Not Fit for Purpose

A critical analysis of the ‘Five Safes’

Dr Chris Culnane & Prof. Ben Rubinstein

University of Melbourne

Overview

1. Setting the Scene
2. Disconnected from Legal Protections
3. Notions of Safety
4. Summary

Setting the Scene

Grown from data protection framework popular in population data research into cornerstone of data sharing and access policies and legislation

- Recent legislation
  - Data Sharing & Access Legislation (Australia)
  - Digital Economy Act (UK)
- Initially only ‘Four Safes’
  - ‘Safe Data’ added to handle ‘de-identified’ data contexts
- Despite its growing popularity there is little rigorous analysis of its capabilities or suitability

Critical Analysis

Our critical analysis[5] concludes that the ‘Five Safes’ is

- Disconnected from existing legal protections
- Appropriates notions of safety without requiring the necessary strong technical measures
- Views disclosure risk as static in time

Legal Protections

Legal protections increasingly focus on the notion of personal data being about an identifiable or reasonably identifiable individual

- Use of personal data requires consent
    - Some secondary uses are permitted as well as certain exemptions on grounds of public safety
- Only data that is no longer considered personal information can be shared or processed without consent (i.e. ‘Safe Data’)
    - Typically this is achieved through de-identification - although this is a troubled concept
- These clear protections are increasingly being undermined by the sharing of personal data by government departments, using the ‘Five Safes’ as an enabler

Notions of Safety

The ‘Five Safes’ framework evaluates safety not risk and is inconsistent with best practice in the wider information security field

- There is no such thing as a 100% secure system
- We identify vulnerabilities and quantify the risk of exploitation
- Only as strong as the weakest link
- Cryptographic key sizes being a good example
  - Key sizes are evaluated on the basis of when they could feasibly be broken by different types of attacker
  - We evaluate who they are at risk from, not who they are safe from

Notions of Safety

Safety is not an absolute position. It is a position on an unspecified and undefined risk continuum.

- Even if such a safety scale were to exist it is unrealistic to expect it to be consistently applied or comparable across organisations
- The language of the 'Five Safes' creates an impression of safety even where one does not exist
- Would have far greater value as the 'Five Risks'
  - Resolves the language issue
  - Maintains focus on the potential adversary

Safe People

- Establishing ‘Safe People’ has previously been the preserve of defence, intelligence, national security and law enforcement agencies
  - Each of whom dedicate significant resources to “vetting” their staff
  - Yet examples of this vetting failing are numerous
- A safety evaluation of people is inherently dynamic as people’s situation changes
- An ‘unsafe’ person may not be a sophisticated adversary or someone intent on harm
  - Curiosity, or just carelessness, can render someone ‘unsafe’

Safe Data

The concept of ‘Safe Data’ is a misnomer

- If data was safe, what is the relevance of the remaining four safes?
- Privacy Law (EU/UK/Aus) sets the threshold for safe data at de-identification
    - The definition of which is problematic and prone to failure
- Mechanisms for the safe release of data do exist, e.g. Differential Privacy, but the ‘Five Safes’ provides no function to favour or define their application
- An organisation can declare data ‘safe’ with no technical proof or evaluation of re-identification risk

Safe Projects

- Assumes all current and future intentions of the project initiator can be determined at the start
- Dependent on ‘Safe People’ on account of the self-reporting of the projects goals, data uses and plans
    - Counter to this assumes that there are ‘Safe People’ willing to propose ‘unsafe’ projects, which in turn would surely render the person ‘unsafe’
- Requires ongoing inspection & auditing, as well as serious consequences for deviation from a ‘Safe Project’

Safe Environments

The notion of ‘Safer’ environments has merits, but the ‘Five Safes’ provides no guidelines as to how to implement such an environment

- Physical Secure Research Environments, with access and technological controls, as well as statistical disclosure measures have real benefit
    - However, they don’t scale, are expensive, time consuming and are difficult to administer
    - Vulnerabilities in protection methods can still exist within the ‘Safe Environment’[3]
- Remote secure environments are sometimes viewed as equivalent when they actually offer materially less protection
    - Dependent on ‘Safe People’ since they are only secure if the users abide by the restrictions and do not actively seek to circumvent the protections

Safe Outputs

- Dependence on de-identification
- If performed by the project and not the owning authority is dependent on ‘Safe People’ and potentially ‘Safe Environments’ as well

Not as safe as thought

- Australian Department of Heatlh MBS/PBS Release - 2016[2]
    - 10% sample of the population's medical and pharmaceutical billing records (medical: 1984-2014, pharmaceutical: 2003-2014)
    - We re-identified providers through vulnerabilities in the "encrypted" supplier ID
    - Subsequently demonstrated patient re-identification risk as well
    - Prior to its public release it had been shared with select groups in the same form
- Public Transport Victoria Myki Release - 2018[4]
    - 3 years of touch-on and touch-off data for 15 million Myki travel cards
    - Over 2 billion touch-on/off events
    - We re-identified ourselves, co-travellers, and a State MP

In all these cases the releasing party thought the data was 'Safe'. The capability to accurately determine re-identification risk remains low.[1]

Summary of Issues

- Language shifts focus from risks and adversaries to misappropriated notions of safety
    - No such thing as ‘Safe Data’ or ‘Safe People’ only some degree of ‘safety’
- The ‘Safes’ are not independent, they are in fact dependent
    - ‘Safe People’ having the most dependants
    - There is no “Graphics Equaliser” concept where strength in one area can cover for weakness in another
- The risks captured by the ‘Safes’ are dynamic, not static
    - Requires ongoing evaluation
    - Must be accompanied by severe consequences for failure to comply
- Not necessarily compatible with Privacy Law
    - Privacy Law seeks to define ‘Safe Data’ only, does not permit sharing outside of that

Questions?

cculnane@unimelb.edu.au

benjamin.rubinstein@unimelb.edu.au

References i

C. Culnane and K. Leins.
Misconceptions in privacy protection and regulation.
Law in Context. A Socio-legal Journal, 36(2):49–60, Apr. 2020.

C. Culnane, B. I. P. Rubinstein, and V. Teague.
Health data in an open world.
CoRR, abs/1712.05627, 2017.

C. Culnane, B. I. P. Rubinstein, and V. Teague.
Vulnerabilities in the use of similarity tables in combination with pseudonymisation to preserve data privacy in the UK office for national statistics' privacy-preserving record linkage.
CoRR, abs/1712.00871, 2017.

References ii

C. Culnane, B. I. P. Rubinstein, and V. Teague.
Stop the open data bus, we want to get off.
CoRR, abs/1908.05004, 2019.

C. Culnane, B. I. P. Rubinstein, and D. Watts.
Not fit for purpose: A critical analysis of the ’five safes’.
CoRR, abs/2011.02142, 2020.