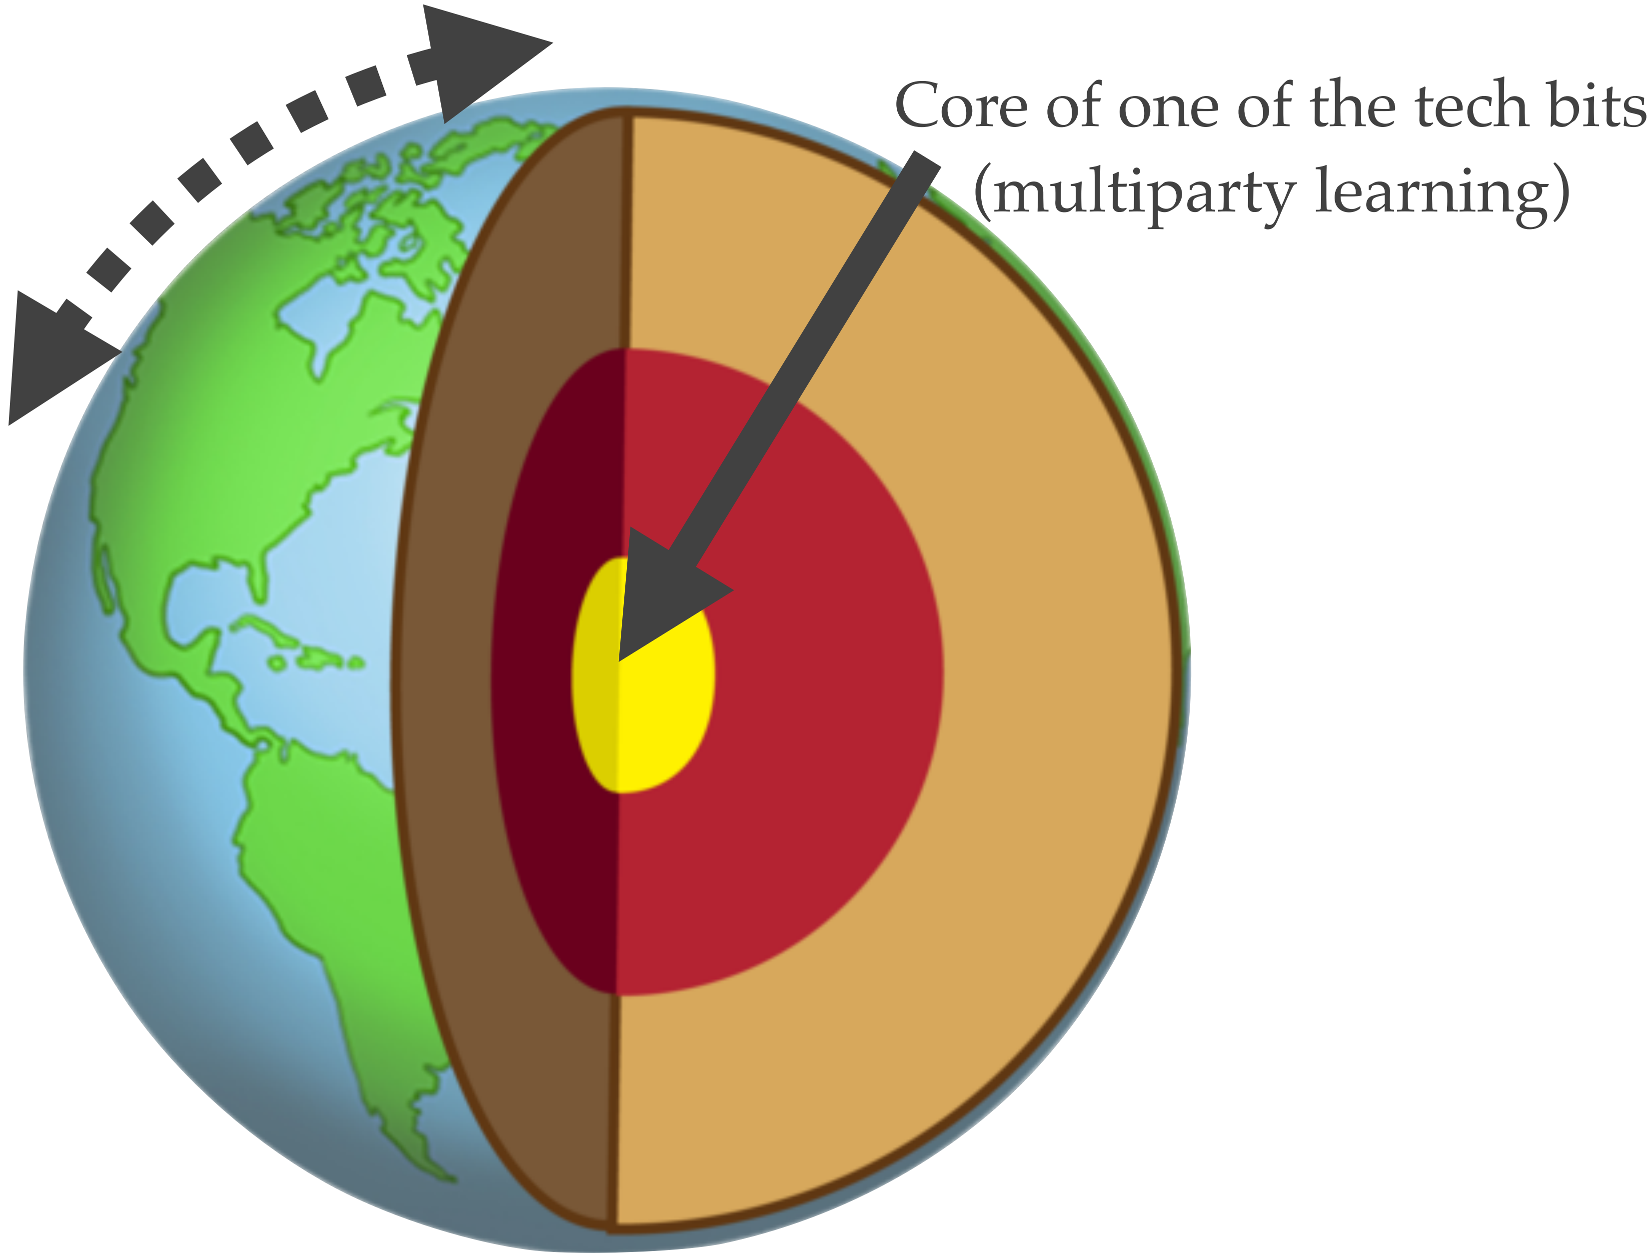

Outline

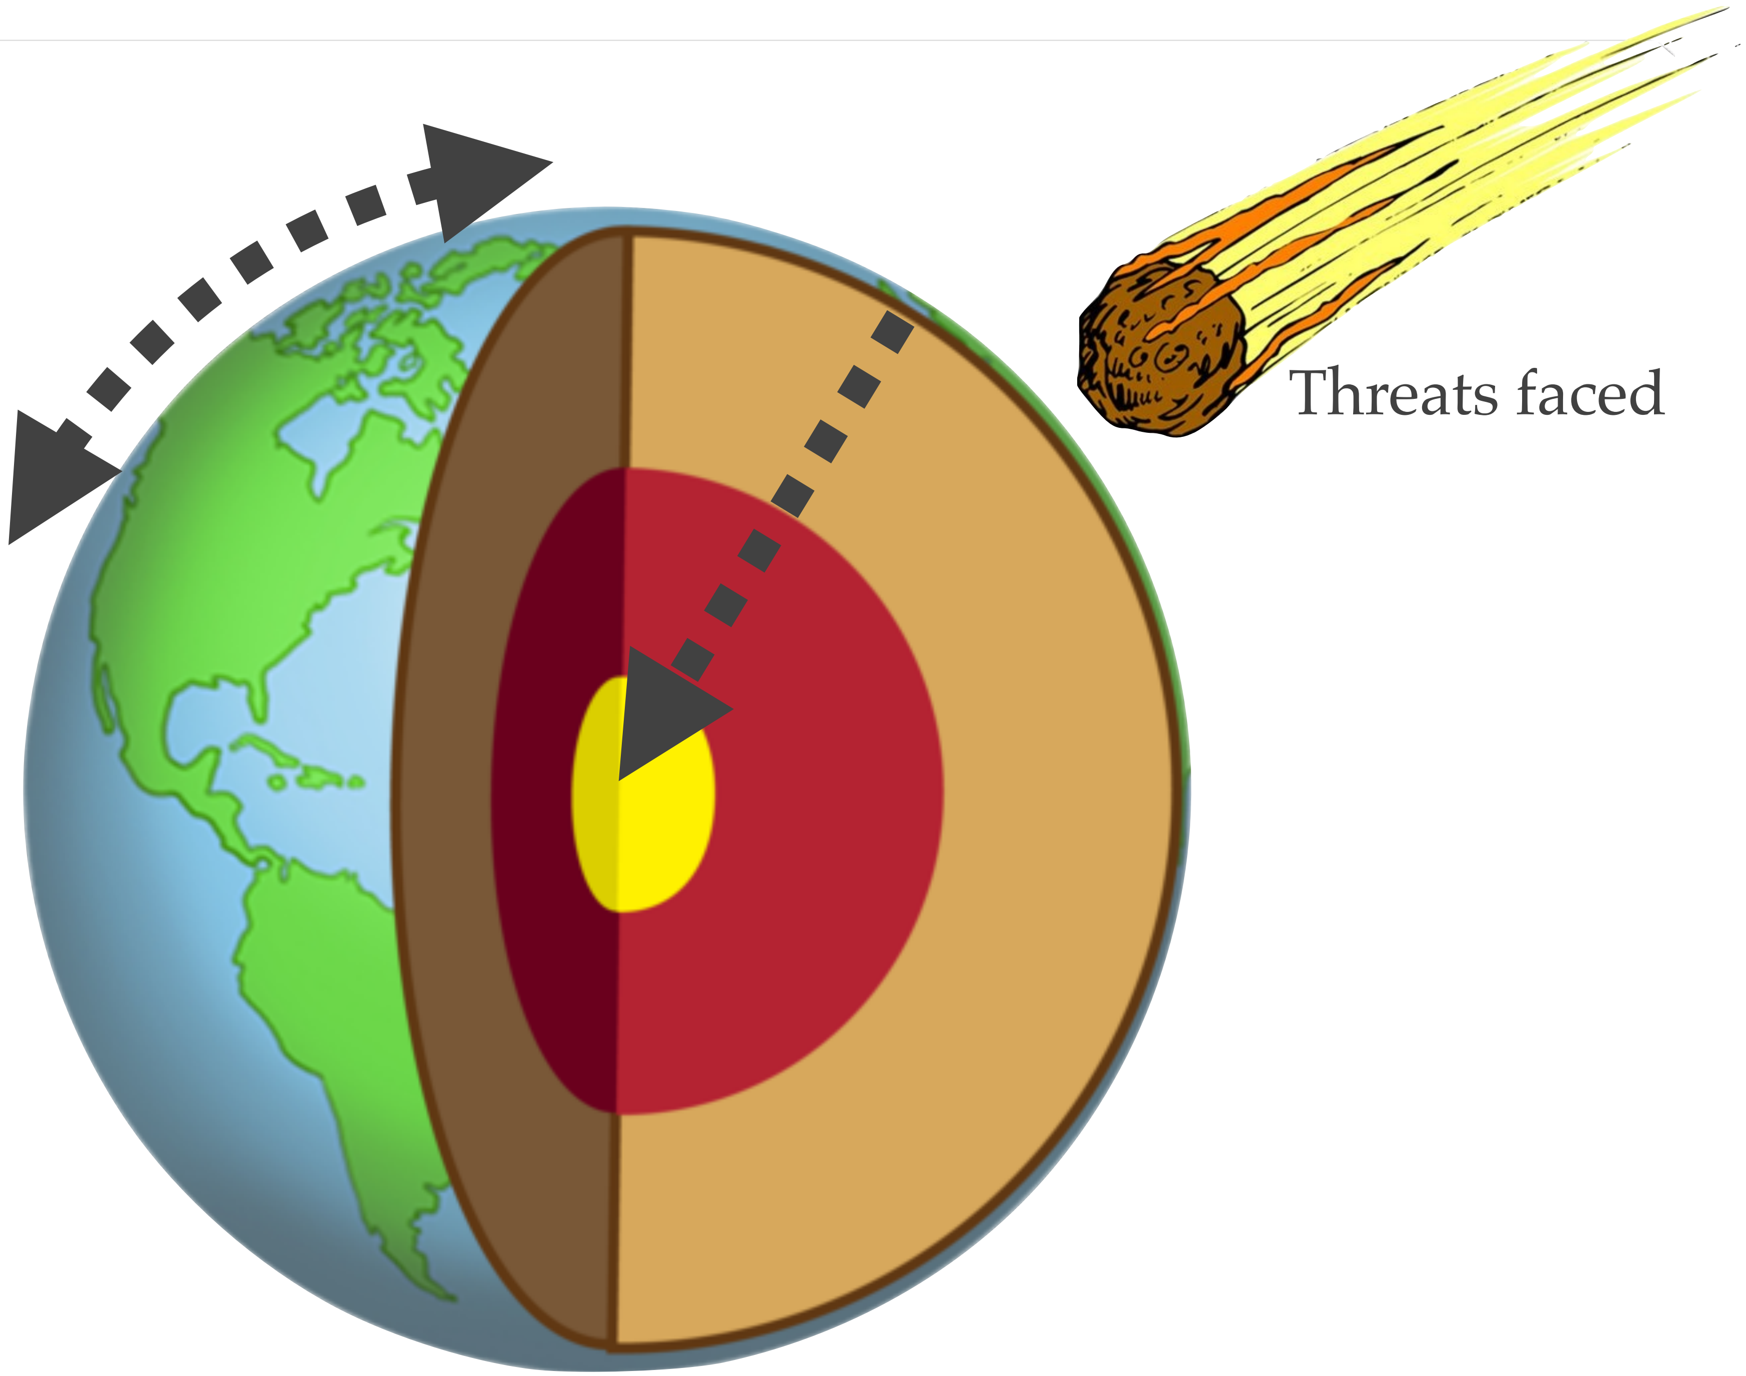

Threats

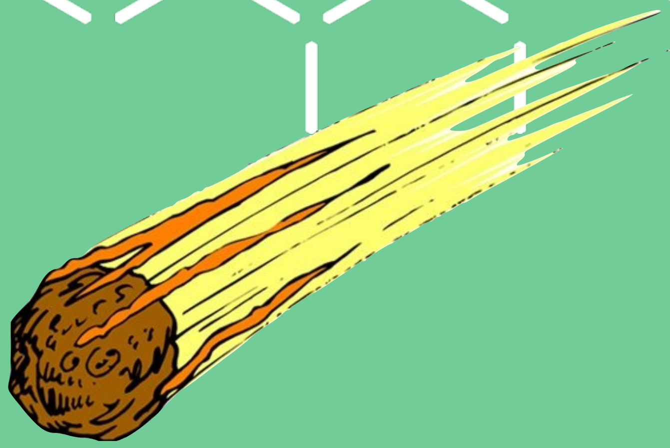

Making "protected" data public…

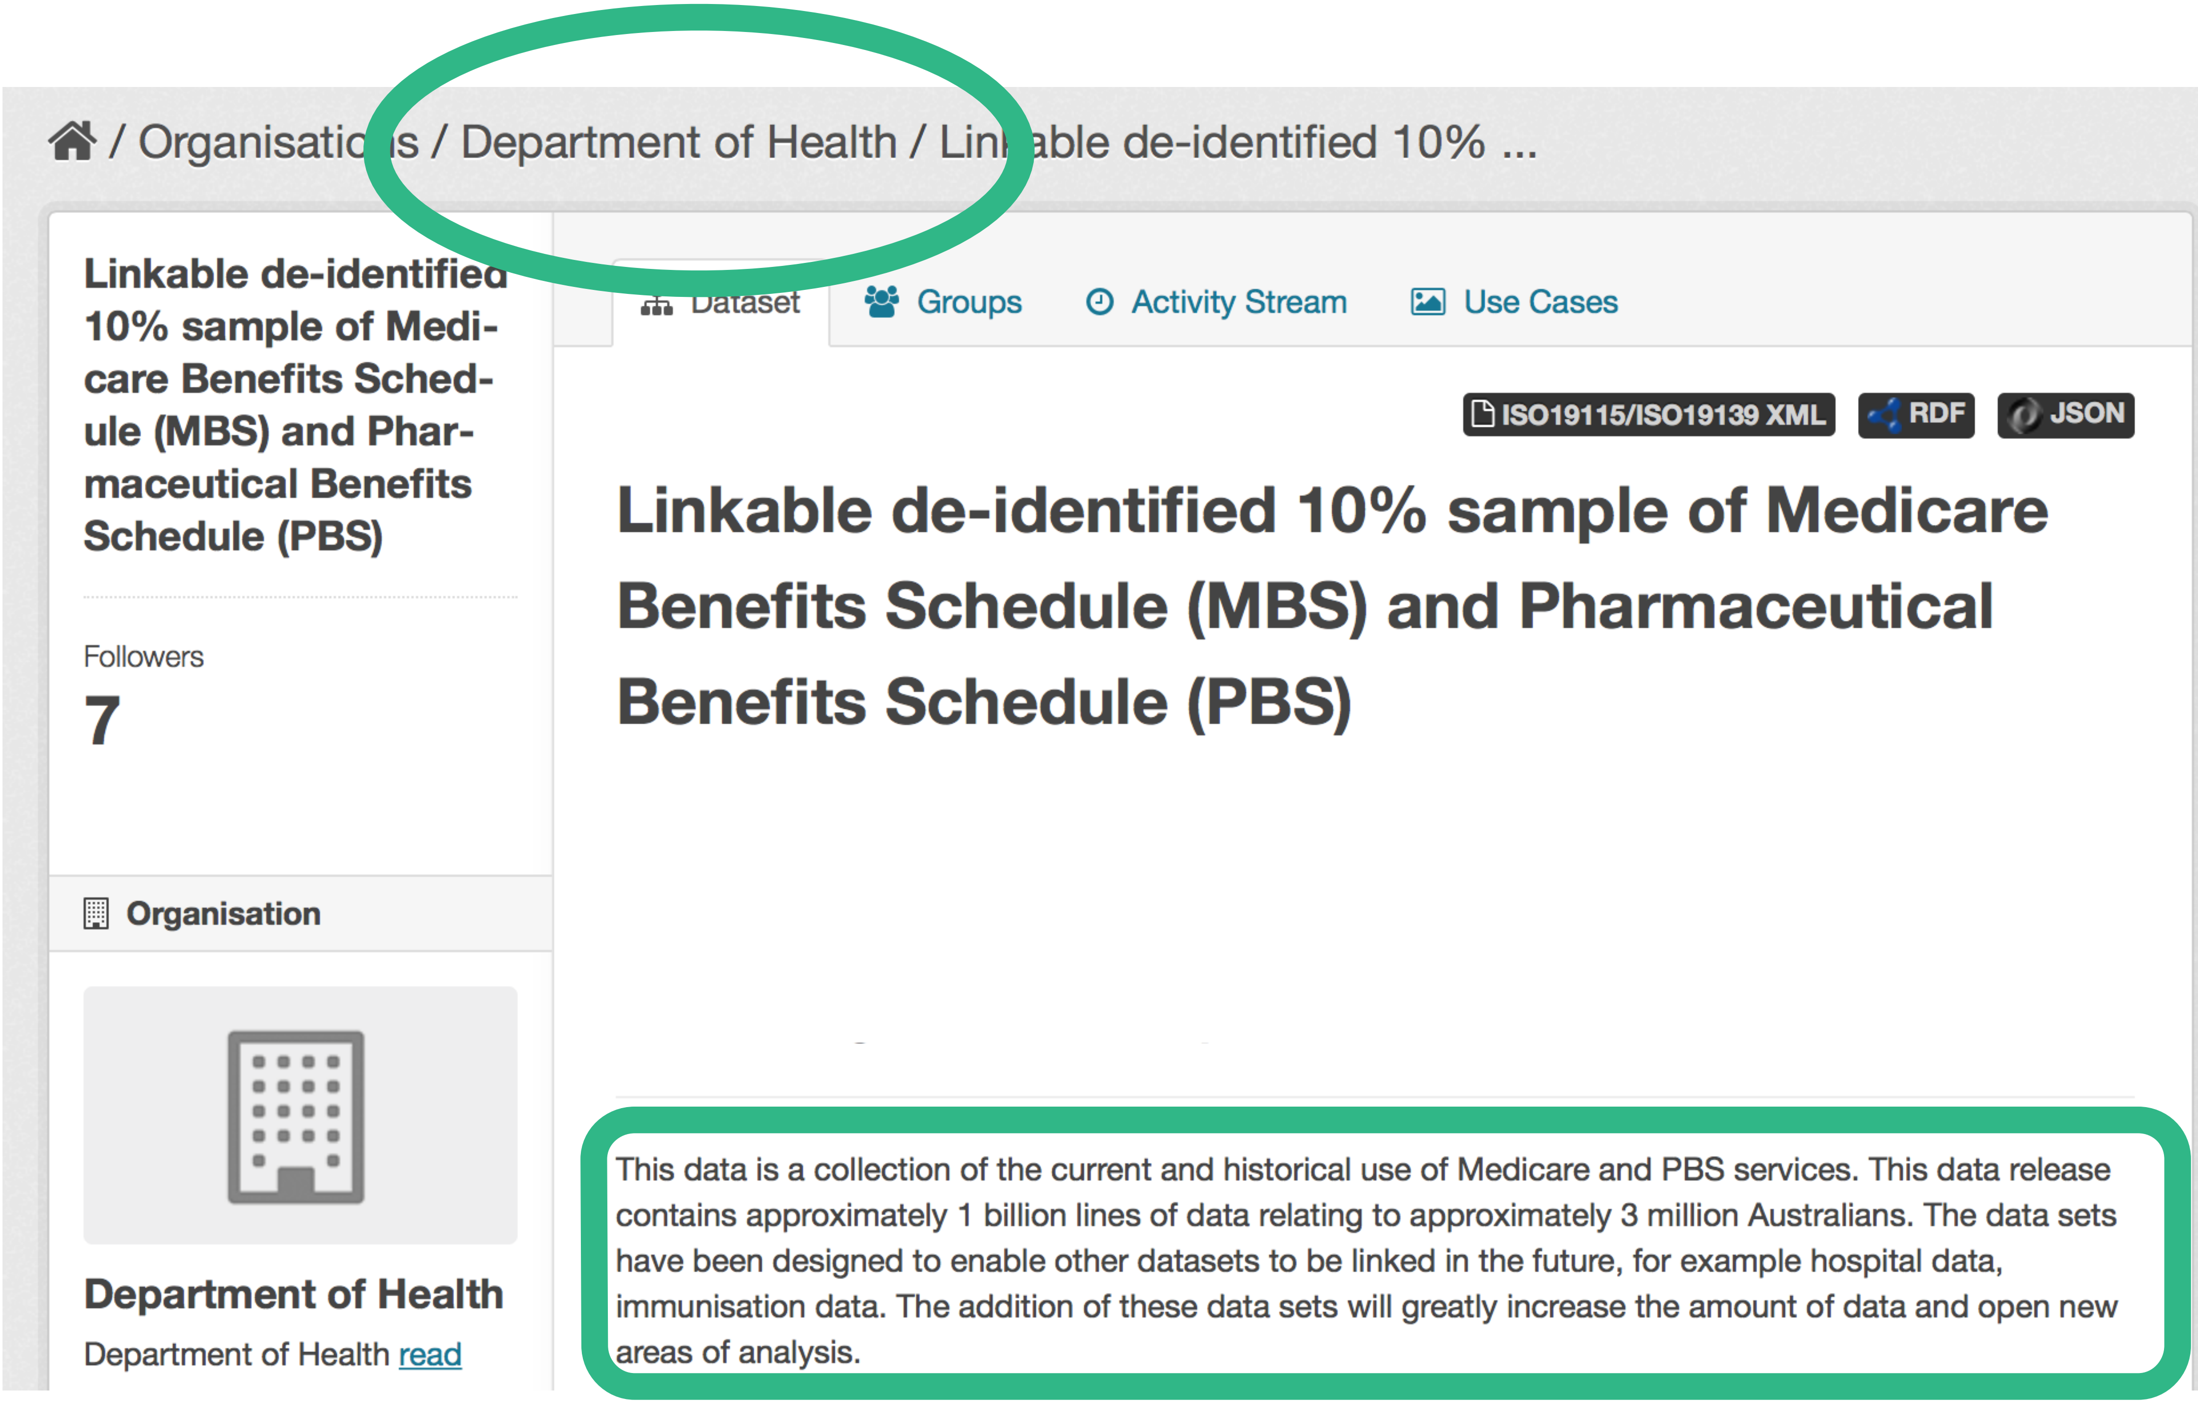

…even with "safe" techniques…

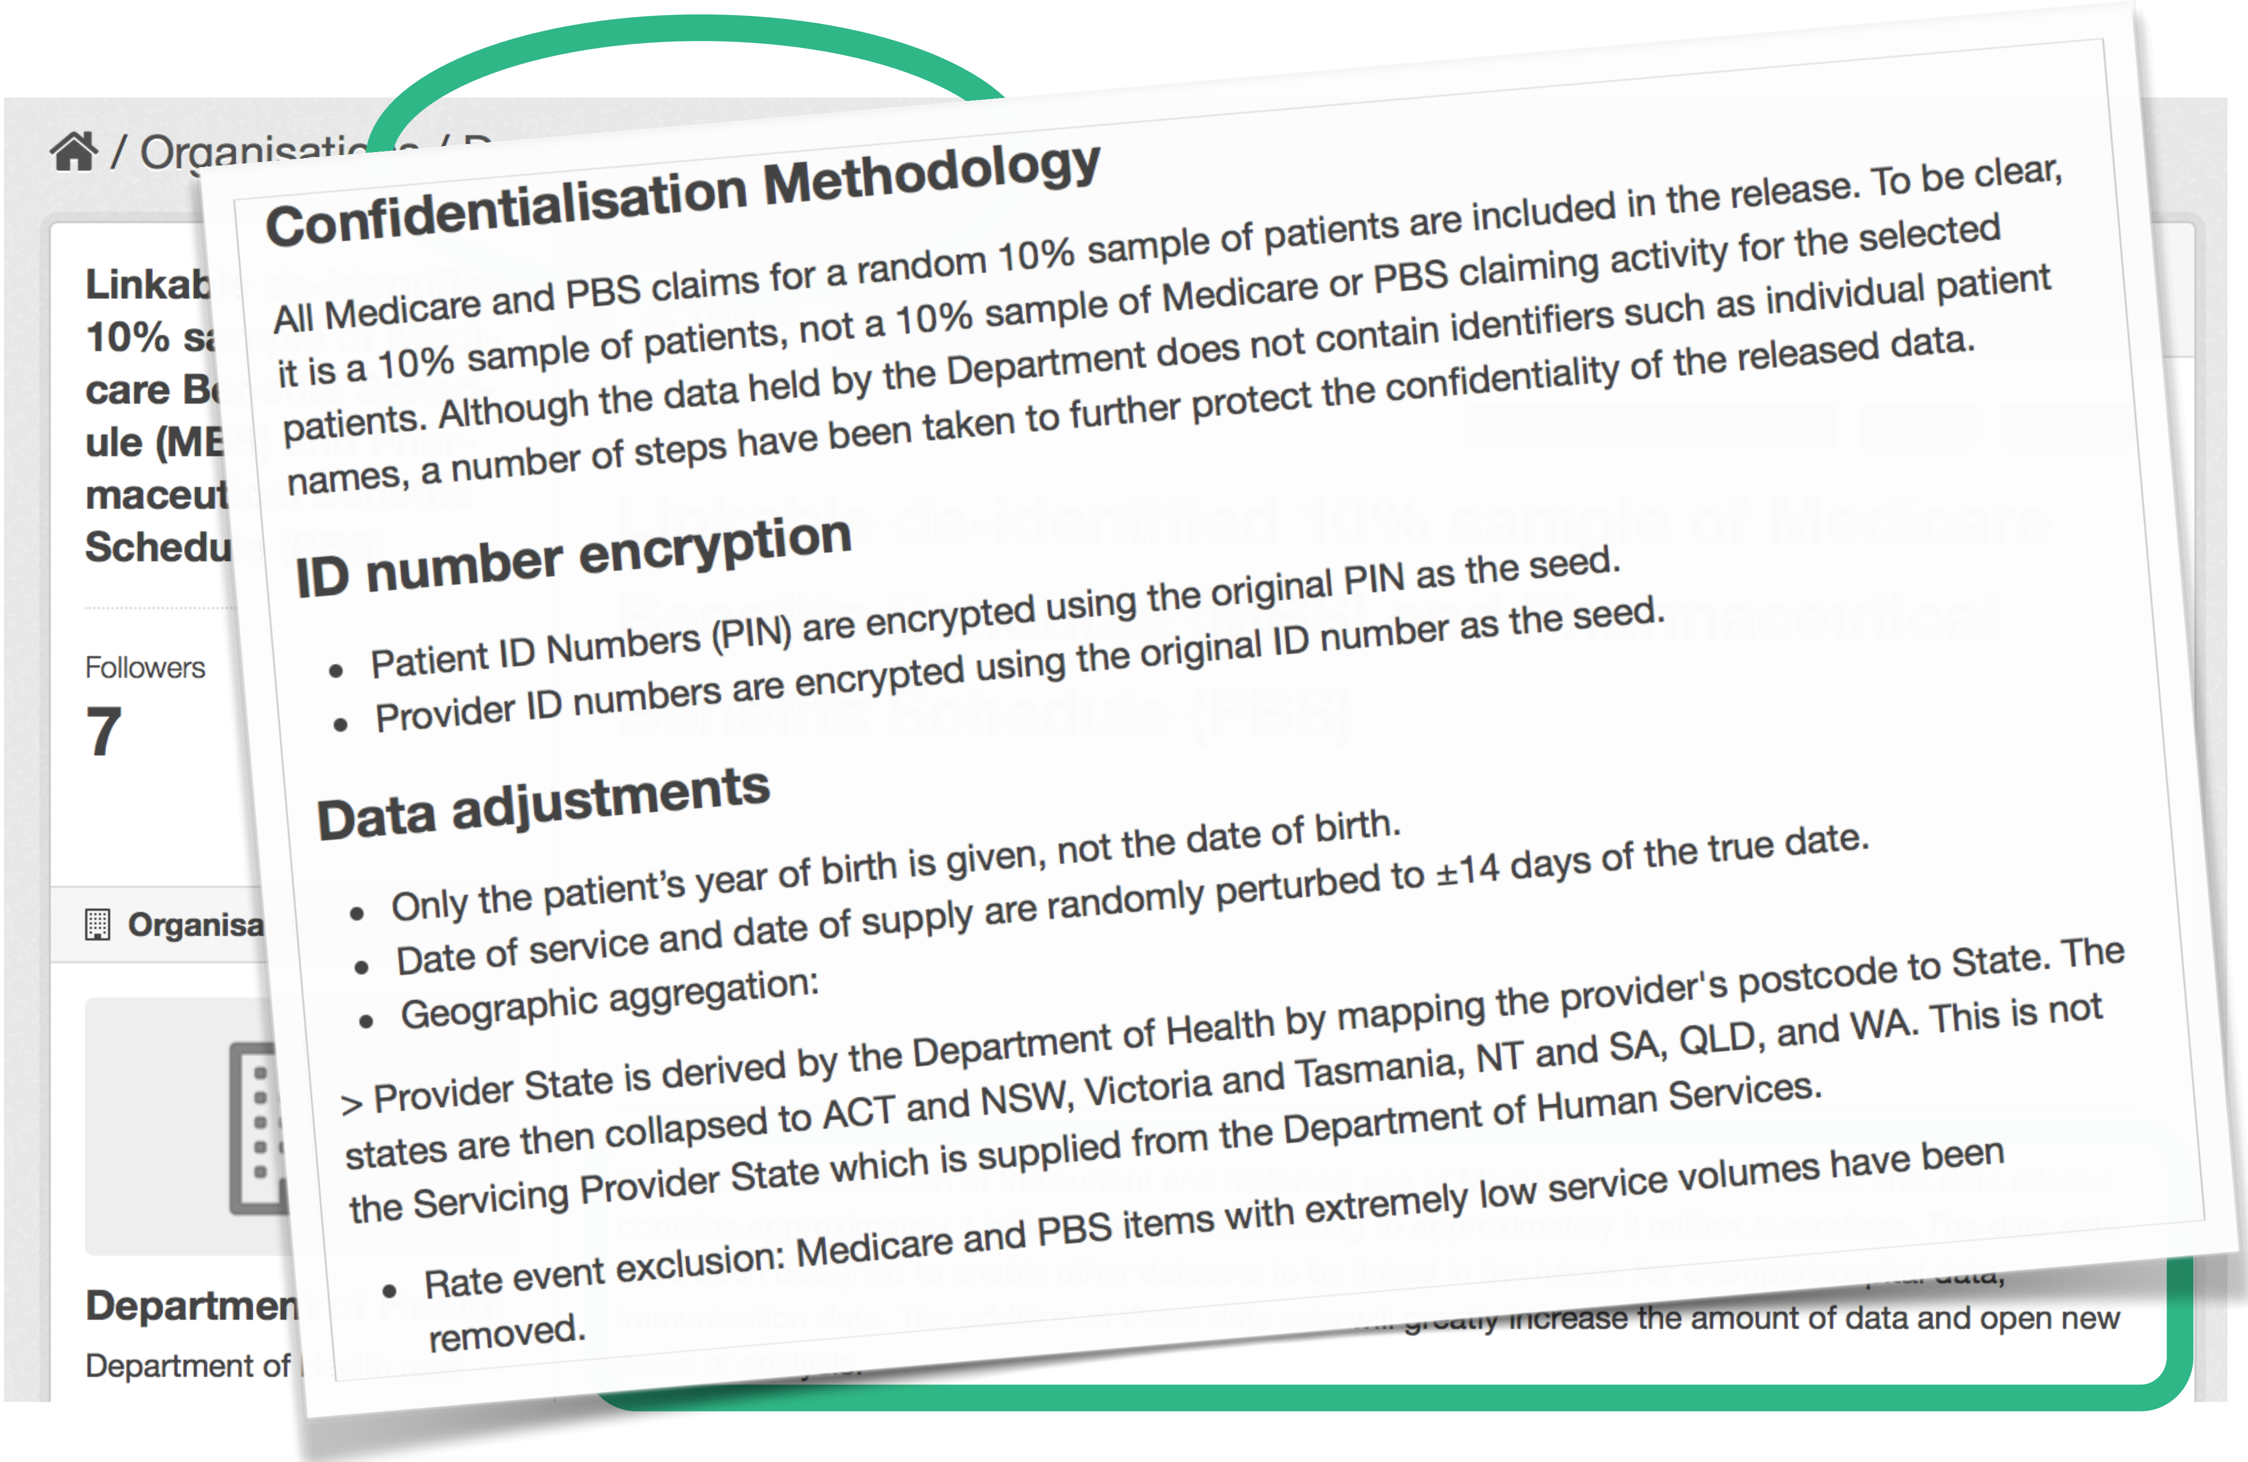

Confidentialisation Methodology

All Medicare and PBS claims for a random 10% sample of patients are included in the release. To be clear, it is a 10% sample of patients, not a 10% sample of Medicare or PBS claiming activity for the selected patients. Although the data held by the Department does not contain identifiers such as individual patient names, a number of steps have been taken to further protect the confidentiality of the released data.

ID number encryption

- Patient ID Numbers (PIN) are encrypted using the original PIN as the seed.
- Provider ID numbers are encrypted using the original ID number as the seed.

Data adjustments

- Only the patient's year of birth is given, not the date of birth.
- Date of service and date of supply are randomly perturbed to ±14 days of the true date.
- Geographic aggregation:

> Provider State is derived by the Department of Health by mapping the provider's postcode to State. The states are then collapsed to ACT and NSW, Victoria and Tasmania, NT and SA, QLD, and WA. This is not the Servicing Provider State which is supplied from the Department of Human Services.

- Rate event exclusion: Medicare and PBS items with extremely low service volumes have been removed.

…may lead to problems…

/ Organisations / Department of Health / Linkable de-identified 10% ...

Linkable de-identified 10% sample of Medicare Benefits Schedule (MBS) and Pharmaceutical Benefits Schedule (PBS)

Followers

7

Organisation

Department of Health

Department of Health read

Dataset | Groups | Activity Stream | Use Cases

ISO19115/ISO19139 XML | RDF | JSON

Linkable de-identified 10% sample of Medicare Benefits Schedule (MBS) and Pharmaceutical Benefits Schedule (PBS)

This data is temporarily unavailable. The Department of Health is currently working on the dataset and hope to have it restored and available again as soon as possible.

This data is a collection of the current and historical use of Medicare and PBS services. This data release contains approximately 1 billion lines of data relating to approximately 3 million Australians. The data sets have been designed to enable other datasets to be linked in the future, for example hospital data, immunisation data. The addition of these data sets will greatly increase the amount of data and open new areas of analysis.

…without extra care…

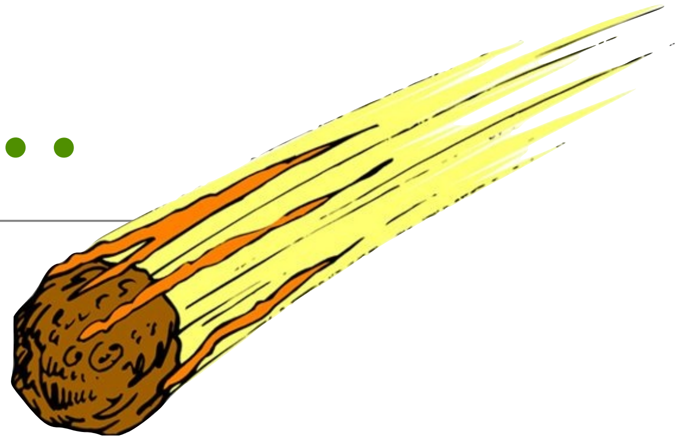

UNDERSTANDING THE MATHS IS CRUCIAL FOR PROTECTING PRIVACY

Publishing data can bring benefits, but it also can be a great risk to privacy

By Dr Chris Culnane, Dr Benjamin Rubinstein and Dr Vanessa Teague, Department of Computing and Information Systems, University of Melbourne

…on the possible attacks…

UNDERSTANDING THE MATHS IS CRUCIAL FOR PROTECTING PRIVACY

Publishing data can bring benefits,
priva

By Dr Chris Culnane, Dr Benjamin Rubinstein and Dr
Information Systems, Univ

Linkage attacks use the unencrypted data to identify people by linking the record with other known information; and

Cryptographic attacks reverse the encryption algorithm to recover encrypted data.

…and then comes (bad) buzz

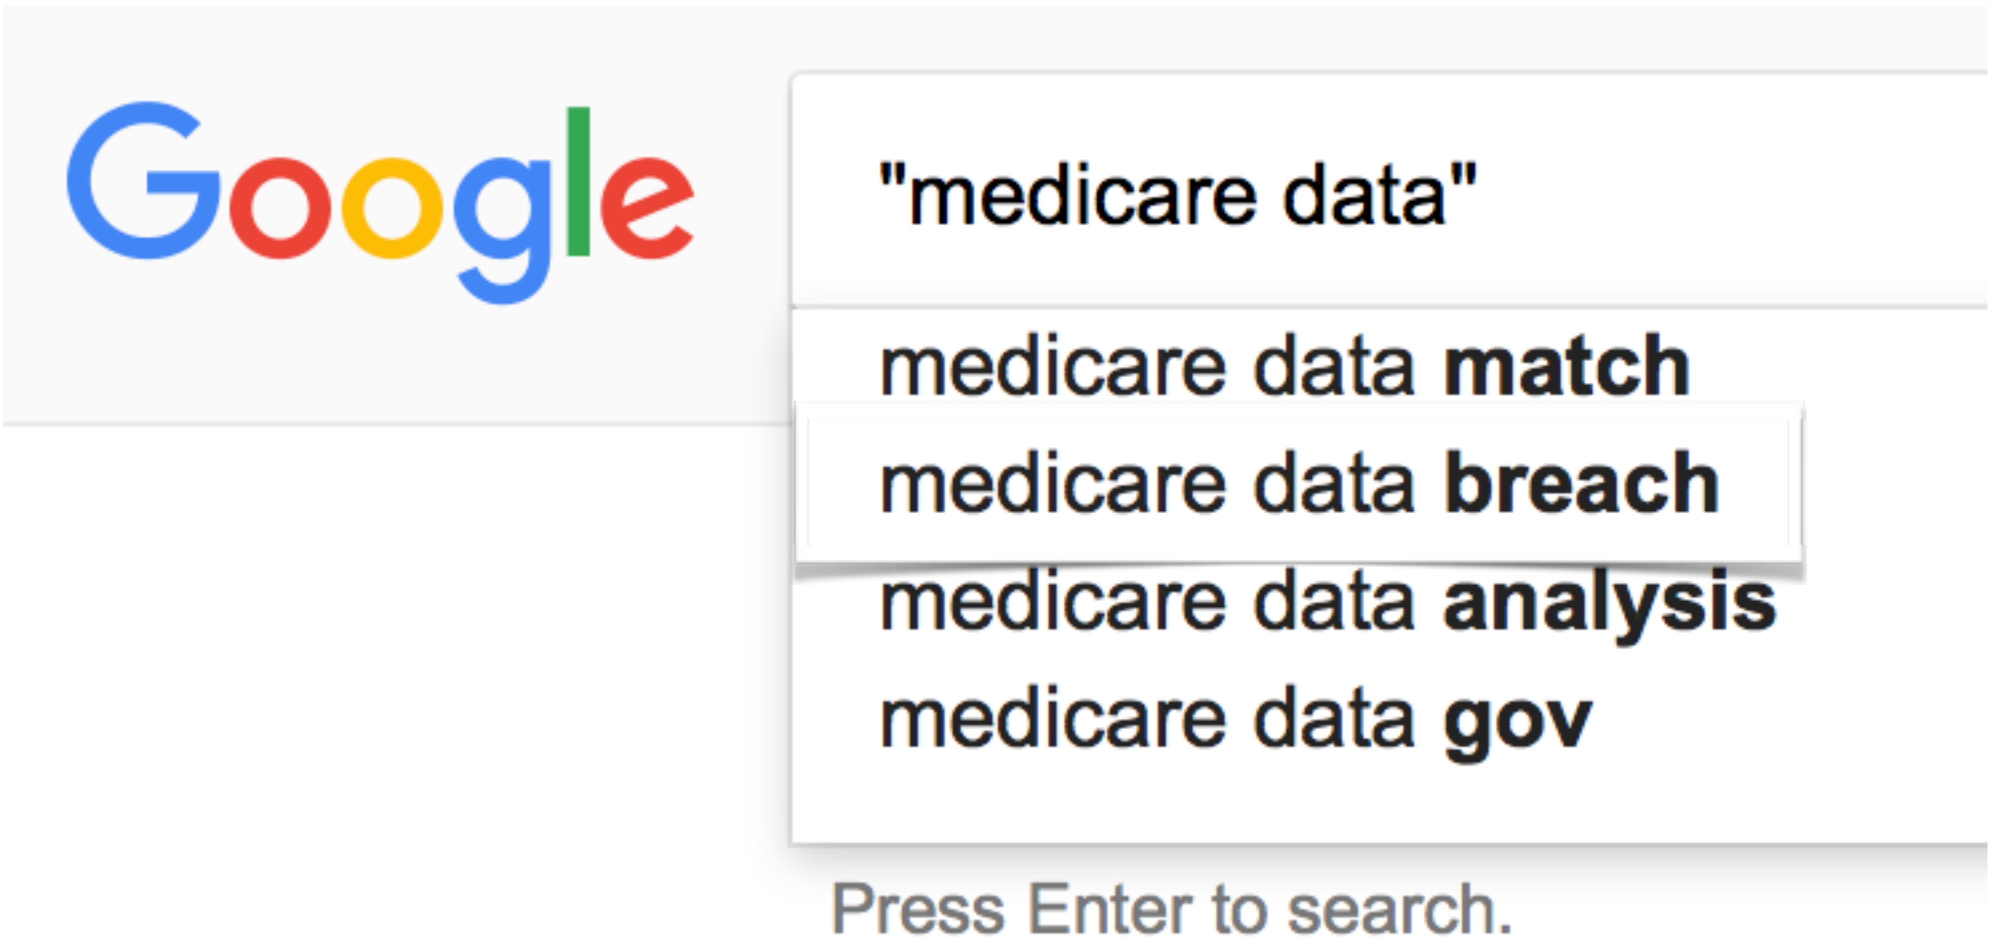

…and then comes (bad) buzz

NEWS

Just In | US Election | Australia | World | Business | Sport | Arts | Analysis & Opinion | Progr

Print | Email | Facebook | Twitter | More

Medicare dataset pulled after academics find breach of doctor details possible

By political reporter Stephanie Anderson

Updated 29 Sep 2016, 6:51am

The Health Department has removed Medicare data from its website amid an investigation into whether personal information has been compromised.

ABC news

…and then comes (bad) buzz


AUSTRALIAN Medicine | News | Health On The Hill | Columns | Opinion | Research | Features | Lifestyle | Letter | Information

MEDICARE DATA BREACH PROMPTS LAW CHANGE

04 Oct 2016

The Federal Government has moved to tighten privacy laws after doctor provider numbers were disclosed in a breach of security around Medicare and Pharmaceutical Benefit Scheme data.

compromised.

The Australian Medical Association

…and then comes (bad) buzz

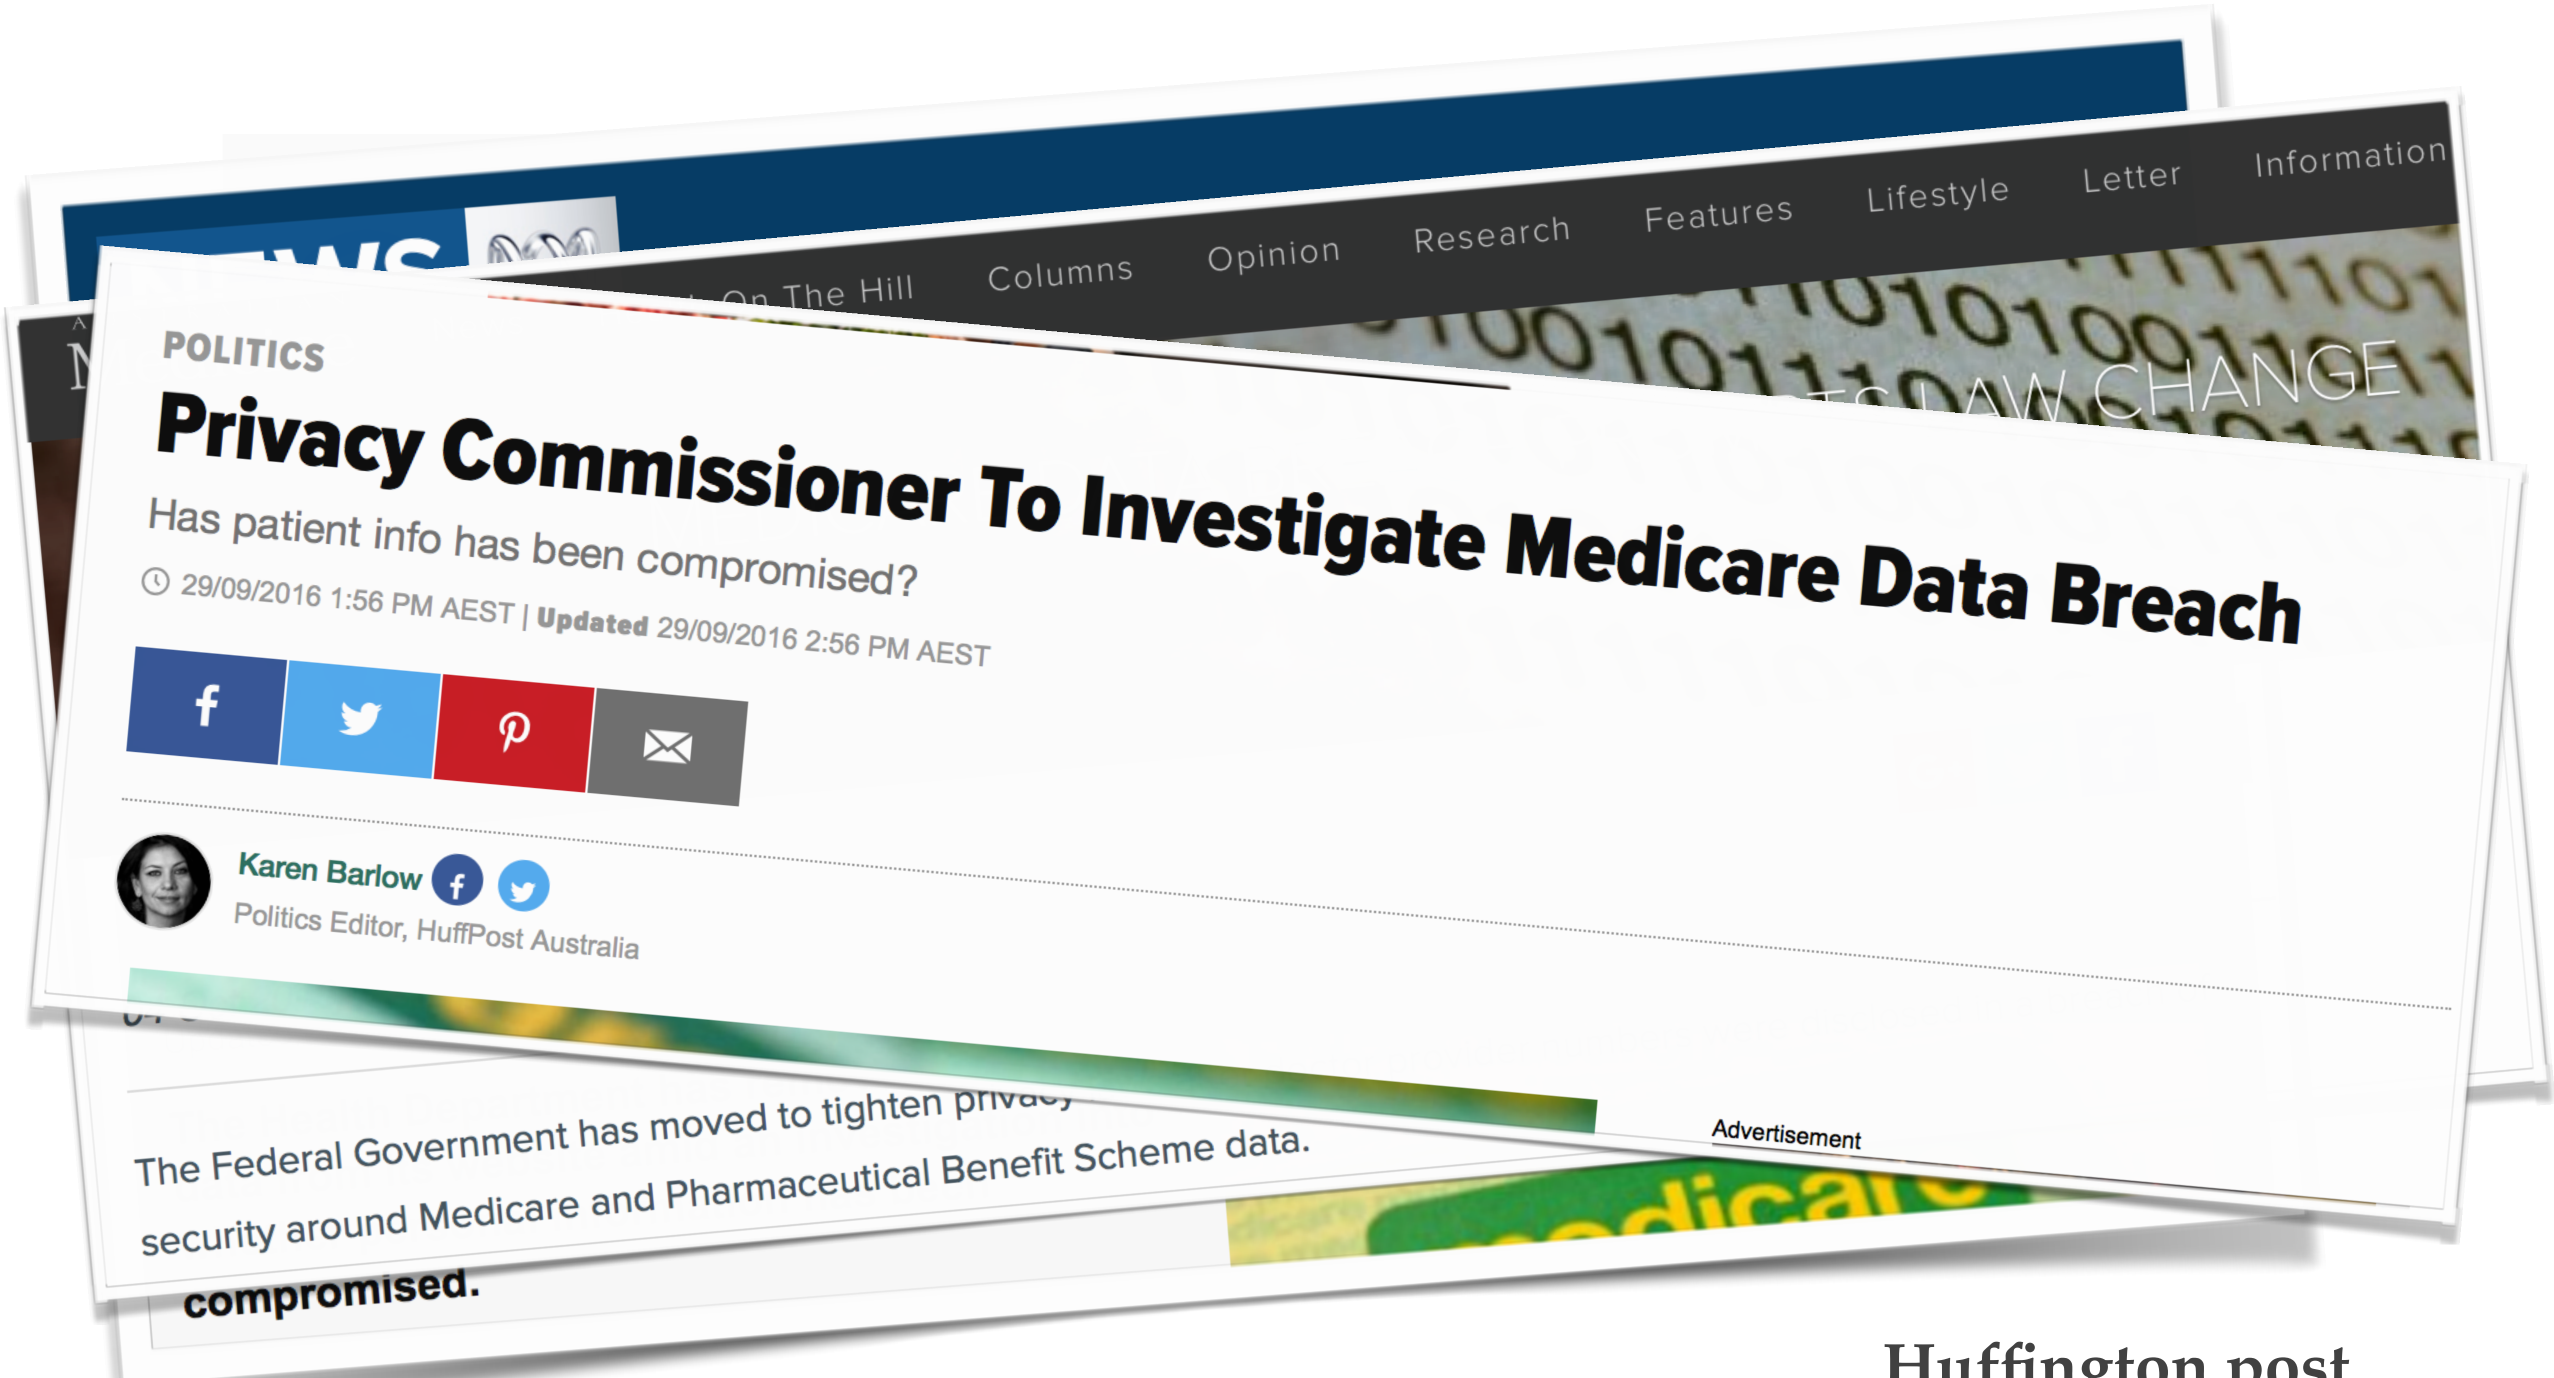

Huffington post

…and then comes (bad) buzz

ITnews

…and then comes (bad) buzz

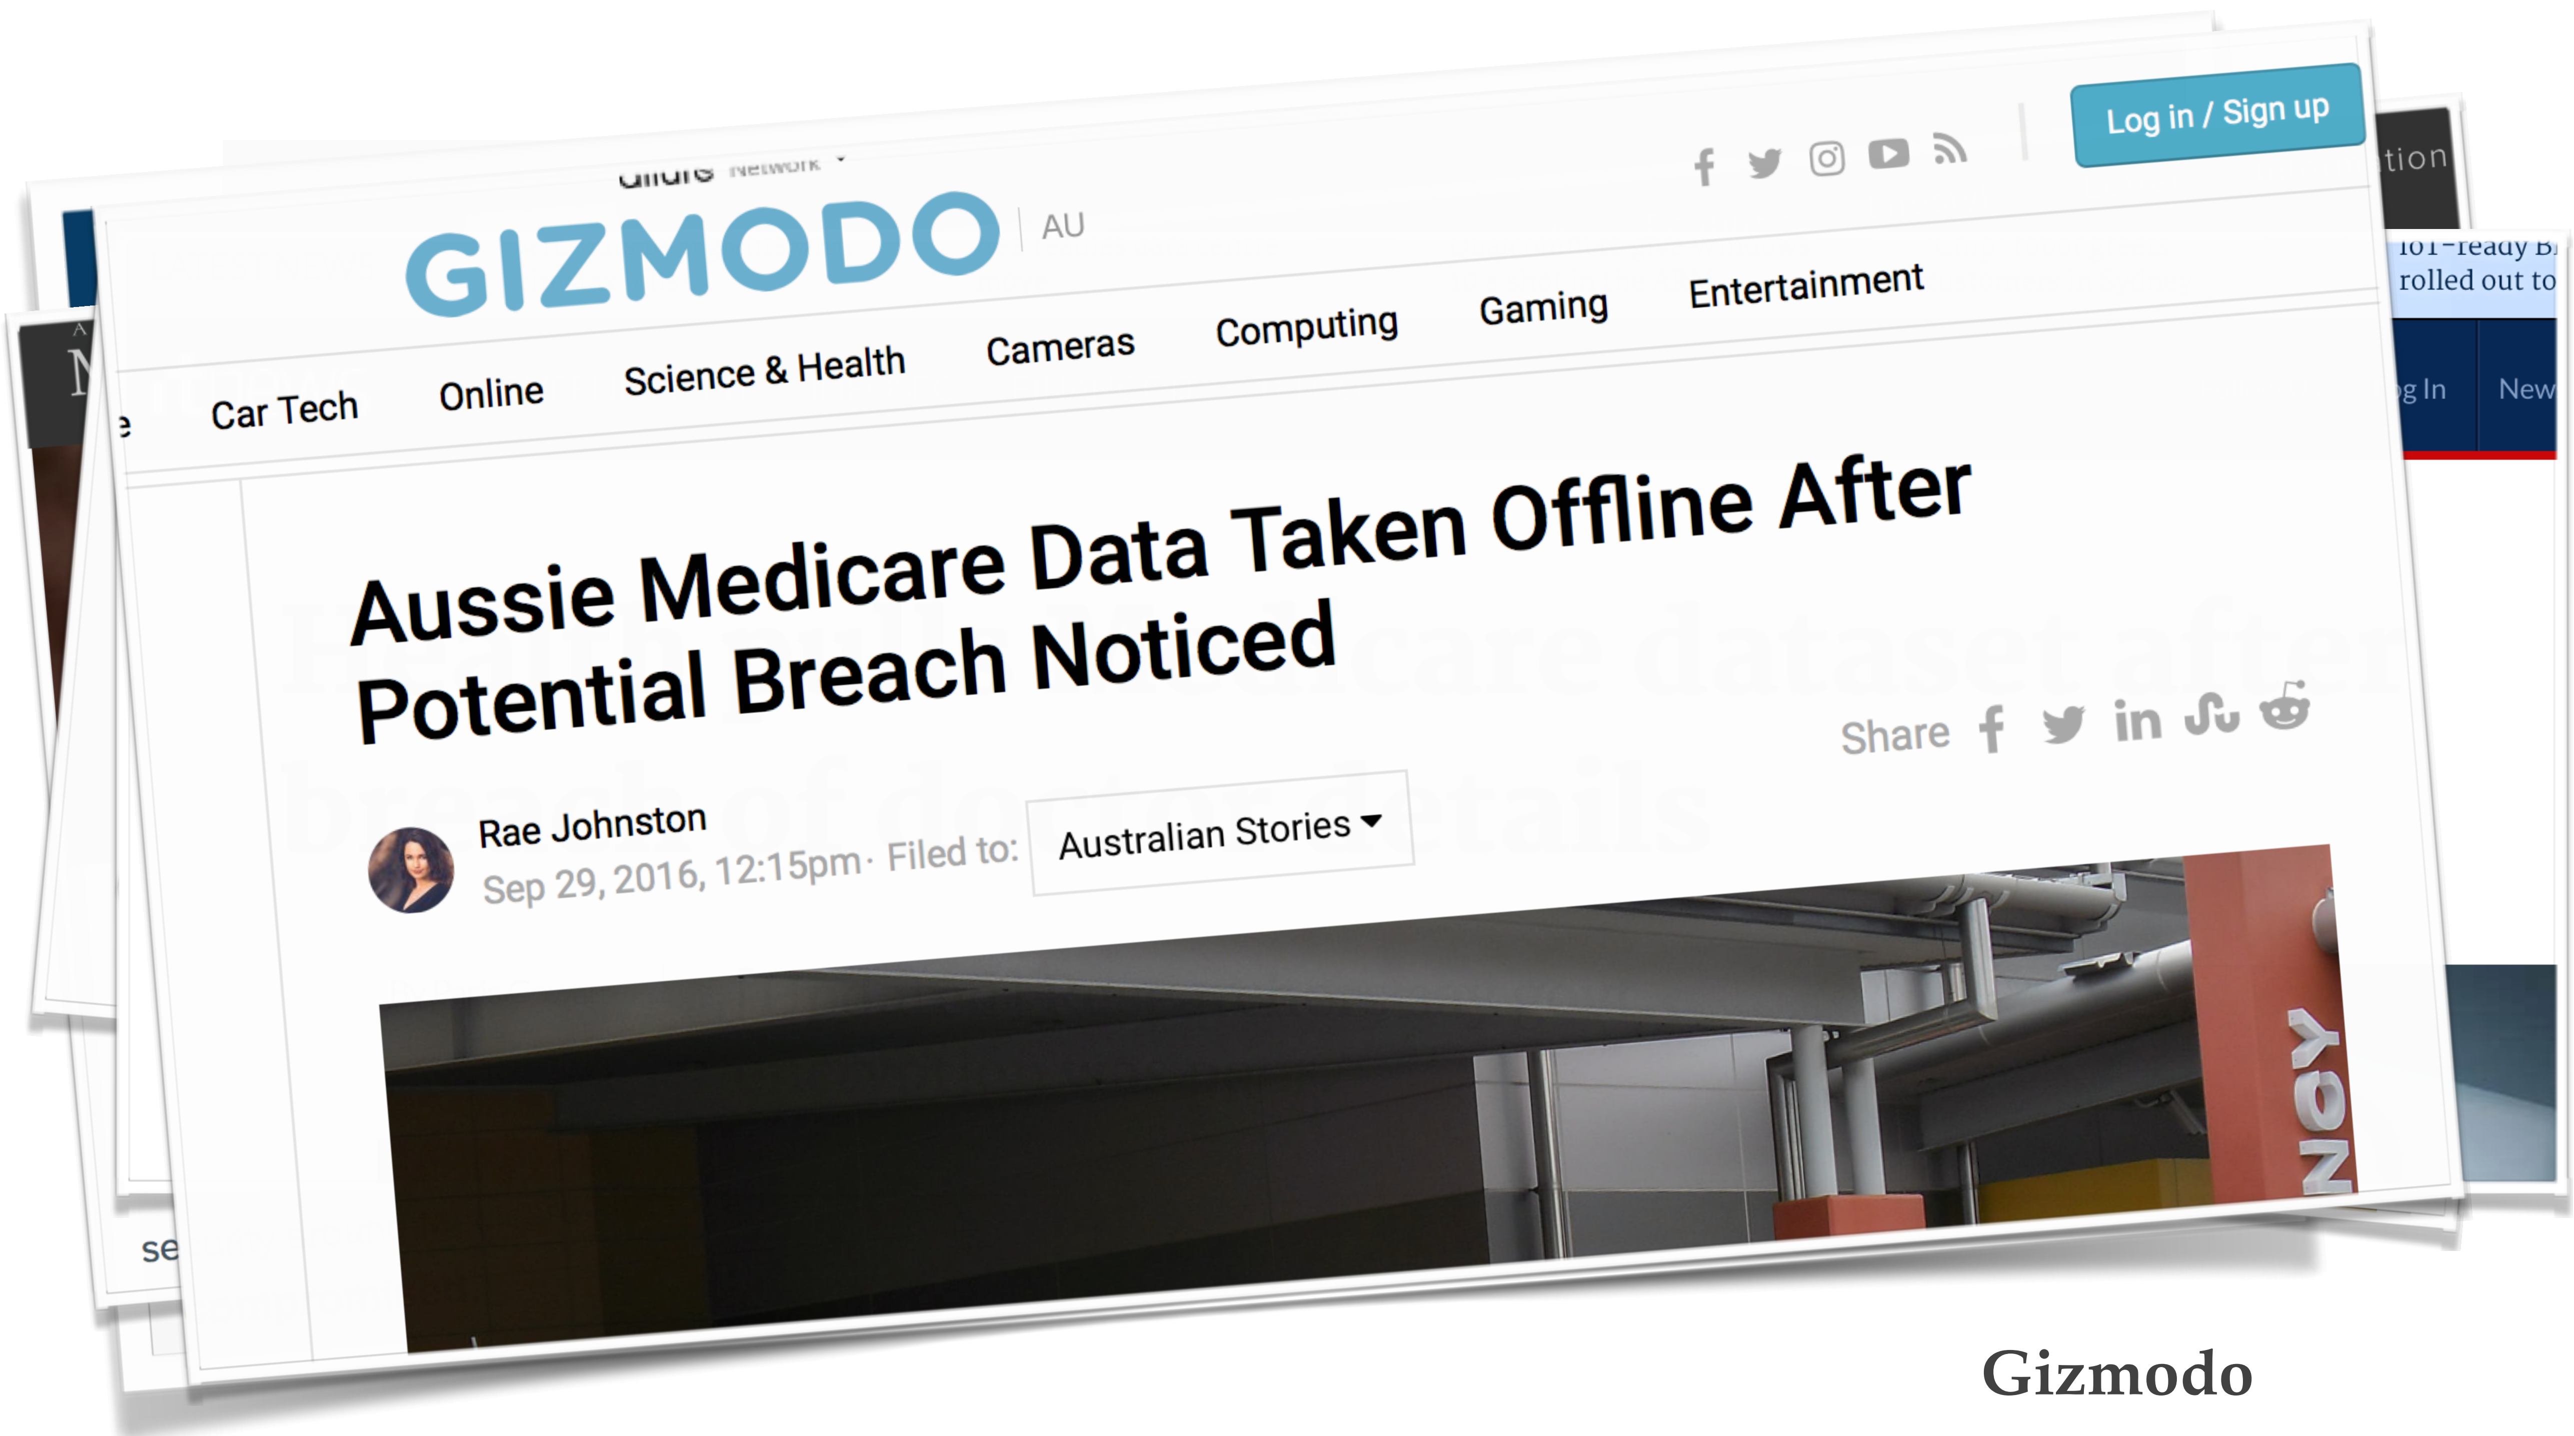

Gizmodo

…and then comes (bad) buzz

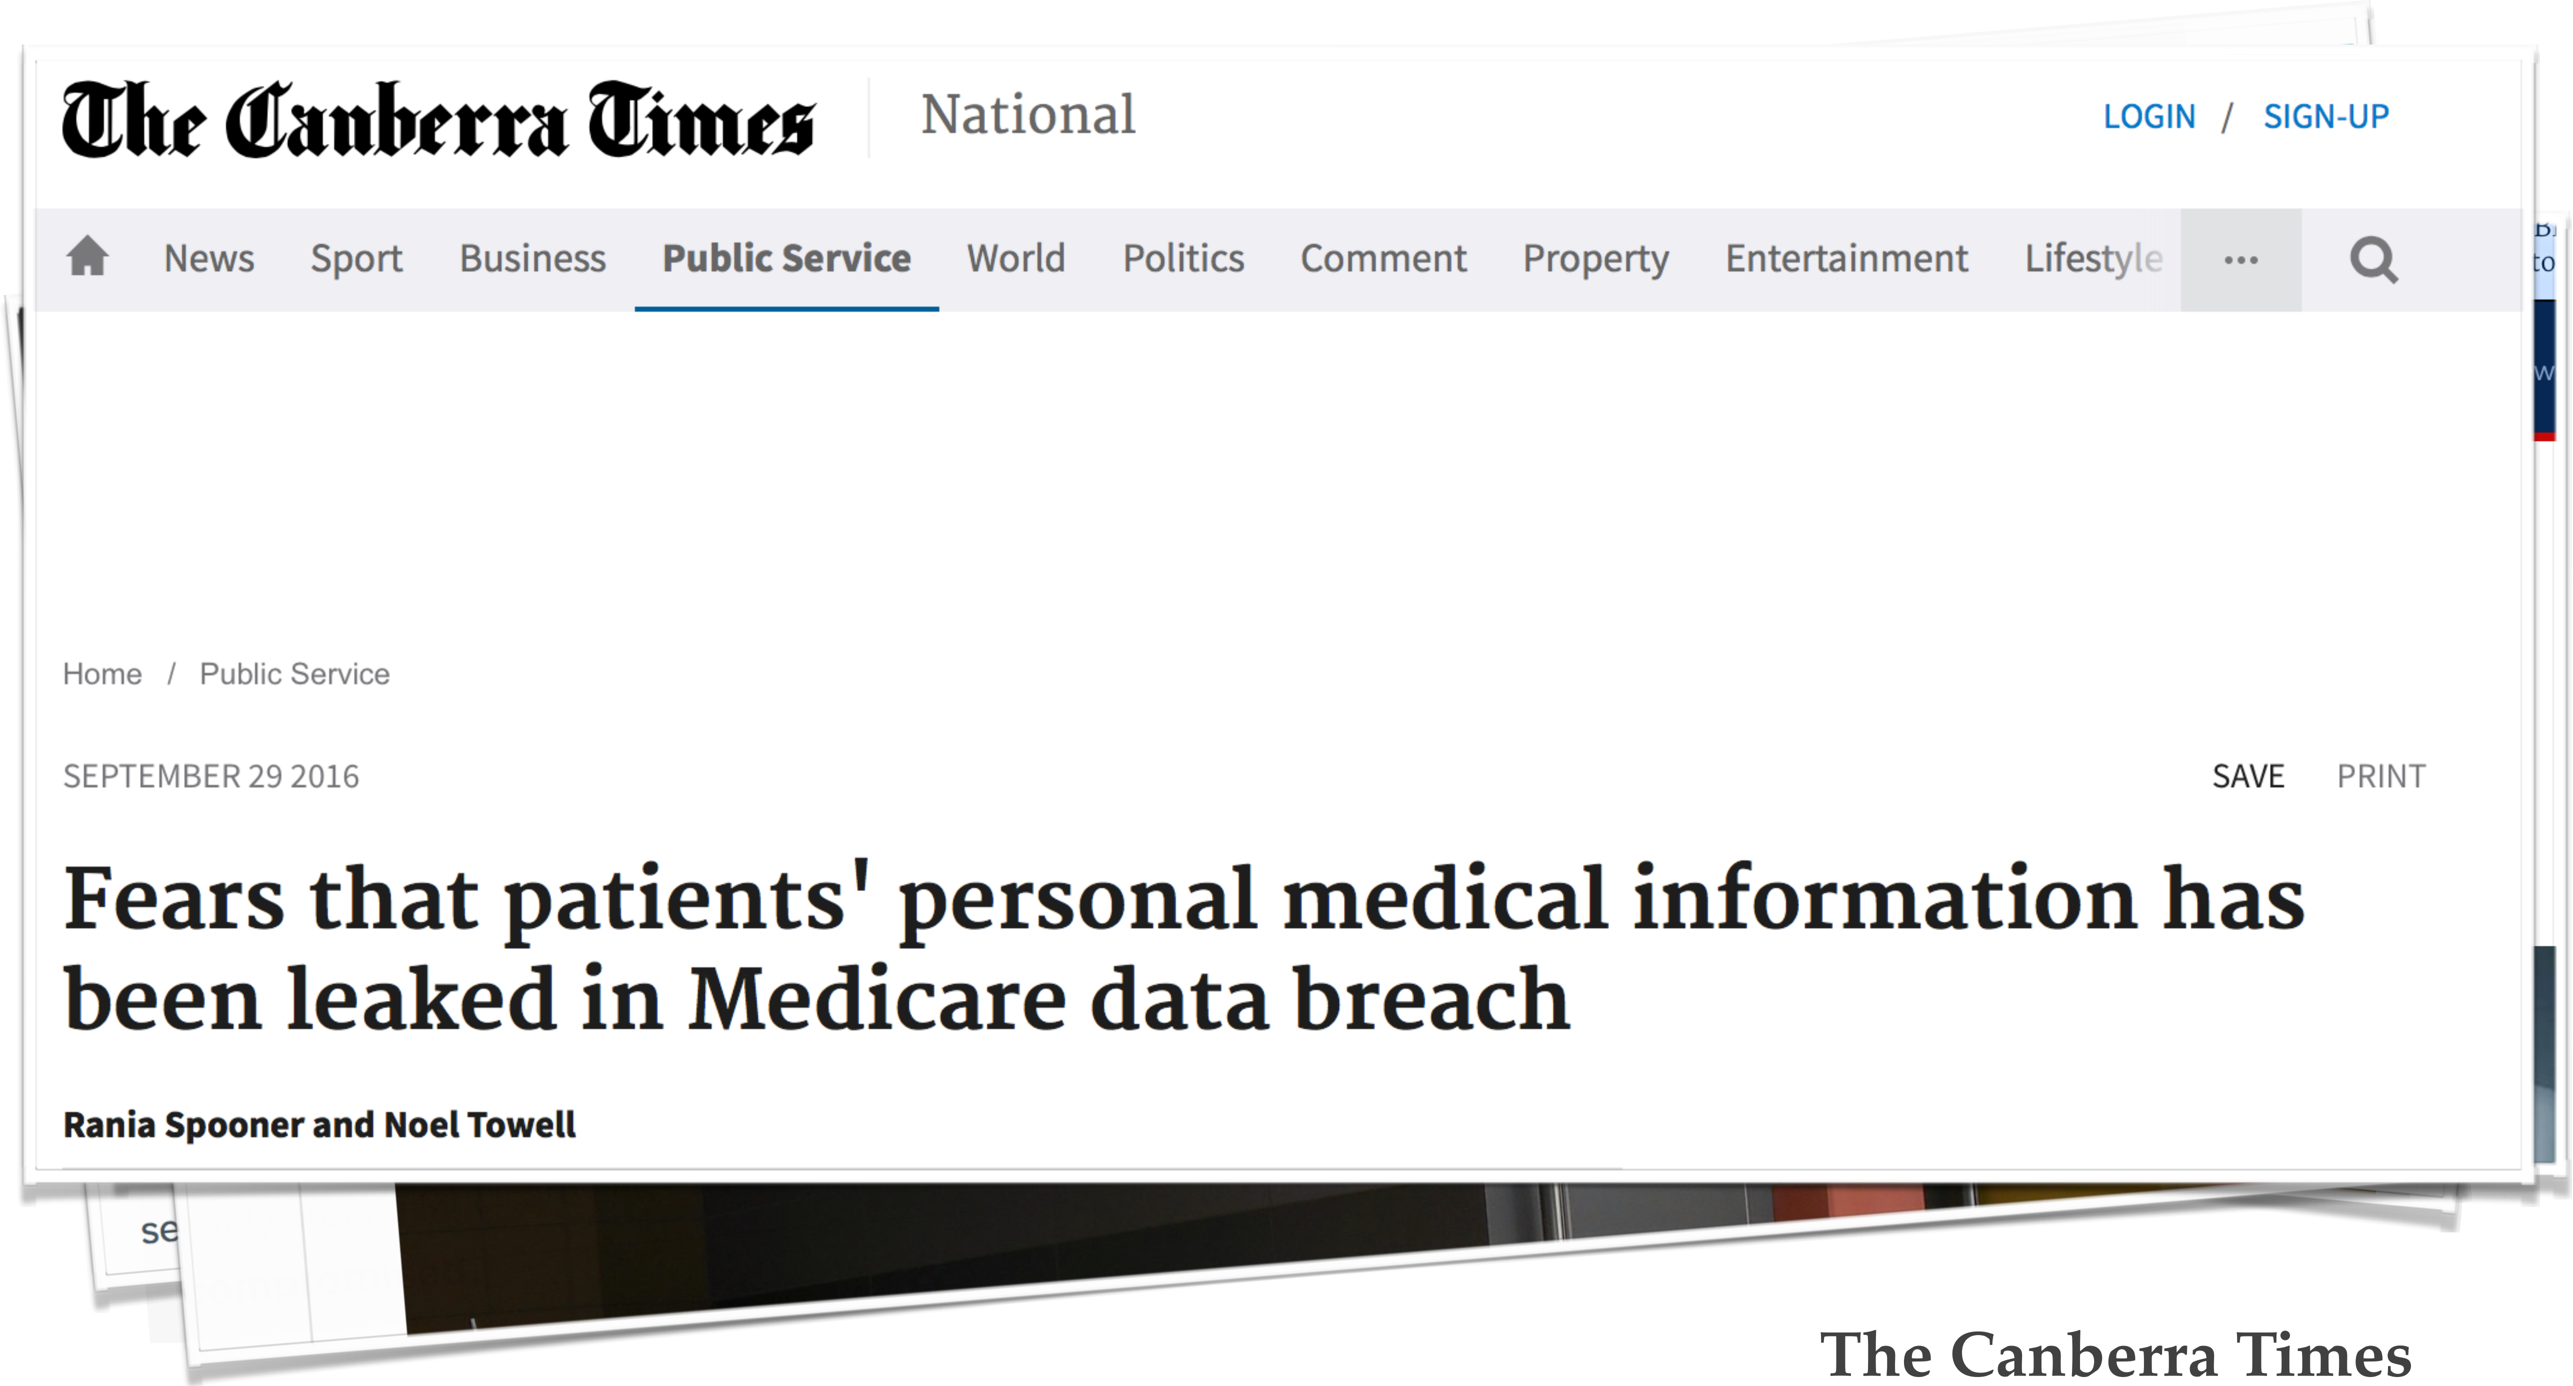
The Canberra Times | National | LOGIN / SIGN-UP

News | Sport | Business | Public Service | World | Politics | Comment | Property | Entertainment | Lifestyle

Home / Public Service

SEPTEMBER 29 2016 | SAVE | PRINT

Fears that patients' personal medical information has been leaked in Medicare data breach

Rania Spooner and Noel Towell

The Canberra Times

Collateral damages

CyberSecurity Online

Key points of the attack

- Questionable choice of ground techniques for the protection, but more importantly
- Attack tackles **bad implementation design** (parameters)
- Attack with **side information** (attacker)

(apologies to my colleagues for depicting them this way)

Lesson

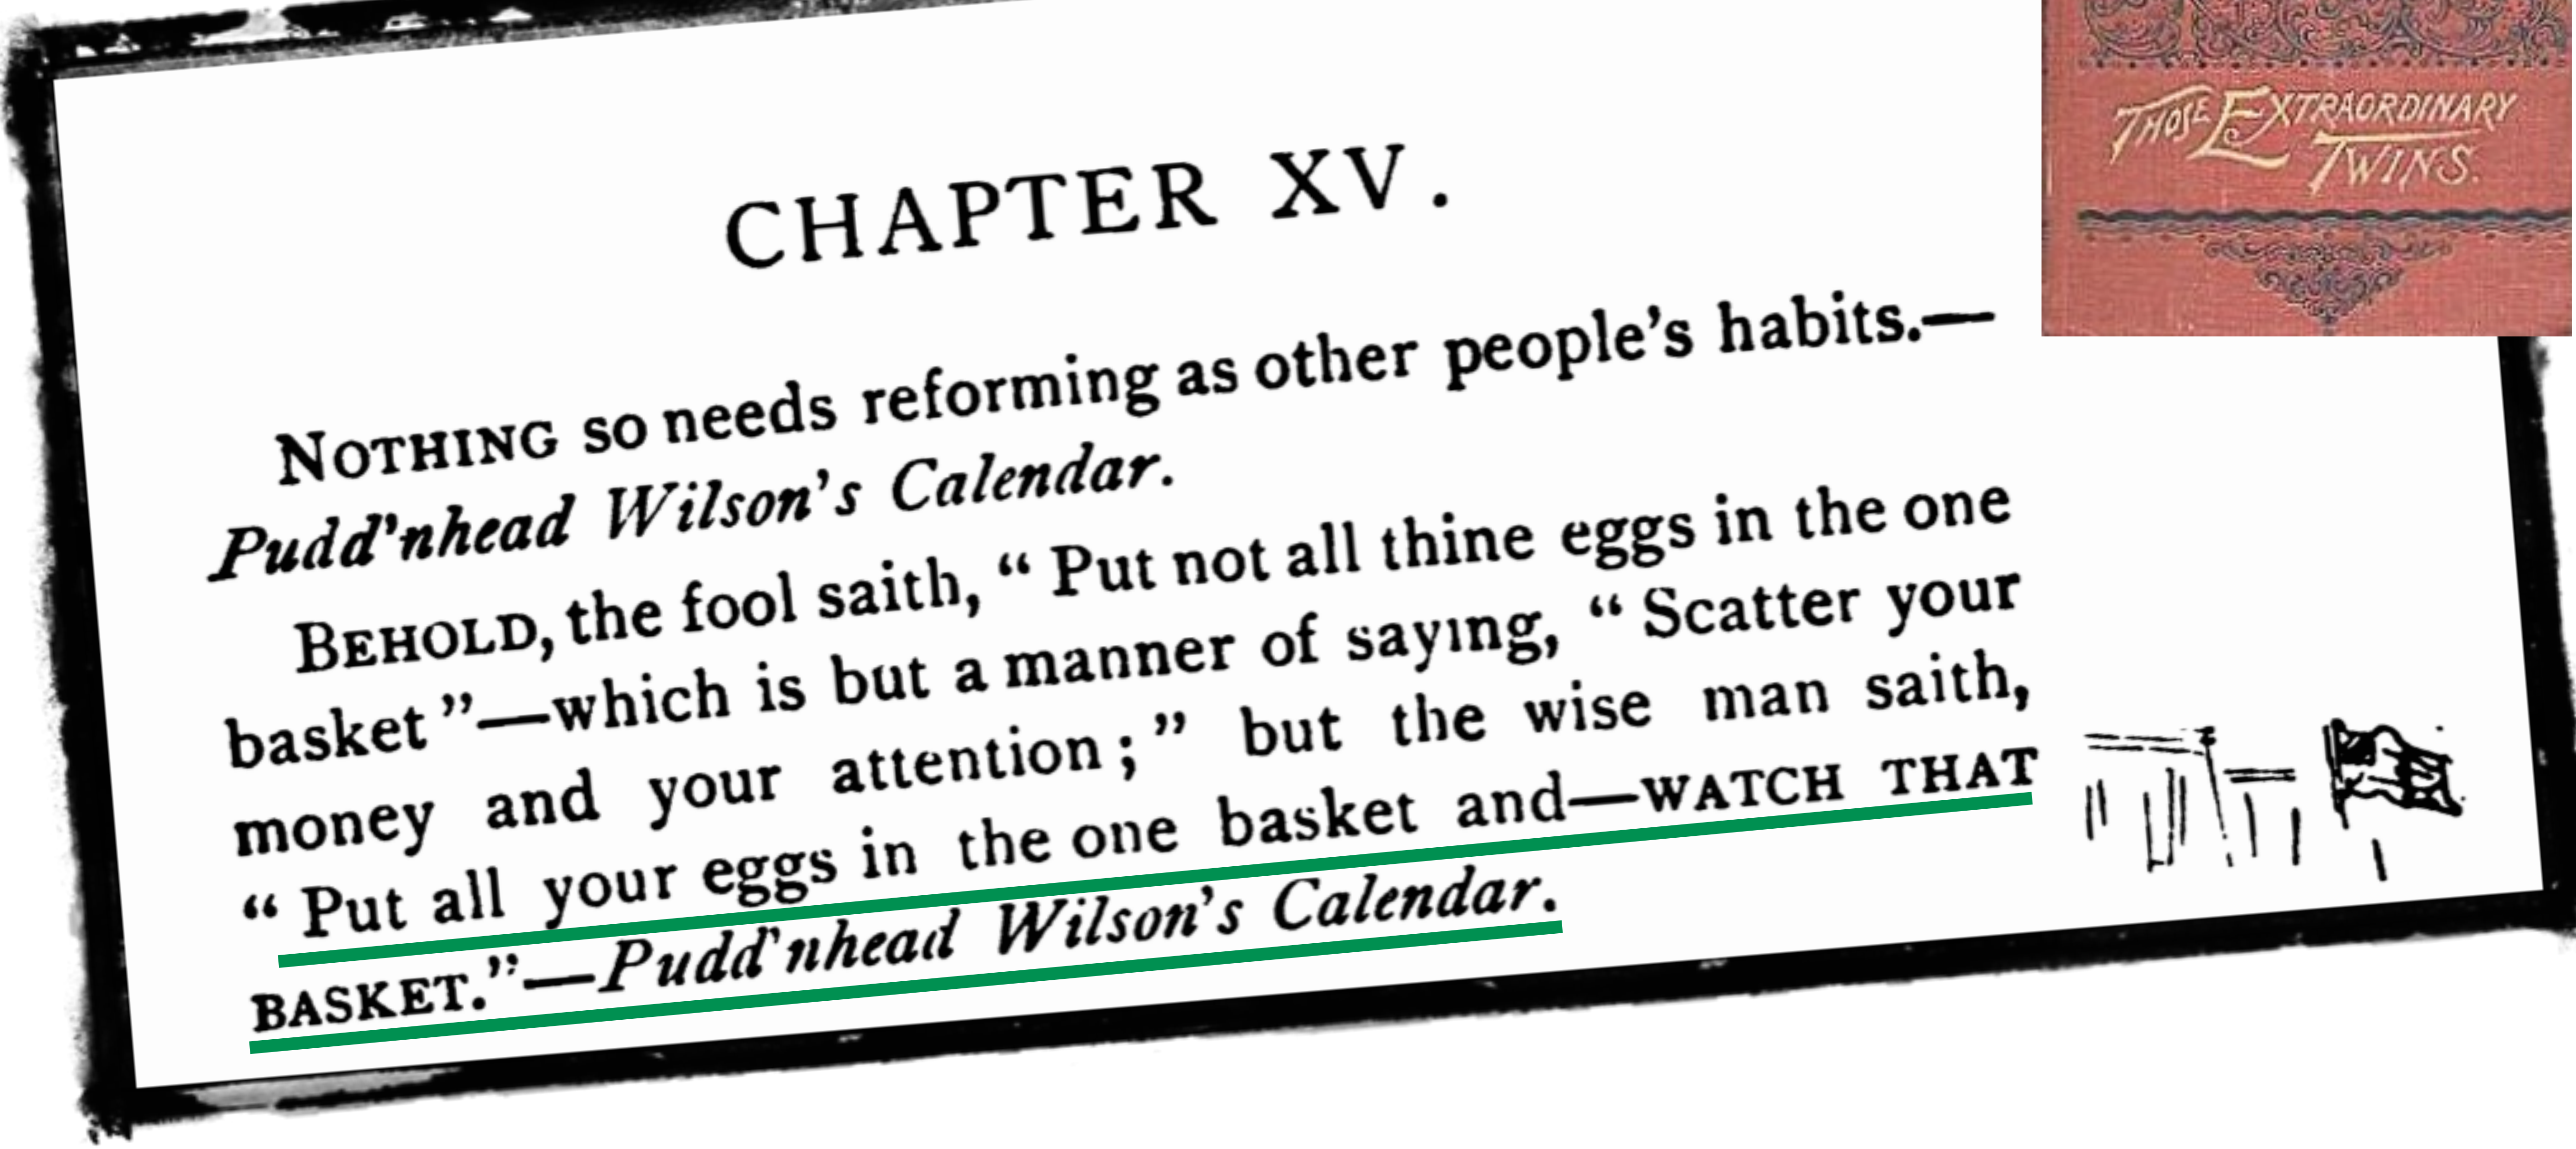

CHAPTER XV.

NOTHING so needs reforming as other people's habits.—*Pudd'nhead Wilson's Calendar.*

BEHOLD, the fool saith, "Put not all thine eggs in the one basket"—which is but a manner of saying, "Scatter your money and your attention;" but the wise man saith, "Put all your eggs in the one basket and—WATCH THAT BASKET."—*Pudd'nhead Wilson's Calendar.*

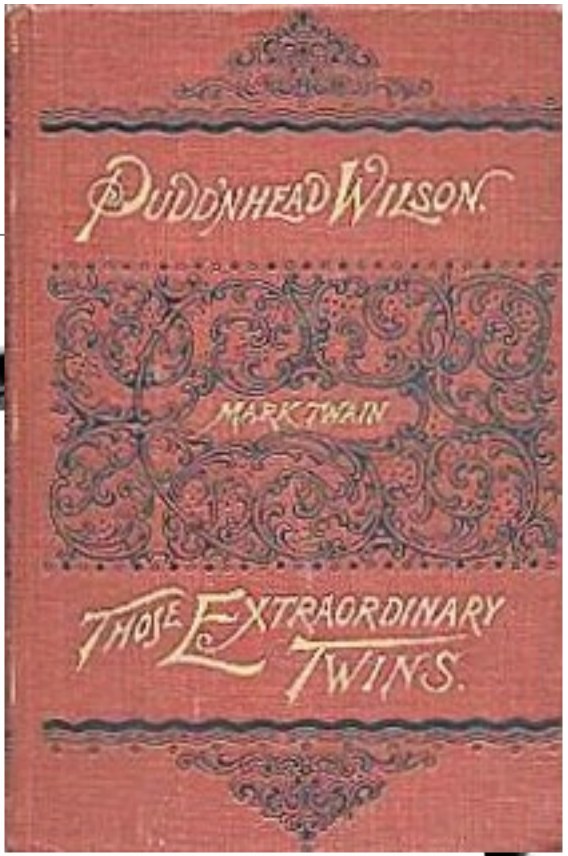

Confidential computing overview & targeted problems

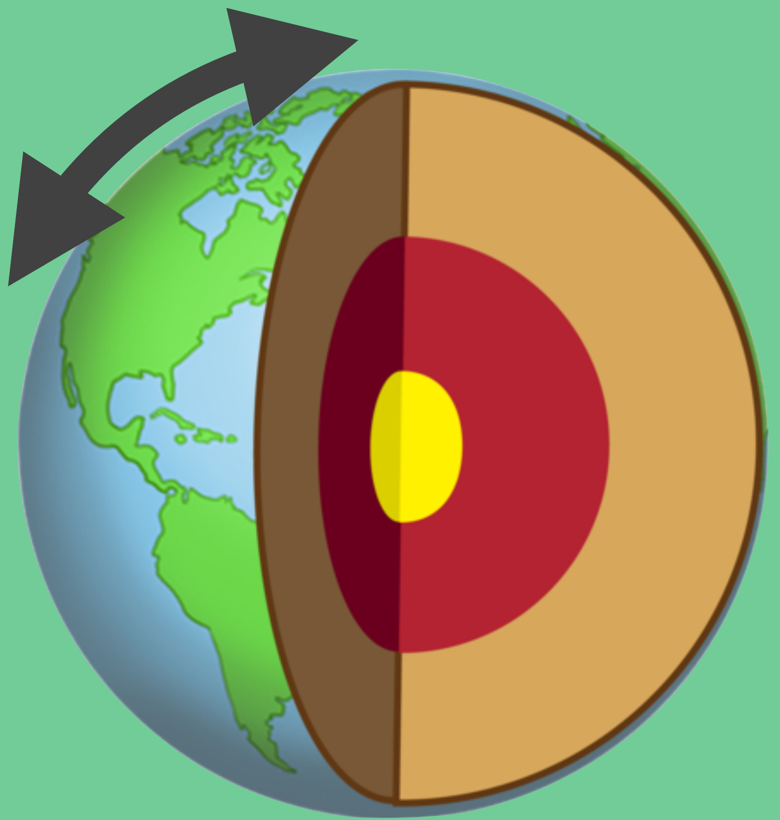

Future value of data

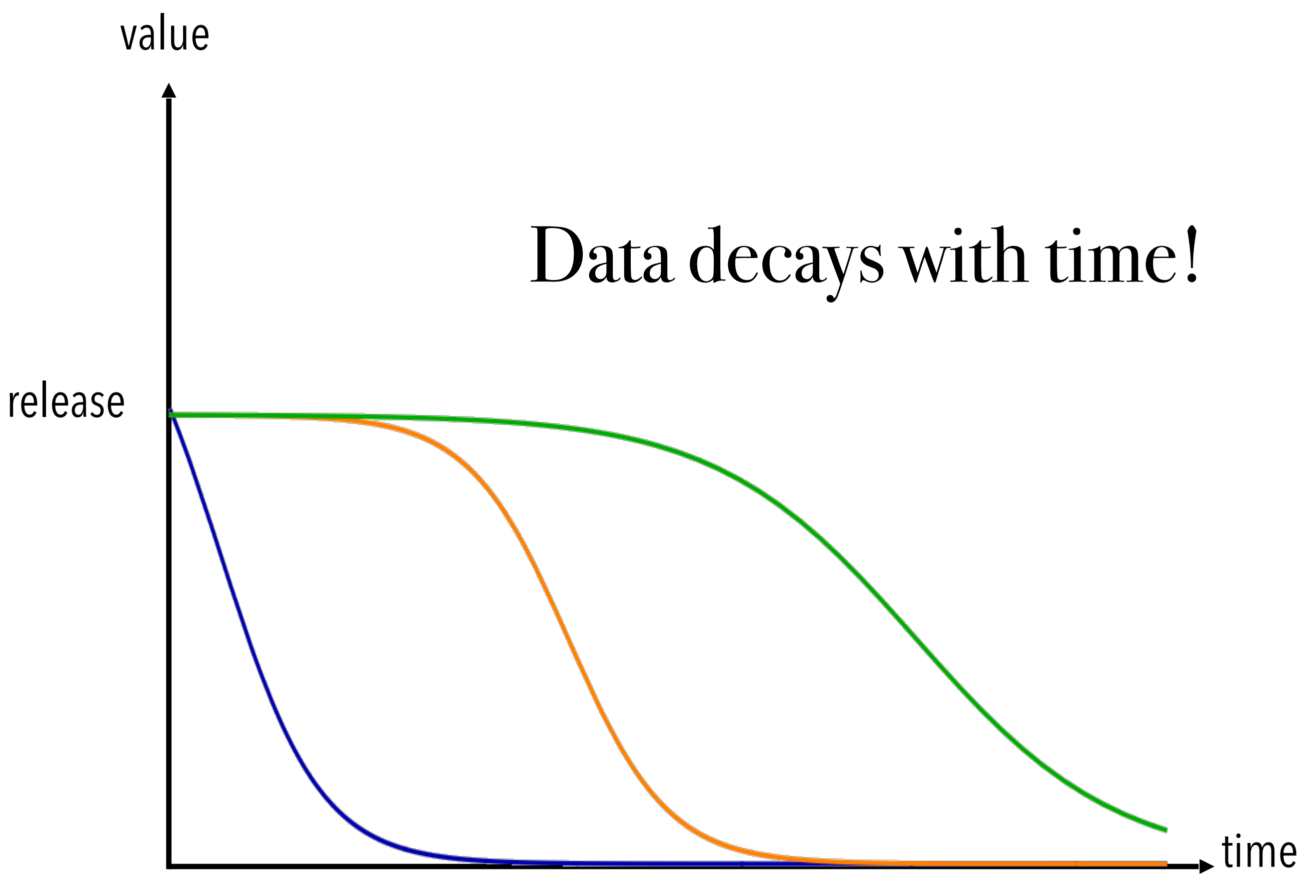

Future value of data

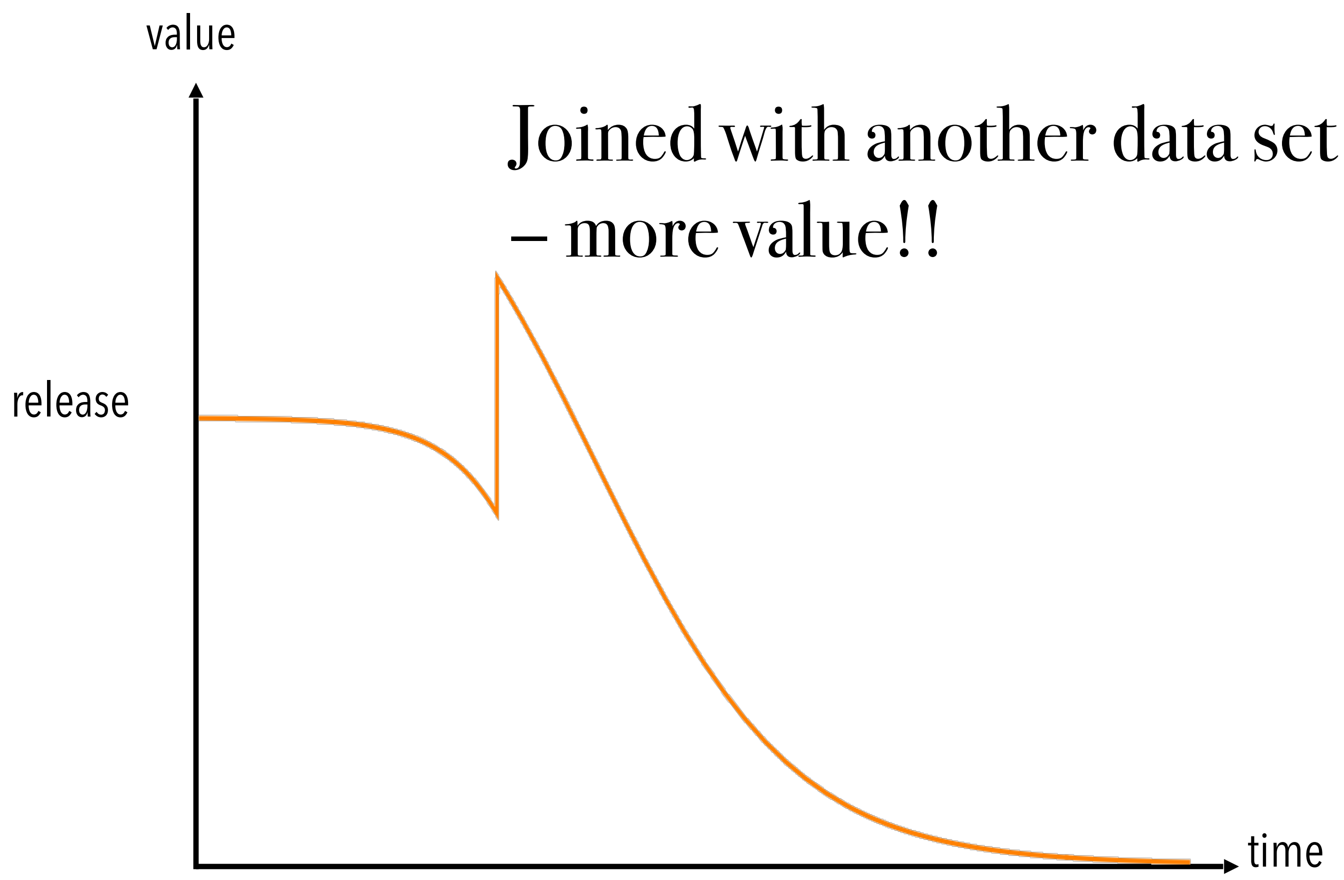

Joined with another data set – more value!!

Future value of data

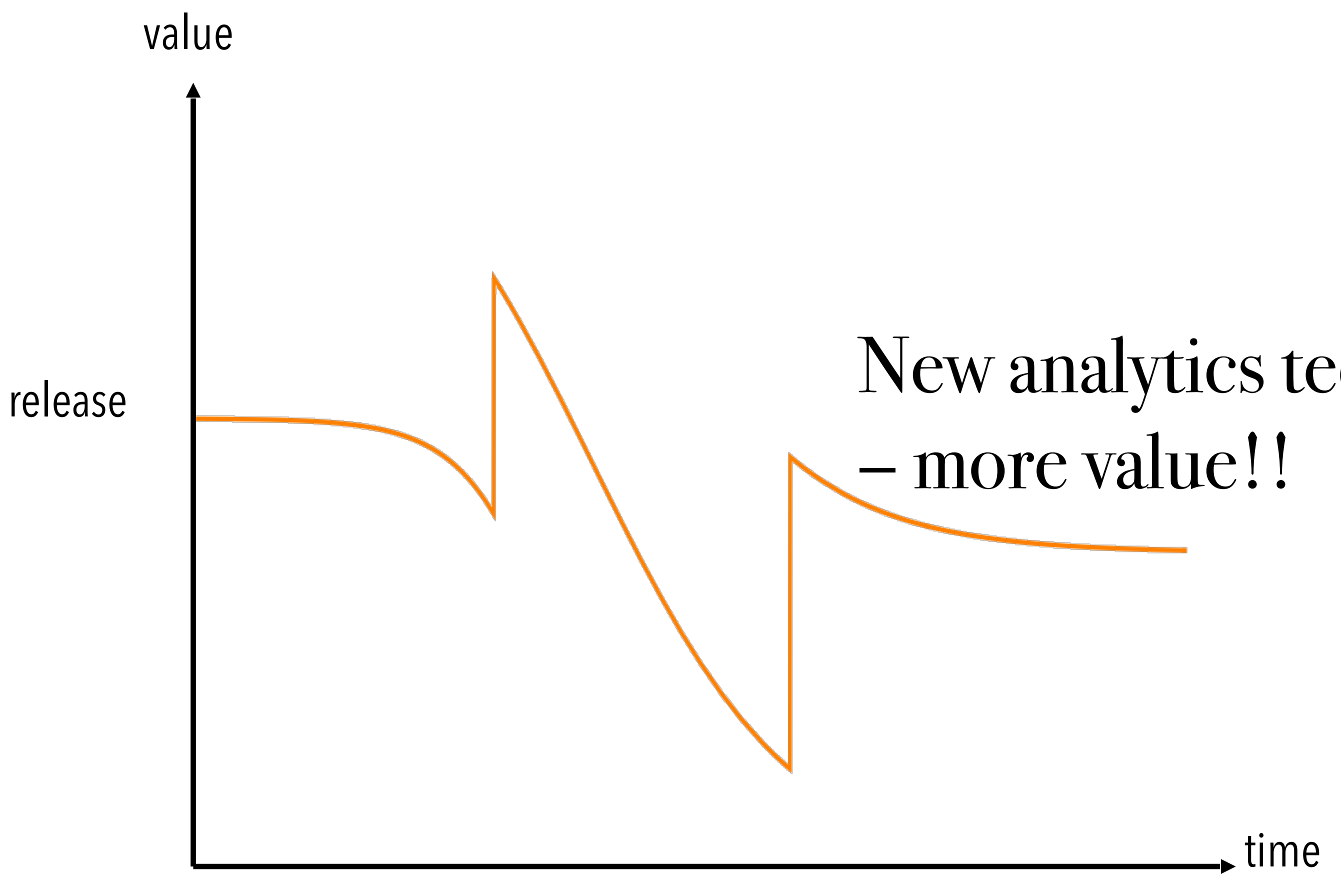

New analytics techniques
– more value!!

Future value of data

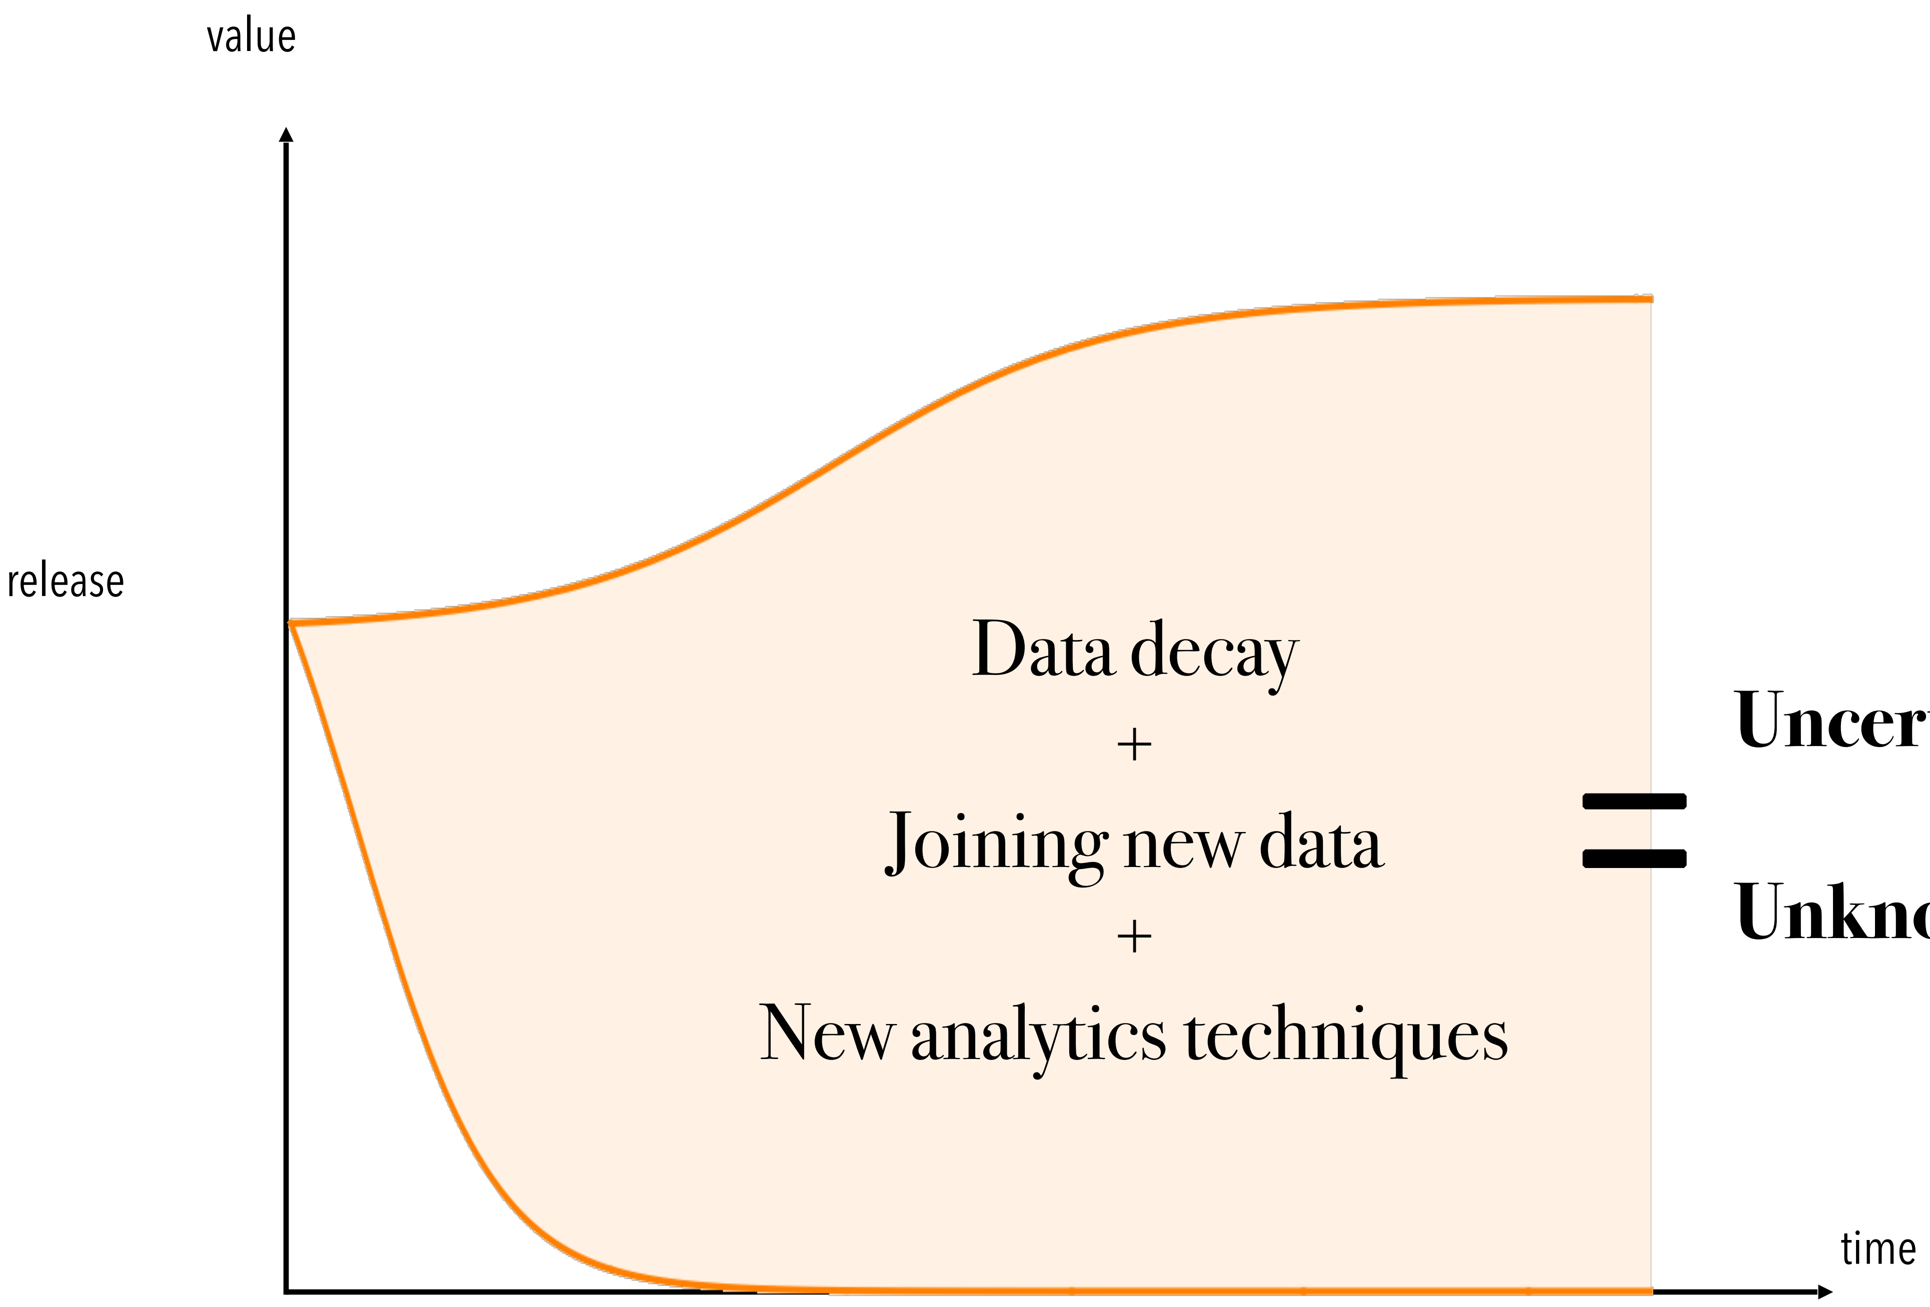

=

Uncertain future value

Unknown future risk

Future value of data

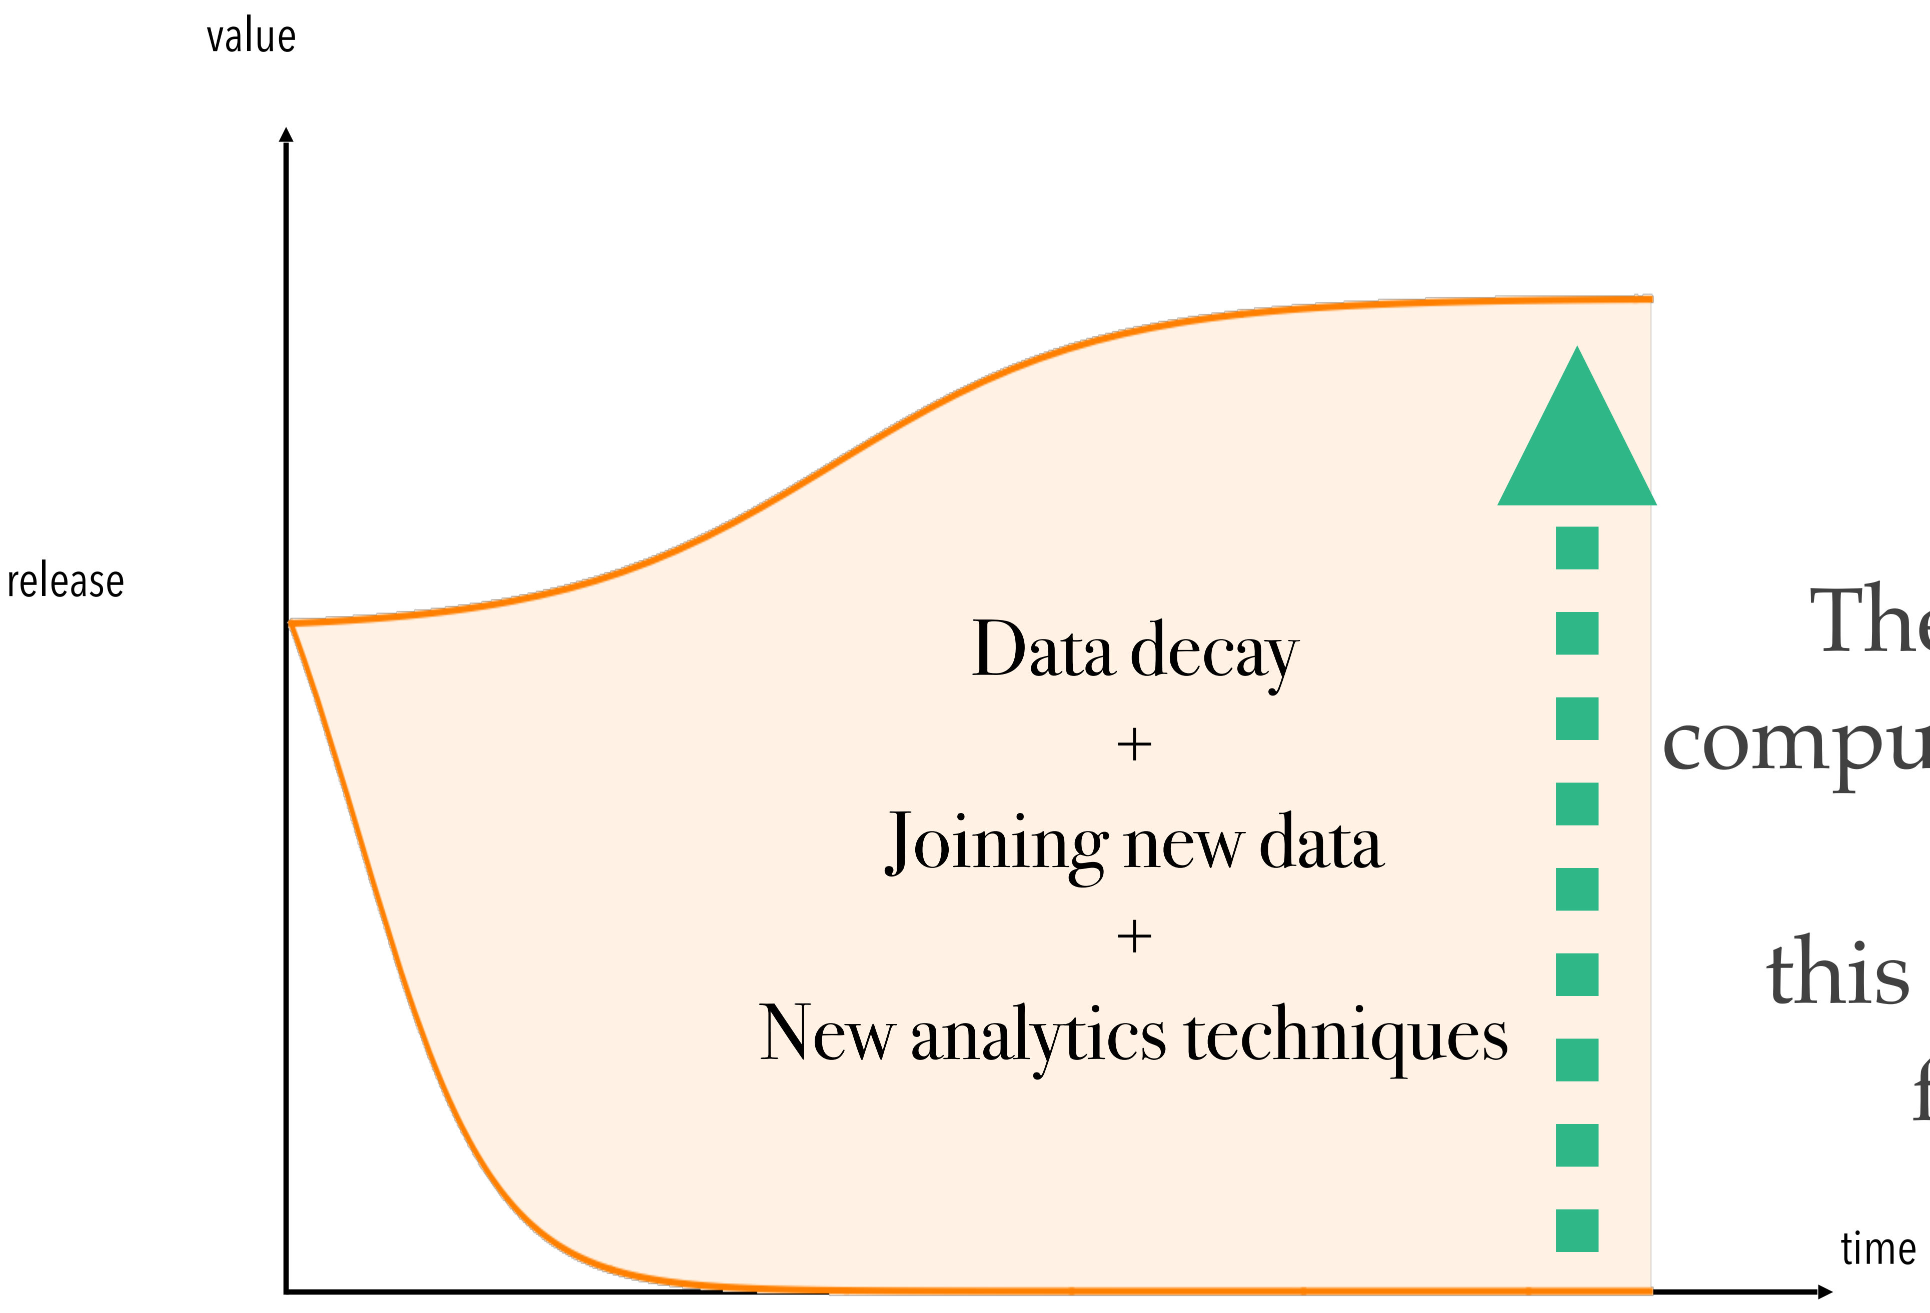

The confidential computing project aims at doing this in a *distributed* framework

Challenge – Summary

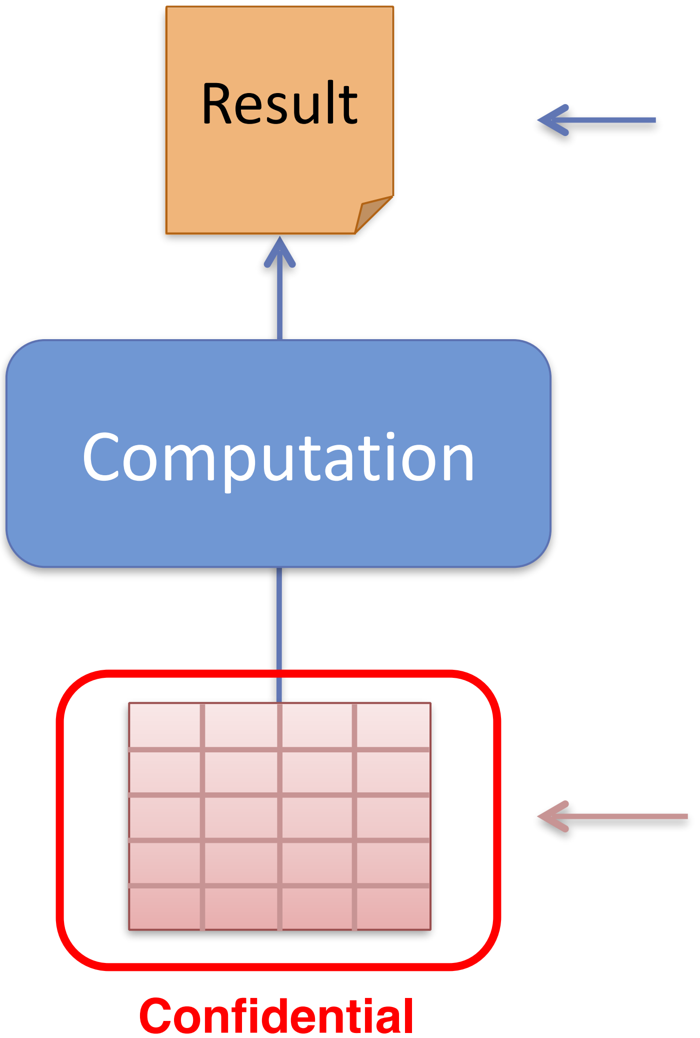

Learn (from) this!

Learn NOTHING

The problem

- How can we learn valuable **insights** from **sensitive** data from **multiple** organisations?

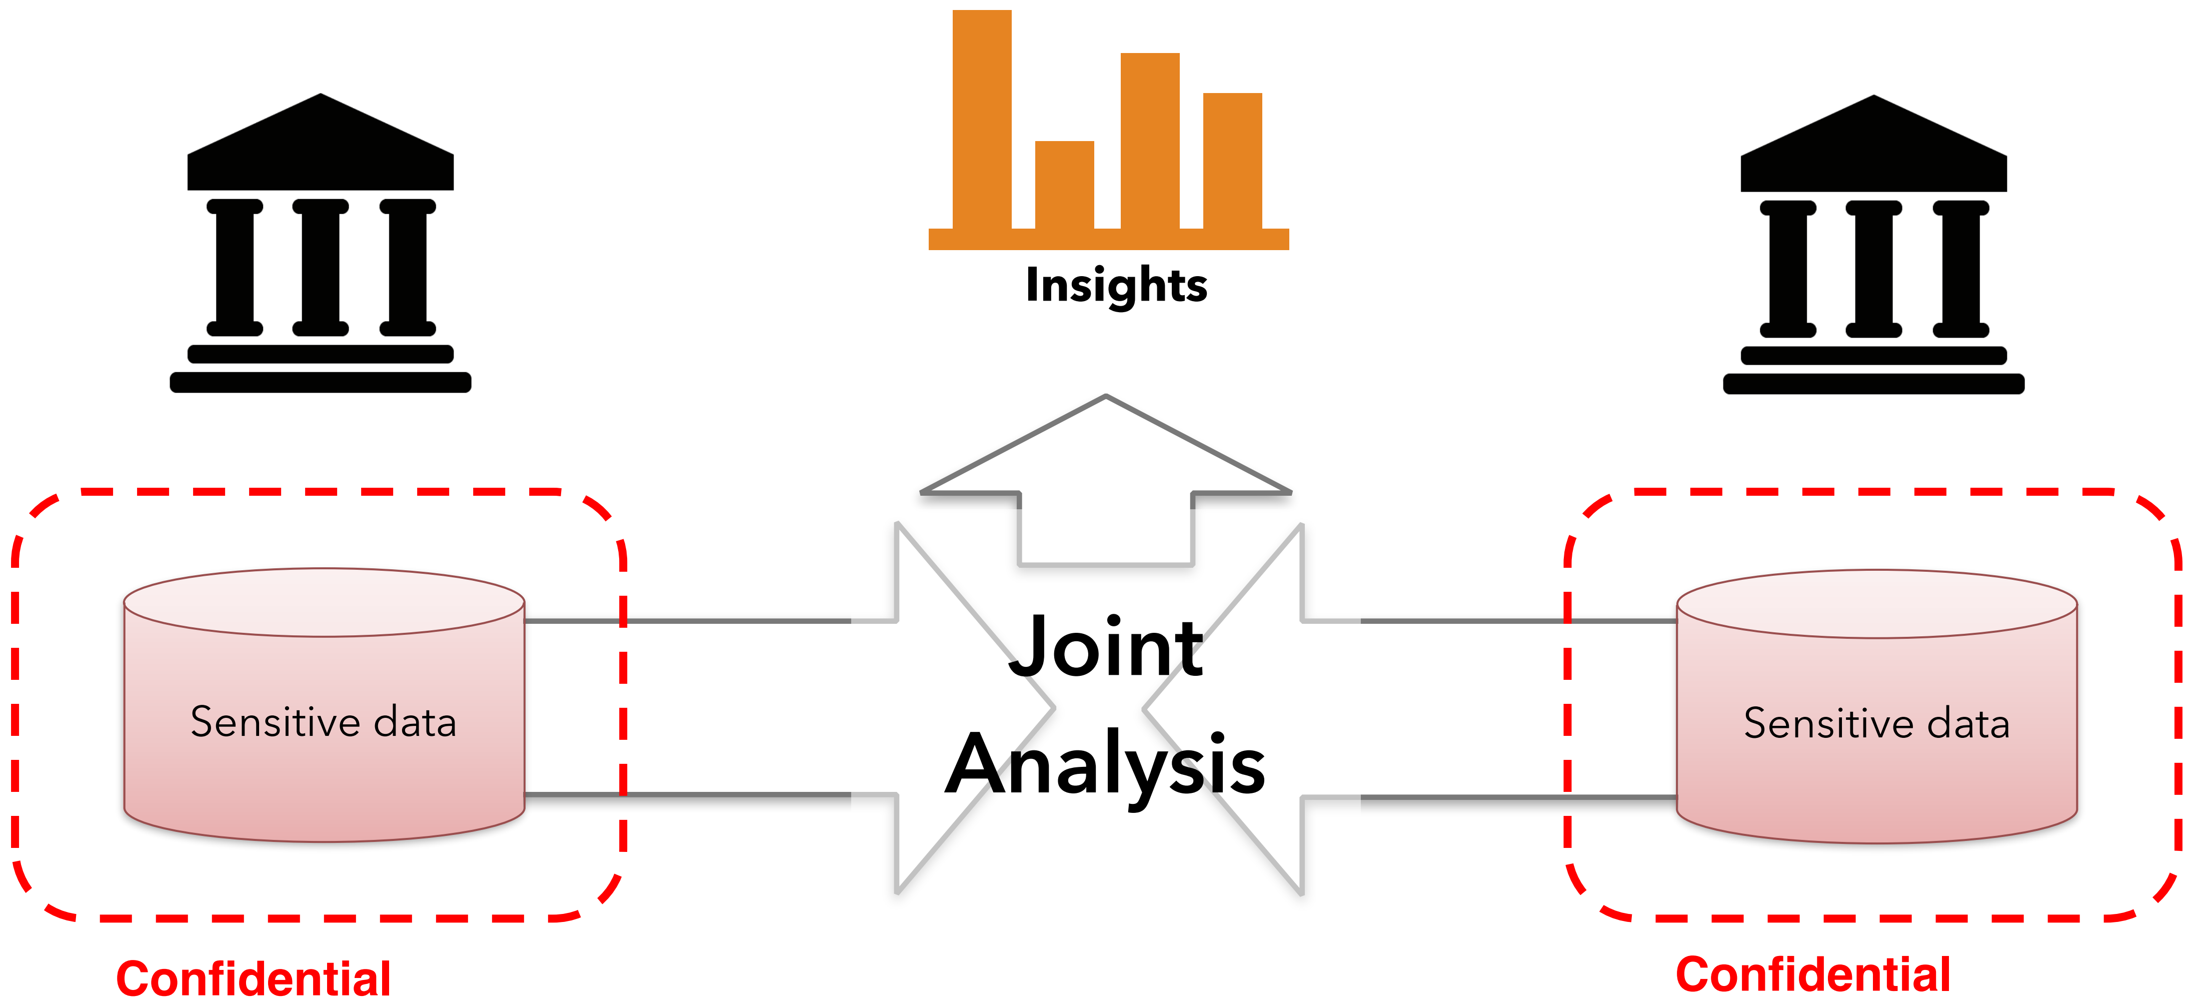

Case studies

Scoring

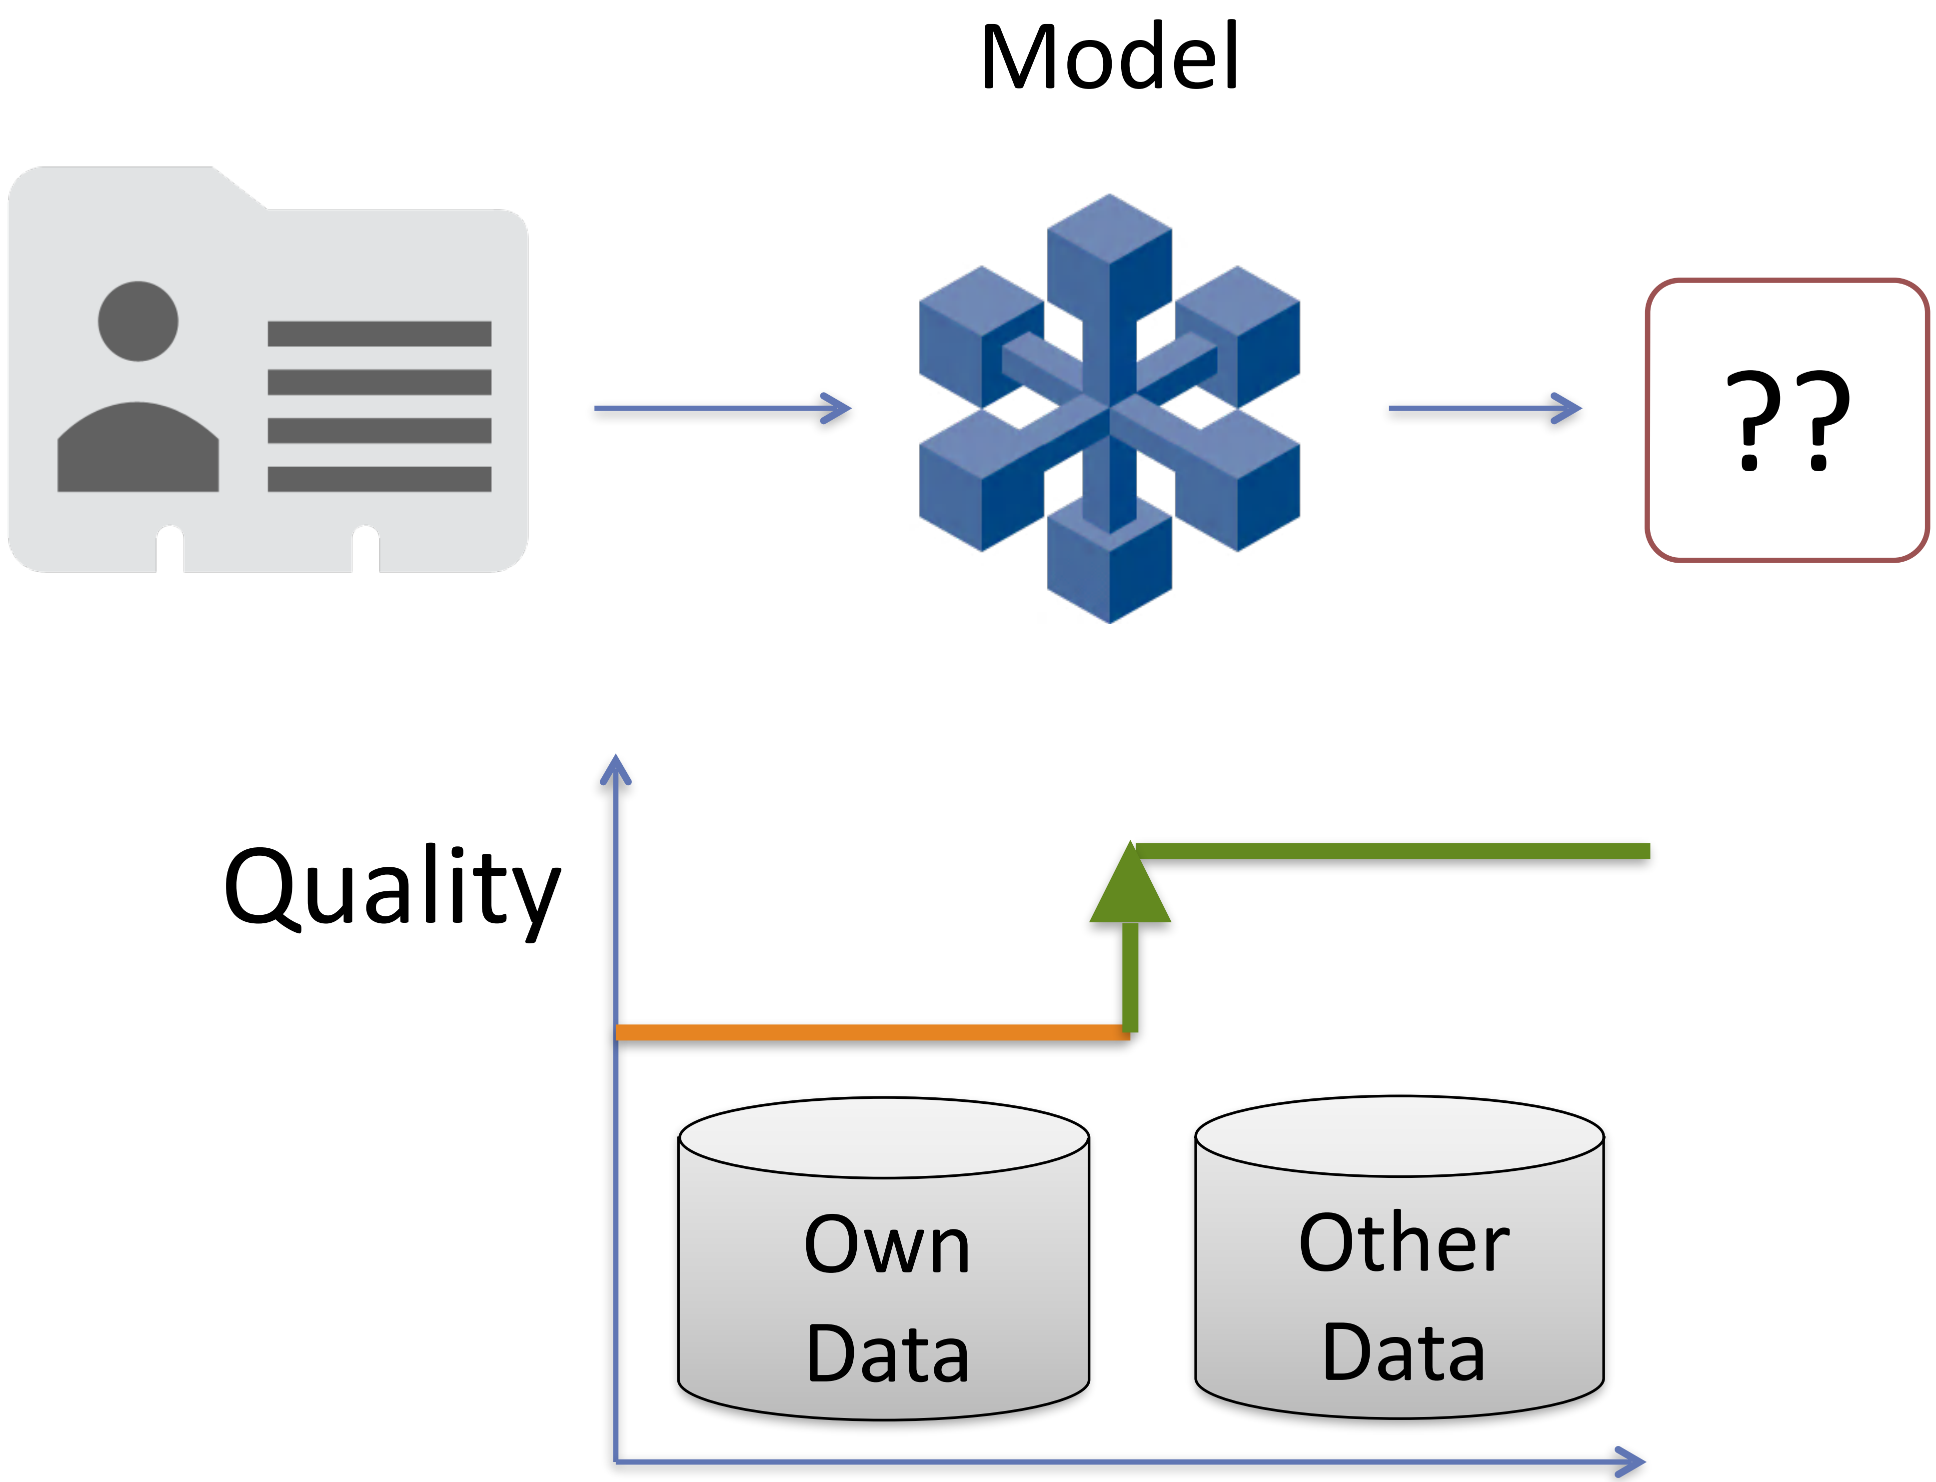

Suspicious activities

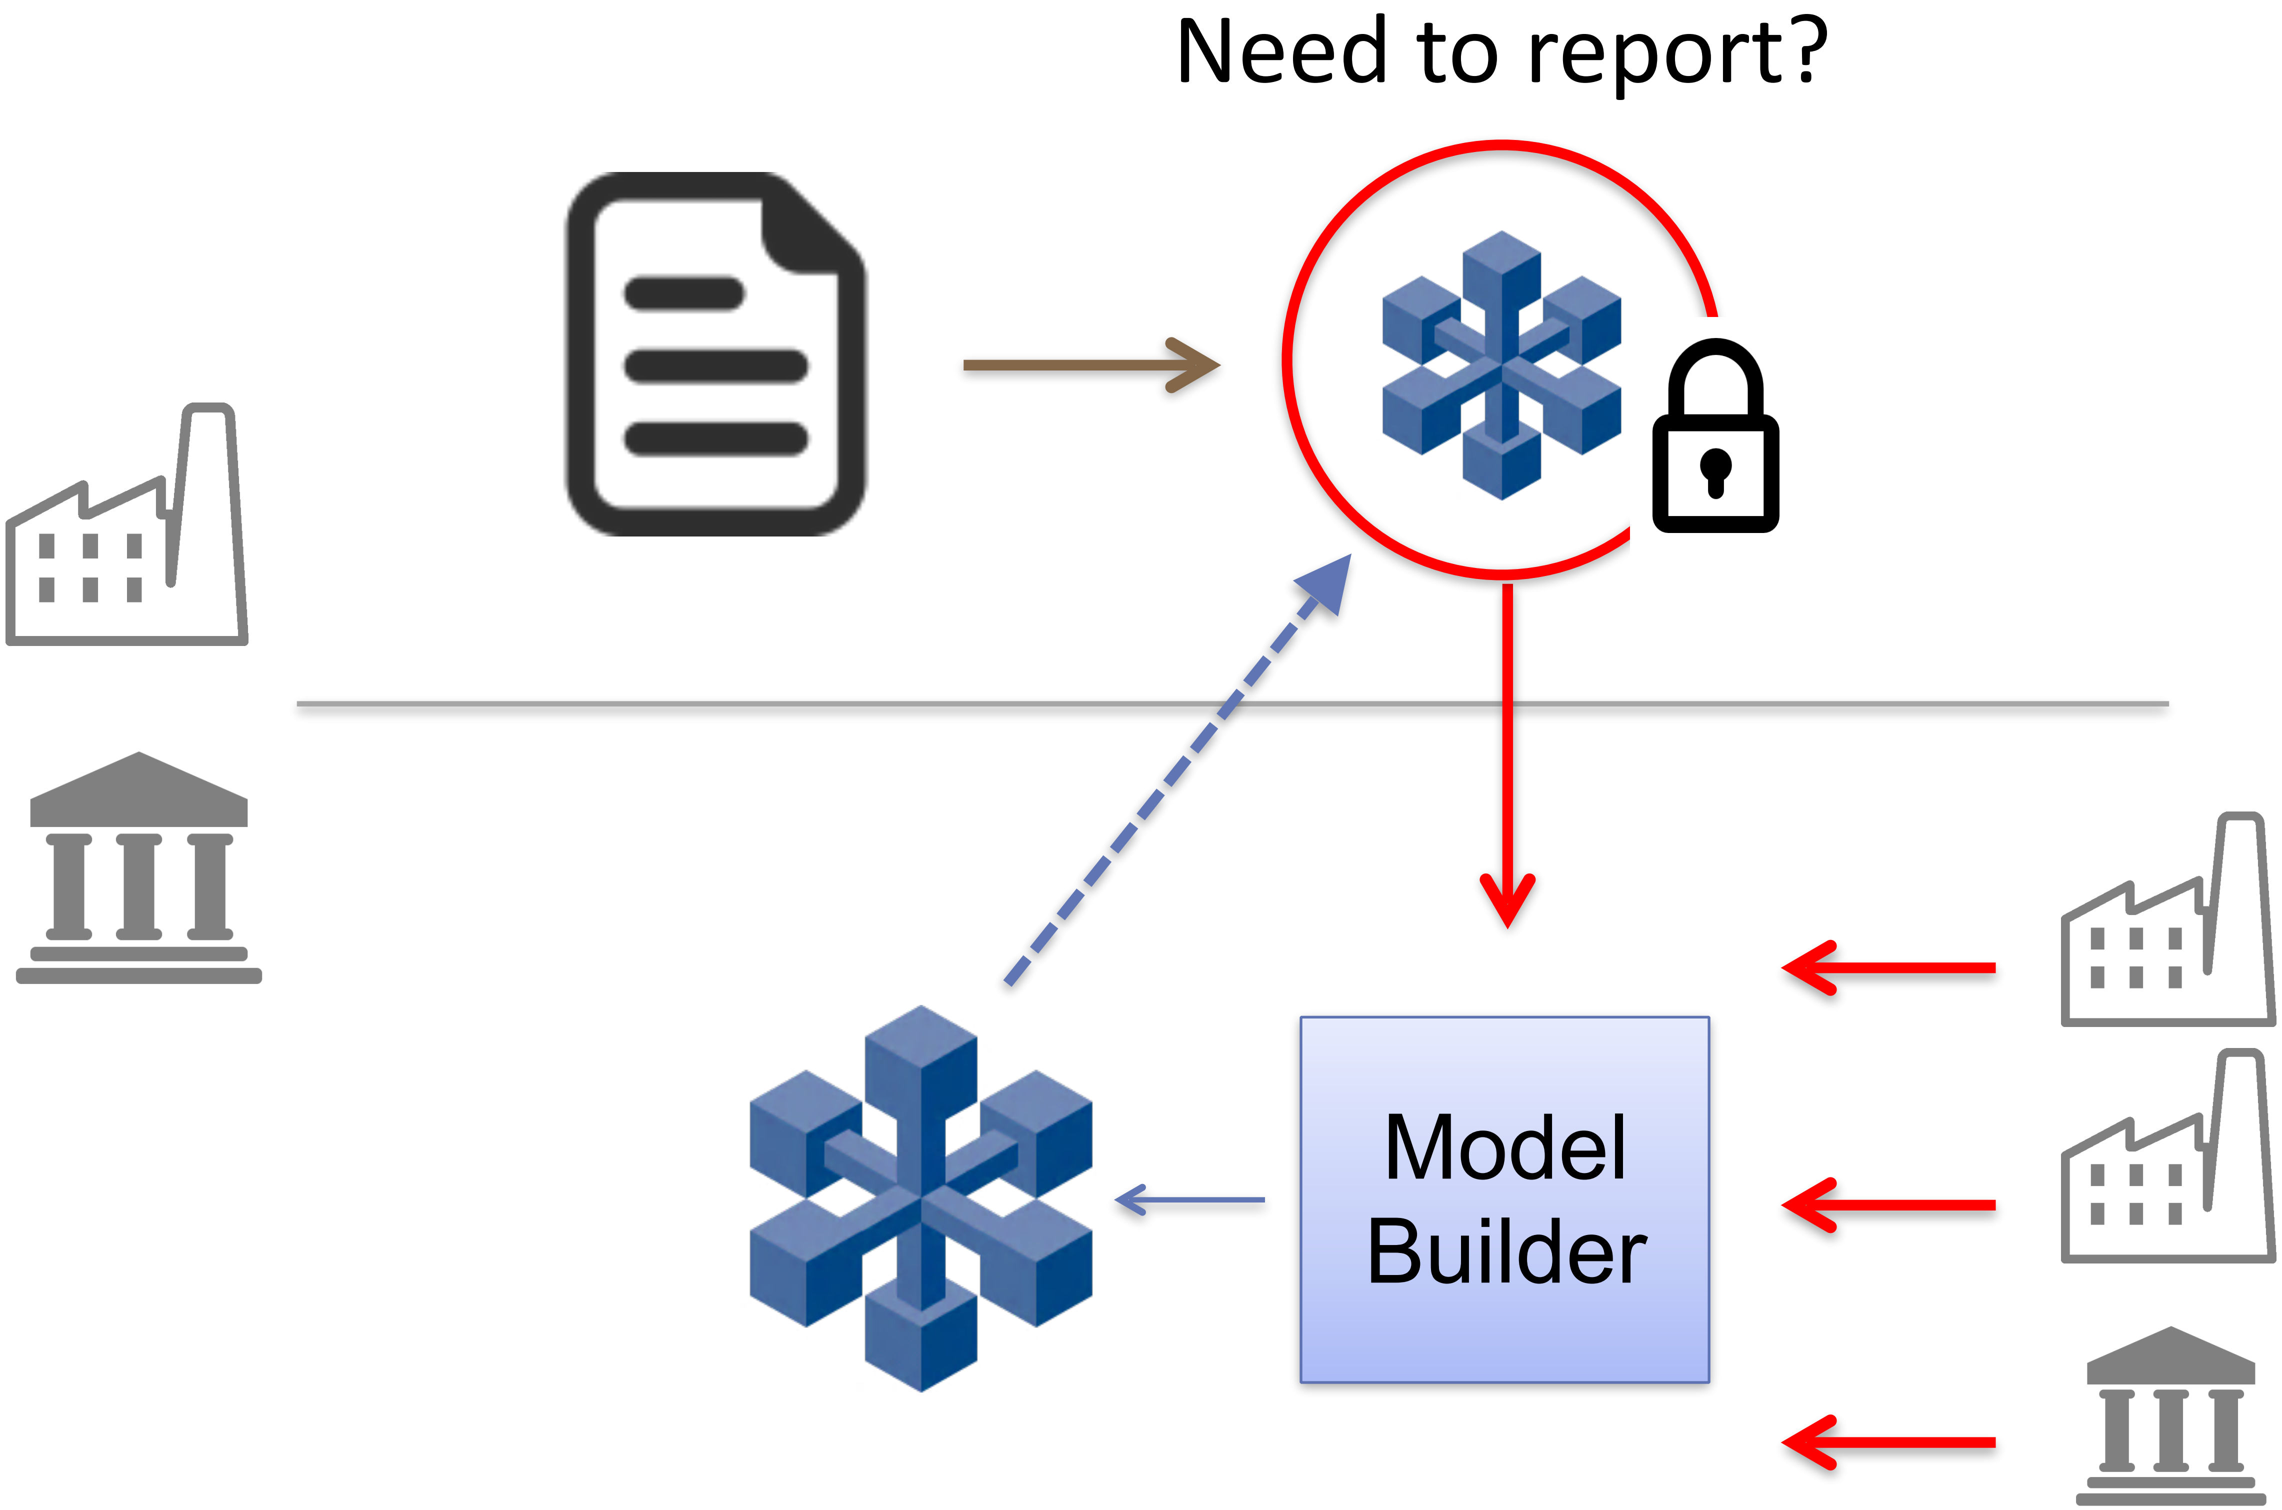

Industry using Gov data

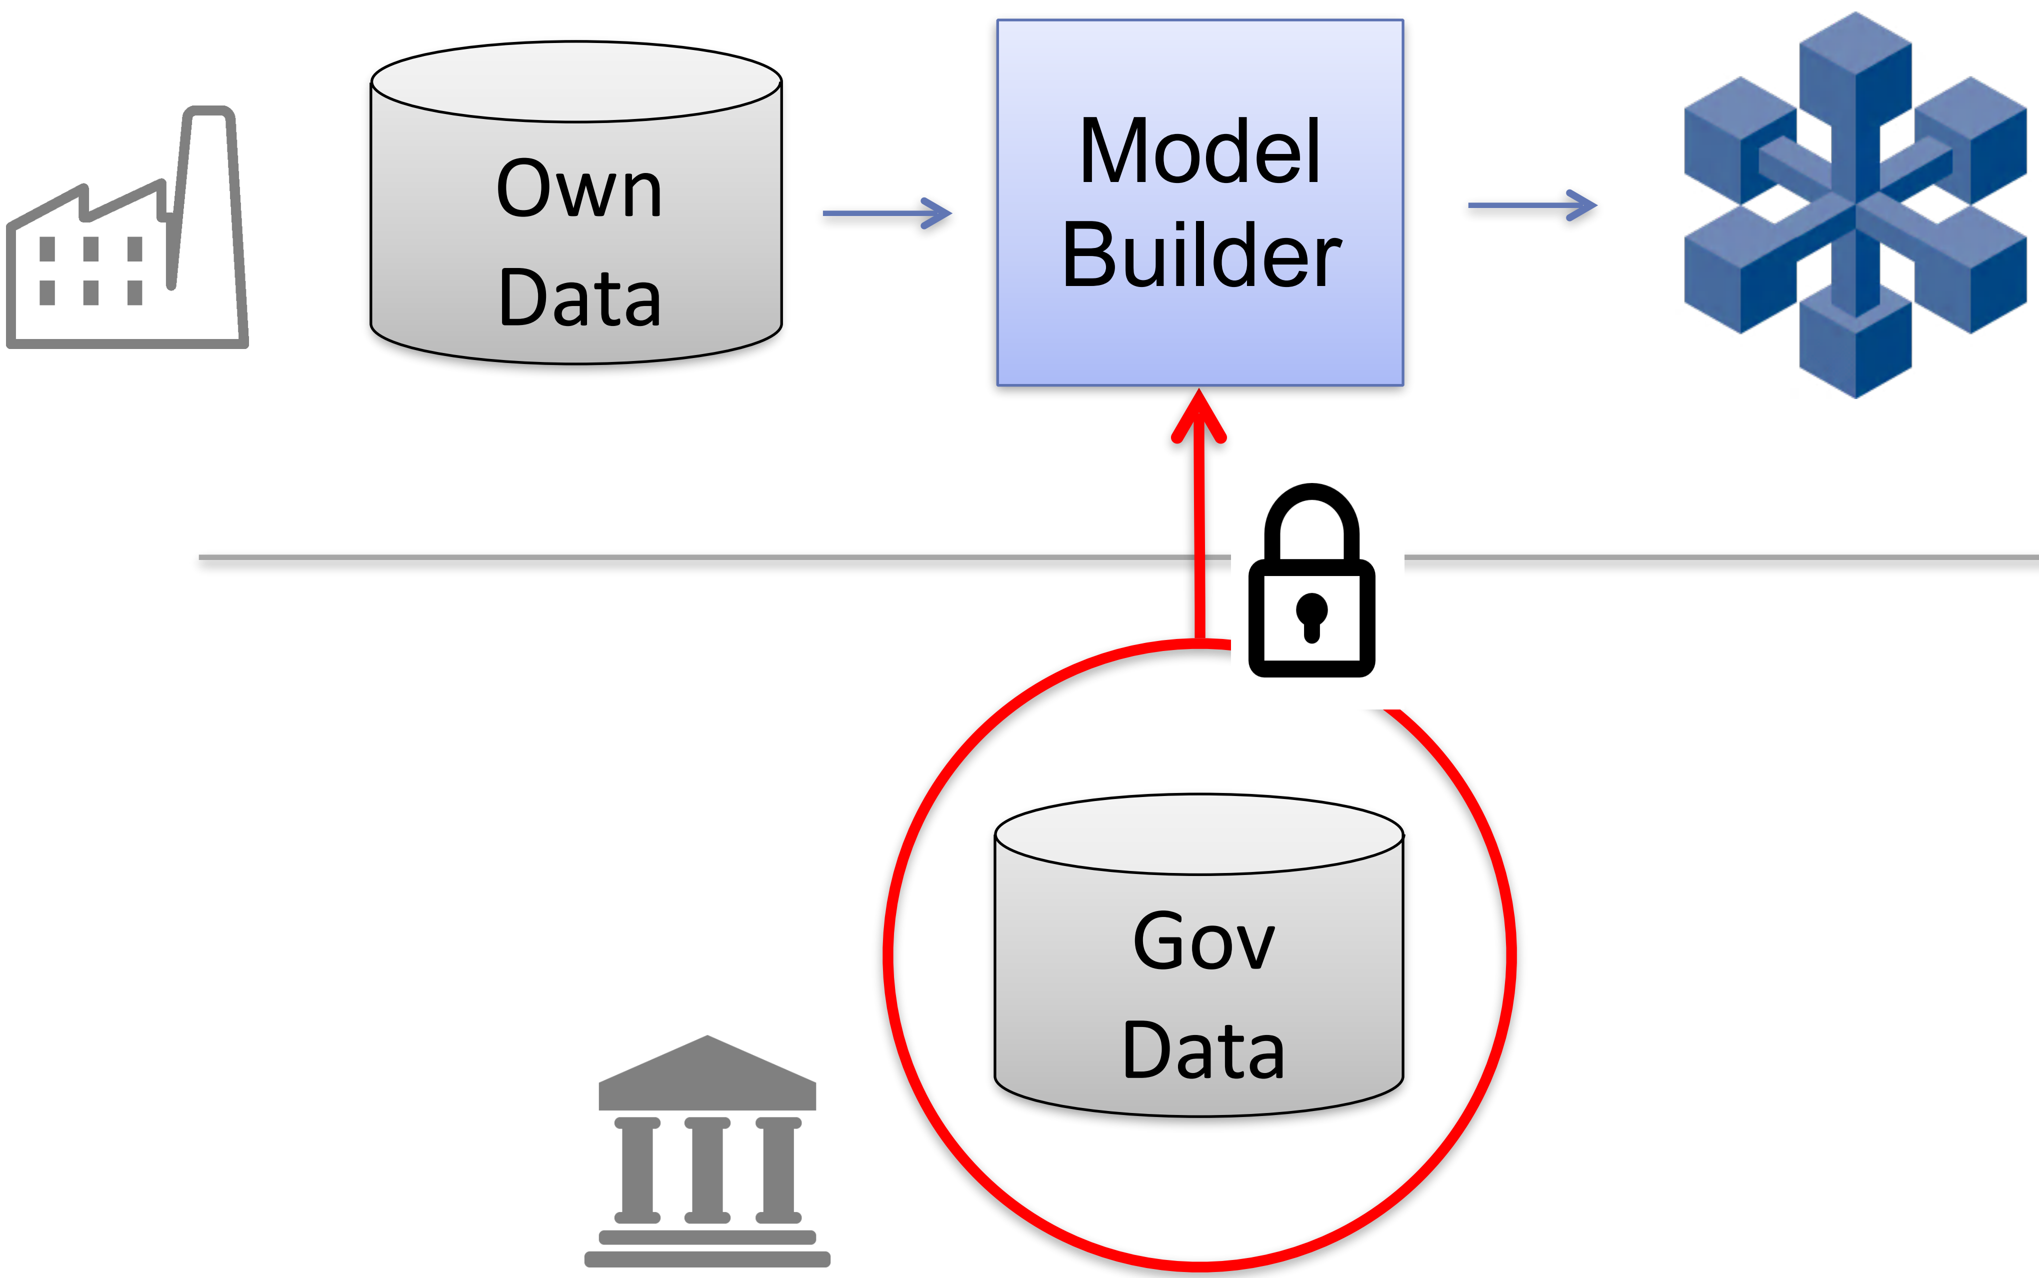

Benchmarking

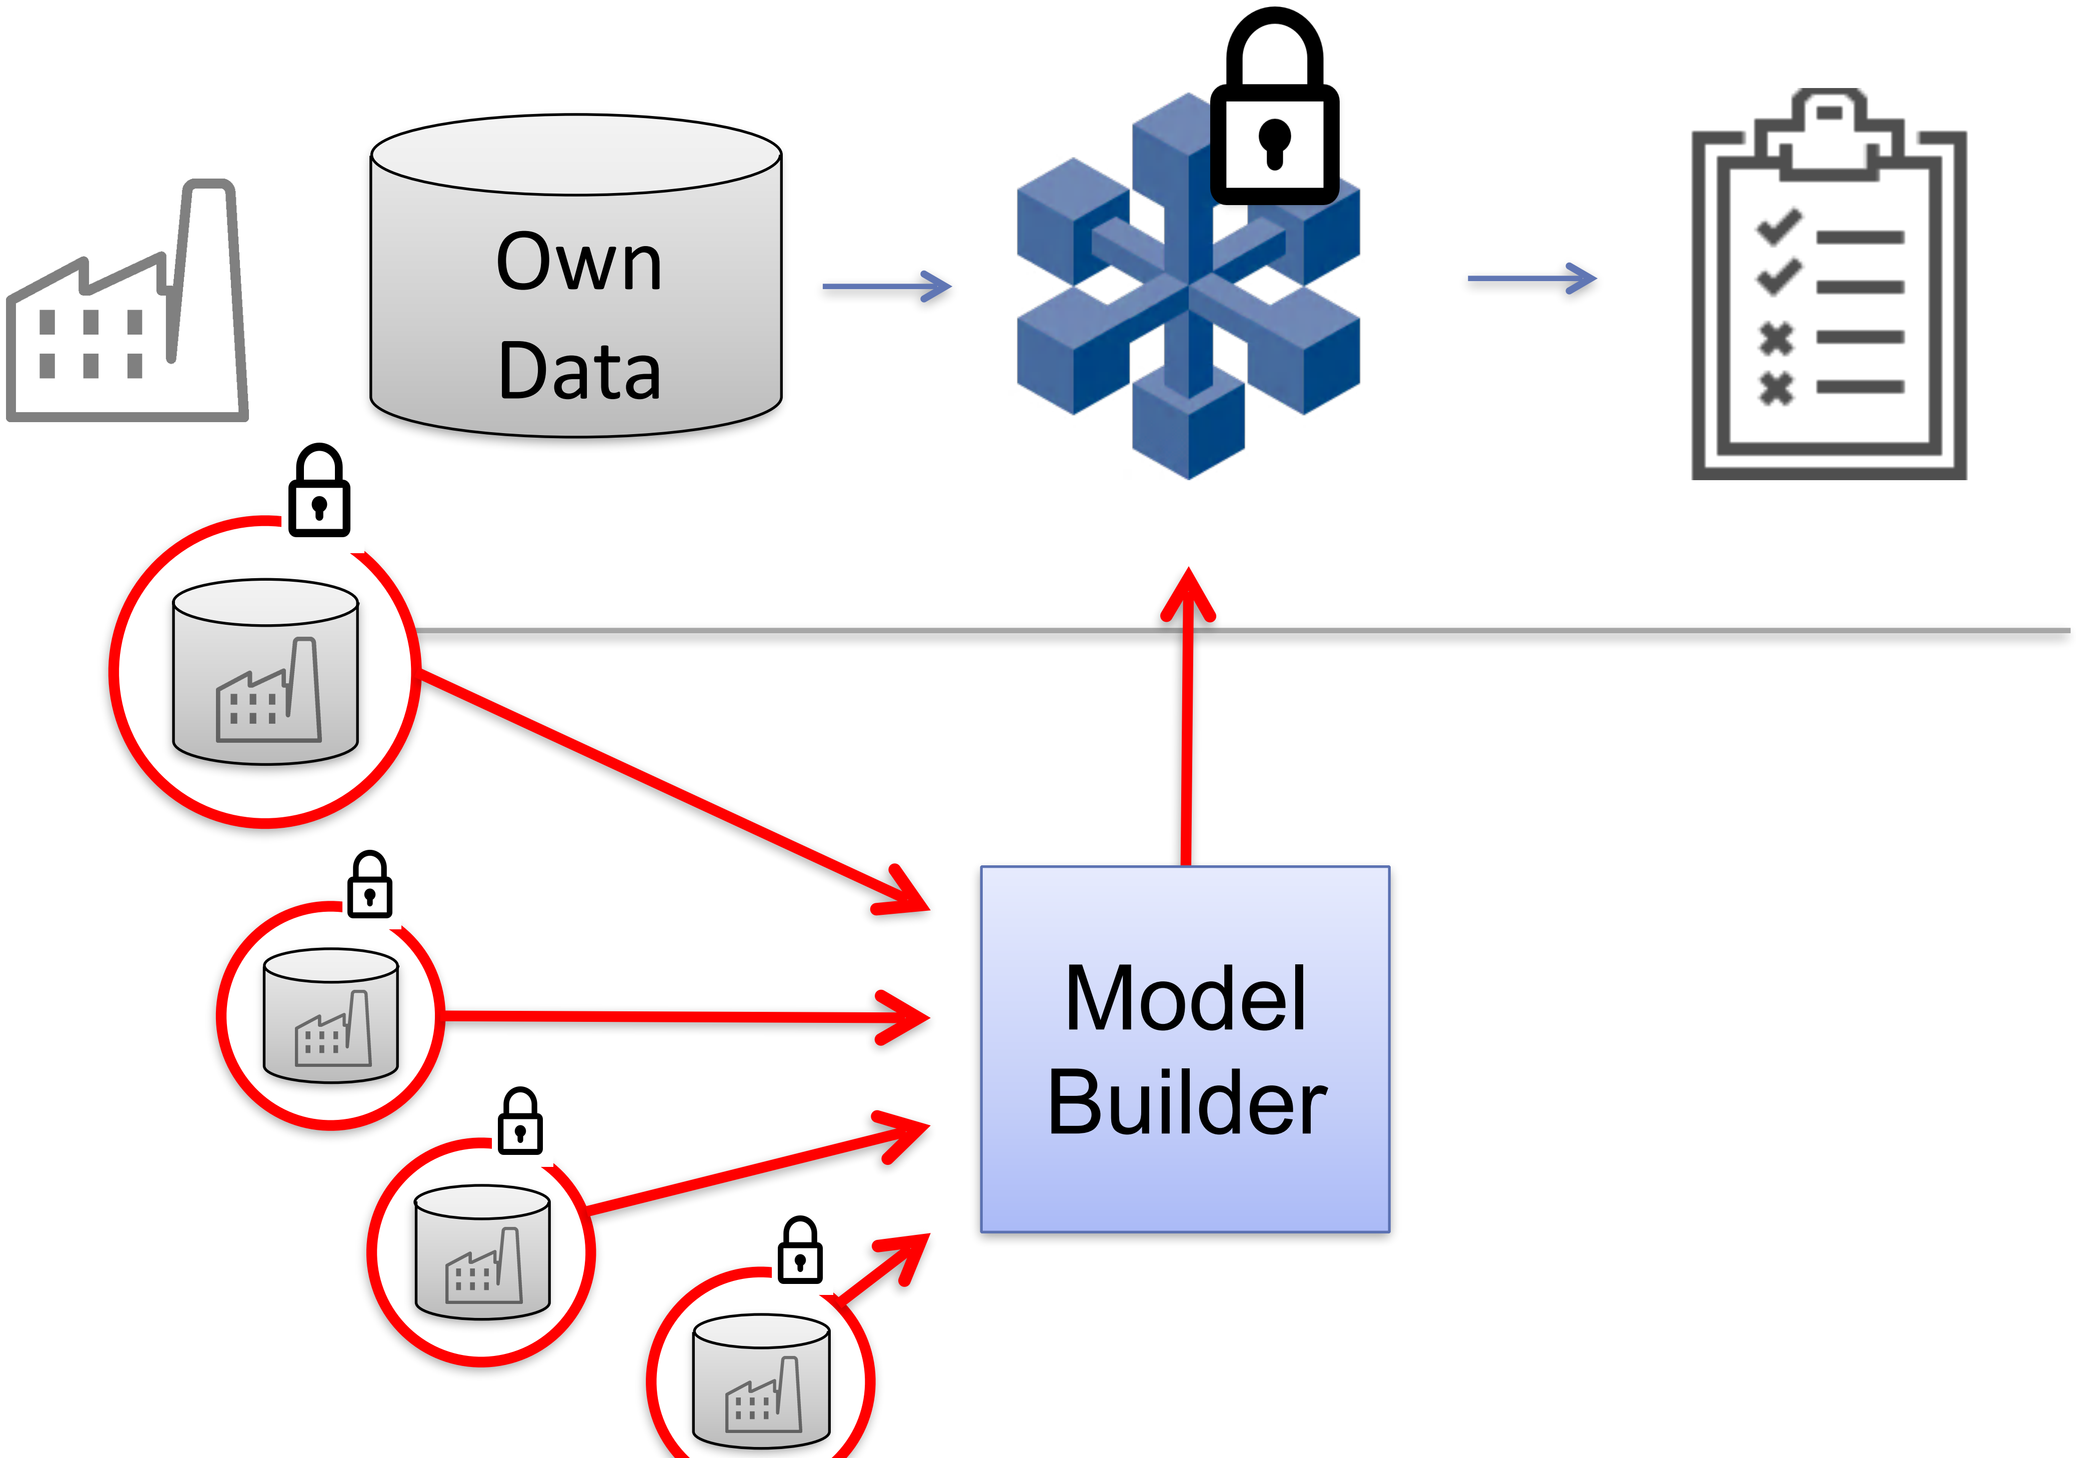

Predictive Maintainance

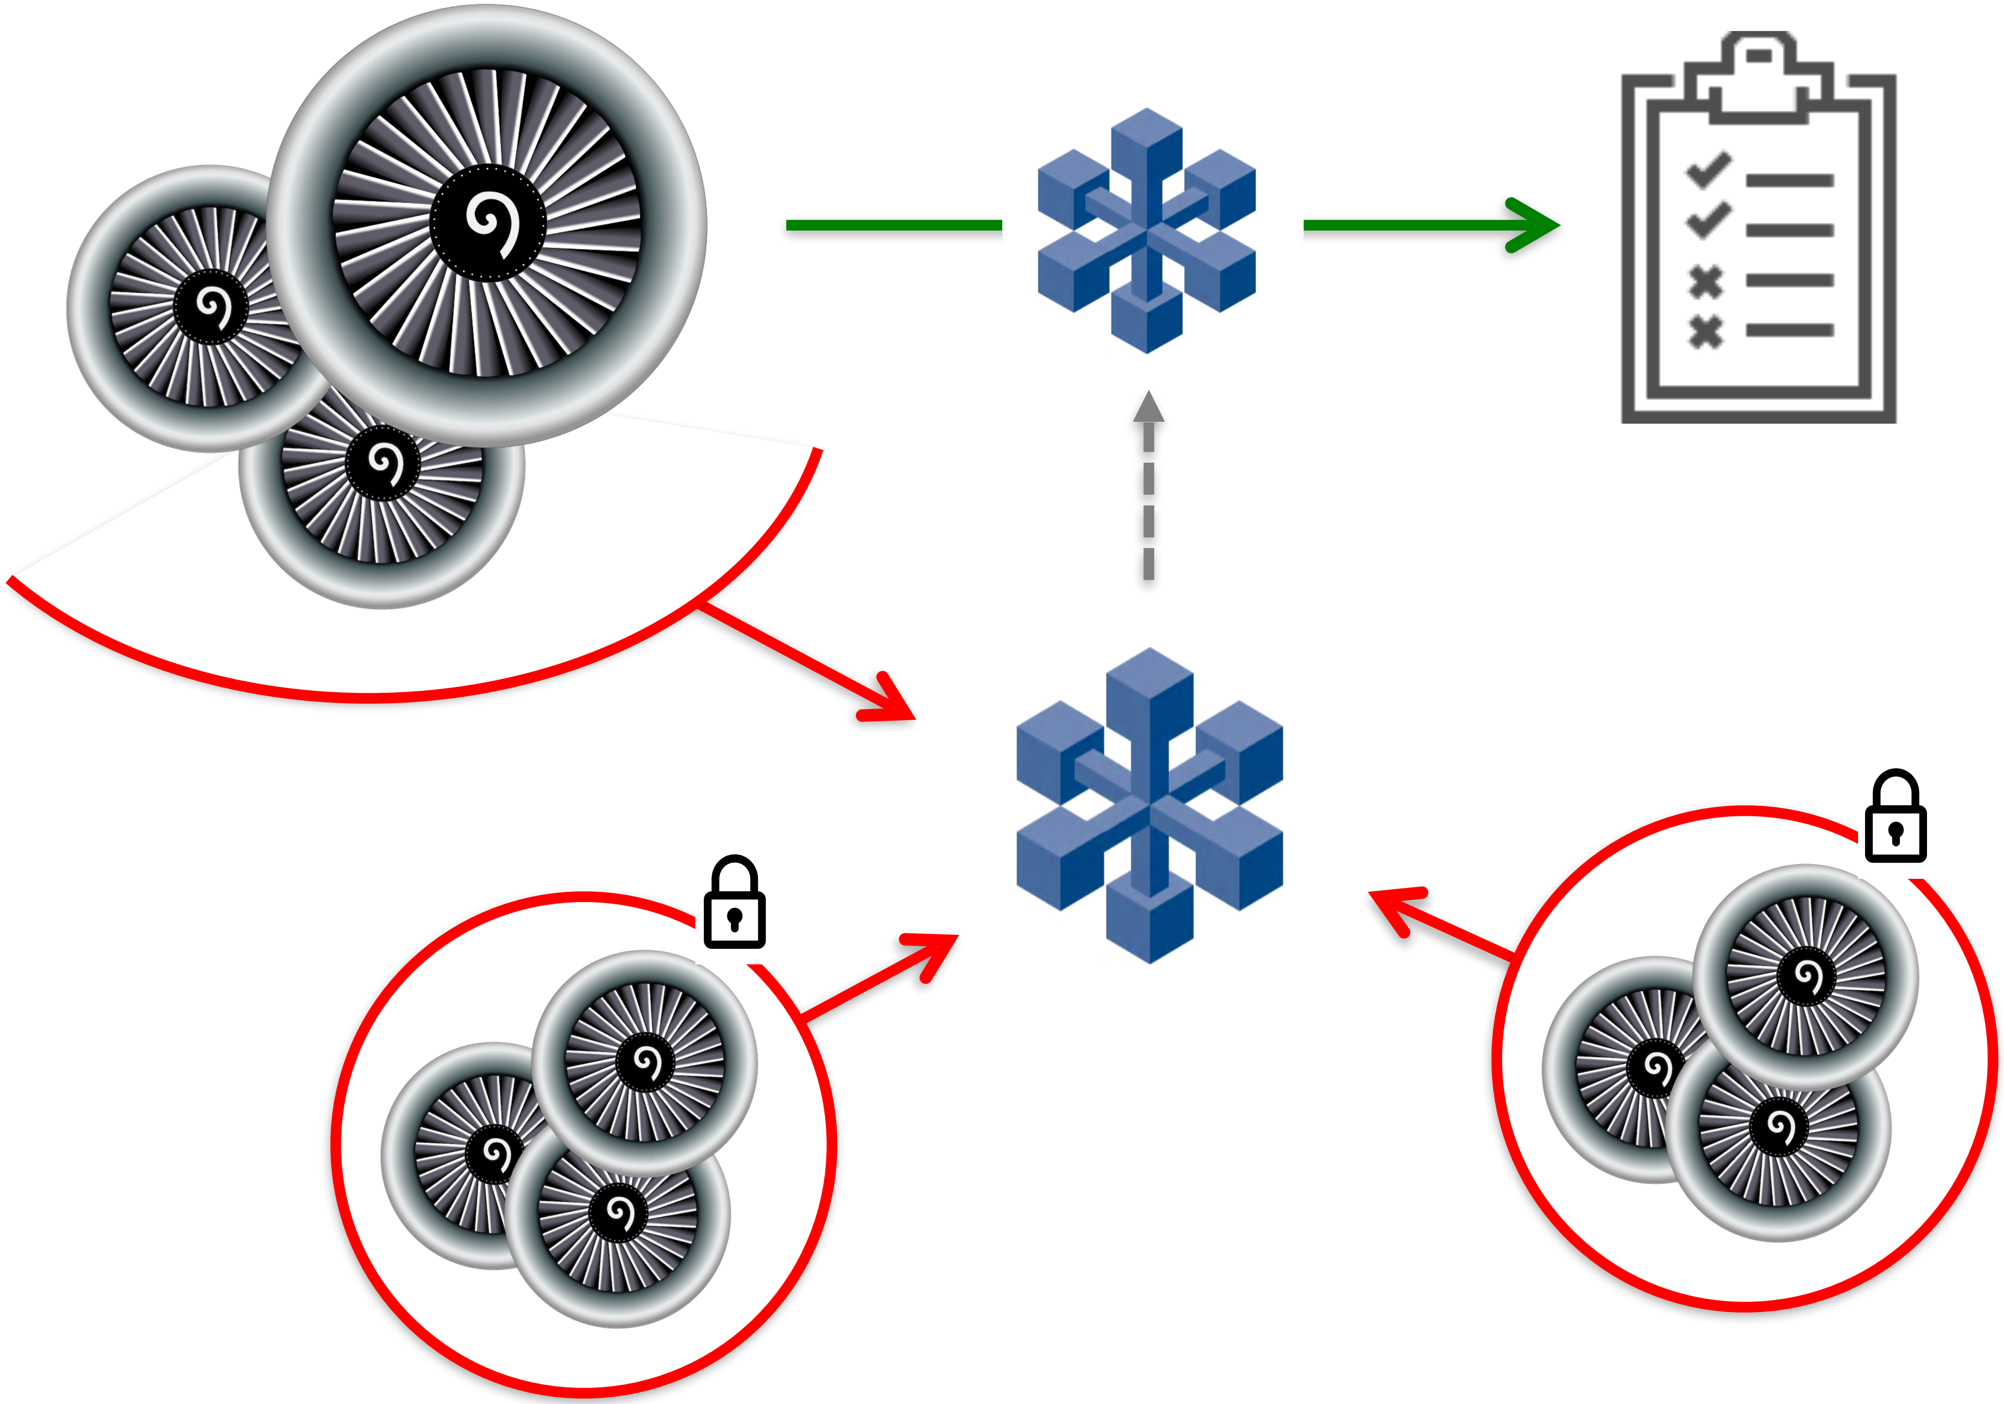

Device analytics

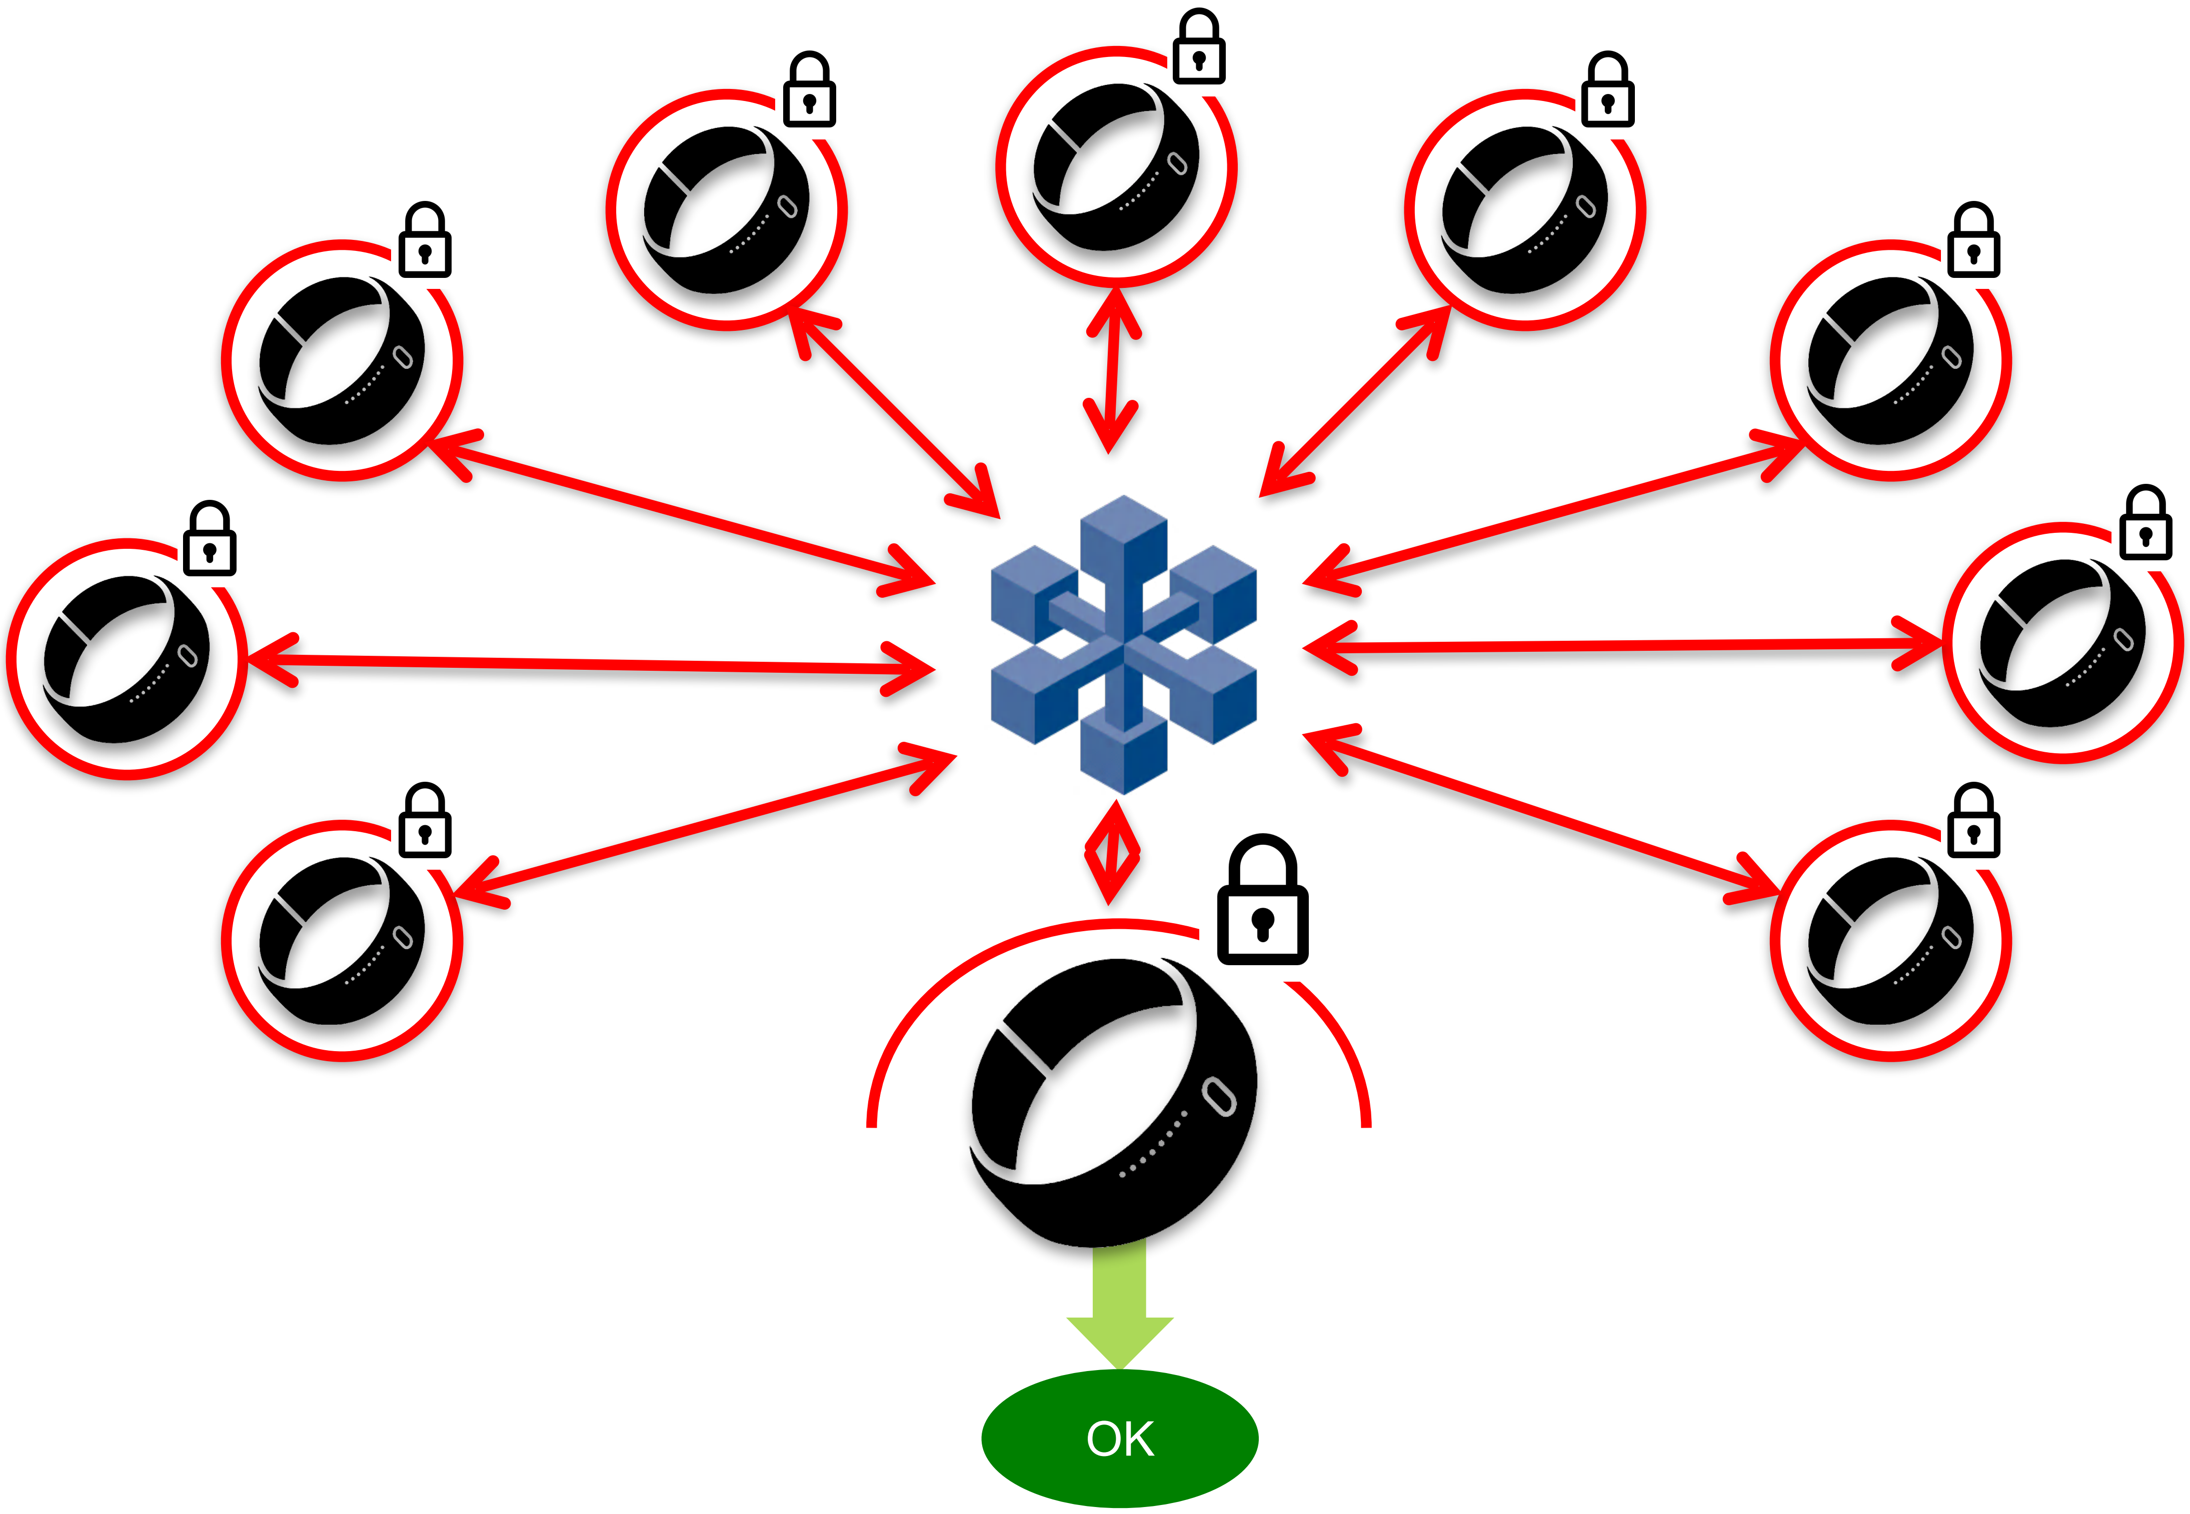

N1 Analytics and an example

N1 Analytics

Platform for federated private analytics
- Automated private record linkage
- Paillier encryption
- Rados
- Web APIs, Java/python Implementation

Standard data analytics techniques on secret data:
- Correlation analysis
- Classification / prediction
- Clustering
- Statistics

Fine grained access control

Scales to millions of records x hundreds of features

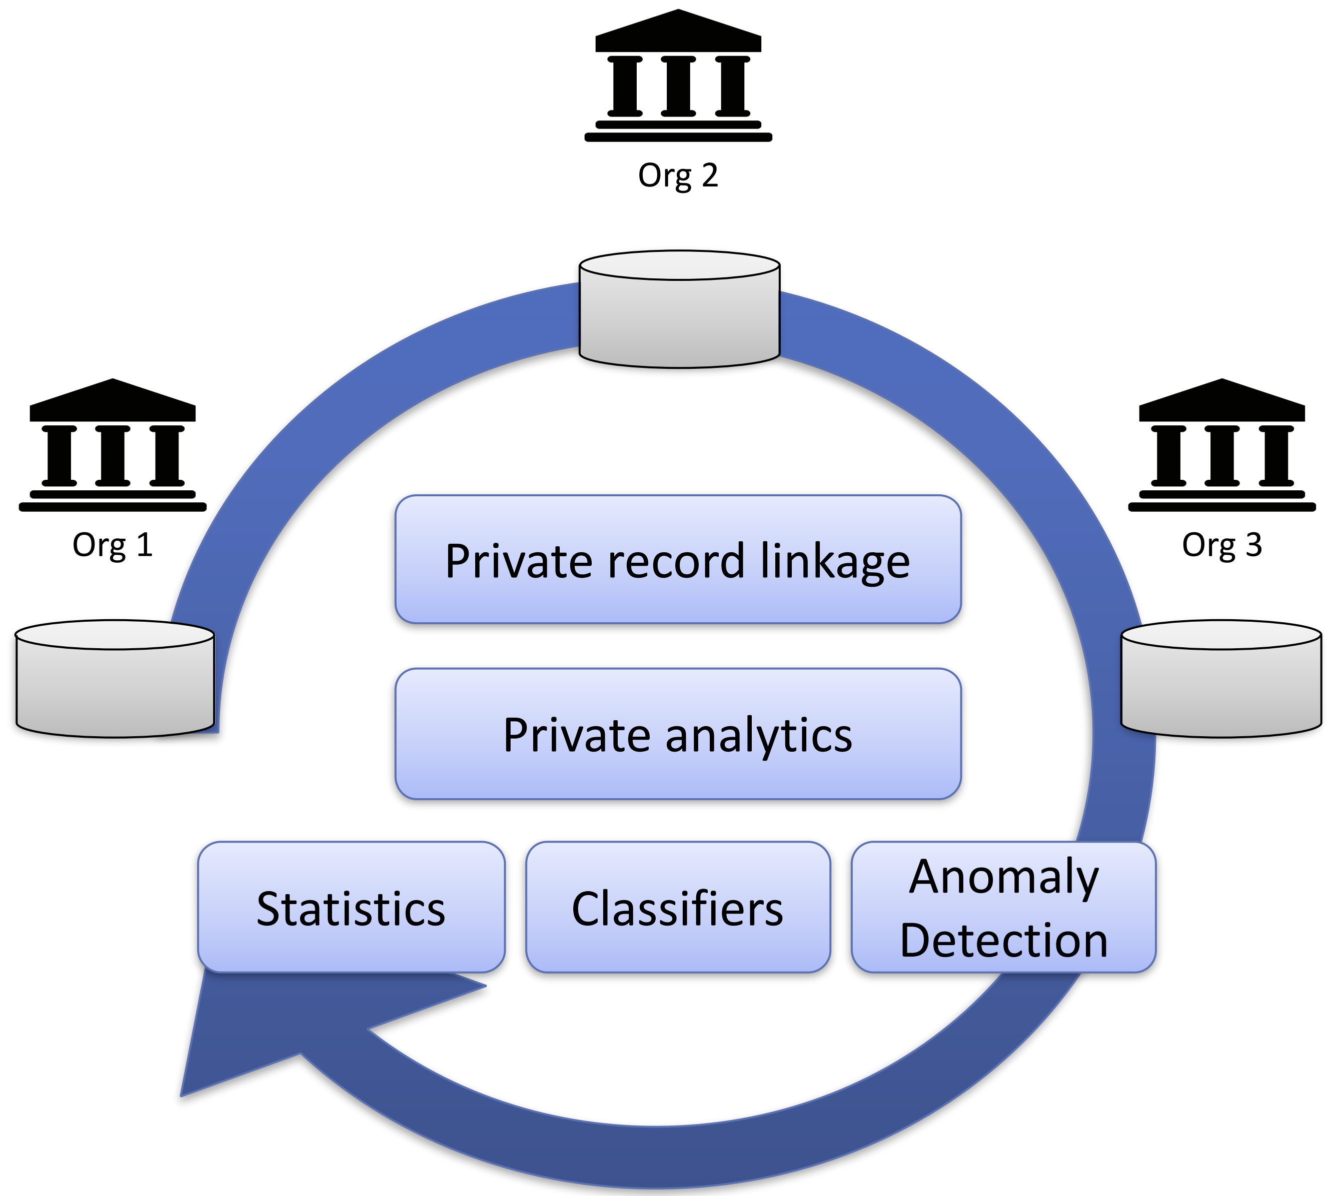

The three basic N1 building blocks

- Private computation
  - Arithmetic on encrypted numbers
- Distributed, confidential analytics
  - Distributed algorithms, computation & protocols
- Private Record Linkage
  - Privacy preserving record level matching

Homomorphic encryption

Partial Homomorphic Encryption	Allows either addition or multiplication of encrypted numbers
Somewhat Homomorphic Encryption	Allows evaluation of low order polynomials
Fully Homomorphic Encryption	Allows evaluation of arbitrary functions

Paillier encryption

Encryption of m: $$c = g^m r^n \mod n^2$$

Addition of encrypted numbers:

$$\mathrm{D}(\mathrm{E}(m_1).\mathrm{E}(m_2) \mod n^2) = m_1 + m_2 \mod n$$

Multiplication of encrypted number by a scalar:

$$\mathrm{D}(\mathrm{E}(m_1)^{m_2} \mod n^2) = m_1 m_2 \mod n$$

Paillier implementation

- Python – open source
  - www.github.com/nicta/python-paillier
- Java – open source
  - www.github.com/nicta/javallier
- Javascript – still under closed development

Distributed, Confidential Analytics

Basic definitions

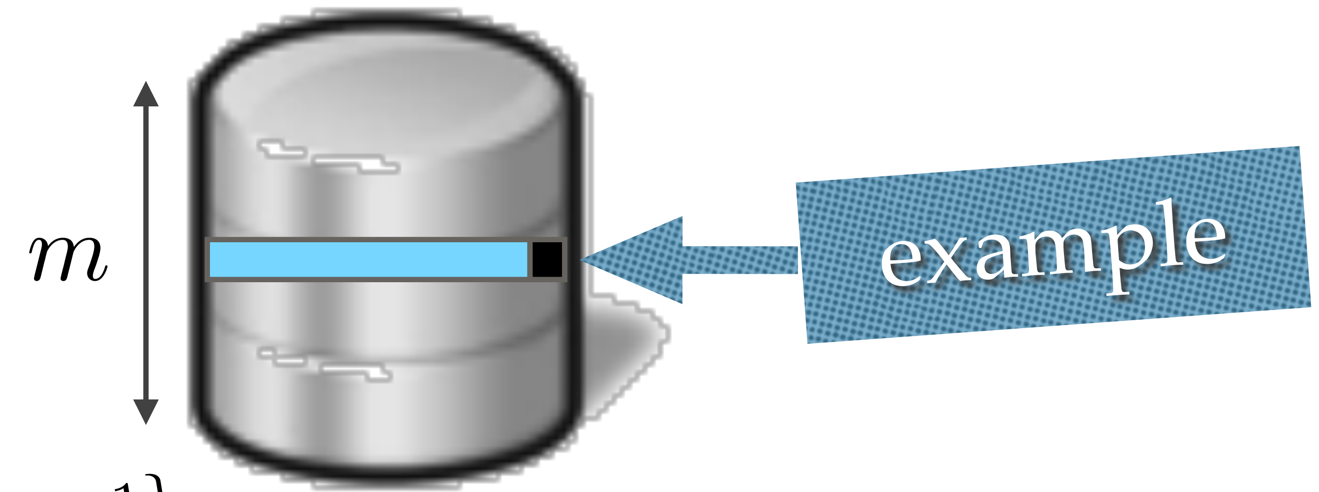

- Input: $\mathcal{S} = \{(\boldsymbol{x}_i, y_i)\}_{i=1}^{m}$ with m examples, $y_i \in \{-1, 1\}$
- Objective: learn safely linear classifier $\boldsymbol{\theta}$ …

Classical technique in the encrypted domain

Minimise for $\boldsymbol{\theta}$:

$$\ell_{\log}(\mathcal{S}, \boldsymbol{\theta}) = \frac{1}{m} \cdot \sum_i y_i \log \hat{p}[\boldsymbol{x}_i; \boldsymbol{\theta}] + (1 - y_i) \log(1 - \hat{p}[\boldsymbol{x}_i; \boldsymbol{\theta}])$$

Log likelihood

Evaluate:

$$\hat{p}[\boldsymbol{x}_i; \boldsymbol{\theta}] = \frac{1}{1+\exp(-\boldsymbol{\theta}^\top \boldsymbol{x}_i)}$$

Logistic function

Requires "secure log" and "secure inverse" protocol using Paillier encryption

Builds on Han et al. 2010 "Privacy Preserving Gradient Descent Methods"

New techniques: public references

- Giorgio Patrini, Richard Nock, Paul Rivera & Tiberio Caetano,
  "(Almost) No label No Cry"
  in ***NIPS 2014***
- Richard Nock, Giorgio Patrini, Arik Friedman,
  "Rademacher Observations, Private Data, and Boosting"
  in ***ICML 2015***
- Giorgio Patrini, Richard Nock, Stephen Hardy, Tiberio Caetano
  "Fast Learning from Distributed Datasets without Entity Resolution"
  in ***IJCAI 2016***
- Richard Nock
  "On Regularizing Rademacher Observation Losses"
  in ***NIPS 2016***

New technique: outline

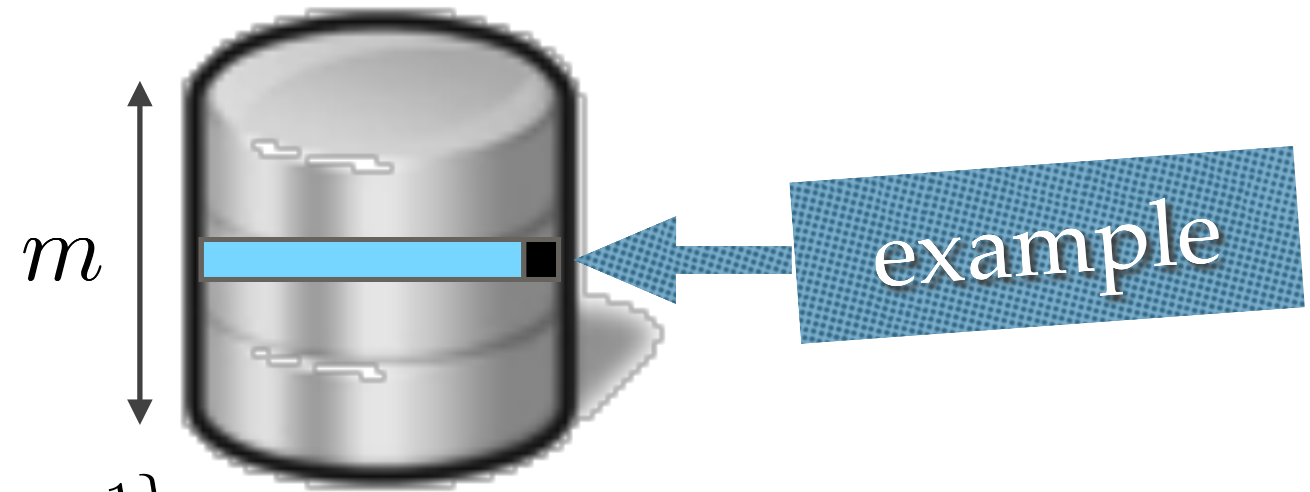

$\in \{-1, 1\}$

- Input: $\mathcal{S} = \{(\boldsymbol{x}_i, y_i)\}_{i=1}^m$ with m examples, Γ sym. pos. def.
- Objective: minimize Ridge regularized square loss for $\boldsymbol{\theta}$:

$$\ell_{\mathrm{sql}}(\mathcal{S}, \boldsymbol{\theta}; \Gamma) \doteq \frac{1}{m} \cdot \sum_i (1 - y_i \boldsymbol{\theta}^\top \boldsymbol{x}_i)^2 + \boldsymbol{\theta}^\top \Gamma \boldsymbol{\theta} \ .$$

linear classifier

Setting: supervised learning

basic

EASY !

- Input: $\mathcal{S} = \{(\boldsymbol{x}_i, y_i)\}_{i=1}^m$ with m examples, Γ sym. pos. def.
- Objective: minimize Ridge regularized square loss for $\boldsymbol{\theta}$:

$$\boldsymbol{\theta}^\star_{\mathrm{ex}} = \left(\mathsf{X}\mathsf{X}^\top + m \cdot \Gamma\right)^{-1} \boldsymbol{\pi}_{\boldsymbol{y}}$$

$$\mathsf{X} \doteq [\boldsymbol{x}_1 | \boldsymbol{x}_2 | \cdots | \boldsymbol{x}_m] \qquad \boldsymbol{\pi}_{\boldsymbol{y}} \doteq \sum_i y_i \cdot \boldsymbol{x}_i$$

rademacher observation (rado)

distributed: supervised learning

- Dataset "vertically" partitioned between 2 peers, P_1 and P_2.
- Have *few* shared features (postcode, gender, etc.)
- And lots of *specific* features (credit history, blood tests, etc.)

distributed: supervised learning

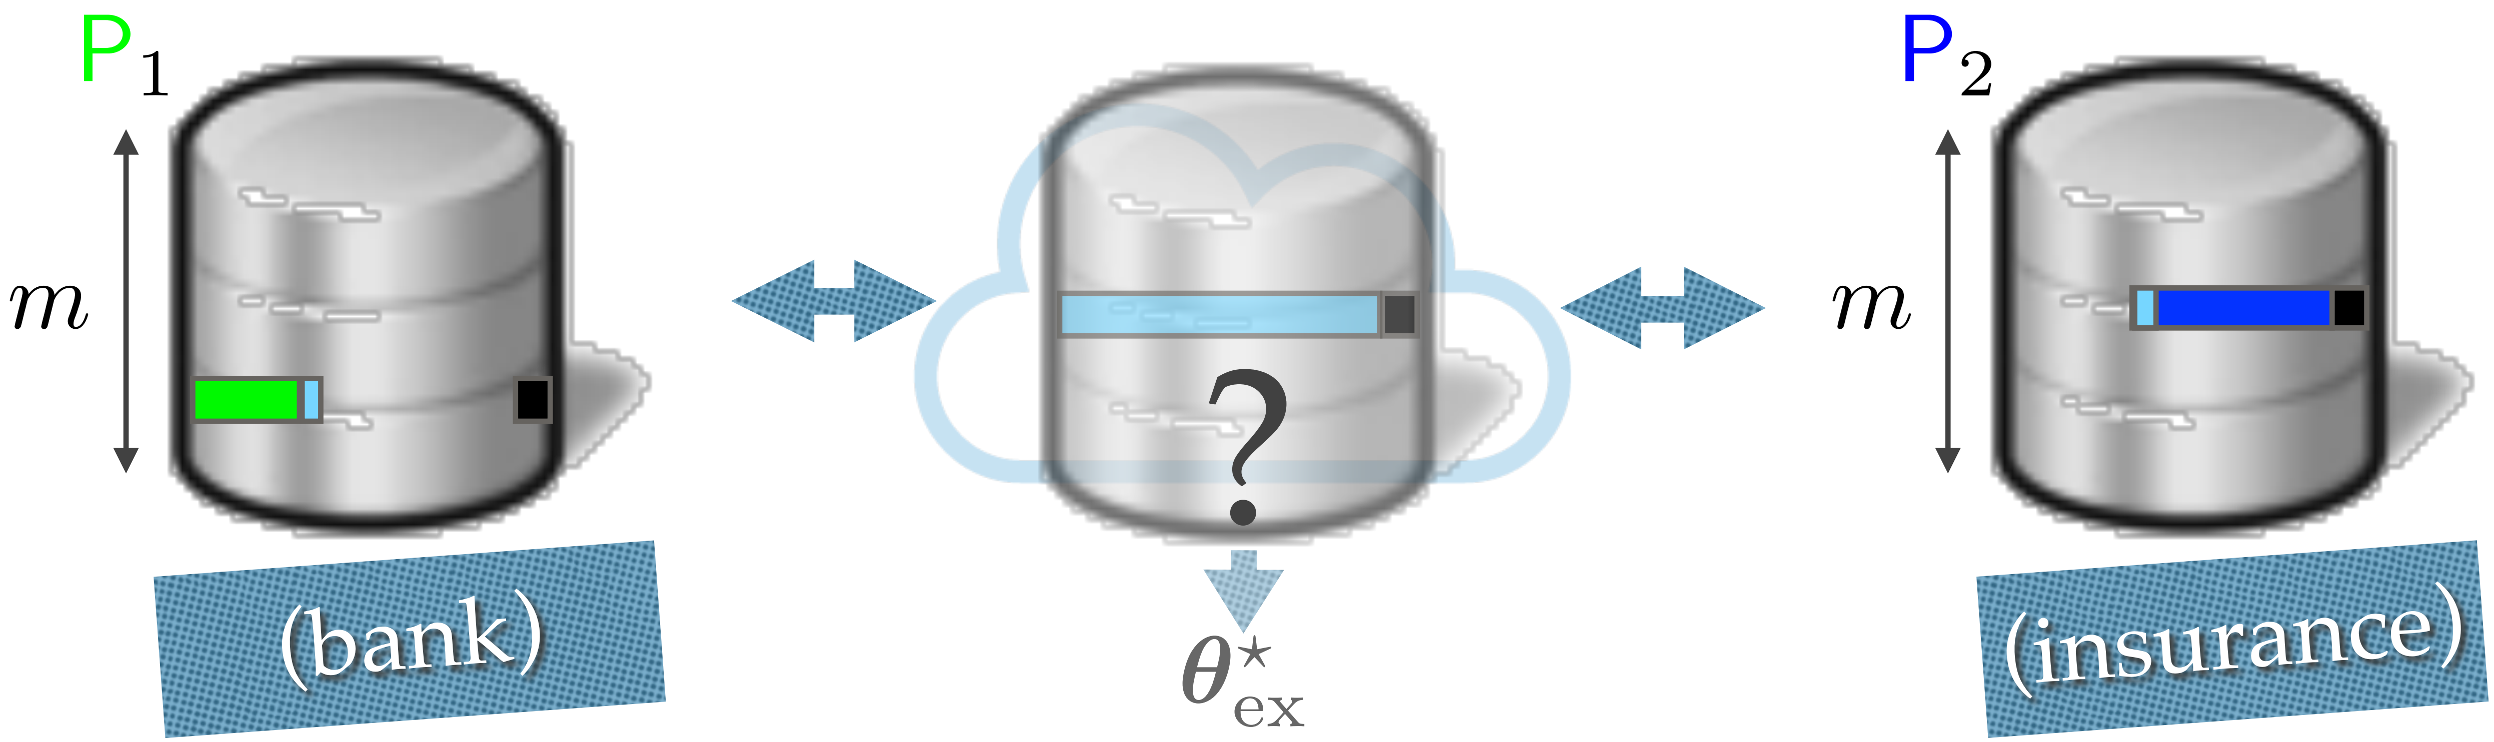

- Dataset "vertically" partitioned between 2 peers, P_1 and P_2.
- Have *few* shared features (postcode, gender, etc.)
- And lots of ***specific*** features (credit history, blood tests, etc.)
- Would like to learn $\theta^\star_{\text{ex}}$ over the union of all features…

~~Setting~~ distributed: supervised learning

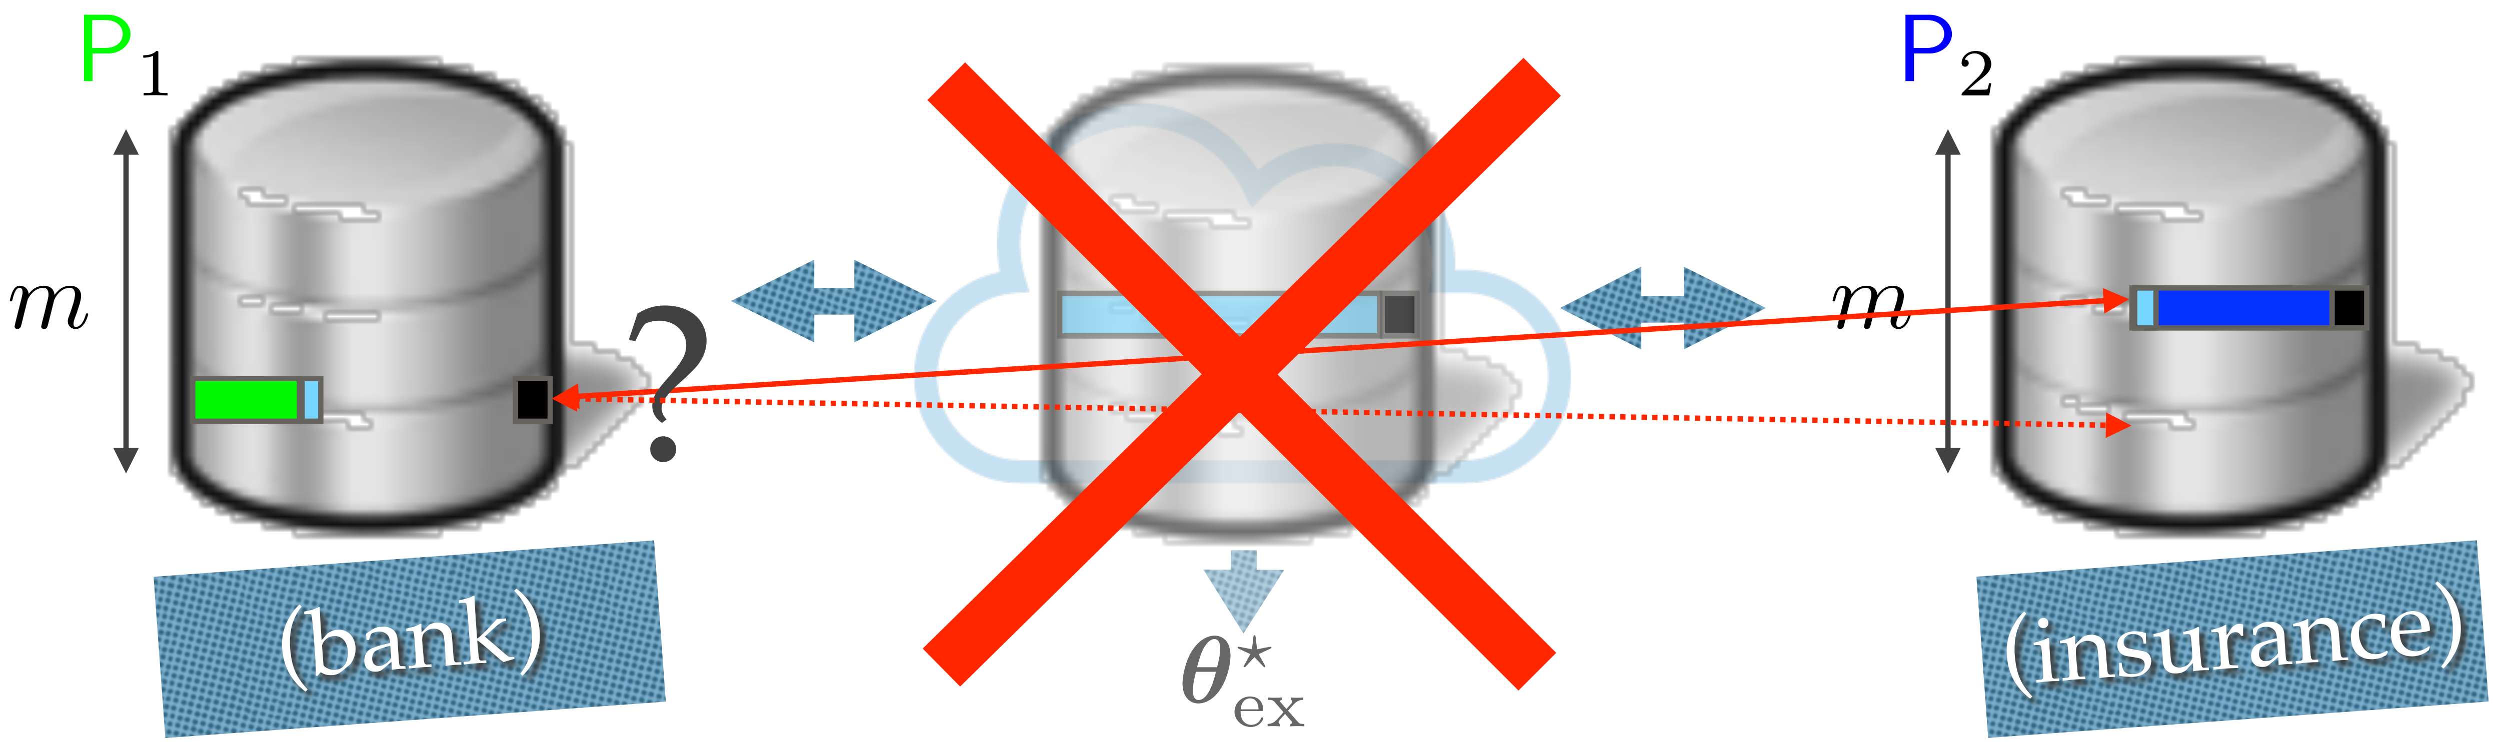

- Dataset "vertically" partitioned between 2 peers, P$_1$ and P$_2$.
- Have *few* shared features (postcode, gender, etc.)
- And lots of ***specific*** features (credit history, blood tests, etc.)
- Would like to learn $\boldsymbol{\theta}^\star_{\text{ex}}$ over the union of all features…
- But no entity matching possible ! (privacy / security)

Let's get more challenging !

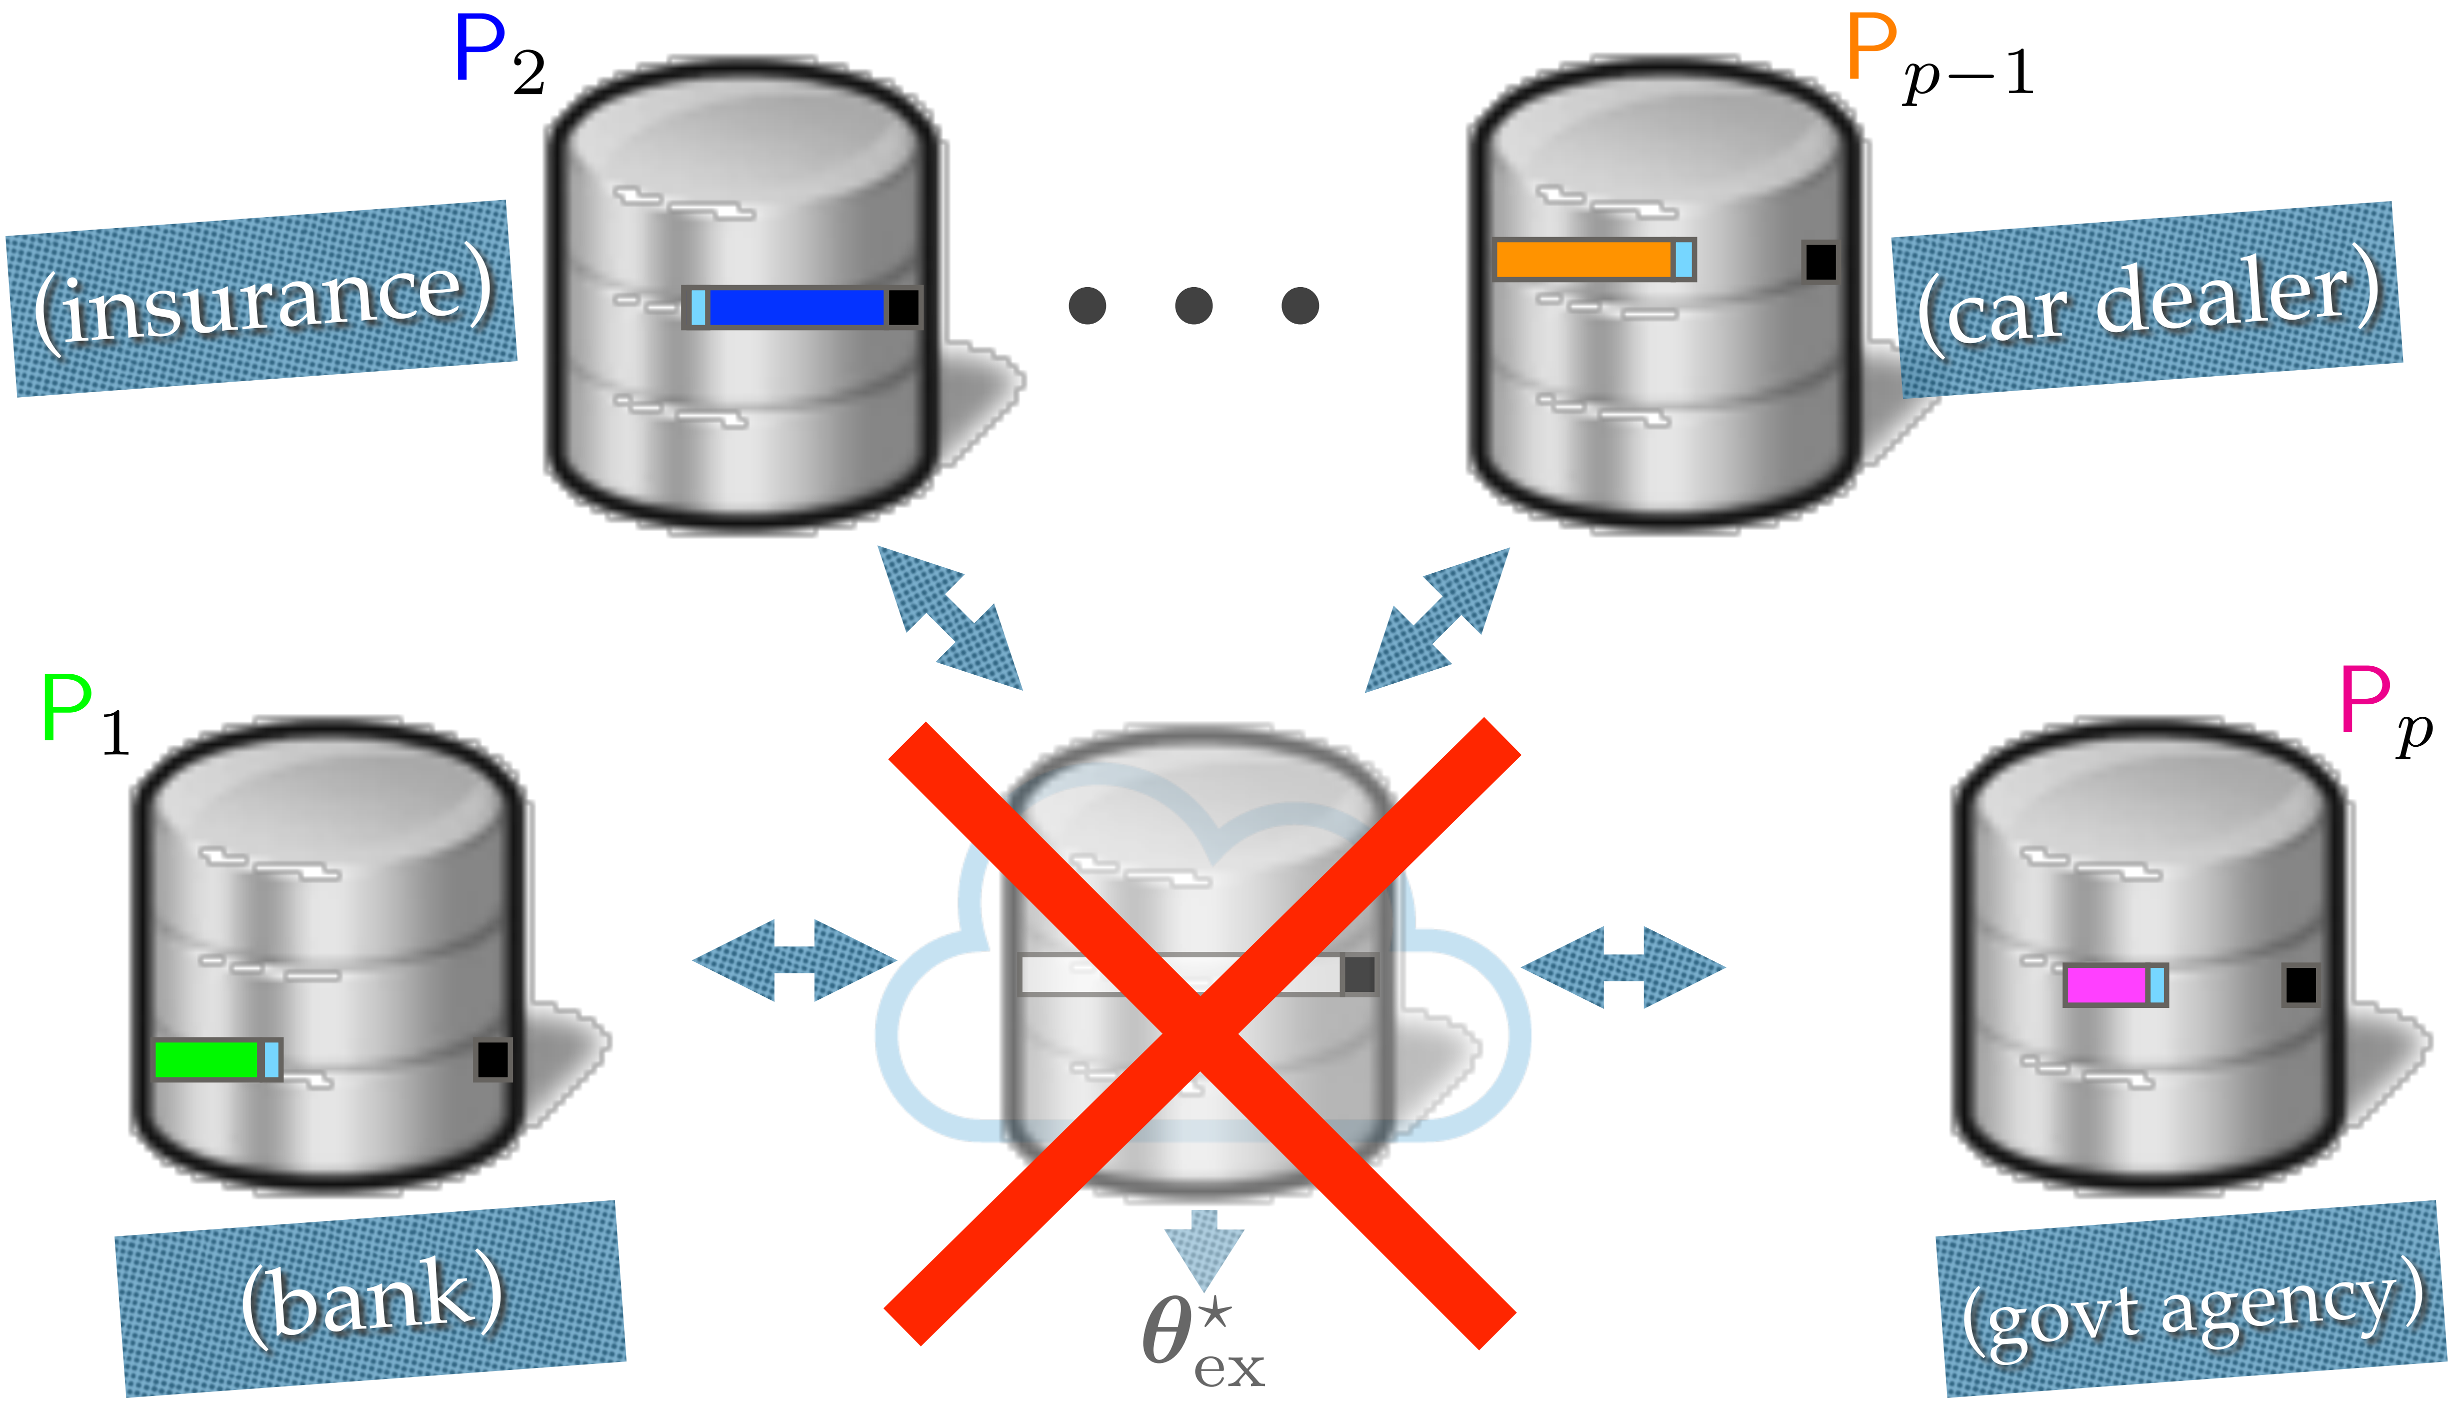

- Same setting & constraint, but arbitrary number of peers

Let's get more challenging !

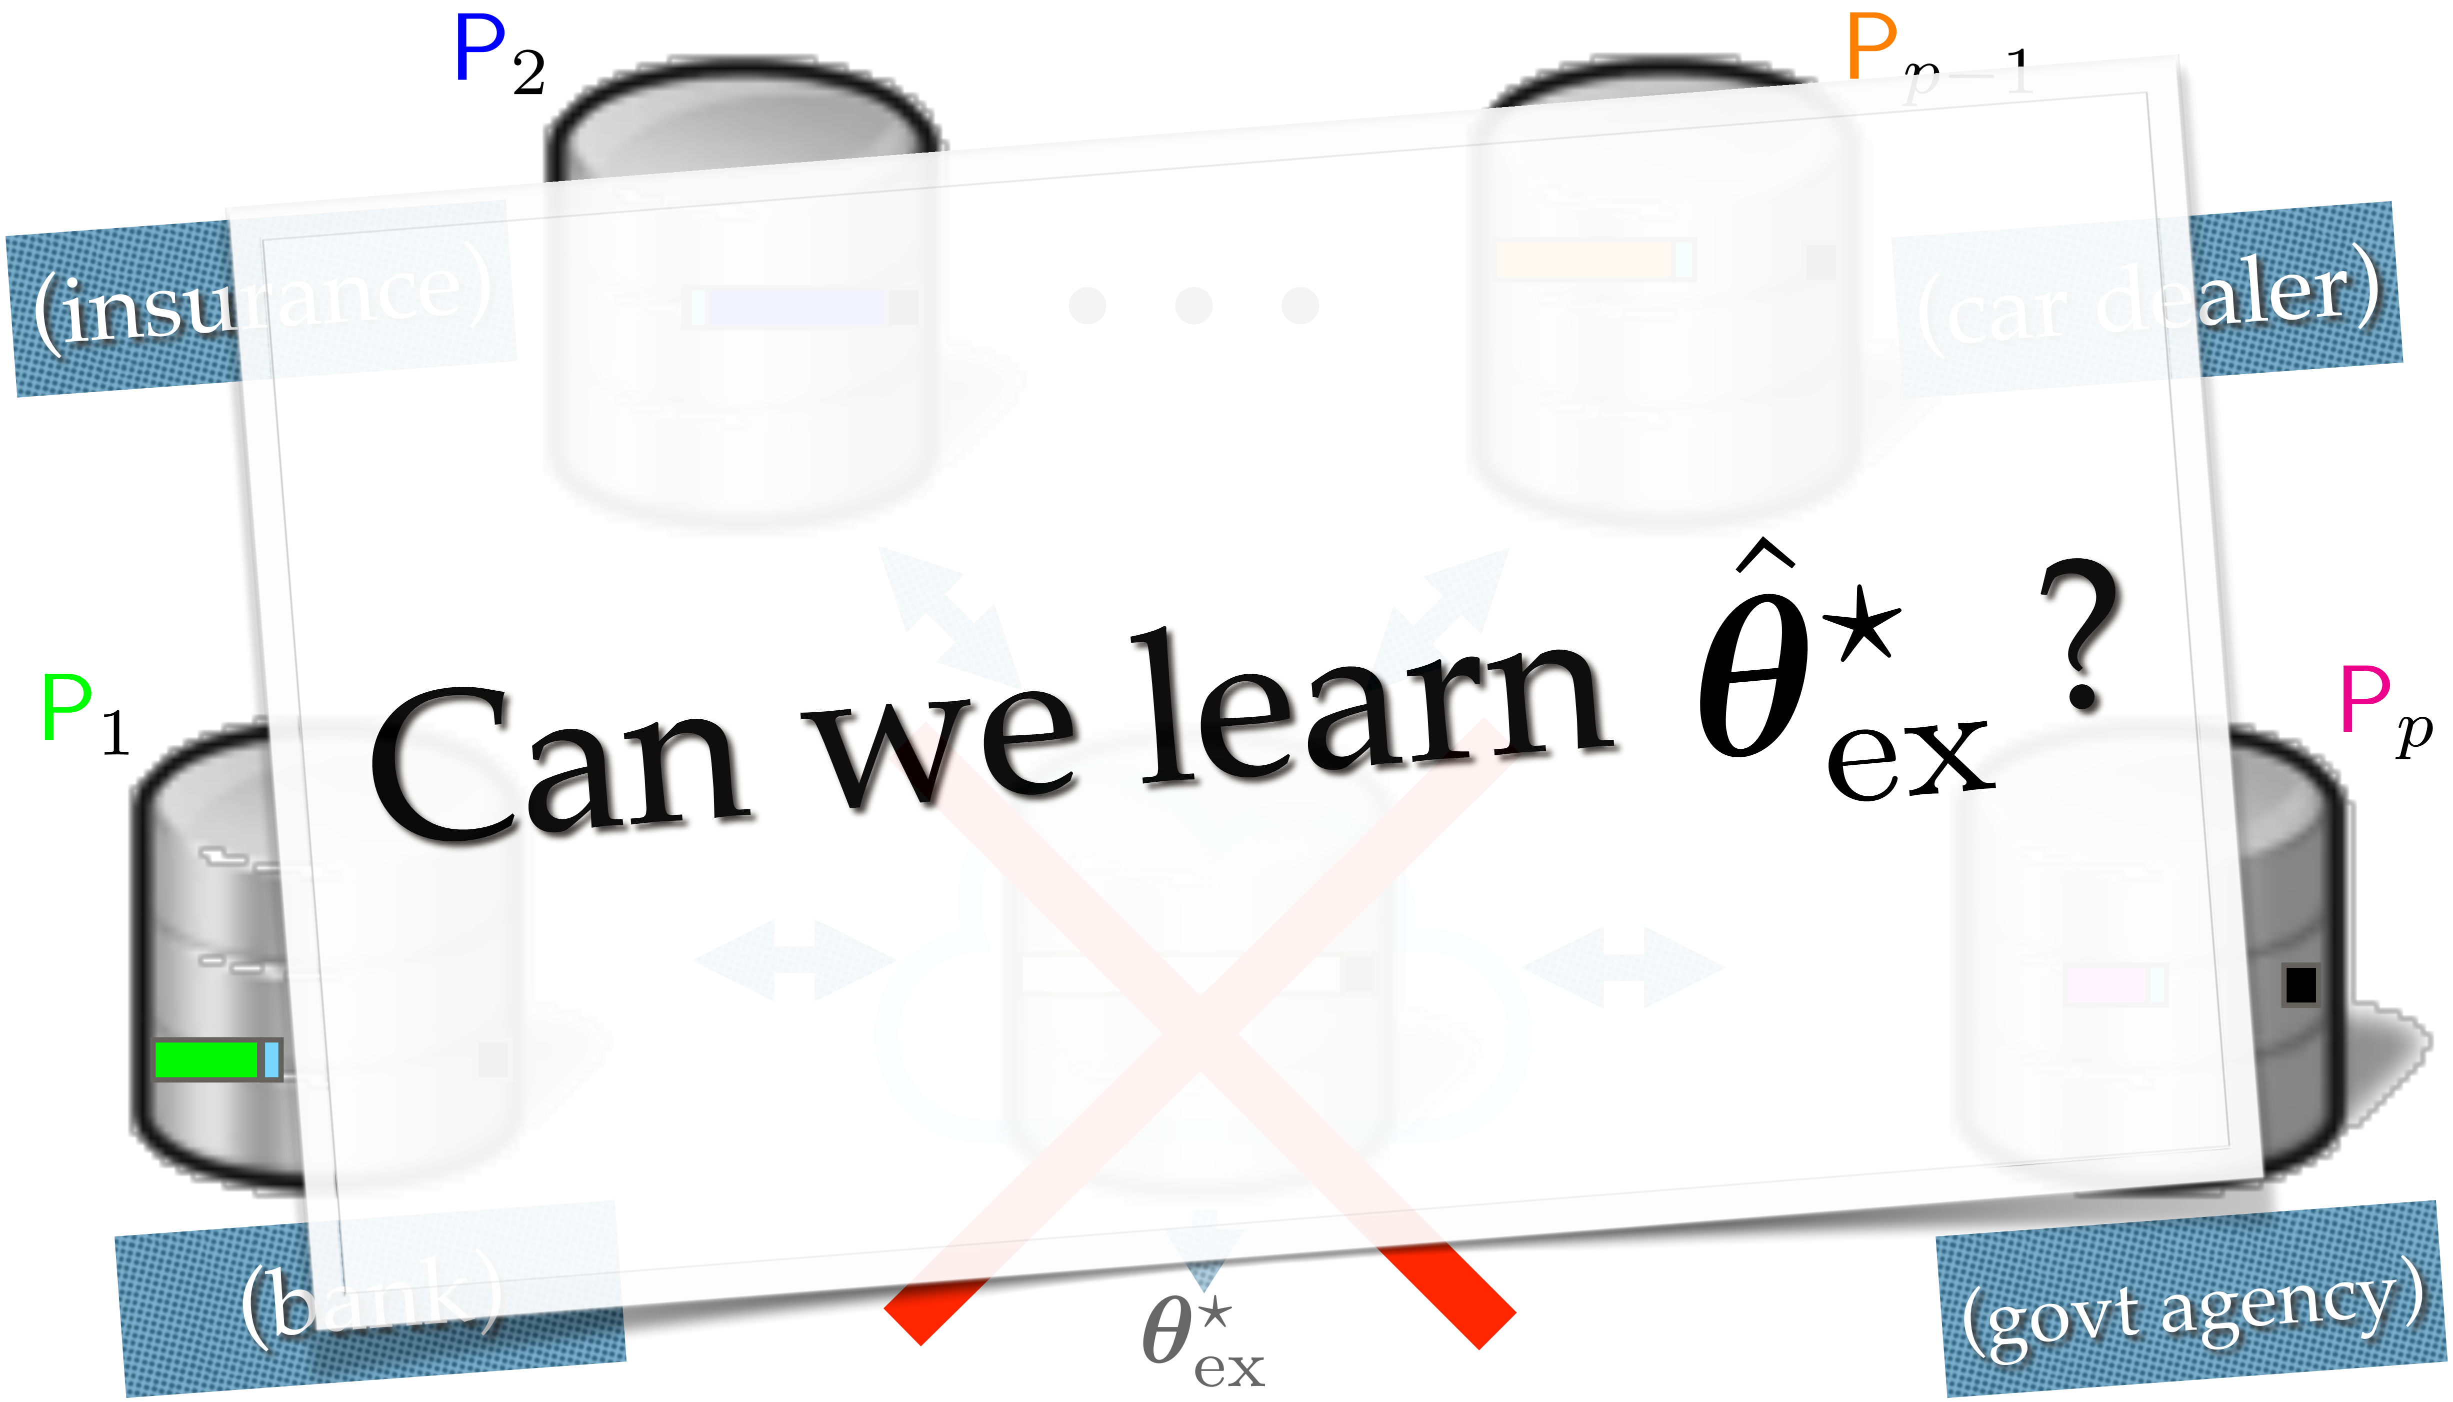

- Same setting & constraint, but arbitrary number of peers

Let's get more challenging !

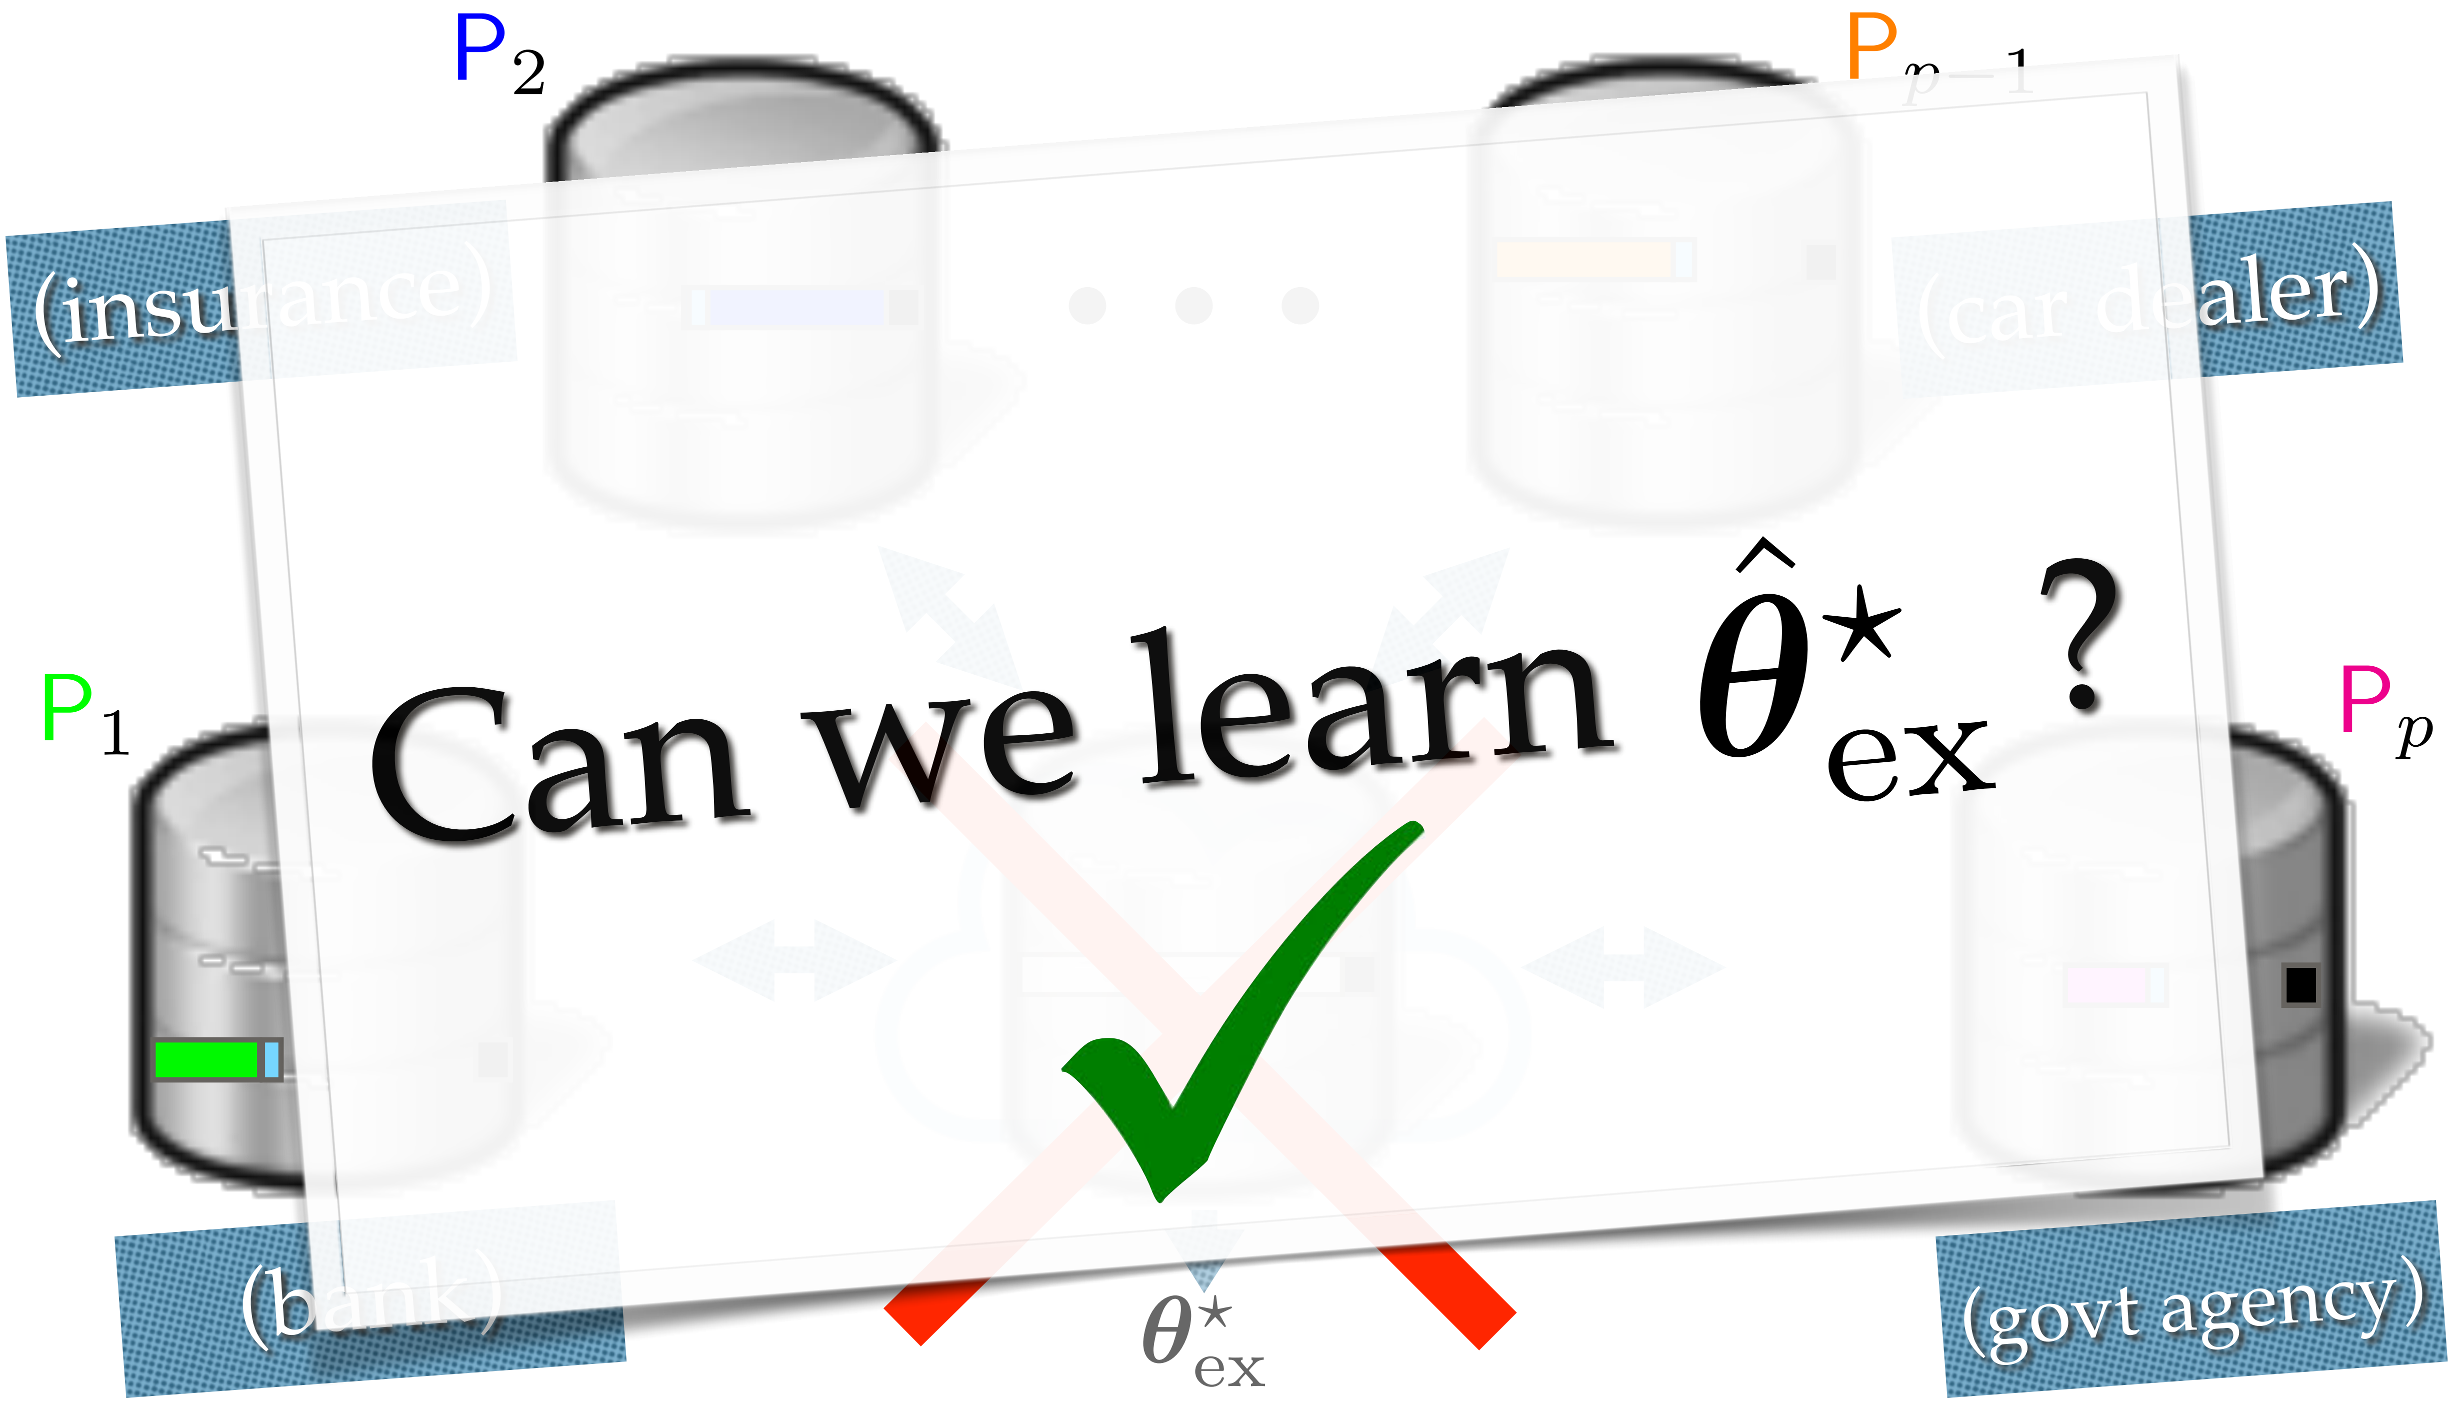

- Same setting & constraint, but arbitrary number of peers

The trick

- Entity matching needed to build complete examples…

The trick

- Entity matching needed to build complete examples… *but* complete examples not needed to learn !

The trick

- Entity matching needed to build complete examples… *but* complete examples not needed to learn !

Bypass the construction of examples, and thereby the need to solve entity matching !

Main Theorem

- Entity matching needed to build complete examples… ***but*** complete examples not needed to learn !
- For ***any*** $\mathcal{S}$ and ***any*** $\boldsymbol{\theta}$,

$$\ell_{\mathrm{sql}}(\mathcal{S}, \boldsymbol{\theta}; \Gamma) = 1 + (4/m) \cdot \ell_{\mathrm{M}}(\mathcal{R}_{\mathcal{S},\Sigma_m}, \boldsymbol{\theta}; \Gamma)$$

Main Theorem

- Entity matching needed to build complete examples… ***but*** complete examples not needed to learn !
- For ***any*** $\mathcal{S}$ and ***any*** $\boldsymbol{\theta}$,

$$\ell_{\mathrm{sql}}(\mathcal{S}, \boldsymbol{\theta}; \Gamma) = 1 + (4/m) \cdot \ell_{\mathrm{M}}(\mathcal{R}_{\mathcal{S}, \Sigma_m}, \boldsymbol{\theta}; \Gamma)$$

Ridge regularized square loss

Loss described using different data: *Rademacher observations.*

(Nock & al., ICML'15)

Main Theorem

- Entity matching needed to build complete examples… ***but*** complete examples not needed to learn !
- For ***any*** $\mathcal{S}$ and ***any*** $\boldsymbol{\theta}$,

$$\Sigma_m = \{-1, 1\}^m$$

$$\ell_{\mathrm{sql}}(\mathcal{S}, \boldsymbol{\theta}; \Gamma) = 1 + (4/m) \cdot \ell_{\mathrm{M}}(\mathcal{R}_{\mathcal{S},\Sigma_m}, \boldsymbol{\theta}; \Gamma)$$

Ridge regularized square loss

Loss described using different data: *Rademacher observations.*

(Nock & al., ICML'15)

All Theorems (almost on 1 slide !)

- Entity matching needed to build complete examples… ***but*** complete examples not needed to learn !
- For ***any*** $\mathcal{S}$ and ***any*** $\boldsymbol{\theta}$,

$$\Sigma_m = \{-1, 1\}^m$$

$$\ell_{\mathrm{sql}}(\mathcal{S}, \boldsymbol{\theta}; \Gamma) = 1 + (4/m) \cdot \ell_{\mathrm{M}}(\mathcal{R}_{\mathcal{S}, \Sigma_m}, \boldsymbol{\theta}; \Gamma)$$

- ***Rado*** set $\mathcal{R}_{\mathcal{S}, \Sigma'} = \{\boldsymbol{\pi_\sigma} \doteq \sum_{y_i = \sigma_i} y_i \cdot \boldsymbol{x_i} : \boldsymbol{\sigma} \in \Sigma'\}$, with $\Sigma' \subseteq \Sigma_m$

All Theorems (almost on 1 slide !)

- Entity matching needed to build complete examples… ***but*** complete examples not needed to learn !
- For ***any*** $\mathcal{S}$ and ***any*** $\boldsymbol{\theta}$, … $\Sigma_m = \{-1, 1\}^m$

$$\ell_{\mathrm{sql}}(\mathcal{S}, \boldsymbol{\theta}; \Gamma) = 1 + (4/m) \cdot \ell_{\mathrm{?}}(\mathcal{R}_{\mathcal{S},\Sigma_m}, \boldsymbol{\theta}; \Gamma)$$

- ***Rado*** s[…] $\{\boldsymbol{\pi}_{\boldsymbol{\sigma}} = \sum_{y_i=\sigma_i} y_i \cdot \boldsymbol{x}_i : \boldsymbol{\sigma} \in \Sigma'\}$, with $\Sigma' \subseteq \Sigma_m$

Reduction trick works for other losses, even *regularised*

All Theorems (almost on 1 slide !)

- Entity matching needed to build complete examples… ***but*** complete examples not needed to learn !
- For ***any*** $\mathcal{S}$ and ***any*** $\boldsymbol{\theta}$,

$$\Sigma_m = \{-1,1\}^m$$

$$\ell_{\mathrm{sql}}(\mathcal{S}, \boldsymbol{\theta}; \Gamma) = 1 + (4/m) \cdot \ell_{\mathrm{M}}(\mathcal{R}_{\mathcal{S},\Sigma_m}, \boldsymbol{\theta}; \Gamma)$$

- ***Rado*** set $\mathcal{R}_{\mathcal{S},\Sigma'} = \{\boldsymbol{\pi}_{\boldsymbol{\sigma}} \doteq \sum_{y_i=\sigma_i} y_i \cdot \boldsymbol{x_i} : \boldsymbol{\sigma} \in \Sigma'\}$, with $\Sigma' \subseteq \Sigma_m$
- A significant subset $\mathcal{R}_{\mathcal{S},\Sigma^\star} \subset \mathcal{R}_{\mathcal{S},\Sigma_m}$ with large size (in m) can be built without knowing entity matching
- classifier $\boldsymbol{\theta}^\star_{\mathrm{rad}} \doteq \arg\min_{\boldsymbol{\theta}} \ell_{\mathrm{M}}(\mathcal{R}_{\mathcal{S},\Sigma'}, \boldsymbol{\theta}; \Gamma)$ is *faster* to build than $\boldsymbol{\theta}^\star_{\mathrm{ex}}$
- …and we ***also*** have $\boldsymbol{\theta}^\star_{\mathrm{rad}} \rightarrow \boldsymbol{\theta}^\star_{\mathrm{ex}}$

All algorithms (on 1 slide !)

- Step 1: build a particular subset of $\mathcal{R} \subset \mathcal{R}_{\mathcal{S},\Sigma^\star}$ with $|\mathcal{R}| \leq m$

- Step 2: build $\boldsymbol{\theta}^\star_{\mathrm{rad}}$: it can be shown that

$$\boldsymbol{\theta}^\star_{\mathrm{rad}} = \left(\mathsf{R}\mathsf{R}^\top + \gamma \cdot \Gamma\right)^{-1} \mathsf{R}\mathbf{1}$$

where R stacks $\mathcal{R}$ in columns and $\gamma \in \mathbb{R}_{+*}$

All algorithms (on 1 slide !)

- Step 1: build a particular subset of $\mathcal{R} \subset \mathcal{R}_{\mathcal{S},\Sigma^\star}$ with $|\mathcal{R}| \leq m$

O(nb_features . m)

- Step 2: build $\boldsymbol{\theta}^\star_{\mathrm{rad}}$: it can be shown that

$$\boldsymbol{\theta}^\star_{\mathrm{rad}} = \left(\mathsf{R}\mathsf{R}^\top + \gamma \cdot \Gamma\right)^{-1} \mathsf{R}\mathbf{1}$$

where R stacks $\mathcal{R}$ in columns and $\gamma \in \mathbb{R}_{+*}$

O(nb_features^2 . m)

Generalisation

- Works for *any* number of peers
- Works outside the vertical partition assumption

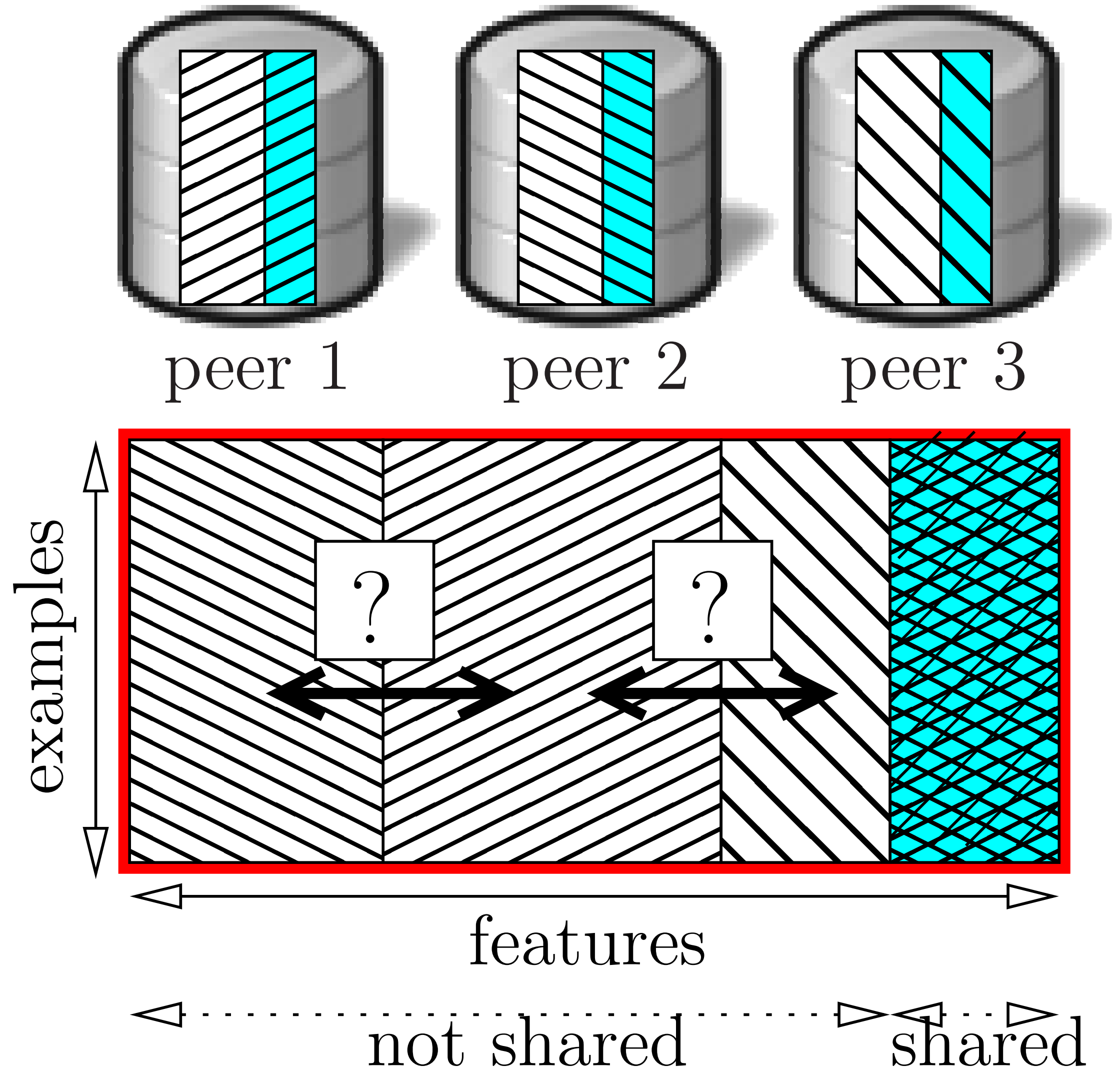

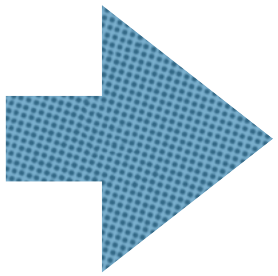

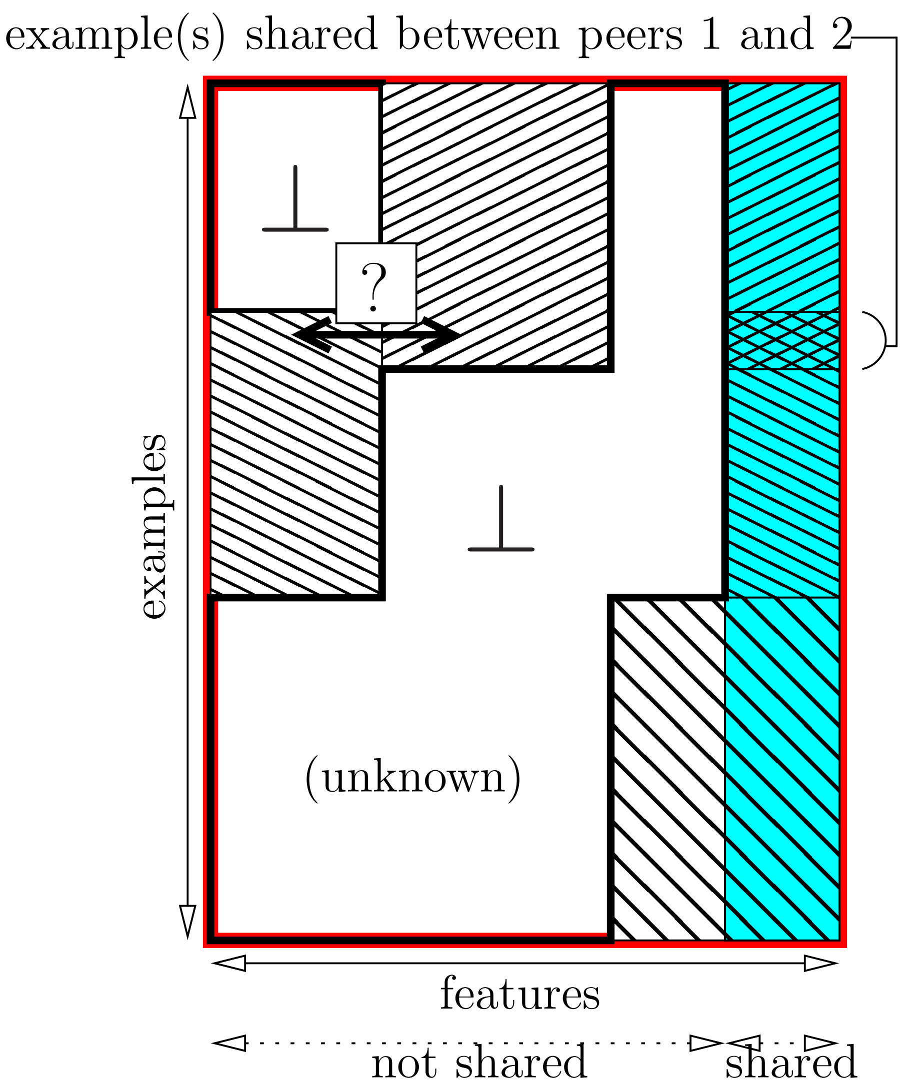

Experiments

- simulation: split datasets — vary #peers, #shared(features), #bins, #joint_examples
- Little experimental influence of #bins (in range 2-5)
- Tested no #joint_examples (peers see all different examples, harder) + small % of #joint_examples

objective: beat the ***best*** peer in hindsight

Experiments

- vary #peers, #shared(features), #joint_examples = **0**

$$\Delta \doteq \hat{p}_{\mathrm{err}}(\text{our algo}) - \min_j \hat{p}_{\mathrm{err}}(\mathsf{P}_j) \;\; (\in [-1, 1])$$

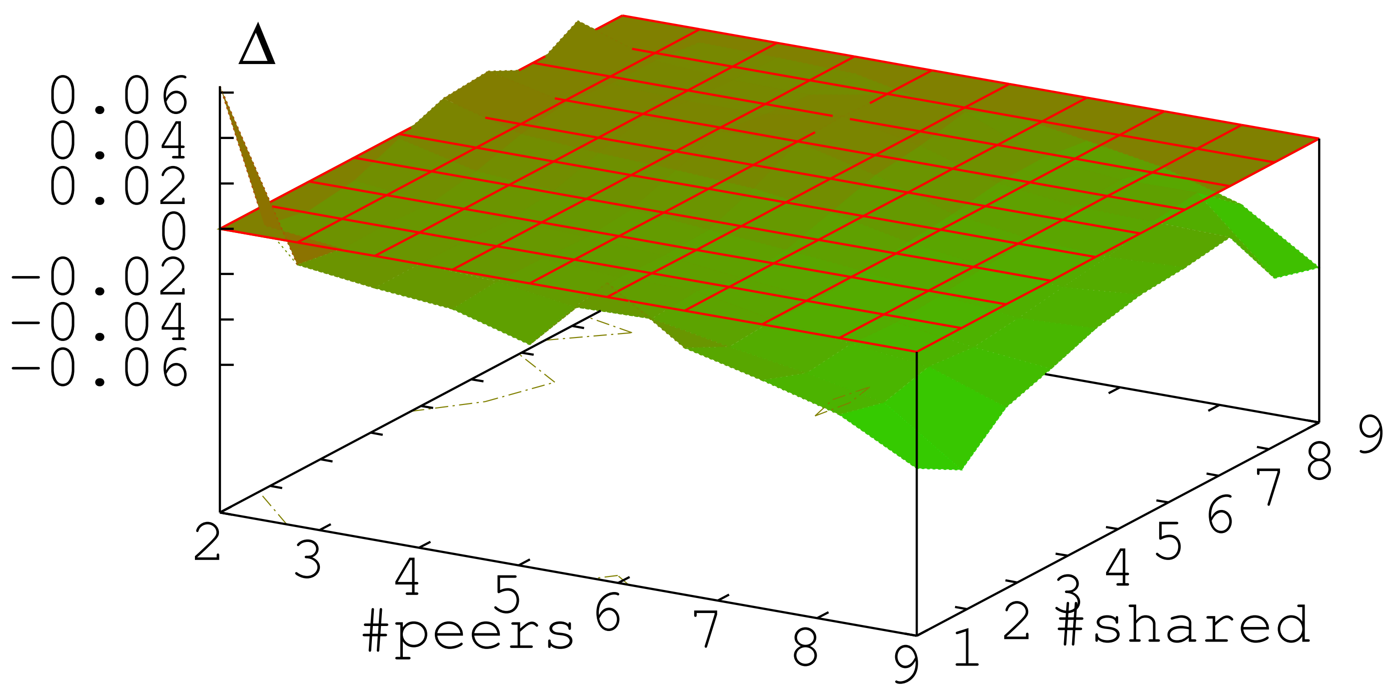

UCI Ionosphere

Experiments

- vary #peers, #shared(features); #joint_examples = **0**

$$\Delta \doteq \hat{p}_{\mathrm{err}}(\text{our algo}) - \min_j \hat{p}_{\mathrm{err}}(\mathsf{P}_j) \ \ (\in [-1,1])$$

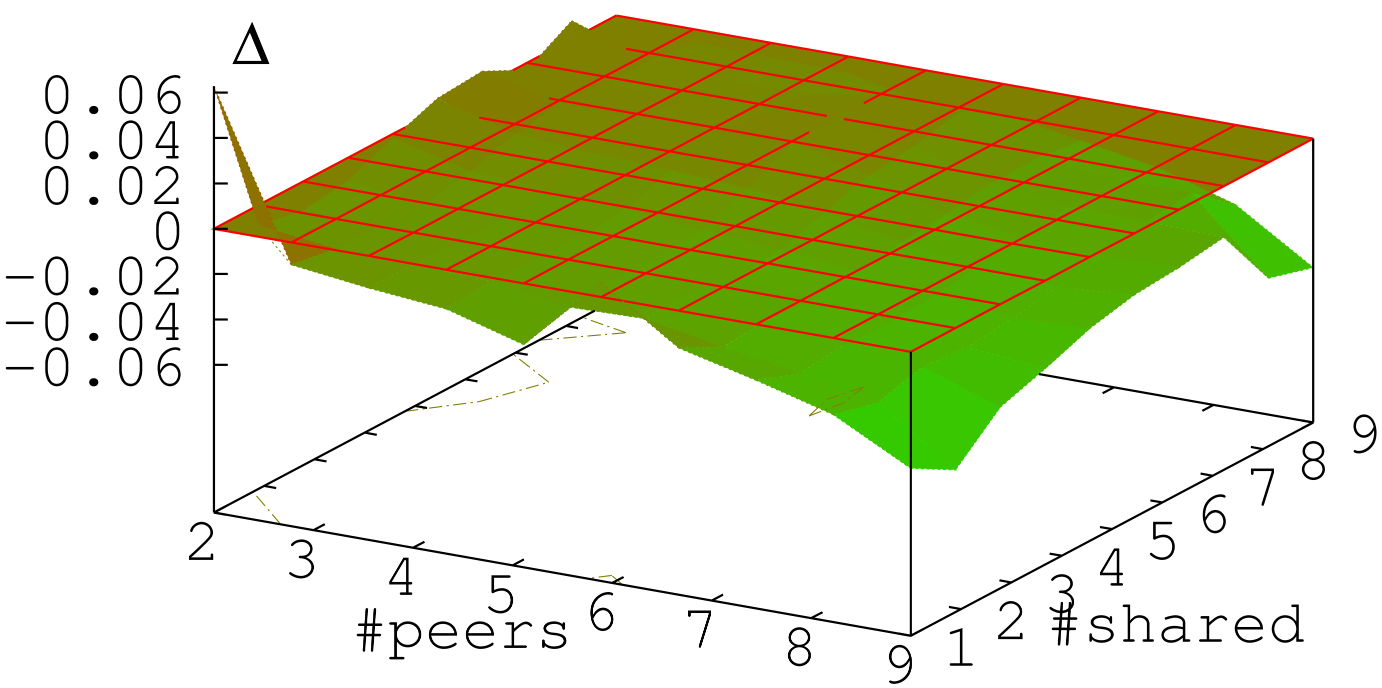

Almost systematically beats all peers

UCI Ionosphere

Experiments

- vary #peers, #shared(features); #joint_examples = **0**

$$q \doteq \text{ proportion of peers } \textit{statistically} \text{ beaten by our algo}$$

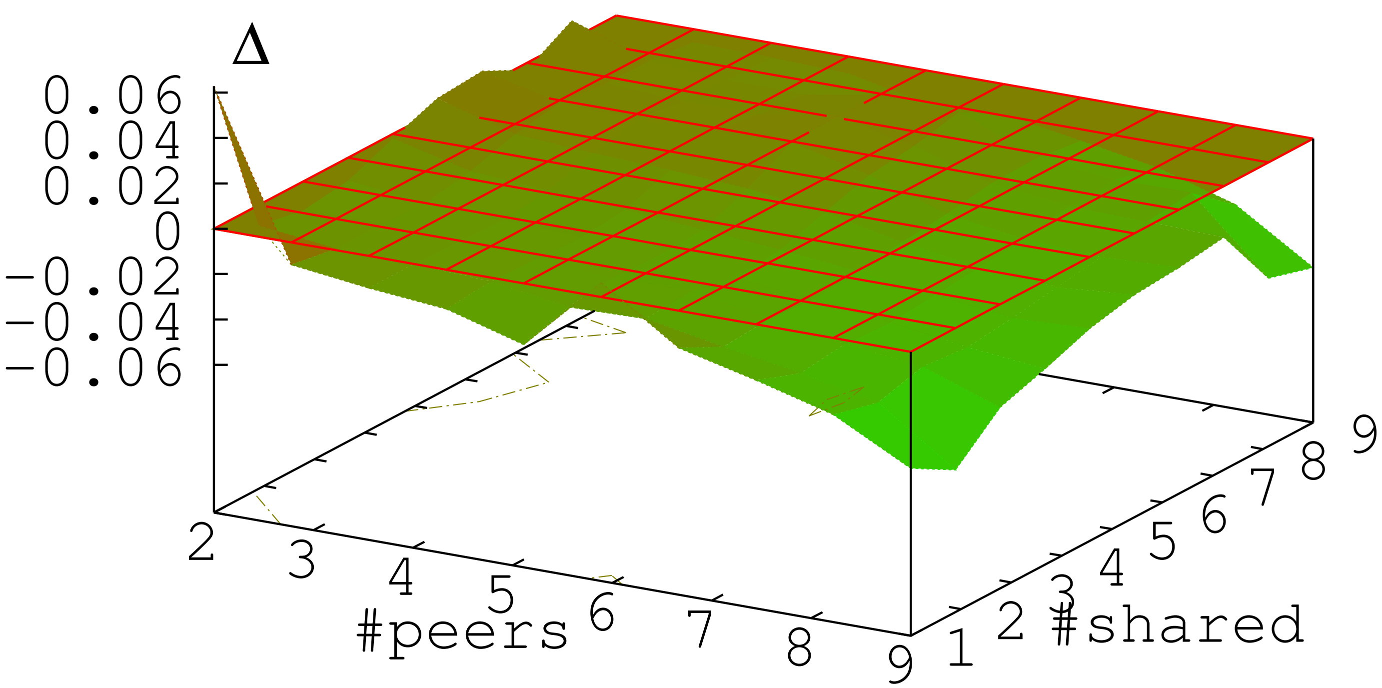

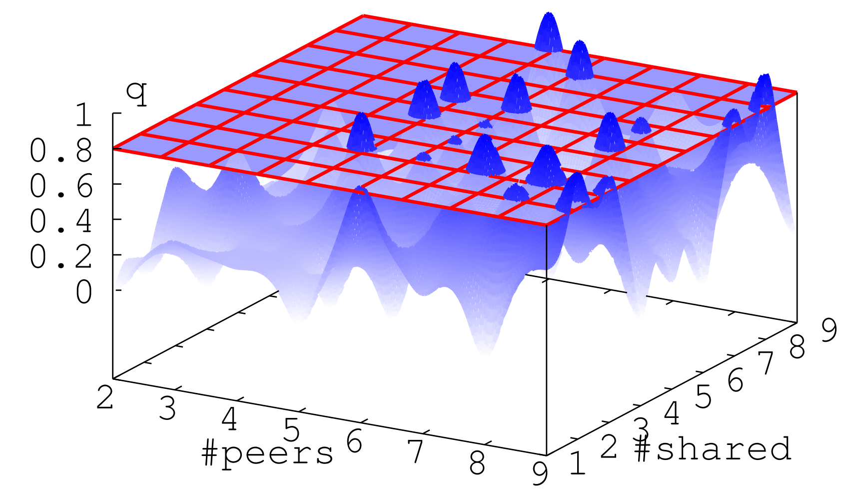

Almost systematically beats all peers… but not always significantly

UCI Ionosphere

Experiments

- See poster, paper & long ArXiv version for more experiments

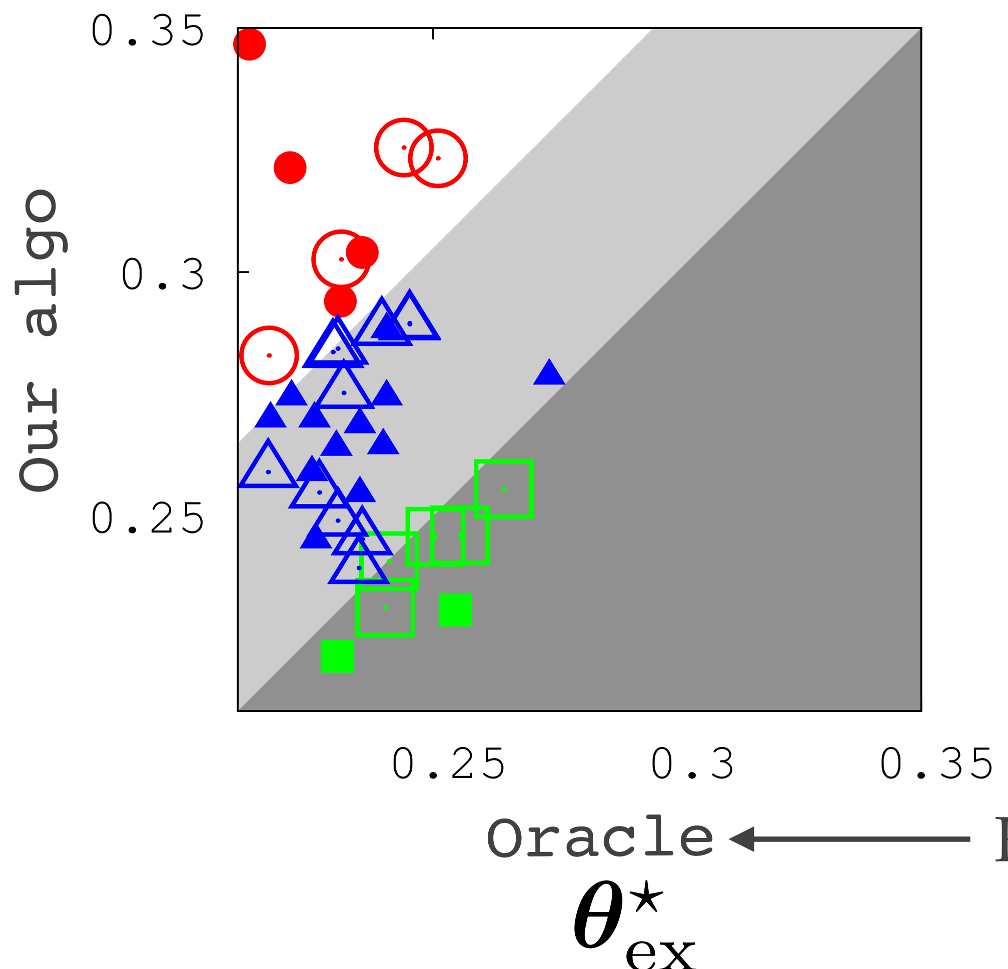

Knows the solution to ER !

Sometimes we are worse (statistically or not)

Sometimes we are better (not statistically) !

UCI Sonar

Radosandprivacy

- Protection guarantees: differential privacy (DP), computational hardness (CH), geometric hardness (GH), algebraic hardness (AH)
  - Crafting of *DP* rados from non-DP examples
  - *CH* of approximate sparse recovery of examples from rados
  - *CH* of pinpointing examples having served to craft rados
  - *GH, AH* of recovering examples from rados
- Crafting of rados from DP (noisified) examples with still ***guaranteed*** convergence rates for boosting over rados

Privacy guarantees ?

Pinpointing examples from rados

- Problem (informal): a super powerful agency $\mathcal{A}$ has a huge database of examples, $\mathcal{S}$. A intercepts a set of rados, $\mathcal{S}^r$. A fixes size m.
- Question: does there exist a subset of $\mathcal{S}$ of size m with which we can *approximately* craft the rados in $\mathcal{S}^r$?

Pinpointing examples from rados

- Problem (informal): a super powerful agency $\mathcal{A}$ has a huge database of examples, $\mathcal{S}$. A intercepts a set of rados, $\mathcal{S}^r$. A fixes size m.
- Question: do there exist a subset of $\mathcal{S}$ of size m with which we can *approximately* craft the rados in $\mathcal{S}^r$?

NP-HARD

Geometric hardness of recovering examples

- Protection guarantees: differential privacy (DP), computational hardness (CH), geometric hardness (GH), algebraic hardness (AH)
  - Crafting of *DP* rados from non-DP examples
  - *CH* of approximate sparse recovery of examples from rados
  - *CH* of pinpointing examples having served to craft rados
  - *GH, AH* of recovering examples from rados
- Crafting of rados from DP (noisified) examples with still ***guaranteed*** convergence rates for RadoBoost

Geometric hardness of recovering examples

- Suppose $\mathcal{A}$ is given *only* a set of rados. $\mathcal{A}$ knows **nothing else** about the examples $\mathcal{S}$, except that all lie in a ball of radius R.
- Then there exists a set of examples $\mathcal{S}'$ with just *one* more example, which produces the *same* set of rados but lies at Hausdorff distance

$$D_H(\mathcal{S}, \mathcal{S}') = \Omega\left(\frac{R \log d}{\sqrt{d} \log m}\right) \qquad (m \geq 2^d)$$

$$D_H(\mathcal{S}, \mathcal{S}') = \Omega\left(\frac{R}{\sqrt{d}}\right) \qquad \text{(Otherwise)}$$

Geometric hardness of recovering examples

- Suppose $\mathcal{A}$ is given *only* a set of rados. $\mathcal{A}$ knows **nothing else** about the examples $\mathcal{S}$, except that all lie in a ball of radius R.
- Then there exists a set of examples $\mathcal{S}'$ with just *one* more example which produces the *same* set of rados but is at Hausdorff [illegible]

$$D_H(\mathcal{S}, \mathcal{S}') = \Omega\left(\frac{R \log d}{\sqrt{d} \log m}\right) \quad (m \geq 2^d)$$

$$D_H(\mathcal{S}, \mathcal{S}') = \Omega\left(\frac{R}{\sqrt{d}}\right) \quad \text{(Otherwise)}$$

Stays as hard if m approximately known

Geometric hardness of recovering examples

- Suppose $\mathcal{A}$ is given *only* a set of rados. $\mathcal{A}$ knows **nothing else** about the examples $\mathcal{S}$, except that all lie in a ball of radius R.

- Then there exists a set of examples $\mathcal{S}'$ with just *one* new example, which produces the *same* set of rados but lies at Hausdorff distance

$$D_H(\mathcal{S}, \mathcal{S}') = \Omega\left(\frac{[illegible]}{\sqrt{\log m}}\right) \quad (m \geq 2^d)$$

$$D_H(\mathcal{S}, \mathcal{S}') = \Omega\left(\frac{R}{\sqrt{d}}\right) \quad \text{(Otherwise)}$$

Hardness does not rely on the computational power at hand

Conclusion: research

- All the results on Rademacher observations rely on the observation that the sufficient statistics for the class is *small* (*one* vector, for *any* symmetric proper scoring rule)
- Therefore, can learn efficiently from weakly labeled data, no-ER data (etc.) as long as it can be reliably estimated

Conclusion: design

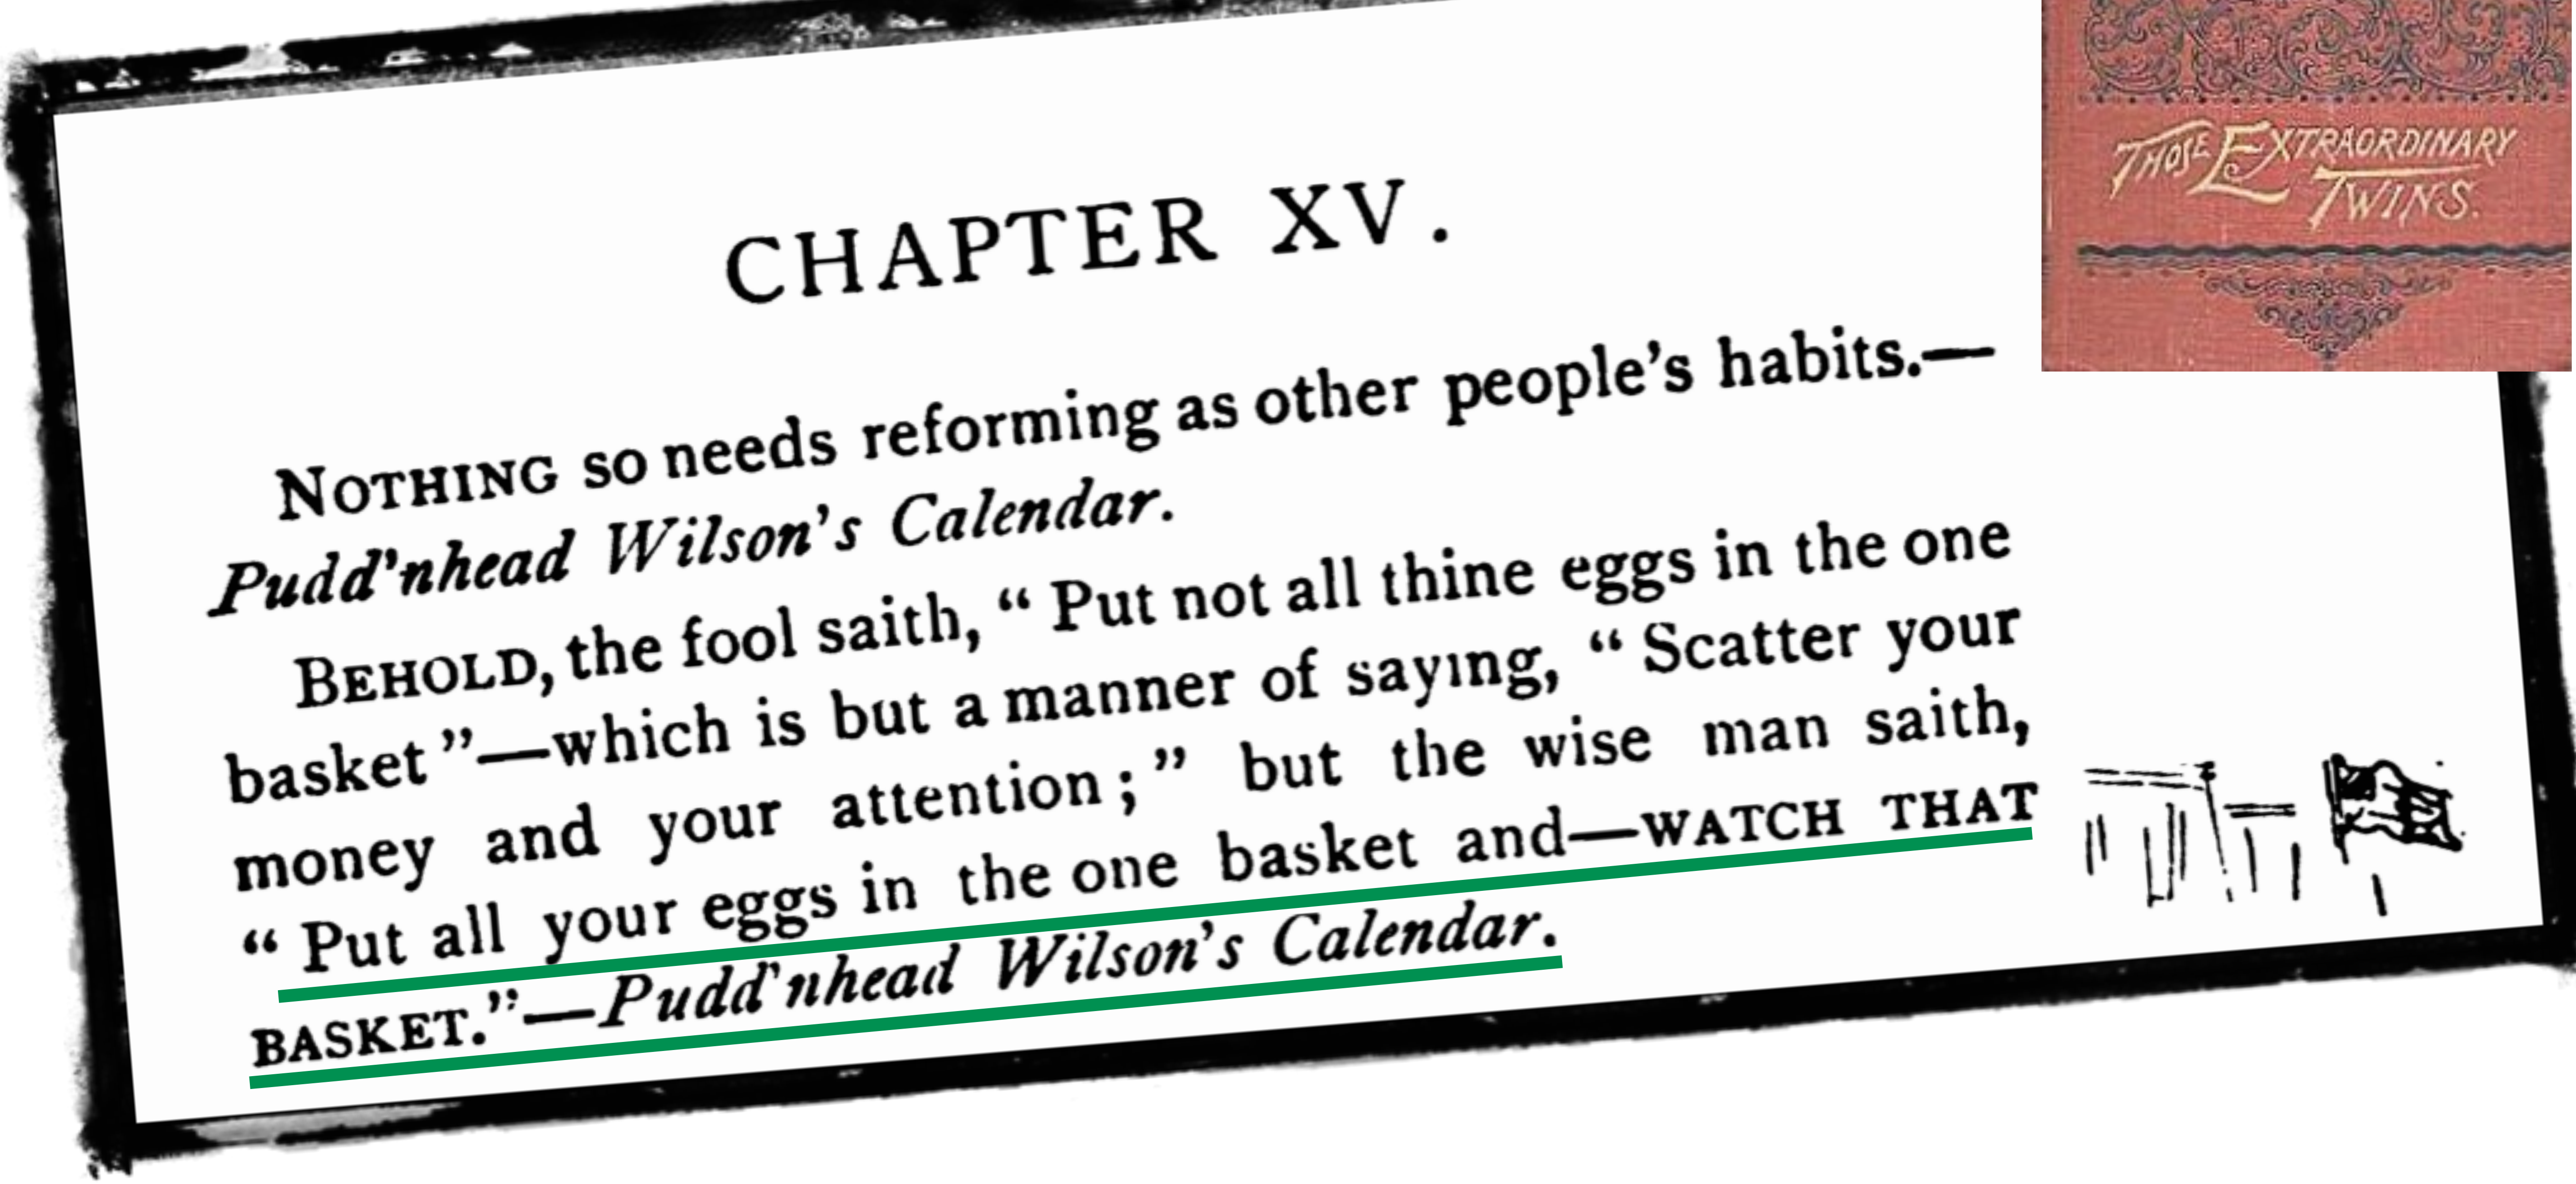
CHAPTER XV.

NOTHING so needs reforming as other people's habits.—*Pudd'nhead Wilson's Calendar.*

BEHOLD, the fool saith, "Put not all thine eggs in the one basket"—which is but a manner of saying, "Scatter your money and your attention;" but the wise man saith, "Put all your eggs in the one basket and—WATCH THAT BASKET."—*Pudd'nhead Wilson's Calendar.*

Thank you!