

UNIVERSITY OF SOUTHAMPTON

FACULTY OF BUSINESS AND LAW

**OPENNESS FOR PRIVACY: APPLYING OPEN APPROACHES TO
PERSONAL DATA CHALLENGES**

by

Reuben Binns

Thesis for the degree of Doctor of Web Science

October 2015

Table of Contents

List of Tables.....	7
List of Figures.....	7
Declaration of Authorship.....	8
Acknowledgements.....	11
Abbreviations.....	13
Foreword.....	15
How to read this PhD.....	16
Part 1: Background.....	18
1.1 Introduction.....	18
1.1.1 Personal data: the view from 10,000 feet.....	18
1.1.2 Privacy, data protection, and social concerns arising from personal data.....	20
1.1.3 Approaches to the policy problem.....	24
1.1.4 Personal data empowerment.....	28
1.2 Privacy and Openness: contradictory or complementary?.....	30
1.2.1 Zero-sum?.....	30
1.2.2 A Middle Ground.....	32
1.2.3 Compatibility, mutual reinforcement.....	33
1.2.4 The Openness Principle's Failings, and Unmet Potential.....	35
1.3 Openness: an overview.....	36
1.3.1 The Origins of Open.....	37
1.3.2 Open Source.....	38
1.3.3 Critiques of Openness.....	42
1.3.4 Towards a definition of openness.....	45
1.4 Openness for Privacy.....	46
1.4.1 Open data for privacy.....	47
1.4.2 Open processing: transparency and modification.....	49
1.4.3 Regulating Privacy with the Open Corporation.....	51
1.4.4 Extending OfP: standards, platforms, collaboration and tools.....	52
1.4.4.1 Open standards and personal data.....	52
1.4.4.2 Open government platforms for privacy.....	53
1.4.4.3 Open collaboration tools.....	54
1.4.4.4 Open source software for privacy management.....	54
1.4.5 Summary of OfP applications.....	55
1.5 Summary.....	57
Part 2: Open Data for Privacy.....	59
2.1 Introduction.....	63
2.2 Background.....	63
2.2.1 Existing transparency mechanisms.....	64
2.2.1.1 Privacy Notices.....	64
2.2.1.2 Public Registers.....	64
2.2.2 Continued emphasis on transparency.....	65
2.2.3 Standardised Formats.....	65
2.2.2 Prior Art.....	66
2.2.2.1 Platform for Privacy Preferences.....	66
2.2.2.2 Collaboration with regulators.....	66
2.2.2.3 A standard in decline.....	67
2.2.2.4 Development of Public Registers.....	68
2.2.2.5 Similarities between P3P and public registers.....	68
2.2.3 Quantifying Privacy Practices.....	69
2.2.3.1 Trading of personal data.....	69
2.2.3.2 Financial Services.....	70
2.2.3.3 Health services.....	70
2.2.3.4 Comprehensive samples for comparison.....	70
2.3 Data Source and Methodology.....	71

2.3.1 Notification Requirements.....	71
2.3.2 Data structure, extraction and selection.....	71
2.3.3 Analysis.....	72
2.4 Results.....	73
2.4.1 Why is data being processed?.....	74
2.4.2 Who is the data about?.....	75
2.4.3 What kind of personal data is used?.....	76
2.4.4 Who has access to the data?.....	77
2.5. Discussion.....	78
2.5.1 Growth in data controllers.....	78
2.5.2 Power law distribution.....	79
2.5.3 Informing public concerns.....	79
2.5.4 Differentiation between practices.....	80
2.5.5 Limitations.....	80
2.6. Recommendations.....	81
2.6.1 Standardisation, Categories and Granularity.....	81
2.6.2 Incentives, monitoring and enforcement.....	82
2.7. Conclusions.....	83
2.8 Epilogue.....	86
Part 3: Open Processing.....	89
3. Abstract.....	91
3.1. Introduction.....	92
3.1.1. Background.....	92
3.1.2. Literature Review.....	95
3.1.3. Aims and Objectives.....	99
3.2. Study Design and Method.....	100
3.3. Analysis and results.....	104
3.4. Discussion and conclusions.....	105
3.4.1 Further research.....	106
3.4.2 Implications for industry and policy.....	107
Part 4: Personal data empowerment.....	110
4.1 Open profiling and the logic of big data.....	112
4.2 The ethics of personal data markets.....	115
4.3 Personal Data Empowerment and the Ideal Observer.....	118
Part 5: Meta-regulating privacy and the open corporation.....	121
5.1. Introduction.....	123
5.2. Privacy Impact Assessments: Background.....	124
5.2.1 Origin of PIAs.....	125
5.2.2 Adoption and implementation of PIAs.....	127
5.3. Regulatory theory of PIAs.....	129
5.3.1 PIAs as self-regulation.....	130
5.3.2 Ensuring implementation through mandatory PIAs.....	131
5.3.3 Mandatory PIAs as legal regulation: would they suffer the drawbacks of 'command and control' regimes?.....	131
5.3.4 PIAs as 'co-regulation'.....	133
5.4. Analysis of mandatory PIAs in the GDPR.....	135
5.4.1 Commission reports prior to the 2012 proposal.....	135
5.4.2 The proposed GDPR.....	136
5.4.2.1 When are PIAs required?.....	136
5.4.2.2 Scope and content of a PIA.....	137
5.4.2.3 Stakeholder consultation.....	138
5.4.2.4 Fines and ongoing compliance.....	138
5.4.3 Summary of the GDPR rationale and provisions.....	138
5.5. Meta-regulation as a model of mandatory PIAs.....	140
5.5.1 Introducing meta-regulation.....	140
5.5.2 PIAs as meta-regulation.....	141
5.6 Evaluating meta-regulation.....	143
5.7. The prospects for PIAs as meta-regulation.....	146
5.7.1 Leveraging regulatees.....	146
5.7.2 Independent scrutiny.....	146

5.7.3 Stability, trust and external support.....	147
5.7.4 Regulatory tiers.....	148
5.7.5 Shaping organisations' compliance.....	148
5.8 Conclusion.....	149
5.9 Epilogue.....	151
Part 6: Conclusion.....	153
6.1 Summary of contributions.....	153
6.1.1 Open Data for Data Protection.....	153
6.1.2 Open Processing.....	153
6.1.3 Meta-regulating privacy and the open corporation.....	154
6.1.4 Summary table.....	155
6.2 Evaluating the Openness-for-Privacy approach.....	156
6.2.2 The promise of OfP.....	156
6.2.3 Limitations and challenges of OfP.....	158
6.2.4 Refining OfP.....	160
6.3 Openness and privacy: mutually supportive principles.....	161
Appendices.....	166
A. Visualisation of international data transfers.....	166
B. Study design flowchart.....	167
C. Study design considerations.....	168
D. Study interface.....	171
E. What's in a name? Privacy Impact Assessments and Data Protection Impact Assessments.....	172
F. PIA Triage Process.....	174
References.....	175

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF BUSINESS AND LAW

Web Science

Doctor of Philosophy

**OPENNESS FOR PRIVACY: APPLYING OPEN APPROACHES TO
PERSONAL DATA CHALLENGES**

by Reuben Binns

This thesis comprises three papers undertaken as part of a PhD by publication or 'Three-Paper PhD', in addition to an introduction and conclusion. The introduction outlines the concept of Openness for Privacy, which describes a class of technological, social and policy approaches for addressing the challenges of personal data. Various manifestations of this concept are investigated in the three papers.

The first paper explores the idea of 'open data for privacy', in particular the potential of machine-readable privacy notices to provide transparency and insight into organisations' uses of personal data. It provides an empirical overview of UK organisations' personal data practices.

The second paper examines services which give individuals transparency and control over their digital profiles, assessing the potential benefits to industry, and the empowering potential for individuals. The first part is a user study, which tests how consumer responses to personalised targeting are affected by the degree of transparency and control they have over their profiles, with implications for digital marketing and advertising. The second part draws from qualitative data, and theoretical perspectives, to develop an account of the empowering potential of these services.

The third paper concerns Privacy Impact Assessments (PIAs), a regulatory tool included in the European Union's proposed general data protection regulation reform. It assesses the potential of PIAs through concepts from regulatory theory, namely, meta-regulation and the open corporation, and outlines implications for regulators, civil society and industry.

List of Tables

1. Main applications of OfP page	58
2. Comparison of DPA, register and P3P fields	75
3. Average Purposes, Classes, Subjects and Recipients,	77
4. Average Recommendation Ratings by source / interface	113
5. Significance tests, SAI vs Behavioural, pure vs misrepresented	114
6. Features supporting classification of PIAs as meta-regulation	153
7. Summary of contributions	168

List of Figures

1. For what purposes is personal data processed?	78
2. Who is the data about?	79
3. What kind of personal data is collected?	80
4. Who has access to the data	82
5. DPIA triage process	148

Declaration of Authorship

I, Reuben Binns, declare that this thesis, titled **Openness for Privacy: Applying Open Approaches to Personal Data Challenges**, and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University;
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- Where I have consulted the published work of others, this is always clearly attributed;
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- Parts of this work have been published as:

Binns, R. (2014a). Personal Data Empowerment and the Ideal Observer. In O'Hara, K., Nguyen, C., & Haynes, P. (eds), *Digital Enlightenment Yearbook 2014 : Social Networks and Social Machines, Surveillance and Empowerment*.

Binns, R. (2014b). Standardised Privacy Policies: A Post-mortem and Promising Developments. In *W3C Privacy Workshop: Privacy and User-Centric Controls*. Berlin.

Binns, R. (2015). Caveat Venditor : Should We Sell Our Own Data ? In *WebSci15: Workshop on the Economics of Surveillance*.

Binns, R., & Lizar, M. (2012). Opening up the online notice infrastructure. In *W3C Privacy Workshop: Do Not Track and Beyond*.

Binns, R., & Matthews, D. (2014). Community Structure for Efficient Information Flow in “ ToS ; DR ”, a Social Machine for Parsing Legalese. In *Proceedings of the companion publication of the 23rd international conference on World Wide Web* (pp. 881–884). Seoul, South Korea.

Binns, R., Millard, D., & Harris, L. (2014). Data Havens, or Privacy Sans Frontières? A Study of International Personal Data Transfers. *Proceedings of the 2014 ACM Conference on Web Science*, 14–15.

Binns, R., Millard, D., & Harris, L. (2015). The Who, What and Why: An Analysis of Personal Data Transparency Notices in the UK. *Journal of Open Access to Law*, 3(1).

Presentations:

- “Remembering Why We Forgot: An Analysis of Wikipedia's

Biography of Living Persons Policy” - Total Archive, CRASSH
2015, Cambridge

- “Is Selling Your Data the Answer to Our Privacy Problems?”
Theorising the Web 2015, New York
- “Privacy and Consumer Markets” - 31st Chaos Communication
Congress (2014), Hamburg

Signed:

Date:

Acknowledgements

First and foremost, I would like to thank my PhD supervisors, Lisa, David, Micheál and Roksana for their guidance, feedback and encouragement.

I would also like to thank the directors, academics, administrators and fellow PhD students at the Web Science DTC and the Web and Internet Science Group for making this PhD possible. I have benefitted enormously from the stimulating and supportive environment which they have worked so hard to foster.

Thanks to the volunteers who patiently tested the initially bug-ridden platform I created for the user study. Thanks to the Information Commissioner's Office, for supplying regular copies, updates and supplementary information regarding the register of data controllers – I will fondly remember my monthly package in the post from Wilmslow.

Thanks to the many interesting people I've had the opportunity to talk to and learn from over the course of this PhD. Thanks in particular to my colleagues at Ctrl-Shift, for giving me the opportunity to gain practical insights into the myriad impacts of the changing personal data landscape.

Last but not least, I'd like to thank my partner Holly, my family, and my friends, for their love, support and encouragement.

Abbreviations

API – Application Programming Interface

CC / BY / NC / ND / SA – Creative Commons / **By** / Non-Commercial / No Derivatives / Share-Alike

CDS – Critical Data Studies

CSR – Corporate Social Responsibility

DPA – Data Protection Act

DPO – Data Protection Officer

DPIA – Data Protection Impact Assessment

EU – European Union

FTC – Federal Trade Commission (US)

FOSS – Free and Open Source Software

GDPR – General Data Protection Regulation

GPG – Gnu Privacy Guard

GPL – General Public License

GNU – Gnu's Not Unix

HTML – Hypertext Markup Language

HTTP – Hypertext Transfer Protocol

IP – Internet Protocol

NGO – Non-Governmental Organisation

OfP – Openness for Privacy

OTR – Off the Record

P3P – Platform for Privacy Preferences

PbD – Privacy by Design

PGP – Pretty Good Privacy

PIA – Privacy Impact Assessment

PIMS – Personal Information Management Services

SAX – Sequential Access Parser

SAI – Self-Authored Interest

TOR – The Onion Router

UK – United Kingdom

US – United States

URL – Uniform Resource Locator

W3C – World Wide Web Consortium

WWW – World Wide Web

XML – eXtensible Markup Language

Foreword

Surfing the web for the first time as a child in the late 1990's, I quickly became awed by its potential. It seemed inevitable that the web would entail the liberation of knowledge, as well as new forms of personal agency, social participation and collaboration. My youthful exuberance continued through the 2000's, as I eagerly read various best-selling popular social science books extolling the virtues of the digital revolution.

But around the turn of the decade, doubts began to grow. By the time I began my PhD in 2011, the web looked quite different. Governments, corporations and venerable institutions (those 'weary giants of flesh and steel'¹) had found their seats at the table, alongside a host of new powerful entities, who provided many of the web's essential services. It had become a place where behaviour is monitored and shaped in opaque ways, where the terms of interaction and information flow are determined by private platforms over which we have no say. Having sipped the web evangelist's Kool Aid, I was feeling the urge to spit it back out and put on a tinfoil hat.

This PhD could therefore be read as an attempt to reconcile this conflict. It explores whether we might use the web's more progressive aspects to address some of its problems; specifically, how various principles of openness might address challenges raised by new uses of personal data. The discussion encompasses more than just what happens on the web, but the web remains a locus throughout.

Personal data has become an essential resource in the modern, data-driven world.² It underlies digital transactions, shapes organisational processes and drives personalised services. But these developments raise some significant concerns and challenges. How can we ensure that this data is used in ways that are compatible with privacy and data protection, that respect ethical and political principles such as autonomy and equality, and that empower rather than undermine the people it relates to?

Almost every set of principles proposed to address these questions appeals to the ideal of *openness*. This PhD explores and expands this notion, looking at the various ways that openness might serve the ends of privacy, data protection and personal data empowerment.

1 From John Perry Barlow's *Declaration of the Independence of Cyberspace* (Barlow, 1996).

2 In what follows, I will depart from Latin grammar in using 'data' as a singular. I hope the reader does not find this too grating.

How to read this PhD

First, some words on the format and style of this document. Unlike a traditional PhD thesis with multiple chapters following one singular narrative, it is comprised of three stand-alone research papers prepended and appended by introductory and concluding chapters (the 'Three-Paper PhD'). The aim of the introductory chapter is to introduce some of the key concepts and central themes that motivate the questions explored in each of the three papers. It also introduces an overarching narrative which connects all three papers. This narrative will be periodically returned to in short epilogues after each paper. The papers themselves are, for the most part, presented in the original format required by the journals for which they were written. Given this, the writing style may change – sometimes significantly – between each section. The concluding section summarizes the findings from each paper and reflects on the overall theme. When read in sequence, these five parts can be read something like a traditional PhD thesis. However, as the three paper format entails, each paper can also be assessed as an independent piece of work.

Part 1: Background

1.1 Introduction

1.1.1 Personal data: the view from 10,000 feet

What is the value of personal data? This seemingly innocent question, frequently asked these days, turns out to be very tricky to answer. It requires a great deal of unpacking, which soon uncovers a host of more fundamental questions. What *kinds* of personal data might we be talking about? What makes data *personal*? Value to *whom*? What *kind* of value are we talking about; value as a tradeable asset, value to an individual, or value to society? Indeed, what determines the value of immaterial stuff, generally speaking, in the 21st century? How is it that intangible assets – such as data, software, intellectual property, algorithms, standards, networks, knowledge – came to occupy such a central position in our economic, social and political relations?

Debates about personal data often look like a microcosm of a much wider debate about the nature of modern economies and the changing face of capitalism. 18th and 19th century theorists, from Adam Smith to Karl Marx, sought to explain *industrial* capitalism, a system characterised by the transformation of raw materials into commodities through a combination of factory technology and physical labour.³ But the 20th century saw a move away from material production, to a system described variously as 'post-Fordism',⁴ the 'knowledge economy',⁵ the 'information age' / 'information society',⁶ or 'cognitive capitalism'.⁷ This change was driven in part by advances in information processing, with the cost of computing power halving every 18-24 months in accordance with Gordon Moore's famous 'law'.

The technology giants of Silicon Valley, promoters of the so-called 'Californian Ideology', exemplify the latest incarnation of this system.⁸ They have 'discovered and invented the new form of value',⁹ their businesses don't rely primarily on extracting value from surplus labor, but on capturing value from the externalities that result from networked digital technologies. For instance, popular social networks aren't sold to consumers as a service. Instead, they derive revenue from new forms of advertising, made possible by the vast amounts of personal data generated from their platforms.

In addition to value-extraction as a by-product of the *free flow* of information, these companies also profit from the *restriction* of certain kinds of information flow. Intellectual property rights (particularly patents,

3 (Marx, 1939), (Smith, 1776). This age was characterised in its latter forms as *Fordism* (Gramsci, 1995).

4 (Amin, 1994)

5 (Drucker, 1969)

6 (Castells, 1999), (Bell, 2007), (Webster, 2014).

7 (Boutang, 2011)

8 (Barbrook & Cameron, 1996)

9 (Boutang, 2011) p49

copyright and trade secrets) allow the fruits of cognitive labour to be restricted, and therefore monetised, through state-backed artificial monopolies. Technical architectures, code and protocols are also deployed to achieve similar ends.¹⁰ These seemingly contradictory strategies lead to counter-intuitive business models, based on striking a delicate balance between opening up and closing down the flow of information and informational goods. On the one hand, data sharing is encouraged between peers, and software source code may be given away for free.¹¹ On the other hand, the data, algorithms and networks that make businesses profitable are often closely guarded. These business models – now no longer the preserve of Silicon Valley - limit external stakeholders' ability to determine how data and code may be used, restricted or modified.

These economic developments have evolved alongside and as a result of significant technological change. In addition to the advances in processing power described by Moore's law, recent innovations have led to a proliferation of data and techniques for analysing it. In the 1990's, the World Wide Web emerged as the primary technical infrastructure for online interaction.¹² Later, the mass adoption of a multitude of personal computing devices, from smartphones to wearable devices, has ensured that a steady flow of data streams emanate from our daily activities. Terms like 'ambient intelligence', 'ubiquitous computing' and the 'internet of things' all attempt to describe the phenomenon of the digital spilling over into the 'real world', sweeping up personal data in the process.¹³

The tools for deriving insight from this data have also changed. So-called 'data science' combines analytical techniques from statistics and computer science to produce insight from multiple large, heterogeneous and unstructured data sources ('big data').¹⁴ A significant aspect of these new data-intensive methods is the extent to which they signal a potential fundamental change in the scientific method. Rather than starting from a hypothesis that predicts a linear relationship between one set of variables and another ('output') variable, data-driven science explores every possible relationship (including highly complex, non-linear functions) between any set of variables. This shift has been described by some as going from data to algorithmic models, from 'model-based' to 'model-free' science, and from parametric to 'non-parametric' modelling.¹⁵

10 Restriction is achieved not only through user agreements (or 'click-wrap' (Murray, 2012)), but through the technical architecture (Lessig, 1999); see also (Zittrain, 2008) on how internet platforms have become 'locked-down'.

11 A competitive business rationale for releasing things for free is to undermine the value of a product / service that competitors are attempting to monetise (Sterling, 2014)

12 (Berners-Lee & Fischetti, 1999)

13 For the implications of personal data and 'ambient intelligence' see e.g. (van Dijk, 2009); (Monteleone, 2011); (de Vries, 2010). For 'ubiquitous computing', see e.g. (Spiekermann & Pallas, 2006); (Langheinrich, 2001). For the 'internet of things', see e.g. (Pepper, 2014).

14 See e.g. (Viktor Mayer-Schonberger, 2013)

15 See e.g. (Russel & Norvig, 2009) chapter 18. For discussion of the claims made regarding the 'model-free' nature of developments in data science, see inter alia (Barnes & Wilson, 2014)

These changing dynamics have given rise to a climate in which data is ever more valuable. According to an oft-repeated phrase, it is the 'new oil' of the digital age, fueling new services and creating billion-dollar industries.¹⁶ But unlike crude oil, personal data is fundamentally about people, and therefore its value is also more complicated.

To add to this complexity, the demarcation of *personal* data from *non-personal* data is much disputed. According to the UK legal definition, personal data must:

‘relate to’ an individual who is alive and is identifiable either from those data or from a combination of those data and other information that [the organisation responsible for it] has or is likely to gain possession of.¹⁷

Where they ought to provide clarity, such definitions have proven slippery in practice, and make the issues even harder to address.¹⁸

These various factors mean that personal data raises some particular social and political problems. The next section outlines the main terms and concepts that have been used to describe the challenges raised by personal data.

1.1.2 Privacy, data protection, and social concerns arising from personal data

Addressing the social and political problems associated with the collection, manipulation and dissemination of personal data has long been recognised as a key challenge for the post-industrial democratic state.¹⁹

These problems are often referred to broadly by the term *privacy*. While the history of the concept can be traced back over millennia, much of modern scholarship – from law to computer science – refers to its establishment in US jurisprudence, in particular to Warren and Brandeis' seminal article 'The Right to Privacy'.²⁰ The authors described a need for a new legal principle to protect individuals in light of threats from new technologies and business models, in particular, instantaneous photography and widespread newspaper

16 The phrase is attributed to Meglena Kuneva, European Consumer Commissioner, in March 2009, in (WEF, 2011)

17 This paraphrasing of the definition from the UK Data Protection Act is borrowed from (Christopher Millard & Kuan Hon, 2012), p 72. The definition is transposed from that of the EU Data Protection Directive (Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the Protection of individuals with regard to the processing of personal data and on the free movement of such data ('Data Protection Directive' or DPD)).

18 These definitions have proven controversial in the courts (see e.g. *Durant v Financial Services Authority* [2003] EWCA Civ 1746, [2004] FSR 573). Subsequent attempts to clarify, e.g. (UK Information Commissioner's Office, 2007), (Article 29 Data Protection Working Party, 2007), have arguably been insufficient, according to (Millard & Church, 2007).

19 See e.g. (Bennett, 1992) p. 1

20 (Warren & Brandeis, 1890). While Warren and Brandeis are commonly taken as a starting point for discussion, the concept can be traced back at least as far as Aristotle (Westin, 1967)

circulation. Privacy was defined here, following Judge Cooley's words, as the individual's 'right to be let alone' from such intrusions.²¹

Later treatments of the concept of privacy have differed somewhat, focusing on an individual's right over the use of personal information relating to them. For instance, Alan Westin described privacy as 'the right of the individual to decide what information about himself should be communicated to others and under what circumstances' (Westin, 1967). In this tradition, privacy is related to the protection of human dignity and individuality (Bloustein, 1964). Related is the idea of *informational self-determination*, a term found in German constitutional law meaning respect for 'the capacity of the individual to determine in principle the disclosure and use of his/her personal information'.²² As Edward Eberle writes, 'in the modern information age... control over personal information is the power to control a measure of one's fate'; necessary for ensuring a personal sphere which allows for the 'freie Entfaltung der Persönlichkeit' ('free unfolding of the personality') (Eberle, 1998, p. 1002). In US jurisprudence, this personal sphere not only protects the development of personality but also independent decision-making; for Louis Henkin, privacy rights protect a 'zone of autonomy', free from government intervention (Henkin, 1974). As these various definitions indicate, privacy harms are, to a significant degree, *subjective* in nature (Prosser, 1960).

Those hoping for a clear, tight definition of privacy may be disheartened by the variety of different articulations it has been given. But this variety does not necessarily make privacy a hopelessly nebulous concept. As Daniel Solove argues, it may simply indicate that privacy is best understood in the Wittgensteinian sense as a 'family resemblance' concept ((Daniel Solove, 2002 p485) citing (Wittgenstein, 1968)). According to this view, the things that privacy refers to may not all share one single essence, but instead share an overlapping set of traits, like members of a family.

These various explications of privacy – particularly those focused on *information* control – capture some of the motivation behind the 'first generation' of information privacy and data protection laws, established in the 1970's and 80's in response to the increasing use of computers in the public and private sector.²³ Privacy now has the status of a fundamental right in many national and transnational legal systems (Bennett & Raab, 2006).

Data protection laws are also widely adopted around the world, and cover a range of prohibitions, rights and obligations relating to the processing of personal data.²⁴ In Europe, the 1995 Data Protection Directive (DPD) was established with the aim of protecting 'the fundamental rights and freedom

21 Ibid, p195, citing (Cooley, 1879)

22 Decision of German Constitutional Court, BVerfG, 1 BvR 518/02 of 4 April 2006, Absatz-Nr. (1-184).

23 See e.g. (Brown & Marsden, 2013) p48. According to (Tene, 2013) the so-called 'first generation' includes, primarily, the 1980 OECD Privacy Guidelines (OECD, 1980), the 1995 EU Directive (European Council, 1995) and various U.S. sector-specific privacy laws dating back to the 1970's.

24 For an overview of the establishment of European data protection laws, see e.g. (Fuster, 2014)

of natural persons, and in particular the right to privacy with respect to the processing of personal data'. As the wording makes clear, the Directive did not aim to protect *personal data* itself, nor did it establish a fundamental right to data protection (since the EU lacks competence to enact fundamental rights legislation). Rather, the Directive aimed to regulate personal data so as to protect the *fundamental rights and freedom* of natural persons, including privacy. An additional aim was to ensure the free flow of personal data within the EU internal market (in fact, this second aim provided the legal justification for the Directive).

These stated aims of the Directive belie a divergence of opinions about what the core aim of data protection actually is.²⁵ The precise nature of the relationship between privacy and data protection is also contested from a conceptual and legal perspective. 'From a conceptual perspective', Gellert and Gutwirth argue, 'data protection is both narrower and broader than privacy', and thus an 'ambiguous relationship' exists between the two concepts (Gellert & Gutwirth, 2012, p. 269-270). From a legal perspective, while data protection and privacy are to some extent recognised as separate legal rights, they are in some contexts treated 'as if they are interchangeable'.²⁶

Despite these ambiguities in the purpose of data protection law, its substantive principles are relatively clear. Most data protection regimes include some variation of the following principles (I refer here to the UK Data Protection Act, for a non-exhaustive illustration of some of the principles common to many other regimes). Personal data must be processed fairly and lawfully (Principle 1). Explicit, legitimate purposes must be identified and specified prior to their collection (Principle 2). Data must be adequate, relevant and not excessive in relation to the specified purpose (Principle 3), as well as accurate and up-to-date (Principle 4). Processing is only legitimate if one of a number of grounds have been met; these include various conditions in which processing is deemed 'necessary', or if the data subject has given their consent (Schedules 2 and 3). The data controller (the entity who decides on the purposes of processing) must inform the data subject of their identity, the purposes of processing and the (types of) recipients of the data. Data subjects have several rights including the right to obtain information about the existence and purposes of processing of their data, and an explanation of the logic of an automated decision based on it.

These principles stem from their expression in guidelines produced by the

25 Fuster and Gutwirth note the 'divergent interpretations of the nature of personal data protection' (Fuster & Gutwirth, 2013 p.1); while Andrew Charlesworth notes divergence of views as to data protection's purpose both between scholars and between European Union member states (Charlesworth, 2006 p.1).

26 For instance the EU Charter of Fundamental Human Rights recognises privacy and data protection as two separate independent legal rights, while Dutch, Spanish and Finnish law, and international human rights texts, regard data protection as a *subset* of privacy. However, they are also sometimes conflated; as Orla Lynskey has argued, the Court of Justice of the European Union has 'treated privacy and data protection as if they are interchangeable' (Lynskey, 2014 p.3). See also (Tzanou, 2013).

OECD in 1980, but the nature of the problem they are expected to deal with has changed quite dramatically.²⁷ The notion that privacy is simply about giving individuals *control* over access to and use of their personal information, is being challenged by a new set of concerns arising from the changing technological and business environment mentioned above (section 1.1.1).²⁸ As a European Commission study argued in 2010, personal data is now increasingly collected and automatically analysed.²⁹ This activity carries 'the risk of individuals becoming mere objects, treated (and even discriminated against) on the basis of computer-generated profiles, probabilities and predictions, with little or no possibility to counter the underlying algorithms' (ibid, p.18).

The report warns that decisions 'will increasingly be taken "because the computer said so" - without even the officials or staff carrying out the decision able to fully explain why'.³⁰ In this environment, the right to limit access to certain types of data by certain types of actors is a blunt tool, and does not go far enough; what's at issue are the judgements formed and decisions taken once data has already been collected.³¹ As Hildebrandt and Gutwirth argue, when we consider the risks of profiling, 'a paradigm shift is needed from privacy and protection of personal data, to discrimination and manipulation and transparency of profiles'.³² Whether these issues (of discrimination, manipulation and transparency) should be classified as conceptually distinct from privacy, as Hildebrandt and Gutwirth's wording suggests, or alternatively as *sub-concepts* under the umbrella term of privacy, is currently an open question.³³ But lack of agreement on this classificatory point has not prevented scholars from elucidating these new concerns.

They include 'social sorting', where 'discrimination and privilege are entrenched through the unplanned consequences of data gathering and analysis'.³⁴ Along these lines, there are growing concerns about the capacity for data-driven systems to erode liberal values like equality and autonomy.³⁵

27 (OECD, 1980). For reflection on their first 30 years, see (OECD, 2011)

28 See also (Allen, 2000), (Dwork & Mulligan, 2013), (Austin, 2014), (Nissenbaum, 2009) chapter 5.

29 (European Commission, 2010a)

30 Ibid p18. See also (Chopra & White, 2011)

31 and in any case, the ability to exercise such rights are diminished in the current environment, where data collection is ubiquitous and often a necessary precondition of receiving various essential services.

32 From (Hildebrandt & Gutwirth, 2008) p2

33 For instance, Dwork and Mulligan examine how issues of discrimination and fairness are frequently conflated with the concept of privacy, noting: 'the ease with which policy and technical proposals revert to solutions focused on individual control over personal information reflects a failure to accurately conceptualize other concerns' (Dwork & Mulligan, 2013) p. 38. Similarly, Raj Patel argues that 'the privacy approach is an inadequate framework for conceptualizing the harms posed by the use of big data' (Patel, 2015) p.1. By contrast, concepts like discrimination and 'decisional interference' have sometimes been classified as sub-concepts under the umbrella term of privacy, e.g. (D. J. Solove, 2006).

34 (I. Brown, 2013 p. 10), citing (Lyon, 2001).

35 See e.g. (Barocas & Selbst, 2014), (Sandvig et al., 2014), (Dwork &

Data-driven discrimination has become, it is claimed, the 'new normal' in the consumer sphere (Turow & McGuigan, 2014). Personal data may be collected in order to learn about people's differences and treat them differently on that basis; but what looks like helpful personalisation in one context might be unfair discrimination in another (Zarsky, 2014). Statistically-generated categories may be used as stereotypes, without any critical evaluation of whether it is just or fair to apply them to individuals (Gandy, 2010). Likewise, there are fears that in monitoring individuals' behavioural quirks and biases, and tailoring their digital environments accordingly, these systems threaten to undermine individual agency and exploit idiosyncratic human vulnerabilities (Calo, 2013b).

The problem is increasingly framed not in terms of *what personal details are collected by whom*, but *how* data can be used by powerful, opaque systems, to *what* significant societal effects. The vocabulary to describe these issues is in flux; the relevant phenomena have been identified as big data, profiling, datafication, dataveillance, and surveillance; while the values they purportedly threaten include civil rights, fairness, non-discrimination, equality, and autonomy.³⁶ For instance, one recent strand of critique fixates on *algorithms*, and is often accompanied by calls for 'algorithmic transparency', 'algorithmic accountability', 'governing algorithms' and 'algorithm auditing'.³⁷ This wave of discussion has even been described, partly in jest, as 'The Great Algorithm Panic of 2015'.³⁸

1.1.3 Approaches to the policy problem

Even if there is agreement on the need to address privacy and personal data concerns, this policy problem has been addressed from many different perspectives. This section briefly describes some of the predominant approaches, questions and substantive disagreements which motivate the wide variety of contributions to this policy debate.

As Colin Bennett notes of this particular policy area:

“A multitude of works exists on this overall subject. There are polemical books designed to alert the general public to the privacy problem; there are legalistic analyses of complicated and esoteric doctrinal and statutory questions; there are philosophical works on the various ethical and moral dimensions of privacy; there are more technical treatments from the computer scientists and information systems experts; there are official and unofficial reports from national commissions, international working parties, civil liberties groups and professional associations. Many conferences, seminars,

Mulligan, 2013)), (Sweeney, 2013), (Frank Pasquale, 2014).

36 See e.g. (The White House, 2014)

37 See e.g. 'Auditing Algorithms' workshop at the International Conference on Web and Social Media (ICWSM) [<https://auditingalgorithms.wordpress.com/submissions/>]; 'Governing Algorithms' and 'Algorithms and Accountability' conferences at the New York University School of Law in 2013 and 2015 [<http://www.law.nyu.edu/centers/ili/AlgorithmsConference>].

38 Professor Kate Crawford, during a presentation on 'Algorithms and Social Control' at Theorising the Web, 2015, New York. As quoted in (kitabet [twitter user], 2015).

and colloquia have been held; and all these accompanied by a steady flow of journalism.” (C. Bennett, 1992, p. vii)

Bennett was writing in 1992, but his remarks remain true; the multitude of works being produced on the topic has continued in the intervening years without abatement.

There are various broad frameworks through which the policy problem might be approached. Economists, for instance, have precise ways to define the ultimate objectives of any policy, in terms of social efficiency or welfare (Bohm, 1987). According to this view, policymakers addressing these challenges should seek to establish conditions which are likely to lead to personal data being used (or prevented from being used) in ways that maximise social efficiency (Brown, 2015, p. 4). Legal scholars, particularly those from the 'law and economics' tradition like Richard Posner, often adopt this position, described as 'welfarist'.³⁹ They also tend to characterise the rationale for policy intervention in terms of one or more of the paradigmatic cases of 'market failure' standardly used to justify state intervention in the market, namely: externalities, public goods, monopolies, and imperfect information (Adler, 2009).

Personal data can give rise to negative externalities, for example, when an organisation collects data for one purpose but sells it to third parties for direct marketing purposes. The organisation gains a benefit without having to compensate the individual for the costs they incur from potentially being subjected to invasive messages. The existence of a negative externality here might justify regulation that limits such transactions (Varian, 1997). Regulators might also step in where there are *information asymmetries* between consumers and companies regarding how data is used, or where there is a *lack of competition* for privacy-preserving services (for instance, where network effects mean that services like search engines and social networks operate as natural monopolies).⁴⁰ In addition to addressing these standard market failures, there is also a growing enthusiasm for regulation that aims to 'correct' common cognitive biases (Sunstein & Thaler, 2008), in so far as they play a role in privacy-related consumer behaviour (Acquisti & Grossklags, 2007).

There are of course other frameworks from which privacy and data protection might be approached, such as human rights or social justice.⁴¹ In these cases, privacy may be seen as just one of many potentially competing values – from free speech, egalitarianism, authority or autonomy – each of which may be more or less fundamental depending on one's political persuasion. The task for those working within such paradigms is to examine how various policies relating to personal data support or conflict with

39 (R. A. Posner, 1973). See (Kaplow & Shavell, 2002) for an assessment of the prevalence of this paradigm in law and economics.

40 Information asymmetries have been identified in e.g. (Romanosky, Acquisti, Hong, Cranor, & Friedman, 2006); (Özpolat, Gao, Jank, & Viswanathan, 2010), and natural monopolies in (J. O. Brown, Broderick, & Lee, 2007).

41 For examples of each (not necessarily representative), see (Banisar, 2000), (Lee Bygrave, 1998) on human rights, (Dwork & Mulligan, 2013) on fairness, (Francis & Francis, 2014) on social justice.

certain rights and social values. Sometimes this will involve explicit appeals to one or more traditions of political philosophy; the value of privacy, and of attempts to protect it, have thus been assessed and contested from within liberal, libertarian, socialist, communitarian and other perspectives.⁴²

Within these distinct political and economic outlooks there is room for significant disagreement, because the same political norms can be interpreted as having different implications for privacy. Consider the following questions, that even those who are committed to the same outlook might reasonably disagree on. Should liberals regard strong data protection as unduly paternalistic (Bergkamp, 2002); (Cavoukian, Dix, & Emam, 2014)), or as an essential pre-requisite for liberal autonomy ((Henkin, 1974); (Allen, 2000))? Should communitarians see privacy and 'big data' as conflicting public and private interests (and therefore allow the former to 'trump' the latter), or rather as two different kinds of public good (O'Hara, 2010)? Should those who are opposed to neoliberalism resist the 'datafication' of everyday life, or instead embrace it as a new opportunity to challenge the capitalist status quo ((Silverman 2015); (McQuillan, 2014); (McQuillan, 2015))? Does faith in free markets mean allowing industry to commoditise personal data without restriction, or does it mean limited government intervention in order to establish personal property rights over personal data, as many have suggested?⁴³ Such questions demonstrate how traditional political and economic perspectives may not provide clear guidance on policy issues relating to personal data. This has consequences for the way policy proposals are assessed; discussion often proceeds on the basis that policy proposals can be assessed in isolation from these more general traditional theories.

Furthermore, as in other areas of technology policy, privacy and data protection raise their own set of questions and dividing lines, owing to the complex interrelations between regulation, technology, business models, consumer behaviour and public attitudes. These issues include whether regulation can 'catch up' with technology;⁴⁴ whether code itself should be considered a form of regulation, and how this might affect the regulatory approach;⁴⁵ whether faith can be placed in technology itself to ensure fairer

42 See e.g. 'communitarian' (O'Hara, 2010), (Etzioni, 1999), 'libertarian' (Block, Whitehead, & Kinsella, 2005), and 'socialist' (Fuchs, 2012) perspectives. Utilitarianism is occasionally explicitly appealed to (e.g. (Alder, Schminke, Noel, & Kuenzi, 2008)), although not frequently – possibly due to the view that utilitarianism is just a philosophical articulation of economic welfarism (see (R. Posner, 1979) for an overview and critique of this view).

43 Personal property rights in personal data are much-discussed: (Spiekermann et al 2015); (P. M. Schwartz, 2004); (Samuelson, 2000); (Lemley, 2000); (Prins, 2006); (Bergelson, 2003); (Litman, 2000); (Murphy, 2012); (Payne & Trumbach, 2009). Some scholars even appear to assume (in my view, incorrectly) that in so far as existing privacy and data protection regimes require consent as a legitimating basis for processing (which is, in fact, rare), they are already 'functionally equivalent' to property rights in personal data; see (Bergkamp, 2002)p36.

44 (Moses, 2007); (Brownsword, 2008)

45 (Lessig, 1999); (Murray, 2007)

outcomes;⁴⁶ the extent to which people do or do not care about privacy;⁴⁷ the potential of new business models and organisational structures,⁴⁸ alternative networks, decentralised technologies and public infrastructure,⁴⁹ or transparent and accountable systems.⁵⁰ These dividing lines are arguably a more important source of disagreement in the policy debate than the philosophical and political divides mentioned above.

In the context of data protection, much of the debate concerns the choice of regulatory approach. One factor is the extent to which regulators ought to pursue traditional legal regulation, or leave industry alone to self-regulate.⁵¹ There are a variety of intermediate approaches that might also be pursued. A sub-question here concerns the balance of *ex post liability* versus *ex ante regulation*.⁵² The former focuses on punishing the misuse of personal data *after* harms arise, by giving data subjects a right of action against data controllers. The latter focuses on preventing potential harms *before* they occur. In practice, privacy and data protection law has combined both approaches, although the balance is subject to change.⁵³

The internationalisation of privacy and data protection creates other complicated dynamics (Greenleaf, 2012a). For example, attempts to harmonise between jurisdictions may lead to either a ratcheting up, or a levelling down of standards.⁵⁴ Much of this discussion focuses on Europe – as the purported 'engine of a global regime' (Birnhack, 2008), and the US.⁵⁵ But other jurisdictional differences are also important to consider.⁵⁶

-
- 46 See e.g. (Mascetti, Ricci, & Ruggieri, 2014), (Dwork, Hardt, Pitassi, Reingold, & Zemel, 2011) on technological means to support fair information processing. See also the proceedings of the Fairness, Accountability and Transparency in Machine Learning (FAT-ML) conference [<http://www.fatml.org/index.html>]
- 47 See e.g. (Barnes, 2006); (Turow, Hennessy, & Draper, 2015)
- 48 See e.g. (Mantelero, 2014); (FA Pasquale, 2010); (Heath, Alexander, & Booth, 2013)
- 49 See e.g. (Narayanan, Toubiana, Barocas, Nissenbaum, & Boneh, 2012) ; (Wendy Seltzer, 2014) ; (Kleek, Smith, & Shadbolt, 2012); (Thiel, Hermann, Heupel, & Bourimi, 2013)
- 50 See e.g. (Butin et al., 2012); (Pearson & Charlesworth, 2009); (Kolovski, Katz, & Hendler, 2005); (Seneviratne & Kagal, 2014a)
- 51 For a discussion of the merits of each approach, see (Tang, Hu, & Smith, 2008)
- 52 See (Kolstad, Ulen, & Johnson, 1990), (Hiriart, Martimort, & Pouyet, 2004), (Innes, 2004). For discussion of the merits of these different approaches in the context of data protection and privacy, see e.g. (Romanosky & Acquisti, 2009), (Grimmelmann, 2010), (B. Koops, 2014)
- 53 For instance, the EU Data Protection Directive 95/46/EC both imposes many legal requirements on data controllers, as well as securing the 'right of every person to a judicial remedy for any breach of the rights guaranteed him by the applicable law' (Article 22, (European Council, 1995)). The proposed General Data Protection Regulation arguably increases both the *ex ante* and *ex post* approaches (B. Koops, 2014) p7.
- 54 (Bennett & Raab, 1997); (Bennett & Raab, 2006); (Binns, Millard, & Harris, 2014)
- 55 e.g. (PM Schwartz, 2013); (P. M. Schwartz & Solove, 2014); (Steinke, 2002); (Cate, 1995))
- 56 (Greenleaf, 2011); (LA Bygrave, 2010)).

This section has given a broad overview of the many different disciplinary approaches, perspectives and concerns involved in this policy area. Contributions can range from abstract philosophical analysis to microeconomic models, from regulatory theory to studies of consumer behaviour. These different levels of analysis relate to each other in complex ways. This variety is understandable, since the issues involved are multi-faceted. In order to come to any substantive conclusions, these different strands need to be given due consideration. This thesis aims to do this, and to do justice to the breadth and complexity of this policy area, by synthesising a range of disciplines and methods, and addressing the topic at multiple levels. In this sense, this work takes on the interdisciplinary perspective of Web Science.⁵⁷

1.1.4 Personal data empowerment

Before beginning the discussion of the over-arching narrative, there is one other aspect of the changing personal data landscape to be introduced. This is what we might call *personal data empowerment*. The term is used here to describe an ideal whereby individuals use their own personal data to serve their own purposes, on their own terms, rather than organisations collecting and using it for their purposes, on their terms. Such purposes include making better decisions, managing ones personal life more effectively, and understanding and shaping ones own behaviour better.

This opportunity can be seen as a continuation of a more general trend in personal information technology. Before the advent of the web, searching large troves of data was confined to big organisations with expensive mainframe computers. Now, many individuals perform such searches in a purely personal capacity, dozens of times a day. Similarly, personal data is currently collected by organisations for their own purposes – such as service provision, customer relationship management, and operational efficiencies. But individuals generally lack equivalent systems to collect and use data for their own purposes.

Computer scientists and designers have long explored the potential for personal computers to help individuals perform personal equivalents of large organisations' computational practices, including in personal information management (PIM) (Jones, 2007). Some of this work is focused primarily on how people already use existing systems to manage their personal information (e.g. (Aviv, Boardman, & Jones, 2004)), while others are more focused on designing new systems, on the basis of explicit design principles.⁵⁸ The aim of the latter is not simply to help people to become more organised, but also to address various concerns such as privacy, consumer exploitation, or centralised control by monopolistic technology companies.

Furthermore, beyond dealing with privacy and other concerns, personal data empowerment provides a vision for how personal data might give

57 (Berners-Lee et al., 2006); (Halford, Pope, & Carr, 2010)

58 E.g. (Kleek et al., 2012); (Kleek, Smith, & Packer, 2013); (Tuffield & Shadbolt, 2008); (Moiso & Minerva, 2012); (Kirkham & Winfield, 2011); (Mun, Hao, Mishra, & Shilton, 2010); (Mortier et al., 2010); (Anciaux, Bouganim, Pucheral, Guo, & Le, 2013).

individuals new capabilities (Binns, 2014a).⁵⁹ The opportunity has been recognised by government and industry in recent years. Initiatives in the UK, US, Canada, France and Finland have been established to encourage innovation in personal information management services (PIMS).⁶⁰ By creating new efficiencies, as well as entirely new kinds of services, PIMS are seen as an opportunity for both economic growth and individual empowerment.⁶¹ They operate in a variety of sectors, ranging from finance (Abiteboul, André, & Kaplan, 2015) to personal healthcare (Ueckert, Goerz, Ataian, Tessmann, & Prokosch., 2003).

The notion of personal data empowerment through PIMS is not without its detractors. Some are sceptical about whether the proposition PIMS offer is genuinely empowering (Milyaeva & Neyland, 2015), while others support their goals but doubt PIMS's technical and economic feasibility (Narayanan et al., 2012).

The privacy and data protection debates referenced above are framed primarily in terms of how to *preserve* or *protect* certain pre-existing individual and social interests in the face of new technologies. If empowerment is mentioned at all in this context, it is usually conceived in terms of giving individuals a choice or a voice in organisations' data processing activities (with some exceptions).⁶² This is understandable, since this discussion is focused on perceived risks and harms, and on regulating organisations' behaviour. However, the potential of personal data empowerment is important to consider alongside discussions of privacy and data protection, for reasons which will become clearer during the remainder of this thesis. Suffice to say, notions of empowerment are intertwined with notions of protection, and any thorough treatment of the issues arising from personal data ought to include both aspects.

This concludes the overview of the challenges which define the scope of this work. The remainder of this introductory chapter introduces a narrative which weaves together the content of the three papers. It outlines a particular approach to addressing the challenges of personal data, which I call Openness for Privacy (OfP). It begins with a familiar debate about the tensions between openness and privacy. It argues that the traditional

59 The role of internet technologies in more general forms of consumer empowerment has been explored in academic literature (e.g. (Pires, Stanton, & Rita, 2006) and (Füllera, Mühlbacherb, Matzlerb, & Jaweckic, 2009)).

60 The UK government's 'midata' initiative seeks to ensure that individuals have a right to a raw data copy of their personal data from any provider ((UK Cabinet Office, 2012); (Shadbolt, 2013)) and similar initiatives exist in the US, Canada (the 'Blue / Green Button' initiatives), France ('MesInfos' [www.mesinfos.fing.org]), and Finland (Poikola, Kuikkaniemi, & Honko, 2015).

61 See e.g. (World Economic Forum, 2011); (WEF, 2012); (WEF, 2013); (Searles, 2013); (Shadbolt, 2013); (Ctrl-Shift, 2014); (Heath et al., 2013)

62 There are some exceptions. For example, Peter Swire defines 'data empowerment' as a state where 'ordinary people can do things with personal data that only large organizations used to be able to do', and argues that this needs to be taken into account in data protection discussions (Swire, 2012)

conception of the principle of openness, as applied to privacy, is too narrow to be effective. However, if we consider broader ideas about what openness means and what it can achieve in a variety of contexts outside of privacy, we can arrive at a more helpful notion of openness *for* privacy.

1.2 Privacy and Openness: contradictory or complementary?

This section considers, at an abstract level, the tensions between openness and privacy, and some of the ways that they might be reconciled. This provides the backdrop for the concept of Openness for Privacy to be introduced in section 1.4.

1.2.1 Zero-sum?

Privacy and openness are often presented beside each other as opposing or competing concepts. *The* question of privacy and openness, we are invited to assume, is which of them should take priority over the other. The terms used to describe this conflict may differ – we might talk about privacy, anonymity or data protection on the one hand, and openness, transparency, or freedom of information on the other – but in case after case, these two clusters of concepts are raised as if there is always necessarily a mutually exclusive choice, trade-off, or balance to be struck between them. This is a *zero-sum paradigm*, in which privacy (or anonymity, or data protection) is pitted against openness (or transparency, or freedom of information), such that more of one necessarily means less of the other, and our main task is strike the right balance between the two.⁶³

Examples of the zero-sum paradigm can often be found in the media, academia, corporate public relations statements and policy discussions over the last decade or more.⁶⁴ Opinion pieces with titles like 'Secrecy vs. Transparency' and 'Openness vs. Privacy' argue that 'modern societies have to find the right balance' between these values.⁶⁵ These tensions came to a head in the early 2000's, when newly digitised records were increasingly made available on the open web. Champions of openness argued that 'access to information... the very lifeblood of self-governance' was in danger of being 'trumped ... by yearnings for privacy'.⁶⁶

In subsequent years, representatives of some technology companies began to argue that 'privacy is dead'.⁶⁷ This bereavement is not a cause for sadness,

63 In game theory this term refers to situations where one agent's gain (or loss) is exactly balanced by another agent's loss (or gain) (see e.g. (Von Neumann, 1953)). Technically, zero-sum games only exist between *agents* rather than *principles* like openness and privacy. I use the term here in a loose, illustrative sense.

64 For academic examples, see (H. R. Anderson, 2011), (Walker, 2000), (Lundblad & Masiello, 2010)

65 See e.g. (Cate, 2001), (Abraham, 2015) ('Privacy vs Transparency')

66 From 'Privacy concerns gone awry' in LJWorld archive, Retrieved from [http://www2.ljworld.com/news/2001/apr/26/privacy_concerns_gone/] in September 2015

67 The phrase is usually attributed to Sun Microsystems CEO Scott Mcnealy (according to (Rauhofer, 2008) p. 1).

they claimed. It is, in fact, the beginning of an alternative world of openness. “If people share more, the world will become more open and connected. And a world that’s more open and connected is a better world”, claimed the CEO of a prominent social network (Zuckerberg, 2010). These claims are echoed by policymakers and political representatives. As a US Congressperson stated in a Senate hearing on consumer privacy:

'What happens when you follow the European privacy model and take information out of the information economy? . . . Revenues fall, innovation stalls and you lose out to innovators who choose to work elsewhere.'⁶⁸

These statements can be distilled into a simplistic argument. According to this approach, openness and privacy are opposing poles of a spectrum on which society must situate itself somewhere. At one pole, we have *total* privacy, a situation in which the value of privacy is maximised at the expense of any other values. Most proponents of this argument appear to have in mind a definition of privacy as personal control over the disclosure and use of personal information. On this definition, *total privacy* is when individuals have absolute say over how their information gets disclosed and used. The implication is that, given total privacy, the disclosure and use of personal information will be significantly limited; many potential beneficial uses will simply not happen.

On the other hand, *total openness* would be when individuals have no such say. At its most extreme, this would mean a free-for-all where everyone could use personal information for any purpose, taking advantage of personal data as a kind of informational good (with its inherent infinite reproducibility) to maximise its uses. A less radical form of openness would be implemented in a limited way, through some form of collective control. Either way, the scenario we are supposed to imagine is one in which personal data flows much more easily and extensively.

Between these two poles, there are intermediate points entailing more or less disclosure of personal information, and a more or less permissive environment for its re-use. We are thus invited to pinpoint where society ought to be along this spectrum.

The zero-sum paradigm is taken to an extreme in David Brin's *The Transparent Society*.⁶⁹ Brin argues that more transparency and less privacy is an inevitable, indeed *desirable* outcome of increasing surveillance, digital recording and storage. When all behaviour, decisions and actions are recorded and stored, most people will know what most other people and institutions are doing, most of the time. While this will involve a loss of privacy, that loss will be outweighed by the benefits of greater accountability and enforcement of the law and social rules in all corners of life, Brin argues. He uses the term *sousveillance* to describe the ability of

68 Hearing of the Subcomm. on Consumer Protection, Product Safety and Insurance of the Senate Commerce, Science and Transportation Committee (Apr. 29, 2010) (statement of Rep. Marsha Blackburn, R-TN)

69 (Brin, 1999).

ordinary members of the public to monitor the world around them.⁷⁰ Unlike traditional forms of transparency, which involve powerful entities disclosing their activities to the public, *sousveillance* allows the public to surveil the powerful using their own devices.

Various technology writers and futurists have supported Brin's views.⁷¹ The basis of their argument is that if surveillance is indeed inevitable, it should be democratised, made available to the weakest so they may document abuses of power. While blanket transparency might give existing powerful entities some more power, in being *sousveilled*, they will be forced to exercise that power more responsibly. Advocates of *sousveillance* have found a growing number of real-life examples to point to in recent years (Mann et al., 2002). For example, human rights organisation *Videre est Credere* ('to see is to believe') have pioneered equipping oppressed communities with cameras to expose human rights abuses against them.⁷² The ubiquity of smartphones with video recording capability has meant instances of police brutality are often captured on camera by bystanders and 'go viral' online.⁷³

However, *sousveillance* alone may not be enough to correct the underlying imbalances of power between citizens and their governments, or consumers and powerful corporations. If it is accompanied by an equivalent rise in surveillance capabilities of the powerful, the overall effect might be worse for the least powerful. And in any case, what good is capturing abuses of power on camera if the apparatus to correct them is absent, or systematically favours the powerful?⁷⁴ For these reasons, Brin's notion that privacy problems may dissolve through greater openness has been dismissed as 'wishful thinking' (Clarke, 1993).

1.2.2 A Middle Ground

Perhaps the middle of the spectrum between openness and privacy might allow for the best of both. A strong commitment to both principles is not uncommon; Justice Brandeis is largely credited as both the inventor of the U.S. right to privacy and a vehement advocate for transparency.⁷⁵ From this perspective, both have their place, and it is up to society to define the borders between them through democratic debate. While transparency may be like sunlight, it need not be indiscriminate in its glow.⁷⁶ We do not need all sections of society and all types of data to be open in order to reap the benefits of openness. Rather, like a series of spotlights, transparency can be

70 The term is originally credited to Steve Mann (Mann, Nolan, & Wellman, 2002)

71 See for instance (K. Kelley, 2014), or (Jarvis, 2011); Jarvis has even been called 'our decade's David Brin' (<https://twitter.com/hoofnagle/status/124698140174594048>). For critique, see (Morozov, 2011)

72 See [www.vedereonline.org]

73 For instance, recent video footage of police violence (Rabinowitz, 2015)

74 See discussions following the Eric Garner case, which suggest such footage has little effect, e.g. (McLaughlin, 2014). See also (Ullrich & Wollinger, 2011)

75 *Ibid* and (Brandeis, 1913)

76 "Sunlight is said to be the best of disinfectants" - (Brandeis, 1913) p2

applied in specific places to different degrees, keeping other areas in the darkness of privacy. Rather than a single spectrum of openness and privacy, we have many spectrums, which can cut across multiple political or social fault lines.

The situation is at its most straightforward when dealing with information that is not privacy-sensitive in the first place. The movement for open data, for example, initially focused on various kinds of non-personal data, such as geospatial data or government spending, which presents no or little risk to privacy.⁷⁷

The waters become muddied, however, by the many examples of data which may be in the public interest, but still contain data identifiable to individuals. For instance, to what extent should the personal lives of elected representatives be subject to public disclosure? Should the recipients of public subsidies be publicly listed so that the fairness of such programmes can be scrutinised? Furthermore, many datasets which do not directly identify individuals might nevertheless compromise privacy if they are subject to re-identification attacks.⁷⁸ Tensions between openness and privacy therefore remain.

One straightforward principle for deciding these matters is succinctly expressed in the mantra of 'privacy for the weak, and transparency for the strong', where privacy and transparency are applied as a means of rebalancing power inequalities.⁷⁹ While Brin may cite the likes of John Locke and Adam Smith as inspiration for his radical transparency, this more nuanced position is perhaps closer to the original enlightenment vision espoused by these classical liberals. Individual citizens can hold the state accountable, but are free from unwarranted interference in their own lives. This simple mantra still leaves room for a great deal of debate. It begs the question; what sort of power does an actor need to have, to what degree, to justify an invasion of their privacy?

1.2.3 Compatibility, mutual reinforcement

Such controversies, recently raised in the context of open data, are part of an old and wide-ranging discussion on the balance between privacy and the public interest (for an overview, see (Janssen & Hugelier, 2013)). There is already a body of work which attempts to reconcile these two principles under one information rights framework.⁸⁰ This involves acknowledging that

77 See, for instance, the UK government's open data portal (data.gov.uk), which states that it is 'only about non-personal, non-sensitive data ...' (<https://data.gov.uk/faq>)

78 See (Narayanan & Shmatikov, 2007), (Ohm, 2010).

79 The exact quote is from (Assange, Appelbaum, Müller-Maguhn, & Zimmermann, 2012), but versions of the same sentiment have often been repeated elsewhere. E.g. 'Privacy protections must be inversely proportionate to power and... transparency requirements should be directly proportionate to power' (Abraham, 2015), or 'Transparency is an opportunity and even obligation for corporations and other institutions. But it is not an opportunity or obligation of individuals. Individuals have the obligation to withhold and protect their personal information' (Tapscott, 2010)

80 E.g (Banisar, 2011), (O'Hara, 2010), (Floridi, 2014)

zero-sum scenarios may exist, but then reconceptualising the debate to show how openness and privacy are not fundamentally opposed. The point is not to pick a side (openness *or* privacy), but to explore how each of them have a role to play in serving a common set of principles for information policy.

In addition to incorporating the two principles in one framework, there may be ways to reconcile them even further. As well as simply being merely compatible, openness and privacy might in some senses be *mutually reinforcing*. One direction of reinforcement is *privacy* in support of *openness*. For instance, one might argue that strong privacy safeguards are necessary to reassure the public about the release of open data comprised of their aggregated personal data.⁸¹ Or one might point to anonymity as a necessary condition to encourage whistleblowing, which is a kind of 'vigilante' openness.⁸² These are the kinds of considerations a privacy advocate might appeal to, to convince an advocate of openness to care about privacy on their own terms.

My interest here is in mutual reinforcement in the *other* direction – what might openness have to offer privacy? A simplistic answer can be sketched on the basis of the arguments above. Openness can help ensure that power is exercised responsibly. If privacy concerns are essentially about the power imbalances that arise when states and private entities use personal data, then openness about their use of that data may help ensure that such power is exercised fairly and responsibly. In other words, openness supports privacy in the same way it can support other values: by allowing society to monitor adherence to it and challenge those who fail to uphold it. This gives us an intuitive sense of how openness might serve privacy and other concerns related to personal data.

This answer is hardly revelatory, however, since the value of openness in this regard is already recognised in privacy and data protection law. Most attempts to create principles for the use of personal data include a principle of openness on these grounds.⁸³ A resolution recently issued at an international meeting of privacy and data protection commissioners reads:

“Openness is a longstanding fair information principle that is reflected in several international instruments... Effective communication of an organisation's policies and practices with respect to personal data is essential to allow individuals to make informed decisions about how their personal data will be used and to take steps to protect their privacy and enforce their rights.”⁸⁴

Openness is seen by regulators as a precondition of *trust* between data subjects and data controllers:

'It is becoming increasingly apparent that the protection of privacy

81 E.g. (O’Hara, 2011), (Jonas & Harper, 2010), (R. Meijer, 2014)

82 E.g. (Moore, Huxford, & Hopper., 2014)

83 See the 'Openness Principle' in (OECD, 1980), which is now 'broadly reflected in data protection and privacy laws around the world' (35th International Conference of Data Protection and Privacy Commissioners, 2013) p. 3

84 (35th International Conference of Data Protection and Privacy Commissioners, 2013)

demands a partnership between individuals and the corporations with which they interact. Like any successful partnership, this must be based on trust *and therefore openness*' [emphasis mine]⁸⁵

As these statements indicate, openness is not a recent addition to data protection policy. The need for transparency over the monitoring of individuals was recognised in the data protection regimes established in Europe since the 1970's. It was arguably a response to the effects of secret registers of personal information (held by both governments and the private sector) during the atrocities of World War II and in oppressive post-war regimes.⁸⁶

A similar dynamic continues to play out in more recent debates. See, for instance, contemporary calls for more openness from government intelligence agencies over their surveillance programs, especially in the wake of the revelations of whistleblower Edward Snowden.⁸⁷ We therefore already have at least some understanding of how openness can serve privacy, in both a government and private sector context. One might therefore wonder whether any more needs to be said about the matter.

1.2.4 The Openness Principle's Failings, and Unmet Potential

The hypothesis explored in the rest of this PhD is that there is a great deal more to be said about how openness can support privacy and data protection. The way openness is currently appealed to in this context is severely limited. It usually means organisations simply documenting their personal data-related practices, perhaps in a privacy policy on their website, or in the small-print on a registration form. In *theory*, individuals will read this information, and be informed about how their data may be used. On the basis of this information, individuals will – again, *theoretically* – make decisions and exercise various rights, giving them a degree of choice over how their data is used.⁸⁸

This system, where organisations disclose what they're doing and individuals choose whether to accept this or not, is described as the 'notice and choice' or 'notice and consent' model.⁸⁹ It is at the heart of many privacy and data protection frameworks, stemming from the OECD principles ((Cate, 2010); (D Solove, 2013)). As the name implies, these rights and powers of data subjects can only be meaningfully exercised if the organisation is open about its activities and policies (i.e. if it provides 'notice').

This model is supposed to empower individuals with the information and rights they need to determine, to some extent, how their data will be used.

85 (Office of the Privacy Commissioner of Canada, 2013)

86 e.g. (Bamberger & Mulligan, 2013)

87 (Stacey & Aglionby, 2013); (Lucas, 2014)

88 Data subjects can be seen as exercising choice when the data controller's legitimating ground for processing is consent. But choice may also be exercised even when the legitimating grounds are not based on consent, for instance if the data subject has the right to object to processing or if they can simply not use the service.

89 E.g. (Calo, 2012), (Calo, 2013a), (L. Cranor, 2012), (Barocas & Nissenbaum, 2009)

But it has many flaws. One is the sheer difficulty of reading and understanding privacy notices, which are usually written in legalese,⁹⁰ and then intelligently weighing up the various benefits and potentially harmful consequences. Another is the lack of meaningful choice within and between organisations' personal data practices. Often, consent to a raft of uses is a precondition for using a service (with no room for negotiation); and there is little difference between the practices of competitors.⁹¹ Furthermore, even when organisations do offer more fine-grained consent controls, individuals are expected to deal with unfamiliar, frequently changing, bespoke settings interfaces (Boyd, 2008). There is no common protocol for individual preferences to be managed in an integrated and efficient way.⁹²

The result is that individuals generally do not even attempt to become informed ((Osman & Rahim, 2011); (Mainier & O'Brien, 2010)). Requiring organisations to be open about their use of personal data does not appear to have resulted in individual empowerment. If anything, it has disempowered individuals by perpetuating the pretence of informed consent. We might therefore conclude that the idea of openness for privacy has been tried – and shown to be a failure.⁹³

There is no denying that openness, as it is currently conceived in the privacy and data protection world, is flawed. But declaring it hopeless would be premature. Perhaps we merely suffer from a lack of imagination about what openness could mean, the various ways it might be implemented and the ends it might serve. If we explore the full scope and potential of openness in this context, we might find it has more to offer.

The practice of using openness to address privacy concerns is therefore already instantiated, albeit in a limited and flawed way. This is Openness for Privacy (OfP) version 1.0, as appealed to by the policymakers quoted above. What might version 2.0 look like? To develop it, I argue we need to look further into what openness could mean, drawing from a wider range of principles, activities and models that have been heralded under the banner of 'open'.

The remainder of this chapter begins that exploration. First, it presents an overview of various strands of work on the idea of openness in general; what it *is*, the various motivations for adopting it, its promises and its failings. This overview will help us to reconsider the role openness could play in managing the challenges of personal data, leading towards a more developed version of Openness for Privacy.

1.3 Openness: an overview

What is 'openness', more broadly construed? While the term may have some

90 e.g. (A. McDonald & Cranor, 2008), (Li, Sarathy, & Xu, 2010), (A. M. McDonald, Reeder, Kelley, & Cranor, 2009))

91 e.g. (Bonneau & Preibusch, 2010), (L. F. Cranor, Idouchi, Leon, Sleeper, & Ur, 2013)

92 (Binns & Lizar, 2012)

93 e.g. (Mantelero, 2014), (Barocas & Nissenbaum, 2009), (Custers & Hof, 2013), (B. Koops, 2014), (Coles-kemp, 2010), (Rainer & Stefan, 2010), (Austin, 2014)

intuitive connotations, it is undeniably vague; what it means to adopt openness as a principle and in practice is equally ambiguous. Even when it is attached to a particular object, for instance *open data*, it remains nebulous, with a diverse range of political, social and economic aims and a wide policy remit (Worthy, 2013).

Arriving at a comprehensive definition of openness that encompasses its many meanings and contexts is a difficult task. For present purposes, we need an account which is general enough to encompass the range of practices which might apply to a wide range of organisational, commercial and public policy challenges in the domains of privacy, data protection, and personal data.

Despite the term's wide-ranging scope and nebulous interpretations, a burgeoning body of multi-disciplinary research explores the connections between 'openness' in a wide variety of contexts. This section will draw on this research to derive a picture of what openness means, its core features, and how it is supposed to operate. Such a picture will allow us to explore in subsequent sections how such practices might work to address the aforementioned challenges raised by personal data.

1.3.1 The Origins of Open

Where does the idea of openness come from? Numerous attempts at tracing a genealogy of the concept begin by citing Karl Popper's *The Open Society and its enemies*.⁹⁴ Popper's philosophy argued against what he termed the 'closed society', in which political programs are driven by a set of unchallengeable received truths. By contrast, an 'open' society is one in which knowledge is persistently open to question and refutation. In such a society, no entity would have a monopoly on truth; instead, dominant beliefs are continually tested against evidence, as a diverse citizenry applies the critical rationalism of science towards the evaluation of political action. Truth, if it can be found, emerges as a result of anyone being able to challenge and potentially refute existing political programs. For Popper, the questions which normally occupied political philosophy were actually less important than the struggle between these two fundamental categories of open and closed. Popper saw widely divergent political ideologies, like communism and fascism, as equivalent in the sense that they are closed.⁹⁵ His defense of liberalism, democracy and the free market was based not on adherence to liberal ideals per se, but rather because he believed that these were the conditions under which openness would thrive. Openness is

94 See e.g. (Krikorian & Kapczynski, 2010), (A Meijer, 2009), (Albert Meijer, 2013), and (Tkacz, 2012). The term *genealogy* is used by these writers in a loose sense to mean an overall picture of the origins of an idea, or a collection of ideas bearing a family resemblance, in light of the wider context in which the idea(s) arose. This is somewhat aligned with the Nietzschean / Foucauldian sense of the term (as described in e.g. (Sherratt, 2006)), but should not be regarded here as an attempt to faithfully adhere to such an approach.

95 It should be noted that Popper's critique of Marx was not aimed at the substantive ideas expressed by Marx himself, but rather the 'prophetic element in his creed' which was 'dominant in the minds of his followers' (Popper, 1945, Chapter 21).

therefore not just one value amongst many, for Popper, but rather the most important precondition for political progress.

Popper's notion of openness is deeper and more politically fundamental than the one employed in discourse on data protection and privacy cited above. Popper was not simply concerned that governments and other powerful entities be transparent about their activities. His conception of openness was one in which questions about knowledge and value are open to contestation and revision by all members of society. For Popper, it is not enough for governments to be transparent about their activity; the principle of openness also encompasses the ability of citizens to challenge it and develop their own alternatives.

Popper's political philosophy contains an implicit claim about knowledge, namely that it is best served in decentralised systems of governance. The idea is that no one political party, no matter how wise its members, could ever know what would be best for a society overall. This notion of decentralisation influenced the economic theories of his student Friedrich Hayek in the 1940's.⁹⁶

Hayek brought Popper's political claim about knowledge into the domain of economics. He noted that the question of what new goods and services a society ought to produce, and who ought to be able to consume them to what degree, is incredibly complex and requires a great deal of knowledge and information. This knowledge is fragmented, scattered between the many different, disparate agents that make up an economy. Even the most well-equipped governments are unlikely to be able to gather all that information and disseminate it appropriately. Instead, the relevant information is communicated through the price signals that arise in a free market. In this way, agents are able to act rationally in response to events that they may have no knowledge of; a shortage of a resource causes producers to charge more, which causes consumers to buy less of it, even if they have no knowledge of the cause of the price change. When knowledge is decentralised in this way, forms of spontaneous order can arise.

1.3.2 Open Source

These political and economic notions of openness may seem quite disconnected from modern discourse on openness, especially as it is used in the context of computer software. However, a line has been drawn from Popper and Hayek's ideas to debates in the 1980's about free and open source computing, which in turn influenced the modern discourse on openness (Tkacz, 2012).

The free and open source software (FOSS) movements grew in the 1980's in reaction to the increase of proprietary software. Richard Stallman, a researcher at MIT, argued that intellectual property (in the form of trade secrets and copyright licences) was being used as a means to control what end-users could do with their computers (Kelty, 2008). Stallman advocated instead for the freedom of end-users to use, inspect, modify and distribute software as they wish, and began work on the GNU operating system which

⁹⁶ See e.g. (Notturmo, 2014), (Hayes, 2008)

would enshrine these freedoms through its more permissive 'copyleft' license (Chopra & Dexter, 2010). The related movement for 'open source' software was centered around the nascent Linux operating system, which eventually fused with Stallman's GNU project. While the open source movement was less concerned with user freedom as an ethical issue, it also advocated making underlying source code open through less restrictive licenses.

These movements re-energised the notion of political openness, and went on to inspire its multiple contemporary incarnations.⁹⁷ A set of core ideas is shared between Popper/Hayek's political/economic openness and that of the FOSS world. One is the premise that no one individual or group is capable of knowing what's best for all people and sections of society. Therefore, the best mode of organisation is one in which ideas and economic actors can compete in a decentralised fashion. Analogous ideas can be found in arguments against closed source computing. The totalitarian thinking of Popper's closed society is replaced, in Stallman's philosophy, by the totalitarian behaviour of proprietary software vendors, enforced through intellectual property laws.

Similar parallels exist between Hayek and the open source community. The latter's celebration of their distributed, non-hierarchical mode of production is arguably a parallel of Hayek's rejection of central planning (Tkacz, 2012). Eric Raymond distinguishes open and proprietary production methods by a now famous analogy of 'the cathedral and the bazaar' (Raymond, 1999). Closed production is akin to the building of a cathedral, 'carefully crafted by individual wizards'. Open production, as exemplified by the Linux development community, 'seemed to resemble a great babbling bazaar of differing agendas and approaches... out of which a coherent and stable system could seemingly emerge' (ibid, p.1). The bazaar, he claims, can be more effective due to the kind of participation it enables. Traditional firms only have access to the sum of their employee's knowledge, so their ability to imagine and build new features or to detect and fix problems are necessarily constrained by their workforce. An open source project, however, has potential access to a wider pool of talent and auditors. As Raymond quips, 'given enough eyeballs, all bugs are shallow' (ibid, p. 6). Open collaboration on digital platforms can also reduce the transaction costs traditionally associated with large organisations, allowing for alternative modes of production between peers ((Benkler, 2002), contra (Coase, 1937)).

The success of FOSS was taken by many as evidence for the viability of a new mode of organisation and production, and thus established a divide between open and closed information technology. The latter restricts the freedom of individuals to use technology, content and immaterial goods for self-devised purposes (Doctorow, 2004). Proprietary or closed models are, it is argued, at odds with the nature of our networked, digital world ((Lessig, 2004); (Boyle, 2002)). Whereas the latter demonstrates that, when the costs of creation and dissemination are drastically reduced, alternative forms of organising production are possible and desirable:

'Assume a random distribution of incentive structures in different

97 (Tkacz, 2012), (Krikorian & Kapczynski, 2010), (A Meijer, 2009)

people, a global network: transmission, information sharing and copying costs that approach zero, and a modular creation process... Under these conditions we will get distributed production without having to rely on the proprietary / exclusion model. The whole enterprise will be much, much, much greater than the sum of the parts' (Boyle, 2002, p. 322)

Advocates of openness therefore see it not only as a backlash against the restrictions of proprietary licensing, but also as an alternative approach to production, collaboration, and participation.

Openness also takes on a further political dimension in debates about information economics. According to Krikorian and Kapczynski, free and open source software can be seen as just part of a wider movement in response to the commodification and privatisation of abstract objects, ideas and methods (Krikorian & Kapczynski, 2010). Some see this kind of commodification as a necessary mechanism for capital accumulation in post-industrial economies (Solow, 1956). But critics argue that it amounts to an 'enclosure' of the 'commons of the mind', depleting the common stock of ideas and methods on which cultural and technological development depends ((Boyle, 2008); (Zittrain, 2008)).

These ideas have been applied in a variety of settings beyond software, and are increasingly described by the prefix of 'open' (Benkler, 2006). These include open government, open innovation, open hardware, open design, open education, open access, open science, open data, among others.⁹⁸ Perhaps the most famous example is Wikipedia, the open encyclopedia, whose content and software is produced collaboratively, released under an open license, and is organised and produced according to similar 'open' principles (Schneider, Passant, & Breslin, 2010).

Openness can be seen not just as a property of a product or mode of production, but also something to be embedded into formats, systems, protocols and, by extension, entire markets (DeNardis, 2011). The world wide web itself is a prime example of this; unlike predecessor hypertext systems, it was based on non-proprietary protocols and standards which allowed anyone to design and run their own servers and clients (Berners-Lee & Fischetti, 1999). It also gave significant powers to the individual user, including the ability to configure many aspects of their experience and to inspect the source code of pages through browsers' 'view source' function. As an open platform, the web allowed for competition between firms, rather than acting as a single monopoly provider of various internet-based services ((Shapiro & Varian, 1998), (Boudreau, 2010)).

The ongoing development of the web's open standards is also an illustration of the open approach to governance, adopted by the World Wide Web Consortium and other key stakeholders in the internet such as the Internet Governance Forum ((Ziewitz & Brown, 2013), (L Bygrave & Bing, 2009), (Brown & Marsden, 2013)). This 'open and collaborative' approach acknowledges that 'the success of the Internet depends on more than the work of one, single organization – no matter how big, diverse, or influential

⁹⁸ See the Peer-To-Peer Foundations' map of 'Open Everything' for a comprehensive list: [http://p2pfoundation.net/Open_Everything]

it may be', and emphasises operating 'collaboratively and inclusively' to reach decisions.⁹⁹ Connections have been drawn between this approach and other, more general governance styles; for instance, the 'Open Method of Co-ordination' in EU policymaking has been described as 'the Linux of EU integration' (Sundholm, 2001).

In recent years, *open data* has become one of the primary examples of openness. Its underlying rationale is that data collected by an organisation for one purpose might also be useful to others outside the organisation for different purposes (Pollock, 2008). Releasing data with a permissive rather than restrictive license could create positive opportunities at little or no cost to the organisation releasing it (Heimstädt, Saunderson, & Heath, 2014). Open data can be seen as a key ingredient of open innovation in the public sector, or 'citizensourcing'.¹⁰⁰

Examples include public transport timetable data being used to create travel advice services, or medical procurement data being analysed by third parties to identify savings for health providers (Shadbolt & O'Hara, 2013). Beyond its more practical uses, such data is also seen as key to 'cement trust between the government and citizens',¹⁰¹ and potentially 're-articulates notions of democracy, participation and journalism' (Baack, 2015). Open data about the private sector also has many uses; for instance, data on registered companies has been used by journalists to identifying corporate hierarchies and potential conflicts of interest (Lindenberg, 2014).

The broad applicability of openness has led some of its proponents to ambitiously claim that it is a foundation for new social and political systems, arguing for a 'Read/Write Society' (Lessig, 2006), or an 'Open Source Democracy' (Rushkoff, 2003). Its ideals are said to 'align ... with the political values of self-determination and autonomy, as well as those of collective governance' (Krikorian & Kapczynski, 2010, p. 36). More recently, open data advocates have placed openness in grand, historical terms, likening it to the translation of the Bible into English during the Reformation (Pollock, 2015).

Various (possibly unintentional) allusions to Popper's open society and Hayek's decentralised planning neatly bring the genealogy around full circle: open data is said to reflect 'a cultural shift to an *open society*'¹⁰², while an 'open source democracy' would work 'not by *central planning*' but through 'participatory, bottom-up and emergent policy' (Rushkoff, 2003, emphasis mine). Even the UK Chancellor, in 2008, proposed 'open source politics' as a way for interested citizens to collaborate on solving problems, instead of relying on politicians and civil servants' 'monopoly of wisdom'.¹⁰³

99 From [<http://www.internetsociety.org/what-we-do/how-we-work>]. Internet governance has even been suggested as a model for general international self-regulation (Mestdagh & Rijgersberg 2015)

100 E.g. (Hilgers & Ihl, 2010); (Kassen, 2013); (Misuraca, Mureddu, & Osimo, 2014)

101 Tim Berners-Lee, as quoted in (Ahmed, 2015). See also (O'Hara, 2012b)

102 Gavin Starks, presentation at OpenTech, 2015 [slides available at: <http://www.slideshare.net/theODI/odi-2015-06-opentech-gavin-starks>] (emphasis mine).

103 In an interview posted on YouTube.com, retrieved from

1.3.3 Critiques of Openness

It is unsurprising, given the many bold claims that have been made about openness, that it has attracted deserved interrogation from many quarters. Before proceeding to attempt to unify these disparate forms of openness under one definition, it is worth addressing some of the criticisms which have been leveled at the notion of openness. If it turns out to be a fundamentally flawed concept, it would be a mistake to continue with it.

The picture of openness which emerges is that of a very broad concept – perhaps too broad. One important criticism is that openness is vague to the point of vacuity. This is sometimes expressed by reference to the ambiguous political alignments of openness. It has been criticised by some for serving a neoliberal agenda ((Bates, 2012); (R. Kitchin, 2014)), and yet praised by others as an alternative to neoliberal paradigms about property (Krikorian & Kapczynski, 2010). Openness is simultaneously a space free *from* certain aspects of the market (for instance, intellectual property disputes), and yet also a space free *for* the market, where new businesses can compete to add value to underlying open information and digital infrastructure. It has been described as post-political or post-ideological, purportedly able to 'subvert the left-right divide' and 'appeal to libertarians, liberals, the postsocialist left, and anarchists' alike (Benkler, 2010). While some advocates see the broad church of openness as a good thing, critics argue that in being all things to all people, openness risks becoming 'dangerously vague' (Morozov, 2013).

This alleged emptiness at the heart of the open paradigm goes back, according to Tkacz, to its genealogical roots in Popper and Hayek's notion of the Open Society. Tkacz argues that their political philosophy suffers from an internal void, which is inherited by modern manifestations of openness as exhibited in activist groups, web entities like Wikipedia and Google, and the open government movement.

For Tkacz, Popper's notion of openness is reactionary; 'it gains meaning largely through a consideration of what it is not' (Tkacz 2012, p.400). The problem with this is that the openness of one era spawns forms of closure in the next. Popper and Hayek's visions of openness, defined in terms of their opposition to Platonic idealism, fascism and communism, were successfully achieved in the form of the capitalist liberal democracies which dominated the latter part of the 20th century. But the success of this vision led to new forms of closure – in the form of neoliberal programs to commodify information and proprietary digital infrastructures – which Popper and Hayek's openness is blind to.

These very closures prompted the second wave of openness of recent decades, according to Tkacz. But just as Popper's openness was defined primarily by its opposition to Plato, fascism and communism, the new wave of openness 'is articulated alongside an entourage of fractal sub-concepts that defer political description: participation, collaboration and transparency' (ibid, p. 403). Tkacz claims that applying open as a political descriptor closes down discussion, and stops the policy or program from being properly interrogated. As a result, each iteration of openness only has the

[<https://www.youtube.com/watch?v=PZwFDKOP9Jo>] in September 2015

conceptual resources to oppose the prevailing forms of closure. If a particular iteration of openness succeeds, it will eventually face new forms of closure and leave us bereft of the means to critique them.

I will not attempt to assess the merits of Tkacz's exegesis of Popper here. The critique is an important challenge to anyone tempted to use the 'open' label to describe their approach, including my own proposal of an 'open' approach to privacy.¹⁰⁴ It is therefore important to respond to this challenge before proceeding.

Let us concede, for now, that openness is most easily and frequently defined in terms of what it is opposed to. The first thing to note is that there is nothing inherently wrong with negative definitions. Concepts as important as liberty, health, and peace can all be usefully defined by what they are not.¹⁰⁵ One reason that openness may be easier to define in negative terms is that it may be a family resemblance concept (Wittgenstein 1968). That is, there are no necessary and sufficient conditions for something to be considered open, but rather, we consider things to be open because they share certain features with each other (even though there is no one single feature that each and every one of them share in common). Negative definitions and family resemblance accounts of concepts may be less satisfying, but that doesn't make them wrong or useless.

Such a definition of openness does have certain risks. Appeals to openness so defined may lack substance, and may risk inadvertently ignoring potential new forms of closure. Any attempt to apply openness in a particular domain will therefore benefit from including specific positive proposals and be mindful of the potential for further forms of closure.

The idea that our notion of openness needs to be consistently re-invented in response to the closures which crop up in new environments is also not necessarily a flaw. The same could be said of many worthwhile political concepts which may be most useful when articulated with a particular context in mind. Rather than attempt to rebut entirely this critique of openness, we may take it as a warning; that openness risks being empty if it is not appropriately contextualised.

What of Tkacz's charge that appeals to openness close down debate rather than foster it? If this were true, then we should not expect to see much debate or disagreement amongst members of 'open' initiatives. But in fact, such communities do appear to have healthy levels of critical self-reflection about what openness means and how it should be practiced.

For instance, the open source community has long debated the merits of different kinds of software licenses (e.g. 'Apache' versus 'GPL') on the basis of competing notions of openness (Rosen 2005). Similar conflicts arise

¹⁰⁴ Tkacz does a good job of extracting some of the positive descriptions of openness to be found scattered around the chapters of the Open Society. But by explicitly excluding many aspects of Popper's political thought, such as its relation to his thoughts on the scientific method and critical rationality, he potentially misses out important material that could be used to construct a more substantive version of Popper's openness.

¹⁰⁵ See e.g. (Berlin 1969).

between advocates of different content licenses, with some arguing that the 'non-commercial' and 'no-derivatives' variants of Creative Commons ought to be discontinued (Pollock 2012). Finally, there is much consternation in the open community regarding so-called 'open-washing'.¹⁰⁶ This term is used by openness advocates to decry those they perceive to be using the label to give their project an undeserved veneer of justification, on the grounds that it is not 'truly' open. These controversies suggest that far from being an empty slogan, openness has a substantive meaning which its advocates are careful to contest, define and defend. Rather than closing down discussion, appeals to openness are in fact critically evaluated by the open community.

Having defended openness against the charge that it is an empty concept which closes down discussion, there are two other major charges we must consider.

One is that openness is not necessarily egalitarian, and worse, may only empower the already empowered.¹⁰⁷ Merely giving everyone permission to reproduce and modify the source code of a computer program does not ensure that everyone will have equal capacity to do so. Open data released by governments may be downloaded by anyone, but the ability to derive meaningful analysis or build profitable services from it is not equally distributed. Inversely, those who are expected to become more open may face different costs in doing so. A mandate of openness might fall harder on those businesses who are less able to derive alternative revenue streams from their software, or on those governments with less technical capacity to publish their data in an appropriate format. In this sense, some of the rhetoric surrounding openness could be said to lack a critical awareness of the socioeconomic conditions underlying it.¹⁰⁸

A final objection is that openness simply neglects the many values of secrecy and partial information. Institutions may operate better if they do not have to disclose everything. After all, they have evolved to operate in an environment where they are not constantly scrutinised; one could therefore make an evolutionary argument for maintaining secrecy so as to prevent them from having to make painful adaptations (Dennett & Roy, 2015). Having to justify every action publicly could hinder government effectiveness, to the extent that the transparency gains do not outweigh the loss of efficiency (Fukuyama, 2014, p. 504). Furthermore, openness may actually inhibit rather than strengthen trust, since having comprehensive information about another's actions means one doesn't have to 'trust' them at all; instead, it may just encourage more elaborate forms of deception (O'Neill, 2002).

Various rebuttals to these general arguments against openness have been

106 The term has been used to describe software which is not seen as truly open (Schestowitz 2015), and in other putatively 'open' initiatives such as science or education ((van der Woert et al 2015), (Tamang & Donavan 2014)).

107 See e.g. (Wright, Glover, Prakash, Abraham, & Shah, 2009); (Bates, 2012); (Tsiavos, Stefanias, & Karounos, 2013); (Longo, 2011); (Slee, 2012).

108 See e.g. (Gurstein, 2011), (Van Dijck & Nieborg, 2009).

made.¹⁰⁹ It is neither necessary, nor within scope of this section, to rehearse them. Suffice to say that these types of objections do not generally point to inherent and fatal problems with the notion of openness. Rather, they point to a set of risks associated with the concept – namely, that it can seem empty if decontextualised, that it may empower unequally, and that there may be advantages to secrecy.

Each of these risks will be taken into account in the remainder of this thesis, as it explores the application of the concept of openness to privacy. By identifying the ambiguities and tensions in the concept of openness, these critiques inject a healthy dose of scepticism into the hubris which surrounds it at times, and provide a useful set of warnings to be heeded in the remainder of the work below.

1.3.4 Towards a definition of openness

This overview of openness hopefully shows that it can be about much more than the straightforward notion appealed to in data protection discourse, which simply consists of organisations disclosing their practices in some format.

Rather, it can be an approach to managing the flow of information and informational goods; a way of collaborating and organising through digital networks; and a means of convening stakeholders around an endeavor, whether they be governments, companies, civil society organisations or engaged individuals. These notions of openness aim to enshrine the freedom of anyone to scrutinise and modify, and to leverage the nature of digital networks to facilitate more efficient forms of decentralised collaboration.

Having acknowledged above that openness is often defined negatively, and also that it may also be a family resemblance concept without necessary and sufficient conditions, the prospects for a satisfyingly universal and comprehensive definition of openness are slim. But this doesn't mean we cannot arrive at a working definition derived from the sections above. Despite their contextual differences, these various forms of openness are motivated by a common core. Openness, in its various guises, embodies a unifying set of principles and core features:

1. It aims to dismantle structures and systems where decisions are made by select entities in a centralised fashion.
2. It is based on the notion that knowledge, expertise and the capacity for innovation are dispersed widely, and are therefore best leveraged in decentralised manner.
3. It aims to give as many people as possible the opportunity to access, re-use and contribute to knowledge and information ecosystems.

¹⁰⁹ For instance, some argue that the problems these detractors point to are not caused by openness, but by more fundamental problems of the state (Bass, Brian, & Eisen, 2014). Others claim that open data's true potential can be realised if it is properly subsumed within a larger framework of information justice (Johnson, 2014), and supported by the right 'participatory mechanisms' (Peixoto, 2013).

1.4 Openness for Privacy

Equipped with this definition of openness, we can now see just how limited its treatment is in the realm of data protection and privacy, where it rarely goes further than a legal requirement on organisations to disclose their practices with regards to personal data.¹¹⁰ As such, it is a decidedly one-way, top-down practice; regulators force organisations to report their activity, which then trickles down to data subjects who consume this information. As we have seen, this limited sense of openness has not resulted in better privacy protection or empowered data subjects.

By contrast, the definition of openness supplied above demonstrates that it is much more closely aligned with bottom up processes. It must support the decentralisation of decision-making, and the ability of individuals to access, re-use, modify and contribute to their information environment. It not only makes organisations practices more transparent, but it also affords the individual data subject more power to manage data on their own behalf. It thus constrains the data controller while increasing the data subject's options. I therefore propose an alternative approach – *Openness for Privacy* (OfP) – which involves a much broader notion of openness, inspired by some of the examples above.¹¹¹

Before we begin to flesh out this approach, however, it is important to consider whether it is needed. Given the many varying interpretations and multiple facets of both openness and privacy, one might be sceptical about the merits of trying to fuse a grand conceptual approach out of both. Abstract concepts and big ideas may just obscure complexity and nuance. Would it not be better to focus on the details of particular systems or policies?

The recent history of privacy and data protection research suggests that both big ideas and detailed analysis are necessary, and the former can act as a catalyst for the latter. This is arguably the case for 'Privacy by Design' (PbD), a term popularised by the privacy commissioner of Ontario, Canada (Cavoukian, 2006).¹¹² It is decidedly simple; in summary, it urges organisations to consider privacy during the design phase of innovation. As a concept, PbD is quite broad and, perhaps, obvious. It doesn't posit any specific hypotheses, nor does it explicitly advocate the use of particular software engineering patterns, encryption methods, standards or user interfaces (instead, it provides a set of general principles). But it has nevertheless generated a rich stream of more detailed research (as well as

¹¹⁰ Some technology companies have gone a little further on their own accord, for instance, producing reports about the number of government requests to access the personal data they hold. See the Electronic Frontier Foundation's annual 'Who Has Your Back' report for an overview.

¹¹¹ A similar term, 'Open Data Protection', is used in (Pagallo & Bassi, 2013). This approach emphasises how techniques like privacy impact assessments and anonymisation can mitigate the tensions between open data and data protection. While valuable, this differs from the approach developed here in that it looks to general strategies to mitigate tensions between the two interests, rather than specifically at strategies in which openness itself reinforces privacy.

¹¹² See also (Langheinrich, 2001) for origins of the term

changing industry practices and regulator focus).¹¹³ These outputs are arguably thanks to the generality of PbD, not in spite of it.

It is in this spirit that I propose the idea of Openness for Privacy. It is envisioned as an approach to addressing a range of personal data challenges, including, but not limited to, privacy – the elision is for brevity's sake.

In so far as it embodies the definition of openness provided above, it can be seen as a particular application of a general 'open' approach towards computation and data in society. In this sense, it can be taken as a normative political principle which policymakers can aspire to.

It can also be seen as an analytical construct which aims to provide clarity in discussions about privacy and data protection policy. It can help by synthesizing a range of otherwise disparate and disconnected concepts in this domain. Its main purpose within the scope of this PhD is to provide a conceptual basis for the specific research questions and applications which are explored in various ways in the three papers.

There are many different ways we might attempt to transpose the principle(s) of openness into the world of privacy. Not all of these will necessarily be a good idea, and some would be downright misguided.¹¹⁴ But it is my hope that at least some permutations of OFP are worth exploring; the remainder of this chapter will introduce a few of the most promising.

1.4.1 Open data for privacy

In 2012, the San Francisco city authority began publishing their restaurant hygiene inspection data in an open format. Previously, restaurants had only been required to display their inspection ratings on-site. For this on-site information to actually have a meaningful impact on a consumer's choice of restaurant, the consumer would need to enter the premises of several restaurants, inspect their walls to discover their ratings, before choosing one of them – an arduous and inconvenient process that few consumers are likely to undertake.

The authority's open data collated all these ratings in one dataset. After the data was made available, ratings website Yelp began including it in their restaurant rankings, so that consumers who care about food hygiene could make more informed decisions about where to eat. Yelp notified a random sample of restaurants about the change. On subsequent inspections, those restaurants tended to clean up their act and get better results, compared to others who weren't informed.¹¹⁵

What does this story have to do with privacy? It is an example of how

113 A search for the term on an online scholarly index (Google Scholar) suggests there are at least 4000 research papers referring to the concept at the time of writing (September 2015).

114 For instance, releasing bulk personal datasets under an open data license would carry great risks and should only be done in exceptional circumstances, if ever. But openness is about more than licenses, and privacy is about more than defining a single set of permissions for personal data.

115 As described on [<http://officialblog.yelp.com/2015/02/yelp-open-data-the-end-of-food-poisoning.html>] Retrieved September 2015.

information needs to be formatted and delivered in the right way to impact consumer behaviour in a market (Helberger, 2013). As we have seen, the existing notice and choice model is based on the idea that there could be a 'market for privacy' for a given type of service. This depends on the following conditions:

- Consumers being aware of the privacy practices of different service providers;
- Consumers being sufficiently motivated by privacy concerns to choose between providers on that basis;
- At least some providers offering privacy as a competitive differentiator.

In theory, this should lead to a positive feedback loop; the more consumers become aware, the more they will be able to choose providers based on their privacy credentials, and the more providers will compete for privacy-conscious consumers by changing their practices. But if any one of these conditions is missing, a functioning privacy market is unlikely to emerge.¹¹⁶

There is some evidence that consumers do care enough to switch to privacy-preserving products if it is easy to do so, and some providers tout their privacy credentials (Özpolat et al., 2010). But neither of these things matter if it is impractical for consumers to factor privacy-relevant information into their decision-making.

As we have seen, the practice of publishing lengthy privacy policies has not led to the level of awareness that would be necessary to kick-start this kind of virtuous circle. Each policy is long and unique to the provider, and the task of reading and comparing them to each other is laborious. Like hygiene ratings displayed on restaurant walls, they are not available in a format which allows them to be aggregated and compared independently of the vendor; in this sense, we might call them *proprietary* or *closed*.

What made the difference to San Francisco's restaurant industry was having the data in an open, aggregated form which allowed for easy re-use by the third-party rating site. Likewise, practices and policies regarding personal data could be represented in standardised vocabulary and made available as machine-readable data. This data could be used by third parties in various ways, including helping consumers make more informed decisions. Rather than users having to visit each provider's website and read their privacy policy, information about privacy practices could be provided independently and figure into consumers' decisions without requiring them to read it.

Exactly how this might work in practice is a matter for design innovation and empirical research.¹¹⁷ Individual companies can attempt to innovate by simplifying their privacy notice systems, but unless they all move in tandem, so the policies can be assembled together in a common data format,

¹¹⁶ 'If consumers have little reason to know about or believe good privacy practices, no firm has an incentive to follow them' (I. Brown, 2015), discussing (Greenstadt & Smith, 2005)

¹¹⁷ See e.g. (Ackerman & Cranor, 1999), (Byers, Cranor, & Kormann, 2003), (Balebako & Leon, 2011), (König & Hansen, 2012).

consumers will have a hard time making comparisons and innovative solutions are unlikely to scale. If a market for privacy has any chance of becoming a reality, this kind of data might be a necessary step.

Even if the idea of a market for privacy is more fundamentally flawed – i.e. if consumers simply aren't sufficiently motivated by privacy – this data could still prove useful in a variety of other ways.¹¹⁸ Policy-makers and civil society organisations could use it to monitor trends and activities of data controllers, and target their work accordingly. Intermediaries could use it to assess an individual's privacy exposure risk and develop targeted forms of protection.¹¹⁹ Companies could use it to benchmark against their competitors, potentially driving up standards. Organisations might use it to assess the suitability of potential outsourcing providers (a form of preliminary due diligence), or for other business-to-business interactions involving personal data. Knowing what data is held by other organisations can also dictate whether a given dataset can safely be made publicly available and in what form, because the existence of auxiliary datasets is a key risk factor for re-identification attacks (Narayanan, Huey, & Felten, 2016).

The idea of representing privacy practices in an open data format is not new. There is a long history of initiatives attempting to create such a system, with mixed results (for an overview, see (Binns, 2014b)). Several research projects have attempted to standardise large volumes of privacy notices, so that the aggregated data can be used to analyse the practices of data controllers (see, for instance, (L. F. Cranor et al., 2013), (Mary J. Culnan, 2000)). However, this work has so far been limited by the barriers associated with manually encoding policies into a standard data format.

The first of the three papers presented below is a contribution to this stream of research (Binns, Millard, & Harris, 2015). It attempts to overcome some of the traditional limitations in this field, by using a large novel source of standardised data from the UK regulator. It contributes to both the ongoing empirical research into the trends of data use, and to the development of design requirements for standardised privacy notice systems.

Open data about organisational uses of personal data is one important part of the OfP approach. But it is not the only part; there are multiple other ways that openness might play a role in privacy and data protection.

1.4.2 Open processing: transparency and modification

The OfP approach can also take inspiration from the notion of freedom in free and open source software (FOSS). For free software advocates in particular, the 'four freedoms' - to use, study, distribute and modify – are key (Stallman, 2002). Their purpose is to ensure that users of software remain in

¹¹⁸ The idea that consumers don't care enough is widespread, but see e.g. (Turow et al., 2015) for evidence to the contrary.

¹¹⁹ The introduction of intermediaries into the equation could bring its own problems; including the question of trust (see the notion of 'agency costs' (Jensen & Meckling, 1976)). These might be mitigated if the intermediary's incentives are aligned with the user's interests, for instance if it is a non-profit organisation.

control and do not become constrained by the software vendor's restrictions.

Take the the freedom to study, i.e. the ability to read the software's source code. This is seen by FOSS advocates as necessary for assurance that the code doesn't surreptitiously run any processes that might be counter to the user's interests. Reading the source code can also help aid independent investigation of why software behaves as it does, and discovery of security flaws, without having to rely on the vendor's own activity and reporting.

FOSS also aims to allow independent developers to modify software to suit particular purposes and circumstances. For instance, they might adapt it for use with assistive technology for the sensory-impaired. These modifications may end up in a future version of the original software, or be released separately. This ensures that niche users are not reliant on a single software vendor creating the modifications they need; they can challenge the assumptions and reshape the affordances embedded in the standard product.

Even if the average individual doesn't exercise these freedoms, all users can in theory benefit because of the potential for improved security and more diverse functionality. The general ability to scrutinise and modify without restriction are key elements of a general open approach, which goes beyond software to include content, protocols and data.¹²⁰

The OfP approach could seek a parallel kind of empowerment in relation to processes involving personal data. Individuals – or perhaps third parties acting on their behalf – could scrutinise and modify the ways their personal data is used. Processes involving personal data can be regarded as 'open' in this sense if they are open to scrutiny and modification, to independent evaluation and challenge. Like in the FOSS example, the average individual doesn't need to pro-actively exercise these freedoms in order to benefit from the actions of others who do. A small number of dedicated individuals or representative groups can create positive outcomes on behalf of a wider user base.

While this notion of *open* processing of personal data takes inspiration from the FOSS paradigm, it also maps on to various existing concepts and approaches in privacy and data protection. For instance, the ability to scrutinise how one's data is used could be seen as another form of transparency, albeit individualised. Unlike the approach described in the section above which involves generic, ex-ante data on organisations' general privacy practices, in this case transparency means ex-post, individual-level reports on data use (Hildebrandt, 2013).

Various techniques have been proposed to enable this form of transparency. They generally aim to allow users themselves, or independent third parties acting on their behalf, to access verifiable records of the processing of their personal data. These often make use of cryptographic protocols and decentralised networks of trusted peers (e.g. (Seneviratne & Kagal, 2014b)),

¹²⁰ The emphasis on scrutiny can be seen in, for instance, open data advocacy around government spending data; while the freedom to modify content without infringing copyright is a key motivation behind some of the Creative Commons suite of licenses (this is true for the CC-BY/SA licenses, but not the ND or NC variants). See www.creativecommons.org/licenses

verifiable server logs (Butin et al., 2012), and third party certifiers.¹²¹ These forms of transparency are often advocated as a means to support accountability ((Article 29 Working Party, 2010), (Gellert & Gutwirth, 2012)), and controlling downstream data uses ((Seneviratne & Kagal, 2014b), (Kolovski et al., 2005)).

In addition to these individualised, ex-post transparency mechanisms, there are also many proposals which would allow individuals (or third parties acting on their behalf) to modify, shape or otherwise influence the processing of their personal data. The ability to modify and challenge processing of personal data is described by Danezis and Domingo-Ferrer as 'intervenability' (Danezis & Domingo-Ferrer, 2015). It 'encompasses control by the user, but also control by responsible entities over contractors performing data processing on their behalf.' Intervenability is seen not only as technical but also social, since 'many processes of our democratic society and in particular of the juridical systems contribute to effective intervenability' (Danezis & Domingo-Ferrer, 2015, p. 53).

'Scrutability' is a related concept from the field of computer-human interaction, which combines both transparency and intervenability ((Kay, 1994), (Wasinger et al., 2006)). A scrutable system reveals to the user how it personalises their experience using a profile (or 'user model'). Users can understand and control what goes into their personal user model, how it is maintained and what services it is shared with (Kay & Kummerfeld, 2012).

These examples all involve organisations adopting 'open' approaches to their personal data processing activities, namely by opening them up to scrutiny and relinquishing some control. To this extent, the techniques and tools they advocate can be seen as manifestations of the Openness for Privacy approach. Just as the FOSS paradigm values the capacity for individuals and independent third-parties to study and modify software, the OfP approach values equivalent abilities in the specific context of personal data.

This idea raises some key questions. What incentives might organisations have for opening up their personal data processing to data subjects in this way? Why would individuals want to engage (or enlist intermediaries to do so on their behalf)? How might individuals seek to reconfigure their profiles?

These questions are explored in various ways in the second paper presented below. It focuses on the particular context of consumer profiles in digital marketing, an area of increasing interest in industry. Several new businesses have emerged which offer greater control to individuals over their profiles. They aim to provide a win-win proposition for both business and individuals.

1.4.3 Regulating Privacy with the Open Corporation

There are also ways that the traditional relationships and processes of regulation and governance of privacy could be made more 'open'. For instance, 'open policy making', where multiple stakeholders convene to have

¹²¹ See e.g. (Pearson & Charlesworth, 2009), (Mont, Sharma, & Pearson, 2012), (B. Koops, 2013)

an input into new government policies (for instance, formulating white papers), is increasingly seen as an important part of civil service reform (UK Cabinet Office, 2015).¹²² Where policies have implications for privacy and personal data, open policy making could become an important avenue for addressing challenges. A recent example is the UK government's data sharing initiative, which intended to 'support civil society organisations, independent experts, and government departments to explore the benefits, risks, limitations and governance for sharing personal data within government' (Involve UK, 2014).

Recent developments in regulatory practice suggest that openness could also apply to the relationship between regulators, regulatees and stakeholders. For instance, in *The Open Corporation*, Christine Parker outlines an ideal form of regulation in which organisations are made open or 'permeable' to influence from external stakeholders (Parker, 2002). The approach, called 'meta-regulation', has been studied in various contexts, from food and workplace safety to nanotechnology. It provides a compelling vision as to how a form of openness could define more effective interactions between regulators, regulatees and stakeholders in the context of privacy and data protection.

This possibility is explored in the third paper, which focuses on new requirements for Privacy Impact Assessments (PIAs) in the EU's proposed General Data Protection Regulation (GDPR). It is argued that PIAs can be regarded as an attempt by the European Commission to incorporate aspects of meta-regulation into data protection regulation. This points to a positive opportunity to bring Parker's ideal of the Open Corporation to bear on issues of privacy, data protection and personal data empowerment.

1.4.4 Extending OfP: standards, platforms, collaboration and tools

Open data, open processing and the *open corporation* are the subjects of the three papers comprising this PhD. But they are just three possible interpretations of OfP. There are many other potential avenues for exploration. This section briefly introduces a few more examples, as a way to flesh out the OfP approach, before moving on to the papers themselves. As before, what unites these examples is the use of 'open' principles and processes to achieve the aims of privacy, data protection and personal data empowerment.

1.4.4.1 Open standards and personal data

At the heart of the web and other open technologies are open standards. The open data community, for instance, have sought to standardise the formats and procedures for sharing and re-using data (Berners-Lee, 2006). Open standards and rights may play a parallel role with regards to personal data. If individuals are to re-use their own data for their own purposes, various open standards may be required, including the ability to export one's own data from a system and re-use it in another context, or at least to access it via an open API (Binns, 2013a). This is the rationale behind the principle of data

¹²² For further examples, see <http://www.involve.org.uk/blog/tag/open-policy-making/>

portability, which is seen as an antidote to 'vendor lock-in' which limits competition between internet services.¹²³ This capacity – for individuals to access and re-use their own data, for their own purposes – is a foundation for *personal data empowerment* as introduced above.

Open standards also allow for computation using distributed data sources. Since privacy problems are often the result of data being spread across multiple resources, mechanisms to address those problems may benefit from such standards. For instance, linked data and the semantic web might help in managing distributed privacy problems, like the so-called 'Right to Be Forgotten' (O'Hara, 2012a). In addition to standards for personal data, technical standards and protocols in general can have strong implications for privacy ((DeNardis, 2011), (Winn, 2009)).¹²⁴

1.4.4.2 Open government platforms for privacy

There are also opportunities for data protection and privacy regulators themselves to proactively pursue openness, beyond the open data referred to above. This could include more effectively sharing the results of investigations and enforcement actions against data controllers (Geist, 2012). Lists of addresses for data controllers are another example of the kind of basic information infrastructure that regulators could openly provide; in some jurisdictions where they do, external developers have built applications which use them to help data subjects make subject access requests.¹²⁵ In addition, regulators could provide open software and tools to help organisations manage their obligations, as in the case of the New Zealand privacy commissioner's free 'privacy statement generator' tool.¹²⁶

The desire for more openness from privacy regulators is evidenced by a number requests made by organisations (from both civil society and industry) under freedom of information laws.¹²⁷ The aims of these groups range from identifying business opportunities (to provide privacy and security consulting) to political advocacy. These examples suggest how government agencies might act as 'platforms' rather than simply as providers of services or agents of regulation (along the lines of the 'government as a platform' approach (O'Reilly, 2011)).

123 See e.g. (Bühler, Dewenter, & Haucap, 2006), (Hoofnagle, 2009), (Moura, 2014), (Open Identity Exchange, 2014)

124 For example, DeNardis cites how the IPv6 standard faced the design decision of incorporating a physical address in a virtual internet address, thereby indicating the location of an internet user.

125 See the 'Privacy Inzage Machine' tool developed by Bits of Freedom, a Dutch digital rights advocacy group, available at [<https://pim.bof.nl/>], retrieved September 2015.

126 See [<https://www.privacy.org.nz/further-resources/privacy-statement-generator/>], retrieved September 2015

127 See e.g. Egress Software Technologies, who sought data from the UK Information Commissioner's Office on the number of law firms who were investigated for breaches of the Data Protection Act 1998 in 2014 [<https://www.scl.org/site.aspx?i=ne42176>]. Or MedConfidential, a patient privacy advocacy group, who requested background communications behind the National Health Service's controversial care.data programme [https://www.whatdotheyknow.com/request/caredata_programme_board_papers/]

1.4.4.3 Open collaboration tools

Individuals face a bewildering array of choices regarding their privacy. One solution, intimated above, would be intermediary organisations who can research and manage these decisions on the individual's behalf.¹²⁸ But a collaborative, peer-to-peer network approach described above might also be useful or complementary. A peer-to-peer approach was tried with the Platform for Privacy Preferences (P3P), an initiative started in the late 1990's as a way for web users to indicate their privacy preferences to websites in an automated way (L. F. Cranor, 2013). The architects of P3P anticipated that many users wouldn't have well-formed privacy preferences and might wish to defer to the better judgement of their more informed peers. The system therefore enabled users to share their preferences with others, so that a wider pool of users could benefit from their judgement. This model of delegated decision-making has similarities with online 'delegative democracy' platforms (Kling, Kunegis, Hartmann, Strohmaier, & Staab, 2015).

Despite the eventual decline of P3P, there may still be potential in collaborative peer-to-peer approaches to privacy decisions. *Terms of Service; Didn't Read* is an initiative to crowd-source summaries and ratings of the user agreements and privacy policies of popular websites.¹²⁹ The system combines automated and human processes to scrutinise the small print and flag up salient points, which are aggregated and made available to consumers in an easily digestible summary form (Binns & Matthews, 2014). Other examples include privacy protection tools which block harmful entities on the web according to crowd-sourced blacklists.¹³⁰

Online collaboration tools could also be used by data controllers themselves, to pool their resources to drive compliance and best practice. Online crowdsourcing tools can be used to help organisations explore and understand their obligations. Examples include ThinkData, where organisations pool knowledge on data protection compliance through sharing stories (Morin & Glassey, 2012), and Law Stack Exchange, a question and answer forum for technologists seeking advice on compliance with technology law (including privacy).¹³¹ The premise behind these initiatives is that asking any one individual to manage their own privacy, or expecting any one organisation to be capable of manage their compliance on their own, is simply too demanding. Like creating a 4.9 million page encyclopedia, making informed privacy decisions may only be possible through open collaboration.

1.4.4.4 Open source software for privacy management

Last but not least, perhaps the most obvious way openness can support

¹²⁸ This is also the expected outcome under Coase's analysis.

¹²⁹ See www.tosdr.org

¹³⁰ One example is the web cookie blocking tool Privacy Badger, which maintains a blacklist of domains which can be contributed to by volunteers (see www.eff.org/privacybadger). A related example is blocktogether, a tool that allows twitter users to share lists of abusive users (www.blocktogether.org). A more ambitious system for crowd-sourced privacy threats has been proposed in (Narayanan, 2014).

¹³¹ See [thinkdata.ch] and [<http://law.stackexchange.com/>]

privacy is through FOSS privacy-enhancing technologies, including encrypted communication tools like PGP and OTR, anonymous networks like TOR, and tracking protection browser plugins. Beyond FOSS privacy tools for individuals, there might also be scope for FOSS in helping organisations manage their own compliance. These include tools which help organisations track the provenance of data (Perez & Moreau, 2008), and ontologies to describe the compliance-relevant features of data ((Casellas, Nieto, Meroño, & Roig, 2006), (Kost, Freytag, Kargl, & Kung, 2011)).

1.4.5 Summary of OfP applications

The reader may now be feeling overwhelmed by the variety of ways that principles of openness might be applied to issues of personal data. As mentioned above, the idea of openness is open to many interpretations, and this is no different when it is applied to the challenges of personal data. The aim here is to provide a high-level conceptual framework, which can provide new perspectives and stimulate further research. Having defined this new approach, my aim is not to dogmatically defend it, but to critically assess its merits and shortcomings in various contexts.

The following table summarises the main applications of OfP.

Form of Openness	Example applications to personal data
Open Data	Open data on privacy practices (e.g. ICO register of data controllers)
Open Processing	'Scrutable' user models (Kay & Kummerfeld, 2012) 'Intervenability' in processing (Danezis & Domingo-Ferrer, 2015) Data use logging / auditing (Butin, Chicote, & Métayer, 2012), (Seneviratne & Kagal, 2014a)
Regulation through the Open Corporation	UK Government Data Sharing Policymaking process (Involve UK, 2014) Open corporate regulation of privacy and data protection (Privacy Impact Assessments as 'meta-regulation') (Parker, 2002)
Open Standards	Standardised privacy policies and privacy negotiation (e.g. P3P) Open standards for personal data empowerment (e.g. Midata) Privacy within other standards (e.g. IPv6)
Open Government / Gov. as a Platform	Infrastructure and platforms to help data subjects and data controllers manage rights and obligations (e.g. subject access request tools, privacy statement generator)
Open Collaboration Tools	Crowdsourcing privacy intelligence (e.g. Terms of Service; Didn't Read) Sharing compliance knowledge (e.g. Law Stack Exchange)
Free and Open Source Software for Privacy Management	Individual privacy tools (e.g. GPG, OTR, TOR) Organisation compliance software (Provenance-tracking, compliance ontologies)

Table 1. Main Applications of OfP

1.5 Summary

This introductory section has covered a great deal of ground. It began by noting how the economic, technological, and legal environment has changed in recent decades, giving rise to the current concerns about privacy and data protection, and a set of ethical quandaries relating to new data-driven socio-technical processes. It outlined the complex mixture of disciplinary perspectives and the dividing lines which animate this policy area. It also introduced the notion of personal data empowerment; the potential for people to use their own data for their own purposes.

It then moved on to a discussion of the relationship between openness and privacy. It was argued that the dominant narrative, in which openness and privacy are pitted against each other, belies more nuanced attempts to reconcile the two values under one information rights framework. Within this reconciliatory approach lies an under-explored possibility; that openness might in fact directly support privacy.

While there are already appeals to openness in the privacy and data protection world, they are stuck in a pre-digital age. The current approach to openness about privacy practices is limited, confined mostly to organisations publishing some information about what they do. It wrongly assumes that people have the time, skill and will to read and process such information.

In order to flesh out a more ambitious, alternative role for openness in this context, broader notions of the concept have been appealed to. These stem from FOSS and open data, but potentially extend much further, to include notions of open collaboration, regulation and governance. These varieties of openness are not without their problems, but they provide ample inspiration for potential applications to privacy.

These potential applications are explored in the remaining papers comprising this PhD 'by publication'. Therefore, while these papers can be read as stand-alone pieces, they are tied together by the concept of Openness for Privacy. This concept will be periodically returned to in the short prologues / epilogues in between each paper, and finally in the conclusion.

Part 2: Open Data for Privacy

The Who, What and Why: An Analysis of Personal Data Transparency
Notices in the UK

Abstract:

Data protection laws require organisations to be transparent about how they use personal data. This article explores the potential of machine-readable privacy notices to address this transparency challenge. We analyse a large source of open data comprised of semi-structured privacy notifications from hundreds of thousands of organisations in the UK, to investigate the reasons for data collection, the types of personal data collected and from whom, and the types of recipients who have access to the data. We analyse three specific sectors in detail; health, finance, and data brokerage. Finally, we draw recommendations for possible future applications of open data to privacy policies and transparency notices.

Keywords: privacy, data protection, personal data, transparency, web, open data

2.1 Introduction

The use of personal data has become one of the most important issues of the digital age. Regulators, policy-makers and consumer advocates have long argued for transparency from the public and private entities who gather and use this data. Transparency is a core principle at the heart of several foundational privacy and data protection frameworks, and continues to inform new regulations and international instruments.¹³² Although transparency alone may be insufficient, it is seen as a necessary precondition to achieving privacy goals (Gutwirth & DeHert, 2006). In theory, it helps regulators, advocates, researchers and others to monitor and analyse privacy-related practices, guiding their strategy and further action. Ultimately, transparency also aims to empower privacy-conscious individuals, whether directly or through an intermediary, to make more informed choices about whom to trust with their data (Egelman & Tsai, 2006). Transparency is therefore a prerequisite for a functioning 'market for privacy', where companies compete on their privacy credentials in order to attract privacy-sensitive consumers (Bonneau & Preibusch, 2010).

But despite broad support for the principle and purpose of transparency, there has been less agreement on how best to achieve it. Effectively and efficiently recording and publishing what organisations do with personal data, and why, has proven difficult. A significant body of research has built up around the design and testing of improved transparency mechanisms.¹³³ At the same time, many studies attempt to use existing mechanisms – which principally come in the form of 'privacy notices' – to analyse the policies and practices of organisations regarding personal data.¹³⁴ But these studies face significant barriers which limit their potential depth and scope.

This paper addresses both the design of privacy notice transparency systems, and the analysis of their content. We present an analysis of a previously unstudied source of standardised privacy notifications, the UK Register of Data Controllers¹³⁵, which contains notifications made to the UK Information Commissioners Office (ICO) by around 350,000 organisations over an 18 month period. Our aims are to generate a broad overview of the landscape of personal data use by UK organisations, and bring new evidence to bear on some particular topics of pressing public concern, namely the collection and use of personal data by health providers, financial services, and data brokers. The results of the analysis are followed by considerations and recommendations for the creation of such transparency systems.

2.2 Background

In order to provide context, the remainder of this introduction briefly outlines the history and current status of privacy notice transparency systems, some developments which have been proposed, as well as related

¹³² See, for instance, the 'Openness Principle' in (OECD, 1980), and its evolution over the following 30 years (OECD, 2011)

¹³³ For an overview, see (Binns, 2014b)

¹³⁴ E.g. (M. J. Culnan & Armstrong, 1999)

¹³⁵ Available to search at [<http://ico.org.uk/esdwebpages/search>]

examples of mass analysis of privacy notices.

2.2.1 Existing transparency mechanisms

2.2.1.1 Privacy Notices

Transparency has so far in practice been implemented through the use of notices, often published by organisations as 'privacy policies' to be included alongside their terms-of-service and end-user license agreements. The common practice of producing lengthy and legalistic documents means that few consumers read or understand these policies. A study of online privacy policies estimates that the average U.S. Internet user would have to spend 244 hours per year reading the privacy policies of all the websites they visit during that year, suggesting that the cost of being informed may well be too high for any individual (McDonald & Cranor, 2008). The length of these documents is also a barrier to academic research and regulator investigations into organisations' stated privacy practices. Several commercial and non-profit organisations have attempted similar work, classifying and rating privacy policies on behalf of consumers.¹³⁶ But manually parsing the mass of policies is a time-consuming task, which has limited the coverage and effectiveness of these efforts.

2.2.1.2 Public Registers

While privacy notices have received most of the attention in discussions about transparency in this context, certain jurisdictions also maintain an alternative scheme of public registers¹³⁷. This approach involves mandatory disclosures by organisations to a regulatory authority, detailing what data they collect, who they share it with, and why. This information is then gathered in a national register of organisations' personal data practices, which is made available to the public. This system – implemented in most EU member states – is generally held in low regard, with the EU Commission describing it as an 'unnecessary administrative requirement'.¹³⁸ At the time of writing, only eight of those member states with national registers appear to have public websites from which they can be searched, which are of varying quality and usability. Anecdotal evidence suggests that public awareness of these public registers is limited to a small number of data protection specialists, and those who do attempt to use them for transparency purposes find they have low usability and are inconvenient.¹³⁹ Given their perceived lack of utility, it is unsurprising that the new draft proposal for a General Data Protection Regulation (henceforth 'GDPR')¹⁴⁰

¹³⁶ See (Binns, 2014b) for an overview.

¹³⁷ Most E.U. member states have such registers, but exceptions include Germany and Sweden.

¹³⁸ 'Commission proposes a comprehensive reform of data protection rules to increase users' control of their data and to cut costs for businesses' – European Commission, press release available from http://europa.eu/rapid/press-release_IP-12-46_en.htm

¹³⁹ See documented complaints made by an individual attempting to use the UK's online register:

https://www.whatdotheyknow.com/request/non_notification_team

¹⁴⁰ Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on the protection of individuals with regard

dropped mention of such registers altogether.

2.2.2 Continued emphasis on transparency

Both privacy notices and national registers fall short of the kind of transparency system that would be required for meaningful oversight, monitoring and analysis by regulators and researchers, let alone the average consumer. But despite the problems with the existing measures, policy-makers continue to emphasise the need for transparency more than ever. At their 2013 international meeting, privacy and data protection commissioners from around the world released a statement on transparency, recognising that:

“Effective communication of an organisation’s policies and practices with respect to personal data is essential to allow individuals to make informed decisions about how their personal data will be used and to take steps to protect their privacy and enforce their rights.”¹⁴¹

With the desire for transparency greater than ever, there is continued enthusiasm for new approaches to effectively record and publish organisations' policies.

2.2.3 Standardised Formats

Some have propose standardised, short, simplified and / or graphical notices as a solution to this problem. The U.S. Department of Commerce has proposed guidelines for short form notices to describe third party data sharing (National Telecommunications and Information Administration (NTIA), 2013). Similarly, a 'nutrition label' style approach has been discussed in the context of the GDPR, which would complement traditional notices with standardised and required fields represented by a set of common simple visual icons that would become familiar to consumers over time¹⁴². This approach has also been explored by a number of non-profits and consumer-oriented companies,¹⁴³ and more recently was the subject of an initiative by the ICO to develop 'privacy seals'¹⁴⁴. Similar initiatives aim to encode privacy policies into machine-readable XML formats¹⁴⁵. A

to the processing of personal data and on the free movement of such data (General Data Protection Regulation)

141 Recorded in (35th International Conference of Data Protection and Privacy Commissioners, 2013)

142 See article 13(a) of the Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation).

143 See for instance, Mozilla Icons project (https://wiki.mozilla.org/Drumbeat/Challenges/Privacy_Icons), ToS-DR.com, PrivacyScore.com

144 “ICO to launch privacy seals scheme 'within the year'”, DataGuidance 27/03/2014
http://www.dataguidance.com/dataguidance_privacy_this_week.asp?id=2258

145 See proposals from TRUSTe, a web certification body (<http://www.truste.com/blog/2010/09/14/more-on-the-problem-with->

common theme in all of these proposals is the idea that by standardising and digitising organisations' disclosures of how they collect and use personal data, this information can be aggregated, accessed, compared and analysed en mass by regulators, consumer advocacy groups, intermediaries, or indeed, individuals themselves.

2.2.2 Prior Art

As has been noted elsewhere (L. Cranor, 2012), this is not an entirely new idea. The current policy proposals have strong parallels with ambitious efforts in previous decades to create large-scale systems for the transparent use of personal data. Amidst enthusiasm for new measures, there is a danger of reinventing old (failed) solutions, unless we learn from past attempts. These previous initiatives, their promises and failures, could be instructive in setting the context and guiding the development of new transparency systems.

2.2.2.1 Platform for Privacy Preferences

Perhaps the most significant long-term effort in this vein came from a series of initiatives which began with the Platform for Privacy Preferences (P3P)¹⁴⁶ in the mid-1990's. Early proponents described a system whereby privacy policies could be encoded as structured data. This data could be 'understood' by web browsers and other software agents, which could then automatically negotiate with websites on behalf of users according to their privacy preferences – ultimately creating a market for privacy. By the time the World Wide Web Consortium (W3C) specification for P3P was approved in 2002, the negotiation features had been dropped, but a standard for rendering privacy policies in machine-readable XML format remained. The standard had a number of early adopters including news websites, search engines, ad networks, retailers, telecommunications companies and government agencies (L. F. Cranor, 2013). The hope within the web standards community, in particular amongst proponents of the 'semantic web'¹⁴⁷, was that a significant proportion of organisations would independently adopt the standard, thus creating a decentralised database of organisation's privacy practices. If successful, such a system could be intelligently queried and analysed *en mass*, thus helping the activities of regulators, intermediaries and consumers.

2.2.2.2 Collaboration with regulators

Perhaps inspired by this vision, the standard was spurred on by regulators in the U.S., principally the FTC. The standard was initially envisioned as a framework for consumer-focused tools, but the FTC also noted the potential of P3P for use in their own investigations and enforcement actions. In 2001, the FTC incorporated P3P data into their annually commissioned surveys of

p3p/), and the Internet Advertising Bureau's CLEAR Ad Notice project (<http://www.iab.net/clear>).

¹⁴⁶ <http://www.w3.org/P3P/>

¹⁴⁷ The semantic web vision is to turn the human-readable content of the existing world wide web into a machine-readable 'consistent, logical web of data' (Berners-Lee, 2004). For an example application of P3P in the semantic web, see (Gandon & Sadeh, 2003).

website privacy policies (Milne & Culnan, 2002), which were conducted in order to investigate organisations' adherence to the FTC's 'Fair Information Practices'. One of the policy recommendations arising out of these studies was to encourage businesses to adopt the emerging standard for their websites. This would make future longitudinal analysis of privacy practices more effective and comprehensive due to the potential for automated analysis. The gradual adoption of this technology by websites in the following years did result in such work. The first detailed and large-scale analysis of the policies of P3P-enabled websites was subsequently conducted in 2003. It investigated the types of data collected, the uses to which it was put, and the types of recipients the data is shared with (Byers et al., 2003).

2.2.2.3. A standard in decline

Unfortunately, further studies like this were hampered by the decline of the standard. When a modified version of the (Byers et al., 2003) study was repeated in 2006, it was found that the proportion of P3P-enabled policies containing errors had increased. Despite evidence of their increased usability (P. Kelley, Cesca, Bresee, & Cranor, 2010), and backing from regulators, the use of standardised P3P privacy notices began to decline. A 2007 study indicated that the level of P3P adoption in 2005 was low (8.4%), and showed that adoption had remained stagnant since 2003 (Beatty, Reay, Dick, & Miller, 2007). Development of the standard was permanently suspended that year, after the W3C failed to reach consensus on a second version. By 2010, P3P 'compact policies' (shortened versions of full P3P-enabled privacy policies) were even found being used to mislead rather than inform users (P. Leon & Cranor, 2010). As of April 2014, support for the standard has been dropped by all the major web browsers apart from Microsoft Internet Explorer¹⁴⁸. Suggested reasons for its failure include: that it was too complex for websites to translate their privacy policy into the P3P format (A. Schwartz, 2009); that effective user interfaces were too difficult to design (Brown & Marsden 2013, p. 54); and that it had insufficient support from privacy advocates, who were concerned it would become a replacement for existing, more enforceable rights (Electronic Privacy Information Center, 2000).

Other policy languages have been designed to supersede P3P, but none have achieved significant adoption as yet.¹⁴⁹ They may face a 'network effects' problem, in that the positive effects of standardisation only emerge once a significant portion of organisations/websites have adopted the standard (Tsai, Egelman, Cranor, & Acquisti, 2010). Therefore, the initiatives struggle to get off the ground as their full benefits are hard to demonstrate. Similar studies of privacy policies by academics and regulators have

148 IE blocks third-party browser cookies by default if they do not have P3P policies (see <http://www.techrepublic.com/blog/software-engineer/craft-a-p3p-policy-to-make-ie-behave/>)

149 See, for instance the Primelife Policy Language (Vimercati, Paraboschi, & Pedrini, 2009). Also the 'Do Not Track' standard – in which a preference/policy regarding online 'tracking' can be communicated between a client and a server – can be seen as a (minimally expressive) descendent of the P3P standard (see [<http://www.donottrack.us/>]).

continued in the absence of P3P or other standards,¹⁵⁰ but their scale and reach is limited by the fact that policies must be parsed manually before any analysis can be done.

2.2.2.4 Development of Public Registers

Meanwhile, the alternative transparency system of public registers has received far less attention than P3P and its various relatives. Perhaps the earliest reference to the public register model in international privacy and data protection frameworks can be found in the 1980 OECD privacy guidelines, in the detailed comments elaborating on the 'Openness Principle'.¹⁵¹ The guidelines note that openness is a pre-requisite for individuals to exercise their right to access and challenge personal data. One of the suggested means to achieve such openness is through the 'publication in official registers of descriptions of activities concerned with the processing of personal data'.¹⁵² The OECD guidelines formed the basis for many subsequent national privacy and data protection regulations and frameworks, with the result that requirements for national public registers are in place in many countries, particularly in the European Union (which in turn has had a significant influence on the development of data protection laws elsewhere ((Greenleaf, 2012b); (Birnhack, 2008)).¹⁵³

2.2.2.5 Similarities between P3P and public registers

The idea of a centralised public register of organisations' privacy notifications is comparable to that of a decentralised, machine-readable corpus of privacy notices. These systems evolved separately, developed by different communities, yet there are similarities in their original visions. Both aim to be a comprehensive resource of standardised privacy notifications. Their initial implementations certainly differed, with P3P conceived as a decentralised, data-driven system from the outset, and the public register as a highly centralised, analogue resource, conceived of before personal computing became widespread. But in more recent years, many national public registers have been published online¹⁵⁴, and in some cases made available as machine-readable open data – in a similar format to

¹⁵⁰ See, for example, the Global Privacy Enforcement Network's 'privacy sweep' investigation of website privacy policies, available at [https://www.priv.gc.ca/media/nr-c/2013/bg_130813_e.asp]

¹⁵¹ See the OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (OECD, 1980)

¹⁵² Ibid, 'Paragraph 12: Openness Principle'

¹⁵³ There are similarities between public register schemes in operation in Europe and US proposals for registers of data brokers and their activities. The FTC have encouraged “creating a centralized website where resellers would identify themselves and describe how they collect and use consumer data, and the access rights and other choices that consumers have” (United States Government Accountability Office, 2013). A similar public register system covering the use of personal data by government agencies was part of the U.S. Privacy Act of 1974, but was also later criticised for being under-used by ordinary citizens (U.S. White House Office of Management and Budget, 1983)

¹⁵⁴ As well as the UK register which is the subject of this paper, other jurisdictions with online registers include Austria, Belgium, Czech Republic, Ireland, Poland, Portugal, Spain and Serbia.

the P3P standard¹⁵⁵.

In this form, public registers arguably come closer to the original, semantic web vision for P3P than the P3P initiative itself ever did. They contain highly standardised, complete, machine-readable privacy notices from a wide range of organisations. Indeed, the number of organisations contained within the UK register alone (~350,000) is comparatively far larger than the number of organisations with P3P-enabled privacy policies identified in any previous study.¹⁵⁶ In addition, because inclusion in a public register is generally mandatory and enforced by law, the contents of a register are less likely to be biased towards those organisations who would voluntarily adopt a given standard (something which was likely to be the case for P3P).

Noting their similarities - whilst being mindful of their differences - it is possible to draw parallels between the public register data and the corpus of privacy policy data that is the subject of prior studies. Analysis of the register data can therefore be seen as a continuation of the extensive body of existing research into organisation's privacy policies, with the advantages of automation and magnitude (which is lacking in previous studies predicated on manually parsed policies), and completeness (which the P3P studies lack). At the same time, any design insights derived from this analysis are likely to be highly applicable to the various proposals for transparency systems mentioned above.

2.2.3 Quantifying Privacy Practices

As well as the aforementioned FTC-commissioned research, numerous other studies have aggregated and manually parsed privacy notices to derive quantitative insights into organisational policies and practices regarding personal data. Such studies usually aim to identify trends in organisations' stated practices, and/or evaluate the notification/disclosure process itself. Since a prime motivation for studying (and regulating) the use of personal data is to further the interests of data subjects and society at large, this research is often driven, at least partly, by particular public concerns.

Our general analysis extends this existing research by providing a broad overview of the reported uses of personal data across a comprehensive range of sectors and uses. This is complemented by in-depth analyses of three specific uses – trading of personal data, financial services and health provision – each of which have been the subject of sustained interest among researchers and in the public eye. These include:

2.2.3.1 Trading of personal data:

Organisations have come under increasing scrutiny over the buying and selling of personal data in recent years. The 'data broker' industry – where personal data is collected and re-sold – has been the subject of investigations

¹⁵⁵ Both the UK and Poland have stored their register data as machine-readable XML, with fields that correspond almost exactly to some of the standard P3P fields.

¹⁵⁶ The largest number of P3P-enabled websites found in any of the prior studies identified in our literature search was 14,720 (L. Cranor, Egelman, & Sheng, 2008)

by regulators and the media.¹⁵⁷ This is also a theme arising in multiple studies of consumer concerns, where it is frequently expressed in terms of unknown 'third parties' with whom data may be shared. In a qualitative study of UK citizens, it was found that an 'unspecified reference to 'third parties' unsettled participants and helped feed concerns that after the transaction there would be a number of uses of their information over which they could have no control' (Bradwell, 2010). An E.U.-wide study found that of the 54% of citizens who were aware of organisations selling their personal information to third parties, only 35% found the practice ethically acceptable (Brockdorff & Appleby-arnold, 2013) In the following analysis, we examine the extent and nature of this practice as compared to other practices, using the pre-defined register category of 'Trading / Sharing in Personal Information'.

2.2.3.2 Financial Services

Previous research has examined the extent to which organisations in a particular industry or context actually differ in their practices, in order to assess whether there is the possibility of meaningful consumer choice and a differentiated market for privacy (Bonneau & Preibusch, 2010). Cranor et al took advantage of a widely implemented standard for privacy notices adopted by 3,422 US financial institutions (L. F. Cranor et al., 2013). In this rare instance of a relatively successfully adopted standard notification format, large-scale empirical analysis of privacy practices was possible. The authors found significant variety in bank's practices, as well as some evidence of self-contradiction and non-compliance by some institutions. Using data on UK banks and other organisations providing financial services, we similarly investigate whether there is homogeneity in practices, or the possibility of meaningful choice for UK consumers.

2.2.3.3 Health services

In a qualitative survey of attitudes towards privacy and health information, national health service patients in the UK regarded health data as a special category worthy of particular concern (Wellcome Trust, 2013), a finding that is supported in earlier E.U.-wide quantitative studies (Brockdorff & Appleby-arnold, 2013). In February 2014, UK government proposals to share medical data gathered from general medical practitioners under the *care.data* scheme raised controversy and debate about the risks of sharing health data.¹⁵⁸ We present a profile of data use by organisations engaged in health administration and services, in order to provide context to concerns about these kinds of practices.

2.2.3.4 Comprehensive samples for comparison

These specific analyses are presented alongside the analysis of data collection in general (i.e. for all purposes) for comparison. This shows, for

¹⁵⁷ For instance, the Federal Trade Commission - see (Federal Trade Commission, 2013), and the Wall Street Journal's 'What They Know' series, Retrieved from [online.wsj.com/public/page/what-they-know-digital-privacy.html], September 2015

¹⁵⁸ See 'Care.data: How did it go so wrong?', BBC News, 19 February 2014, [<http://www.bbc.co.uk/news/health-26259101>]

instance, whether different kinds of data are more often collected, or whether certain kinds of data subjects are more often involved, in the context of health, finance, or trading, than in the general case. Previous studies have been unable to provide such a comparison, because they are generally limited to particular sectors (e.g. financial companies or social networking sites), with sample sizes that are both small and unrepresentative. By presenting a comprehensive, representative, cross-industry overview of organisations privacy practices, we aim to situate a particular sector or practice in its broader context.

2.3 Data Source and Methodology

The source of the data in this analysis is the United Kingdom Information Commissioner's Office (ICO) public register of data controllers. Data controllers are defined as 'a person who (either alone or jointly or in common with other persons) determines the purposes for which and the manner in which any personal data are, or are to be, processed' (UK Data Protection Act 1998 [DPA], s.1)¹⁵⁹. 'Data processor' is defined as 'any person (other than an employee of the data controller) who processes the data on behalf of the data controller' (DPA 1998, s.1). Personal data is defined as 'data which relate to a living individual who can be identified (a) from those data, or (b) from those data and other information which is in the possession of, or is likely to come into the possession of, the data controller', while 'data subject' is defined as 'an individual who is the subject of personal data' (DPA 1998, s.1).

2.3.1 Notification Requirements

The DPA states that data controllers must contact their national supervisory authority, notifying them of their name and address, purposes of processing, the categories of data types and subjects to whom they relate, recipients to whom the data may be disclosed, and proposed transfers of the data to third countries (DPA 1998, s.16). Furthermore, these notifications should be compiled into a register of data controllers, made available for inspection by any person (DPA 1998, s.19(6)(a)). In the UK, this register is made available to the public to search on the ICO's website, and a regularly updated version of the whole register is available upon request under an Open Government License in a re-usable, machine-readable format. The latter, gathered over an 18 month period, forms the basis of the following analysis.

2.3.2 Data structure, extraction and selection

The register is made available in a semi-structured standard data format (XML), and contains fields corresponding to a)-e) of the notification requirements in the DPA (section 16.1). The ICO provide a set of standard defined purpose types, subjects, classes, and recipients. Data controllers may also describe their activity in their own terms if it is not captured by these standard definitions. In all, three copies of the register were used, from September 2011, September 2012, and March 2013 (unfortunately, data

¹⁵⁹ Note that 'person' in this context means 'legal person', and as such could be an organisation or a natural person.

from September 2013 is unusable for the purposes of this study due to changes made by the ICO in April that year).

DPA required information (section 16.1)	Human-readable register field(s)	XML Tag	P3P equivalent
(a) his name and address	Data controller name and details	<DATA_CTRL_NAME> <DATA_CTRL_DETAIL>	<ENTITY>
(c) a description of the personal data being or to be processed by or on behalf of the data controller and of the category or categories of data subject to which they relate,	Class, Subject	<CLASS> <SUBJECT>	<DATA-GROUP>
(d) a description of the purpose or purposes for which the data are being or are to be processed	Purpose	<PURPOSE>	<PURPOSE>
(e) a description of any recipient or recipients to whom the data controller intends or may wish to disclose the data	Recipient	<RECIPIENT>	<RECIPIENT>

Table 2. Comparison of DPA, register and P3P fields

Because the original XML file was too large to query directly, we first parsed the data using SAX, an event-based sequential access parser API for XML¹⁶⁰. It was then restructured as an SQL database, composed of separate tables for each of the human-readable register fields, along with unique identifiers for each data controller and each 'purpose' instance. This database was then queried to extract relevant portions for further analysis.

2.3.3 Analysis

The first stage of analysis was to measure the occurrence of different classes in the dataset. Given that the data is exclusively concerned with categories (i.e. nominal data), in order to subject it to quantitative analysis we measured the occurrence of certain classes. Quantifying the remaining categories (purpose, subject, class, and recipient) reveals the extent to which certain arrangements and relationships exist within organisations. We can then derive conclusions about the extent and nature of data sharing between individuals, data controllers, and third parties, and the prevalence of certain types of personal data, data subjects, and recipients.

Categories with similar definitions (e.g. 'Marketing' and 'Marketing, Advertising and Public relations') were aggregated. Any category with less than 50 instances that could not be meaningfully aggregated into a more prevalent category was discarded. By conducting the same operations on each of the datasets (from 2011, 2012, and 2013), we measure differences in practices over time. The three specific analyses followed a similar procedure with some differences. For the analysis of personal data use in financial services, two subsets of the data were isolated and analysed. One large subset consisted of instances where data was collected and used for the

¹⁶⁰ See www.saxproject.org

purposes of providing financial services and advice – this included a wide variety of different organisations, not just banks (37,436 distinct organisations in total). A second, smaller subset consisted of 98 data controllers whom we independently (manually) classified as 'retail banks'. These samples were then analysed to establish which classes of data were used (e.g. 'Personal Details' or 'Employment'), and which categories of recipients had access to this data (e.g. 'Credit Reference Agencies' or 'Regulators').

2.4 Results

We found steady growth in overall data collection, the types of data involved, and the types of entities who have access to the data. Each of these fields exhibit a power law distribution with a few very common categories accounting for the majority of the total. The following figures present the general and specific cases side-by-side for ease of comparison.

The total number of data controllers averaged 358,558 across the time period studied, growing by 6.5%. The number of purposes (which could also be understood as the *total number of distinct reasons for which data is used*) exhibited a similar level of growth of 6.3%. The average number of purposes per controller stayed consistent across the period at an average of 3.72, indicating that while the number of organisations classified as 'data controllers' is increasing, the average number of different types of uses of data per controller remains the same. The standard deviation in number of uses is 1.8, indicating that most data controllers are close to this average. The average number of distinct types of subject, class, and recipients per purpose provide a benchmark for analysis of specific sectors and practices, where averages and spread may differ.

	Average	Standard Deviation
Purposes per data controller	3.7	1.8
Classes per purpose	5.7	2.8
Subjects per purpose	7.5	3.5
Recipients per purpose	3.3	1.7

Table 3. Average Purposes, Classes, Subjects and Recipients

The remaining general analysis is broken down by the five fields of 'Purpose' (i.e. why data is collected / used), 'Subject' (who the data is about), 'Class' (what categories of data are collected / used), and 'Recipient' (who is given access to the data). In addition to the total number of entries per category within a field, the prevalence of each purpose category can be calculated in relation to the total number of 'data controller' instances in the entire register, indicating the proportion of organisations engaging in that practice. Similarly classes, subjects, and recipients are expressed as a percentage of the total number of uses (or 'purposes') in the register. This provides a more natural measure, expressing how often a given category appears as a proportion of the total number of data controllers or uses (figures 1-5 express this as percentages).

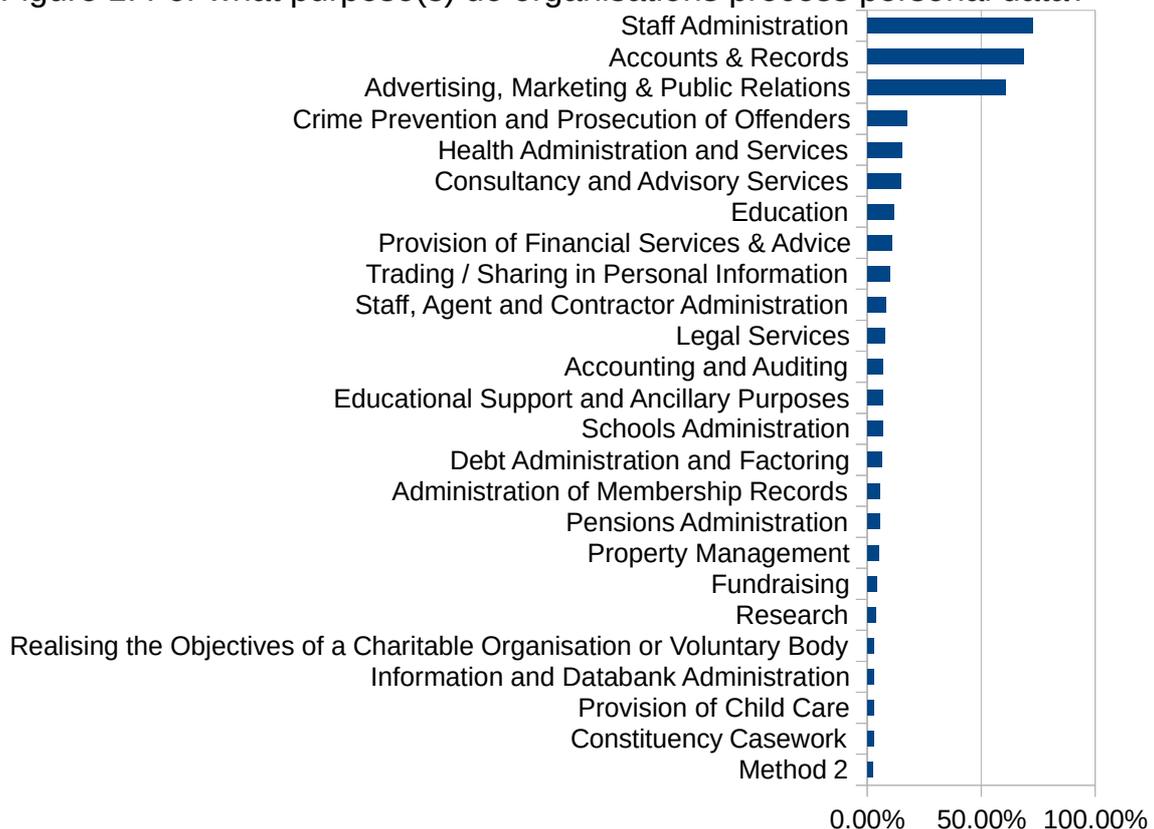
Across all fields, the registrar-defined standard descriptions (i.e. those which are explicitly defined in the ICO's notification handbook given to registrees) featured more heavily than registree-defined descriptions (i.e. those invented

by data controllers themselves). The distribution of entries for each of the categories within a field tended to follow a power law distribution, with a few very prominent categories having a high number of entries, and a 'long tail' of more obscure (mostly registree-defined) categories.

2.4.1 Why is data being processed?

The three most common categories in the 'Purpose' field (namely 'Staff Administration', 'Accounts & Records', and 'Advertising, Marketing & Public Relations') accounted for 54% of the total on average across the period, while the bottom 14 categories accounted for just 13%. Changes between the number of entries for a given purpose category were measured in the 18 month period, with a mean growth of 6% across all categories. While the top 5 categories ('Staff Administration', 'Accounts & Records', 'Advertising, Marketing & Public Relations', 'Crime Prevention and Prosecution of Offenders', and 'Health Administration and Services') grew between 5-10%, the most significant growth was found in more obscure, registree-defined categories such as 'Provision of Childcare' and 'Provision of Investment Management and Advice'. However, this apparently large

Figure 1. For what purpose(s) do organisations process personal data?



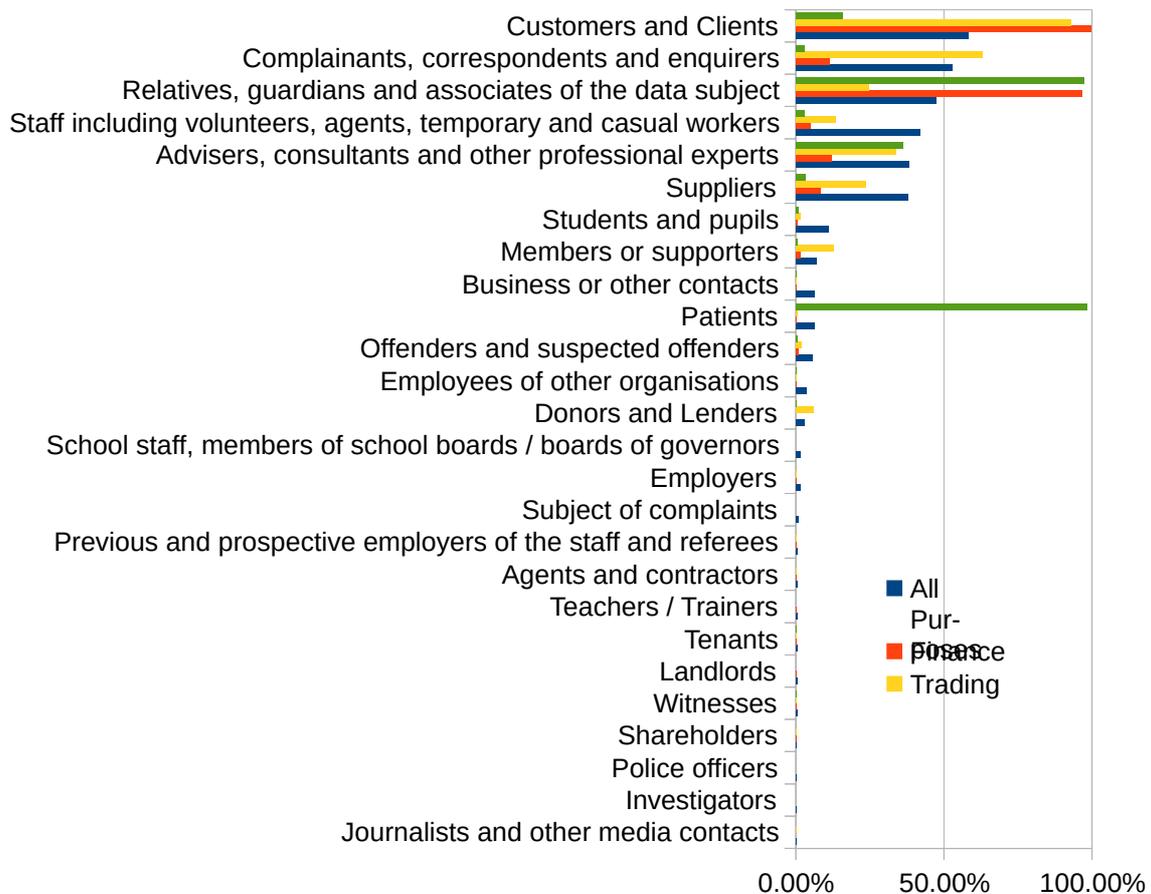
change is likely to be amplified due to the relatively small size of the obscure categories as a proportion of the total.

The three purposes which have been selected for further analysis account for a significant minority of all uses. The use of personal data for 'Health administration and services' was listed by 15% of all data controllers on average across the time period (the fifth most common purpose). 'Provision of financial services and advice' and 'Trading / Sharing in Personal

Information' were both listed by around 10% of controllers, and were (respectively) the eighth and ninth most commonly listed uses of data.

2.4.2 Who is the data about?

The five most common types of data subjects accounted for 85% of the entire field, while the 14 least common accounted for just 4%. The growth for all types of subject was similar to the overall growth in purposes (6.5%). The two categories with the biggest growth were 'Subjects of complaints' and 'Landlords'. In comparing health data to the general case, we find, unsurprisingly, that personal data is far more likely to be about patients and relatives, and far less likely to be about customers and complainants. The



trading of personal data appears to mirror this in reverse, involving relatively few patients and relatively many customers.

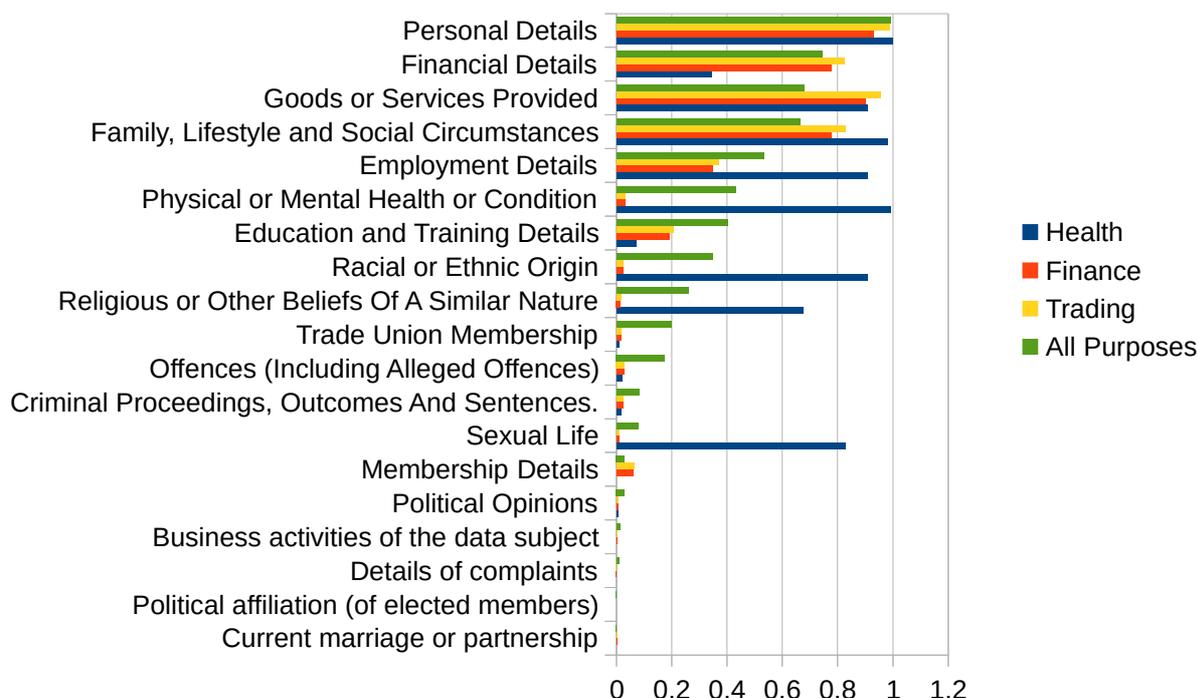
2.4.3 What kind of personal data is used?

The registrar-defined data classes constituted the majority of the categories in the 'class' field, with only five

registree-defined classes achieving more than 50 entries. In cases where data is used for healthcare purposes, this often includes sensitive data, for

What kind of personal data is used?

Data class types as a percentage of health, finance, trading and all uses



instance about an individuals sexual life, which is collected in over 80% of

cases compared to just 7.7% across all purposes. The classes of data collected for the purposes of trading include personal details (99%), goods provided (96%), family and lifestyle (83%), and financial details (82%).

These are the kinds of personal information one might expect to be traded, given that they may pertain to commercially useful knowledge like the kinds of goods people might buy or their creditworthiness. However, a small proportion of instances where personal data was traded / shared involved more sensitive kinds of personal data. The ICO lists 8 classes of 'sensitive' data, all of which were 'traded' in the following percentage of cases:

- Physical or Mental Health or Condition (10%)
- Racial or Ethnic Origin (8%)
- Religious or Other Beliefs Of A Similar Nature (6%)
- Trade Union Membership (4.6%)
- Offences (Including Alleged Offences) (4%)
- Criminal Proceedings, Outcomes And Sentences (1.9%)
- Sexual Life (1.8%)

- Political Opinions (0.7%)

Amongst both retail banks and providers of financial services more generally, the use of personal and financial details, and information about goods and services provided, is almost ubiquitous – at least 97% of entries stated using this information. However, there was more variation in practice concerning other types of data. For example, a quarter of retail banks did not list 'Employment Details', and only half listed 'Education and Training Details'.

Growth across all categories was uniform at around +5%, with the exception of 'Details of complaints' reaching a high of +18% growth.

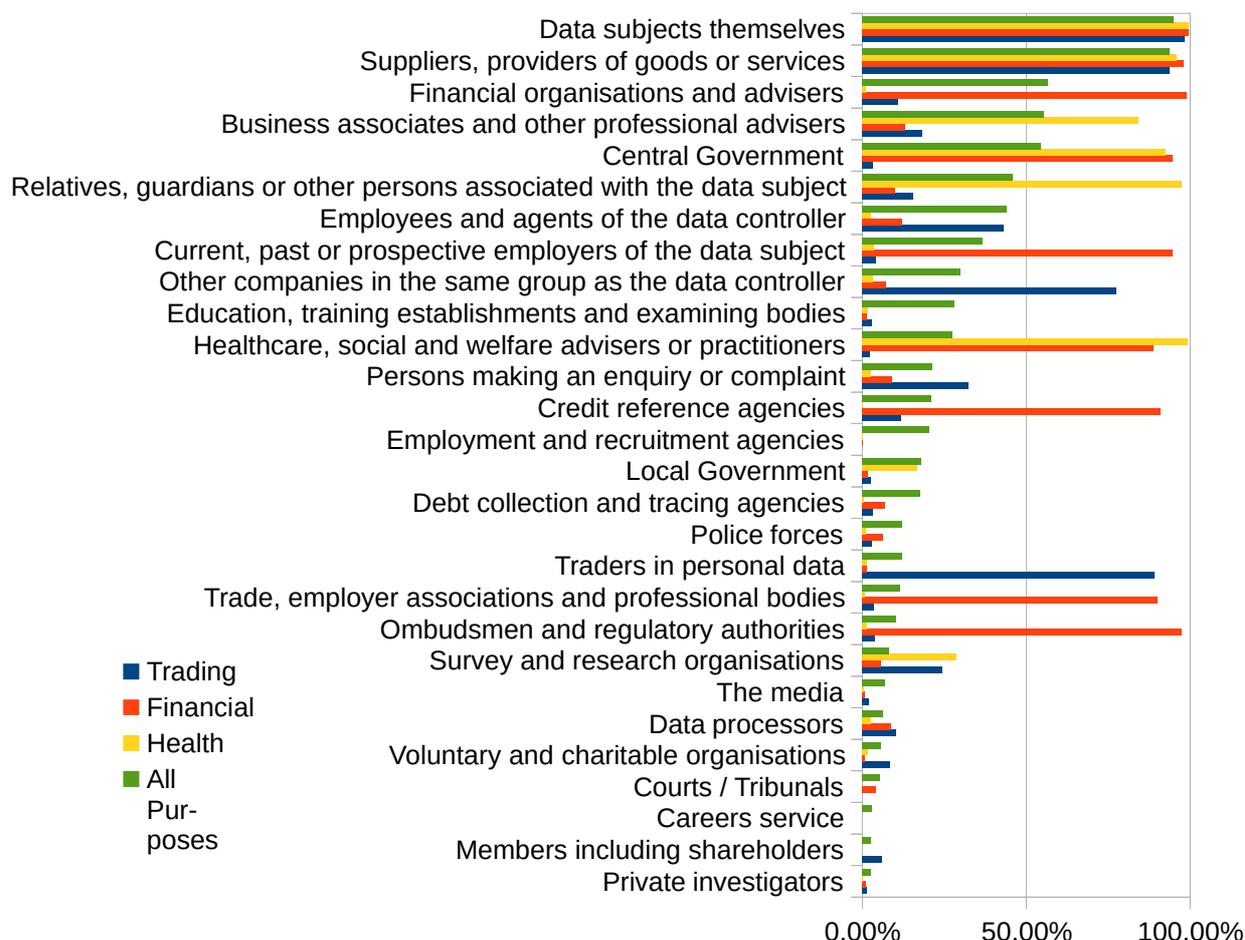
2.4.4 Who has access to the data?

By far the most common potential recipient of personal data is the data subject themselves; in the vast majority (92%) of cases, the data subject themselves is given access. This is probably due to the fact that under UK data protection law, in most cases, data subjects have the right to request a copy of data held about them (exceptions apply in some cases such as criminal investigations). Average growth across all categories was 5.8%.

Overall, in situations where data controllers were 'trading / sharing in personal data', the average number of subjects is 2.7 – much less than the general average of 7.5. However, the average number of recipients was 5.7 – higher than the general average of 3.3. This indicates that when personal data is traded or shared, it is likely to involve only relatively select types of data subject, but the data will then likely be shared with a broader than average range of recipients. These include 'Suppliers, providers of goods and services' (96%) and unsurprisingly, 'Traders in Personal Data' (90%).

Interestingly, the aforementioned classes of 'sensitive' personal data were also being traded/shared with a wide range of recipients. We further investigated the use of sensitive data classes in trading/sharing, finding 134 organisations who state that they trade/share data about individuals' political opinions with 'credit reference agencies'. Data about individual's sexual lives is reportedly traded/shared with 'Traders in personal data' by 226 organisations. 'Trade, employer associations and professional bodies' reportedly receive data about individuals' trade union membership from 288 organisations, and their racial or ethnic origin from 182 organisations.

In the case of access to data collected for provision of financial services, the kinds of entities who have access to this data (i.e. those listed as 'recipients') exhibited a similar pattern. Both general providers of financial services



(displayed in figure 7) and retail banks in particular almost always gave access to a certain familiar list of entities, such as 'Data Subjects themselves' (as is normally required by law), 'Employees of the data controller', and 'Suppliers and providers of services'. However, there were also some differences between the practices of retail banks and general financial service providers. For instance, while 72% of the former shared this data with 'Data Processors', only 9% of the latter did so. Similarly, giving 'traders in personal data' access to the data was more prevalent amongst retail banks than financial service providers (22% versus 1% respectively). Perhaps surprisingly, this trend appears to reverse when it comes to sharing data with credit referencing agencies, where only 52% of banks share, compared to 90% of financial services providers generally.

2.5. Discussion

2.5.1 Growth in data controllers

The general analysis here supports the perhaps unsurprising hypothesis that the use of personal data is increasing, in so far as the total number of entries

in all fields is growing. However, this conclusion should be accompanied with the following considerations. First, since entries in the register describe certain data collection, usage and sharing arrangements, the existence of more entries should not be confused with other measures of data storage and use, such as data points, number of database queries, or volume in bits. Second, the growth rate seems to be driven by new data controllers, rather than an increase in the overall counts of purposes, data classes, subjects, or recipients per purpose. In other words, the number and range of uses of personal data by individual organisations does not appear to be increasing, but the total number of organisations registered as data controllers is.

2.5.2 Power law distribution

The distribution we observed in each field, where a few highly common categories account for the majority of the entries in a given field (a power law distribution), is in keeping with previous research. A similar distribution was observed in classes of personal data collected by US banks (L. F. Cranor et al., 2013) and websites (Milne & Culnan, 2002), where a relatively small number of classes account for the majority, with a 'long tail' of less common classes. Like these studies, we find that it is often the 'long tail' of categories which contain the more interesting and controversial practices (for instance, the use of 'sensitive' data classes in the trading of personal data) which are commonly the focus of media attention and public concerns.

2.5.3 Informing public concerns

Our analysis appears to have revealed a number of uses of data that correspond to the public concerns identified above, such data being sold to third parties. We found that 'trading / sharing in personal information' is prevalent; ten percent of data controllers use data in this way. Furthermore, the personal data being traded is not just data classes like personal details and purchase histories, but also 'sensitive' personal data. For each type of sensitive data, we found at least 200 organisations trading it with third parties (about a third of whom were 'Traders in Personal Information'). Although this is only a small minority of cases (3%), and may even simply be erroneous, it combines two particular public concerns – data being sold to data brokers, and the sharing of sensitive data in particular.

For example, 840 data controllers claimed to be sharing data collected for health administration purposes with traders in personal data. This despite apparently widespread public opposition in the UK to the sale of health data.¹⁶¹ It should be noted that selling sensitive personal data to data brokers does not necessarily contravene data protection law. According to the Data Protection Act, sensitive personal data may be processed if certain additional strict conditions are met.¹⁶² It may be that such conditions are indeed met in

¹⁶¹ According to a poll by Yougov in 2014, 65% of UK adults do not want their medical data to be used by commercial companies (as reported in (White, 2014)).

¹⁶² These include at least one of the following: the data subject has given explicit consent or deliberately made the information public; that processing is necessary for compliance with employment law,

these identified cases. But even if they are met, and are therefore the activity is compliant with the letter of the law, it nevertheless conflicts with widely-expressed consumer expectations.

2.5.4 Differentiation between practices

In contrast to some of the previous research which has shown significant differentiation between company privacy practices (e.g. (Bonneau & Preibusch, 2010), (L. F. Cranor et al., 2013)), we find a lack of variation within each of the three sectors we studied. For instance, where Cranor et al found just 24.4% of US financial institutions shared data with affiliates, we found that 93% of UK financial service providers did so.¹⁶³ Different regulatory environments and other conditions prevent any direct comparison between the UK and US banking sectors, but nevertheless this indicates that UK consumers who prefer financial service providers to not share their data with affiliates have fewer options.

2.5.5 Limitations

Whilst it enables new analysis on an unprecedented scale, this data source is not without its limitations. One is that a large portion of data processing occurs outside its scope. Beyond the 350,000 registered controllers, there may be other liable organisations who have simply failed to comply with their notification requirement. Many companies whose data practices affect UK consumers, such as large international web companies, do not operate their consumer-facing services from UK offices and therefore do not register as UK data controllers in this regard (although this is a complex and changing area).¹⁶⁴ In addition, controllers might not disclose processing of 'anonymous' or 'pseudonymous' data, since according to the ICO these types of data may not be covered by the Act.¹⁶⁵

Another key issue is granularity. Many categories contained in the dataset would be more informative if given separate definitions. For instance, some consumers may perceive a difference between trading personal data for a profit and sharing it for some social purpose (E. A. Bell, Ohno-Machado, & Grando, 2014). As such, the standard description 'Trading / sharing personal information' is too broad.

Furthermore, the data within a purpose entry is not fine-grained enough. For instance, if the data subjects are 'customers' and 'staff', and the data classes are 'financial details' and 'physical and mental health', it matters greatly

protecting individuals' vital interests, administering justice, medical purposes or equality of opportunity. See the ICO's guidance (Information Commissioner's Office, 2015a).

163 In the paper by Cranor et al, affiliates are defined as entities 'related by common ownership or control' to the institution in question), while the register refers to 'Other companies in the same group as the data controller'.

164 Recent CJEU decisions indicate that these companies may indeed count as data controllers in Europe (see Case C-131/12, Google Spain SL and Google Inc. v. Agencia Española de Protección de Datos and Mario Costeja González). See (Caspar, 2015) for commentary.

165 A position reiterated in recent correspondence, as reported in (Burton, 2014)

which of the data classes pertain to which data subject. As it stands, many of the entries in the register give the appearance of potentially unethical or illegal practices because of this lack of differentiation. This has been made worse by the new format which was designed with the aim of making individual entries shorter and more user-friendly.¹⁶⁶ Unfortunately, this has obscured which categories of data subjects, classes, and recipients are associated with which purposes, preventing any meaningful disclosure on a per-purpose basis.

Finally, one might be sceptical about the accuracy of some of the disclosures organisations make. In addition to the possibility of basic administrative mistakes, or failure to reveal certain practices, organisations may also have perverse incentives to state practices that they do not really engage in. This is because there appear to be penalties for not disclosing practices that a controller is later found to be engaging in, but no penalties for listing a practice that a controller does not currently engage in but may at some point in the future. Therefore, a rational controller might be inclined to list as many categories as possible in order to cover themselves and avoid penalties for any activity they later take. This undermines the ability to discriminate between organisations based on their practices. This suggests a kind of transparency paradox; when forced to disclose their activity, organisations over-disclose, undermining the original purpose of transparency (a similar dynamic may exist for privacy policies (P. Leon & Cranor, 2010)).

2.6. Recommendations

With an increasing level of concern from the public and regulators about the collection and use of personal data, calls for more transparency are higher than ever. Researching privacy notifications and the systems and standards that support them is therefore not only of academic interest but also has implications for policy and practice. As evidenced in the background section, a broad swathe of technical and policy proposals have urged the adoption of appropriate standards for privacy-related disclosures.

This paper suggests that there are some challenges that must be overcome before these forms of transparency can feasibly be achieved. The data source studied here is the largest, most comprehensive, most structured store of information on organisations' privacy practices we are aware of. Nevertheless, significant problems prevent it from being a truly informative resource. As mentioned above, the register's format was changed in April 2013, and is likely to be abandoned by the ICO altogether as a result of impending changes to E.U. data protection law. Any recommendations drawn here are therefore aimed at policy-makers and designers of similar transparency systems in other contexts.¹⁶⁷

2.6.1 Standardisation, Categories and Granularity

The analysis above demonstrates the benefits of notifications being made available in a standard, machine-readable format. This significantly reduces

¹⁶⁶ (JISC Legal, 2012)

¹⁶⁷ i.e. those mentioned in the background section.

the barriers to scaling up analysis of organisational privacy practices from individuals to whole sectors or countries. Previous research has had to rely on relatively inefficient methods, from writing special natural language parsing software for standard-form policies, or worse, manual analysis of full-length legal documents. Having machine-readable data to start with (even if it requires additional parsing and processing to be useful), drastically reduces these barriers.

Second, division of privacy-related practices into the categories of purpose, subject, class, and recipient is useful, and all four categories seem necessary in order to derive any kind of meaningful conclusions about an organisations' practices. Any notification system which leaves one or more of these out is likely to prevent meaningful analysis. It is unsurprising, therefore, that other standards for such disclosures, such as P3P and the US financial institution model privacy form (studied in (L. F. Cranor et al., 2013)), have included equivalent categories.

However, the provision of these categories alone is not necessarily sufficient. Most importantly, without further fine-grained differentiation, the notification is likely to leave significant ambiguity. Rather than lists of data subjects, classes, and recipients for one purpose, it would be far more informative to differentiate subjects, classes and recipients individually, rather than aggregating them on a per-purpose basis. This way, the notification would indicate exactly which classes apply to which subjects, and which recipients have access to which classes. This would increase the amount of input involved in each notification, and therefore be more onerous on the organisation making the disclosure. Without such differentiation, the resulting data is far less informative, and potentially misleading.

The aforementioned changes to the notification process for the UK register have unfortunately made the data contained within it even less fine-grained. Instead of differentiation on a per-purpose basis, distinct purposes and associated information about subjects, classes, and recipients, have been amalgamated into one entry. It is no longer possible to ascertain, for instance, the precise purpose or purposes that data about customers are gathered under, and if so, which categories of data are gathered for which purpose. So the problematic lack of granularity encountered in the data prior to April 2013 is now even greater, rendering the resource even less informative than it previously was.

2.6.2 Incentives, monitoring and enforcement

Requiring per-field differentiation would also ideally go hand-in-hand with better incentives for accurate disclosures by organisations. Organisations may assume (correctly or not) that they can reduce their legal liability by simply exaggerating the extent of their actual practices, to 'cover their bases'. There are numerous ways organisations might be encouraged to make more accurate and detailed disclosures, from improved guidance for registration, to mandatory audits. However, one measure would be for regulators to pro-actively monitor the content of each notification using the kinds of techniques explored here, using this as a basis for further

investigation.

Previous research has noted the opportunity this kind of analysis presents for improving regulatory practice. Having found evidence of contradictory, controversial and potentially illegal practices in the disclosures of financial institutions, Cranor et al suggest that failure to identify and act on such evidence is a missed opportunity for the regulator (L. F. Cranor et al., 2013). They ask; 'if we as academics can quickly uncover these issues, why have regulators who are charged with overseeing these financial institutions not already done so?'. The same could be asked of the ICO, in cases where reported practices are controversial. For instance, if an organisation claims to use sensitive personal data to make credit reference decisions, or to sell health data to data brokers, this could prompt a further set of questions to ascertain whether the processing is indeed legitimate, and on what grounds. This kind of targeted action would be too onerous if it involved manually checking 350,000 registrations, and may be beyond the capacity of even a well-resourced regulator. But as demonstrated here, machine-readability means that regulators could easily employ such simple analytical techniques on their own data.

In addition to oversight by regulators, informed decision-making by privacy-conscious consumers is also likely to pressure organisations to make more accurate notifications, and develop more privacy-friendly practices. At present, this is prevented by a lack of useful, usable transparency, and consumer ignorance of organisation's practices. Those consumers who are concerned about their privacy do not have the time or the means to make informed and meaningful choices between service providers. While a better notification system may not in itself change this, it could provide the basis for intermediary services which would rate organisations practices on behalf of privacy-conscious consumers, and in turn provide a commercial incentive for organisations to improve practices.

2.7. Conclusions

Transparency is easy to affirm but hard to achieve in practice. There has been no shortage of enthusiasm for measures which render visible organisations' policies and practices regarding personal data. We are left with a graveyard of incomplete attempts (Binns, 2014b). The data source studied here is arguably the largest and most complete arising from any of them, and therefore provides an instructive case study through which we can assess the viability of this kind of transparency proposal. This kind of resource does enable macro-level conclusions about the types of data used by a range of organisations, the purposes involved, and the types of recipients with access to the data. This information could be an important starting point for more detailed investigations.

However, the data source itself is not designed for a detailed understanding at the level of particular organisations' practices. There are two kinds of problems associated with using these broad, abstract descriptions as a means to assess particular practices: *false negatives*, where the data fails to capture the existence of a practice, and *false positives*, where the data suggests a certain practice is occurring where it is not. The ultimate utility of this

resource may therefore depend on whether there is value in macro-level abstraction despite the strong possibility of these different types of errors. When it comes to describing the use of personal data, there will always be tension between standardisation and nuance, abstraction and detail. Resolving these tensions may be the key to successful transparency systems in this domain.

2.8 Epilogue

This first paper has provided insights into one iteration of Openness for Privacy, based on *open data*. As explained in the introduction, this involves compiling / releasing structured information on the privacy-related practices of data controllers. It is commonly pursued, as in the case of P3P, in the hope of improving the infrastructure for notice and consent. But as hinted in the introduction and this first paper, it also could serve other purposes. These include helping regulators to pursue more data-driven and targeted interventions; helping third parties to evaluate privacy risks on an individuals' behalf; or helping organisations identify their own risks (for instance, discovering auxiliary datasets that might affect the risk of a particular dataset being de-anonymised).

This initial foray suggests this concept has potential, but that it also has some important limitations. In practice, standardised policies like P3P, and the system of national registers, have largely been failures. The register is seen as tedious bureaucracy (Pederson, 2005), and has been abandoned in the proposed data protection reform package because of the 'administrative and financial burdens' it imposed without identifiable benefits.¹⁶⁸

Despite the failures of these existing systems, there remains enthusiasm from many quarters for some kind of standardised data about organisations' privacy credentials (see section 2.3 above). The paper above suggests some key challenges that would need to be overcome. Perhaps most importantly, there is a tension between categorising practices so that they can be compared, and capturing the idiosyncratic contextual factors that might be essential for a meaningful evaluation. Privacy and data protection are complex in ways that simple data summaries and metrics are unlikely to reflect. If such transparency systems ever came to be strongly relied upon, this lack of nuance could be problematic.

If organisations are assessed (whether by regulators, consumers, or others) on the basis of proxy values, they may attempt to optimise the proxy rather than the phenomena it is supposed to measure. Rather than taking action that would actually reduce their privacy harms, data controllers might focus instead on superficial measures that would make them appear less risky within the structure of the notification system. Transparency systems therefore need to be carefully designed to minimise this kind of 'creative accounting' of data.

Nevertheless, standardised categories and metrics still serve important purposes. Perhaps the most important is in improving the empirical investigation of privacy and data protection, to take it beyond individual case studies and unsubstantiated generalisations, towards a more evidence-based policy discussion. While some good empirical work exists (much of which is cited above), more can and should be done.¹⁶⁹

Many commonly held opinions about the actions of data controllers could be

¹⁶⁸ See footnote 283 below.

¹⁶⁹ It should be noted that user behaviour and attitudes regarding privacy are comparatively well studied (see (Patil, 2013)).

tested if there were more of the kind of empirical data I'm advocating. For instance:

- The legal implications of using data for purposes it wasn't originally collected for is a hot topic,¹⁷⁰ but we have little data on the nature and extent of such 're-purposing'. Are organisations really using personal data for new purposes (or, as the paper above suggests, have purpose types actually remained stable in recent years)? If so, what are the new purposes? Does the situation differ by industry?
- Data protection law sets out six lawful bases for processing (including consent, necessity in relation to a contract, protection of an individual's 'vital interests', and others). The relative importance, and the supposed growth or decline of these conditions has been much discussed (e.g. (Zanfir, 2014)), but we have very little empirical data on which lawful bases are actually relied on by data controllers in practice, and whether these are indeed changing.
- With better standardised measures of data use and compliance behaviour, correlations with other organisational measures could be examined.

These examples provide an illustration of how open data about organisations' practices could drive a more empirically informed debate about data protection and privacy. The resource studied here is far from perfect, and cannot answer all these questions, but it demonstrates the potential of a more data-driven approach to policy studies in this area.

In fact, the analysis presented above has already provided empirical input into policy debates. In a short paper based on the same data source, my co-authors and I presented evidence on the extent of cross-border data transfers from the UK ((Binns et al., 2014), see appendix A). This was subsequently cited in a report by the Centre for European Legal Studies on access to data by third-country law enforcement authorities (Carrera, Fuster, Guild, & Mitsilegas, 2015), and in a technical paper on new systems to protect privacy in the cloud (Zeng, Wang, & Feng, 2015). These citations demonstrate the latent demand for more empirical data to inform policy and technical work in this area.

170 (Mantelero, 2014), (Information Commissioner's Office, 2015b), (van der Sloot, 2014)

Part 3: Open Processing

The previous section explored one manifestation of the Openness for Privacy concept. But as I argued in the introduction, openness is not just about organisations stating what data they collect and why; it's also about ongoing personal data processing activities being open to scrutiny, modification, and challenge by individuals. The latter was defined in the introduction as *open processing*, in which the individual can understand and influence the processing of their own personal data in context, in real time. This is explored in the following section.

It focuses on a new service and business model which gives individuals control over the contents and use of their digital profiles. While interest in user-controlled personal data architectures is not new, until recently it has been driven by privacy enthusiasts rather than industry. But as we shall see, marketers and advertisers are beginning to see the potential benefits of this alternative model.

The section seeks to answer two questions. First, is user-controlled profiling a viable option for businesses, or will it negatively impact their revenues in the long term? Second, can it also be a genuinely empowering option for individuals, or are the two mutually exclusive?

The section is divided into two parts corresponding to each question. The first part addresses the first question ('is it viable for business?') through a quantitative user study. This part has been accepted for publication in a forthcoming issue of the International Journal of Internet Marketing and Advertising. The second part addresses whether this can be genuinely empowering for individuals, and synthesises work presented in three previously published papers.

3. Abstract:

A large portion of the content, recommendations and advertisements shown on the web are targeted, based on a profile of an individual user. This paper explores two ways of creating and using such profiles. *Behavioural profiling* – a commonly used technique which makes inferences based on an individual's previous activity – is compared to what I call *self-authored interest* (SAI) profiling, which is based on information explicitly volunteered and controlled by the individual. I present the results of an experimental study comparing the effectiveness of the two systems in generating targeted product recommendations. I find that a) people respond more positively to product recommendations when they are derived from SAI profiles, and b) the mere belief that a recommendation comes from an SAI profile is also associated with more positive responses.

Keywords: digital marketing, profiling, behavioural targeting, privacy, recommender systems, uncanny valley, permission marketing, vendor relationship management, personal data, personalisation

3.1. Introduction

The first part of this paper introduces some background on profiling and personalisation in the context of online marketing. It provides an overview of the industry, recent challenges, and some developing new service models that attempt to give individuals greater control over their own profiles. A review of relevant literature on consumer behaviour, information systems and marketing is presented, which leads to a set of aims and objectives for the study. Two types of profiling – *behavioural* and *self-authored interest* – are defined and compared. Part 2 details the study design and method, followed by results and analysis in part 3. The paper ends with a discussion of the implications of the findings for research and industry.

3.1.1. Background

Since the web was first used as an advertising medium in 1994,¹⁷¹ the digital advertising industry has grown to become a \$137 billion global industry.¹⁷² Unlike traditional television or print advertising, the web enables targeting, whereby particular advertisements can be matched to consumers using data collected about their behaviour. Behavioural targeting takes place within online platforms like social networks, search engines and e-commerce sites, as a means to provide personalised recommendations. Part of the process of targeting is the collection of data about user behaviour, such as the types of websites they have previously visited and the products they have previously bought or looked at. This data is aggregated from thousands of users and used to create predictive models, which allow inferences to be made about what a particular user might be interested in and receptive to. This activity is facilitated by a complex network of intermediaries operating in several sub-markets (H. Kox, 2014). Information about an individual consumer is stored in a profile which is used to personalise content, recommendations and

171 The date is reported by (Edwards, 2013)

172 As reported in (Emarketer, 2014)

advertisements. Personalisation in this context has been defined as 'a process that changes the functionality, interface, information content, or distinctiveness of a system to increase its personal relevance to an individual' (Blom, 2002). The behaviour of many individuals is analysed in order to drive personalisation for one individual (a process sometimes called collaborative filtering (Su & Khoshgoftaar, 2009)).

In recent years, researchers and technology commentators have noted the potential for a consumer backlash against targeting. A parallel has been drawn with the 'uncanny valley' hypothesis in robotics. The hypothesis states that people prefer interacting with robots that have human features; however, if those human features become too realistic, people find them uncanny and cease to enjoy interacting with them (Mori & Minato, 1970). It has been suggested by various industry commenters that a similar phenomenon applies to personalised marketing; that there may be an 'uncanny valley' facing big data marketers.¹⁷³

There are also doubts about the actual success rate (in terms of click-throughs and purchases) of targeted advertising, with warnings that 'peak advertising' is imminent (Hwang & Kamdar, 2013). This is partly due to factors like ad-fraud, but it may also be due in part to negative reactions of consumers.¹⁷⁴

In response to a possible backlash from consumers, some alternative systems have emerged. They offer individuals greater control over the content of their profiles and over what marketing messages they are exposed to. Such controls have already been offered to a limited extent by existing digital advertising networks, who have in recent years introduced account

173 E.g. (Strong, 2014), (McEwan, 2014), (Salmon, 2011), (Bramwell, 2014), (Watson, 2014). This should be separated from other concerns about personalisation, such as objection to personalised pricing (for which, see e.g. (Mikians, Gyarmati, Erramilli, & Laoutaris, 2012); (Miller, 2014); (Acquisti, 2008)).

174 Depending on the type of advertising, between 11% and 52% of impressions are thought to be fraudulent (Association of National Advertisers, 2014)

settings wherein consumers can view and edit the marketing profiles that have been created about them for use in targeted advertising.¹⁷⁵ Browser tools which limit user tracking are also available.¹⁷⁶ But some new companies take a step further, offering consumers the ability to create their own profiles from scratch, select what information they wish to reveal and to whom, and even earn money in return for exposing their profiles to marketers.¹⁷⁷

These include companies like DataCoup, which describes itself as the world's first personal data marketplace, and CitizenMe, which offers to 'unlock the value of your personal data, for you and on your terms'.¹⁷⁸ Autograph 'lets people realise their interests, helping marketers drive response rates', while Handshake, another personal data marketplace, estimates its users could earn between £1,000 - £5,000 (GBP) per year by selling their data.¹⁷⁹ By providing individual control and transparency, these companies aim to make profiling and targeting more acceptable, avoid a consumer backlash and create a mutually beneficial system for both consumers and marketers.

This alternative service model is at odds with current personalisation systems which typically work entirely without any explicit user input. Data is gathered by various tracking technologies, from which statistical models are created and applied to individual profiles, which are used for personalisation and targeting. The individual user does not generally have a say in this process and cannot therefore define their own preferences for themselves.

175 See, for instance, Google's 'Control Your Google Ads' [<https://www.google.com/settings/u/0/ads/authenticated>]

176 Tools such as Ghostery [www.ghostery.com] and PrivacyBadger [www.eff.org/privacybadger].

177 The earliest instance of this kind of service known to the author is RootMarkets, founded in 2006, now defunct. As reported in (R. Hof, 2006).

178 See Datacoup [www.datacoup.com] and CitizenMe [www.citizenme.com]

179 See Autograph [autograph.me] and (Lomas, 2013).

These two kinds of profiles present quite different visions for the marketing and advertising industry. On one hand, there is a purely behavioural model, which has proven successful but gives consumers no control or transparency, and may face a consumer backlash. On the other hand there is a model, largely unproven, which attempts to make targeting more acceptable by giving users control and transparency.

3.1.2. Literature Review

To date, there is relatively little research which directly compares the relative strengths and weaknesses of these two models. This is not surprising given the relative novelty of the user-centric systems.¹⁸⁰ However, a wide range of adjacent research on consumer behaviour, advertising and marketing, and information systems design, provides some relevant insight.

From the perspective of retail marketers, digital profiles exist to drive consumer purchases through personalisation. Consumer purchasing behaviour is well studied in numerous contexts (Solomon, Russell-Bennett, & Previte, 2012), including online. There are numerous factors influencing online shopping, from consumer presence, enjoyment, and attitude to vendor (Neuendorf, Xiong, Blake, & Hudzinski, 2014), to economic gains (Gummerus, Liljander, Weman, & Pihlstrom, 2012), and perceptions of risk (Atorough & Donaldson, 2012). Degree of personalisation is an important factor; personalised recommendations increase online purchasing when compared to non-personalised ones (Senecal & Nantel, 2004).

Behaviourally targeted advertising has been found to increase the click-through rates of advertising by as much as 670% (Yan et al., 2009). It also allows for greater market differentiation and reduced wastage in advertising spending (Iyer, Soberman, & Villas-Boas, 2005).

These results attest to the success of existing targeting systems. But there is also evidence to support the claim that a consumer backlash may ensue. Ur

¹⁸⁰ However, see (Sørensen, Sørensen, & Khajuria, 2015).

et al found a variety of negative attitudes amongst consumers towards profiling and targeted advertising, regarding it as 'inaccurate' and even 'creepy' (Ur, Leon, Cranor, Shay, & Wang, 2012). As the level of personalisation in digital advertising content increases, consumers may pay more attention, but find it less acceptable (Malheiros, Jennett, & Patel, 2012).

These negative consumer attitudes threaten marketers' interests. Privacy concerns have been found to moderate the effectiveness of targeted ads (Alnahdi & Ali, 2014), and reduce consumers intention to purchase ((M. Brown & Muchira, 2004), (Flavián & Guinaliú, 2006), (Valvi & West, 2013)). In particular, where consumers perceive behaviourally targeted ads as creepy and / or threatening, they have been found to lead to a 5% reduction in intention to purchase (Barnard, 2014).

Concerns about privacy and invasiveness of targeted advertising also negatively affect a consumer's attitude towards the advertised brand (Taylor, Lewin, & Strutton, 2011), and reduce their trust in a vendor (McCole, Ramsey, & Williams, 2010). Where trust in a retailer is already low, personalisation drives stronger privacy concerns (Bleier & Eisenbeiss, 2015).

According to Summers et al, when users know an ad is targeted, they perceive it as a social label which may change their self-perception (Summers, Smith, & Reczek, 2014). Their study suggests that 'making consumers aware that they are being targeted can prompt them to perceive an ad as a social label, which they then use to evaluate their own characteristics' (ibid. p1). This also suggests that negative responses to personalisation may also be due to what consumers believe a targeted message says about them, rather than privacy *per se*.

A consumer backlash may result in the avoidance of advertising. While it predates the web (Speck & Elliott, 1997), advertising avoidance has been

observed on the web (Cho & Cheon, 2004), (Duff & Faber, 2011), and social networks (Hadija, Z., Barnes, S. B., & Hair, 2012) It is driven, at least in part, by privacy concerns ((Krasnova, Günther, Spiekermann, & Koroleva, 2009), (Wirtz & Lwin., 2009), (Baek, T. H., & Morimoto, 2012)), as well as high levels of ambivalence and low levels of interactivity (Jin & Villegas, 2007). Avoidance does not just result in the message not being consumed; it is also associated with the consumer forming a negative image of the brand (Alwitt & Abhaker, 1994), (Cho & Cheon, 2004).

Advertising avoidance can be seen as one of the various forms of consumer empowerment enabled by the web ((Schultz, 2006), (Denegri-Knott, 2006)). But avoidance in itself only empowers consumers to mitigate negative aspects of marketing; it doesn't allow them to benefit from the potential positives. It also, of course, runs counter to the objectives of marketers.

Alternative systems attempt to give consumers more control whilst maintaining a viable business proposition for marketers. These systems allow individuals to scrutinise and customise their profiles ((Sundar & Marathe, 2010), (Kay & Kummerfeld, 2012)). One framework for thinking about the alternative is provided by the concept of 'Vendor Relationship Management' (VRM) (Searles, 2013). The term comes from a project originating at Harvard University's Berkman centre and now encompasses a business networking community, which includes some of the organisations mentioned above.¹⁸¹ VRM is a corrolary of customer relationship management (CRM), which helps organisations manage their relationships with customers. Proponents of this concept foresee a wide range of customer-centric technology, tools, and services which help individuals engage with organisations on their own terms (Ctrl-Shift, 2014). When applied to digital marketing profiles, a VRM approach would emphasise giving the individual the ability to create and control their own profile, and

¹⁸¹ For an overview of VRM concepts, by the founder of the project, see (Searles, 2013).

to accept or reject advertising on their own terms.

Some studies suggest that giving such controls to consumers could result in a mutually beneficial position for marketers, by encouraging higher engagement with services. Feeling in control, and knowledge of the algorithm, may encourage consumer engagement with recommendation systems (Blom, 2002). Chauhan and Rathore found that consumers tend to agree to targeting provided they are made more aware of the process (Chauhan & Rathore, 2015).

A potential challenge for this alternative, however, is that consumers may not actually know their own preferences well enough for effective personalisation to work on the basis of their self-authored profiles. Behavioural targeting is based on statistical analysis of many users' behaviour, rather than what an individual thinks would represent them accurately. This focus is perhaps understandable given that an individual's stated preferences are not necessarily accurate or consistent ((Bettman, Luce, & Payne, 2015), (Slovic, 1995)). They may form their preferences on an ad-hoc basis when a decision has to be made, and may be influenced by extraneous factors such as the type of user interface ((Hong, Thong, & Tam, 2004), (Lim & Benbasat, 2000)), and be susceptible to various systematic cognitive biases which influence their reported preferences (Kahneman, 2011). Individuals are often unaware of these factors. For these reasons, an algorithmic system based on behavioural data and predictive models, where users remain passive, may seem preferable to a system which introduces a significant element of error-prone human judgement.

If alternative profiling systems are to be successful, they will need to demonstrate how user-generated profiles can nevertheless create worthwhile opportunities for marketers. At present, there is only limited research comparing user-defined profiles with traditional behavioural targeting. McNee studied a hybrid system which relied partially on user input (McNee,

2003). When users selected the items they wanted to rate, rather than having the system produce a list based on behavioural data, this resulted in an equally accurate user model and also increased consumer loyalty. However, the system this study assessed is still primarily based on behavioural targeting – the user input related only to one aspect of the profile generation. It therefore remains uncertain how recommendations derived from the alternative services outlined above would ultimately compare to mainstream behavioural targeting.

3.1.3. Aims and Objectives

This paper compares mainstream behavioural profiling with a customer-centric alternative, through a user study. For the purposes of consistency and clarity, I introduce the following terms. *Behavioural* profiles are defined as those generated from user behaviour and other surreptitiously gathered information such as browsing history, purchases or search terms, which are analysed to infer user interests. This approach generally precludes individuals from being the primary authors their own profile. On the other hand, *Self-authored interest (SAI)* profiles are generated by individuals themselves, by explicitly stating or selecting their interests (e.g. 'poetry', 'football' or 'DIY'). They do not contain information gathered without the users explicit input. Both behavioural and SAI profiles can be used for various kinds of targeting and personalisation, but for the purposes of this paper I will be focusing on the targeting of consumer products.

The literature review suggests two variables which may affect the success of any system for personalised marketing. One is the extent to which it creates an accurate user model and relevant personalised messages. The other is the consumer's attitudes towards the personalisation system; if they have a negative attitude towards it (e.g. due to the posited 'uncanny valley' effect), this may actually harm the marketer's objectives and reduce the effect of the recommendation. On the other hand, if a consumer understands and feels positively about the targeting process, they may be more likely to respond

more positively to the recommendations that come from it. This suggests that the positive or negative feelings that a consumer has about the process *behind* a recommendation might influence how receptive they are to it, rather than simply being influenced by its accuracy or relevance. If consumers' propensity to purchase is potentially affected by their view on the *process* of targeting, this must also be accounted for in comparing these systems.

This study aims to address both aspects. First, how do consumer responses to recommendations differ between mainstream behavioural targeting and user-centric alternatives?, Second, what is the effect of an individual's *perception of the process* behind the personalisation system on their response?

3.2. Study Design and Method

The experiment placed participants in simulations of behavioural and SAI-based targeting, and asked them to rate a set of product recommendations in terms of how much they would like to buy the recommended product. There were two hypotheses to be tested:

Null Hypothesis A: There is no relationship between consumers' ratings of a recommendation and the type of profiling system (SAI or behavioural) used to derive that recommendation.

Null Hypothesis B: There is no relationship between consumers' ratings of a recommendation, and *their beliefs about* the type of profiling system (SAI or behavioural) used to to derive that recommendation.

For the first hypothesis, the independent variable is the profiling system (including an associated interface) being used to generate a recommendation to the individual, which is either SAI or behavioural. For the second hypothesis, the independent variable is a combination of the type of profiling system actually used to generate the recommendation, and the type

of profiling system that is presented to the user via the interface (which will be different in some conditions). In both cases, the dependent variable is the individual's reported propensity to purchase the recommended item, recorded on a 5 point Likert scale ranging from 5 ('*very likely*') to 1 ('*very unlikely*').

Participants for the study were recruited through advertising on the web and on social networks. I opted for a between-subjects design to avoid contamination by extraneous factors like fatigue. I tested the information displayed in each condition to ensure participants understood the difference between self-authored interest profiles and behavioural profiles.

Manipulation checks were conducted to ensure that the respondents perceived the scenarios they were placed in as I intended.

Behavioural targeting was simulated using an 'advertising API' from a leading international e-commerce site.¹⁸² This system is based on a proprietary recommendation algorithm working from the purchase history of several hundred million active users¹⁸³. Rather than attempt to simulate a product recommendation system from scratch, I chose a real-world system, based on industry-standard algorithms, genuine user data and an extensive product range, to ensure high ecological validity for the experiment.

When queried with a set of previous product purchases as a parameter, the API returns product recommendations based on what customers with similar purchase histories also bought. To create a behavioural profile within the experiment, participants were asked to disclose 5 items they had *recently purchased*. To make this task easier and to aid accurate recall, participants were prompted within the experiment to import recent purchase data from

182 See Amazon Product Advertising API, which provides programmatic access to Amazon's product selection and discovery functionality, including recommendations based on previous purchases
<https://affiliate-program.amazon.com/gp/advertising/api/detail/main.html>

183 While the details of the recommendation algorithm in its current implementation remain proprietary, a paper from 2003 describes how a similar early version worked using 'item-to-item collaborative filtering' (Linden, Smith, & York, 2003)

their online accounts with several popular online retailers.¹⁸⁴ This data populated their behavioural profile, which could then be used as parameters to retrieve recommendations through the API. This ensured that the targeted recommendations participants received in the simulation were very similar to those they would receive on a real e-commerce platform.¹⁸⁵

The self-authored profile simulation used a simple keyword matching system. Participants were asked to enter a set of 5 keyword strings describing various personal interests, which they felt comfortable revealing for the purposes of targeting. These could be specific items (i.e. 'digital camera'), categories (i.e. 'photography' or 'poetry'), or names associated with an interest (such as authors, brands or sports teams). These keywords were used to search for matching products (using product names and descriptions) in the product catalogue.

A number of measures were taken to ensure the processes involved in the creation of the behavioural profile and the SAI profile were equivalent. The same product catalogue was used. Both types of profiles consisted of 5 items (whether prior purchases or self-authored interest keywords), and both took a similar amount of time to create.

In order to test hypothesis B (concerning the relationship between a consumer's rating of a recommendation, and *their beliefs about* the type of profiling system used to derive it), I created control conditions for both systems. In the case of behavioural targeting, participants were induced to believe that all their recommendations were derived from their behavioural data, but a random one out of five were actually, unbeknownst to them,

184 The imported data is supplied in yearly batches; I prompted the user to obtain the latest batch, so that no items in their profile were more than one year old.

185 In this scenario I am considering an e-commerce platform, where a user's entire purchase history may be used for targeting. There may be differences between this and other forms of targeting, such as remarketing, which generally operate on much more recent data. See (Deane, Meuer, & Teets, 2011) for an assessment of the optimal period of behavioural history to include in a profile. I am grateful to an anonymous reviewer for this point.

derived from their self-authored interests. In the case of SAI-based targeting, participants received one random recommendation based on their behavioural data. This allowed me to test whether any difference in ratings between the two systems can be accounted for by the beliefs participants have about those systems rather than any differences in the actual content of the recommendations. See appendix B for a flowchart illustrating the study design.

The observations were taken over two different periods (before and after the UK winter holiday, December 2014 and February 2015), to account for potential differences in consumer behaviour due to the time of year. Half of the participants were recruited through an online platform for conducting user studies, the other half through advertising for volunteers on social and professional networks in my academic institution. In order to determine a sufficient sample size, I conducted a power analysis as outlined in (J. Cohen, 1988). The significance and power were set at the standard levels of 0.05 and 0.8 respectively.¹⁸⁶ To determine an appropriate effect size, I considered the sizes reported in similar academic research,¹⁸⁷ as well as considering what would be seen as a noteworthy effect size in the digital advertising industry.¹⁸⁸ On this basis, an effect size of 0.8 was deemed appropriate. Given an effect size of 0.8, significance of 0.05, and power of 0.8, a two-sample t-test power calculation indicated that a minimum of 25 participants per group would be required in each two-sample test. Given that there were two independent tests to perform, this would require two pairs of such

¹⁸⁶ As outlined in Cohen, (J. Cohen, 1988)

¹⁸⁷ Unfortunately, similar studies of online purchasing often fail to report effect sizes. However, for those who do, a 0.8 effect size is considered large: e.g. (Eslami et al., 2015).

¹⁸⁸ An interesting effect size for industry can be relatively small, especially if the user base is large. For instance, in a study of the effect of emotional posts on Facebook (Kramer, Guillory, & Hancock, 2014), the authors noted that “given the massive scale of social networks such as Facebook, even small effects can have large aggregated consequences... an effect size of $d = 0.001$ at Facebook’s scale is not negligible”. Similarly, (P. G. Leon et al., 2013) argue that “even a small effect size has important practical implications when applied to millions of Internet users” (p9).

groups (4x25), and therefore at least 100 participants in total were required.

115 participants were recruited.

3.3. Analysis and results

Table 4 shows a descriptive statistical summary of the results.

Condition	Source of targeting	Interface presented to user	Average rating
All conditions	N/a	N/a	2.68
Behavioural	Behavioural	Behavioural	2.14
	Behavioural	SAI	2.78
Self-Authored Interests (SAI)	SAI	SAI	2.89
	SAI	Behavioural	2.95

Table 4. Average Recommendation Ratings by source / interface

The average rating for recommendations across all conditions was 2.68.

When participants were shown recommendations based on their prior purchases, (i.e. behavioural targeting) through an interface representing them as such, the average rating was 2.14. However, when behavioural targeting was delivered through an interface that represented the recommendations as if they were SAI-based, the average rating was higher at 2.78. When shown recommendations based on their self-authored interests, through an interface presenting them as such, participants rated them at 2.89 (higher than either of the behavioural targeting conditions). When SAI-based recommendations were represented as if derived from behavioural targeting, they were rated as 2.95 on average.

To test hypotheses A and B, statistical tests were performed on pairs of samples (the Wilcoxon Rank Sum / Mann Whitney U test).¹⁸⁹ The relationship between profiling system and recommendation rating (hypothesis A) was tested by comparing the scores for *'pure behavioural'*

¹⁸⁹ The choice of test followed guidance from (Leeper, 2000). A normality test was performed (using a QQ plot), which revealed the data was not normally distributed, hence a Wilcoxon Rank / Mann Whitney U test was appropriate.

(where the targeting was behaviour-based and the interface faithfully presented it as such) and '*pure SAI*' (where the targeting was SAI-based and the interface faithfully presented it as such). The relationship between consumer's *beliefs about a profiling system* and recommendation rating (hypothesis B) was tested by comparing two pairs of samples:

B1. '*Pure behavioural*' and '*misrepresented behavioural*' (where the targeting was presented as being based on their prior purchasing behaviour, but was actually based on SAI)

B2. '*Pure SAI*' and '*misrepresented SAI*' (where the targeting was presented as SAI-based but was in fact based on their prior purchasing behaviour).

The results of these tests are summarised in table 2. They show that there are significant differences between samples tested in hypothesis A, but mixed results for B, where only one of the two tests produced a significant result.

Test	Averages	Wilcoxon Rank Sum Test
Pure behavioural vs. Pure SAI	2.14, 2.89	W = 1129, p-value = 0.003842
Pure behavioural vs. misrepresented behavioural	2.14, 2.78	W = 950.5, p-value = 0.003562
Pure SAI vs misrepresented SAI	2.89, 2.79	W = 1664.5, p-value = 0.6898

Table 5. Significance tests for SAI and Behavioural, pure vs misrepresented

We can therefore reject A. There are significant differences between the ratings of recommendations from behavioural profiling and those from SAI profiles, the latter being positively associated with relatively higher ratings. However, our investigation of B is inconclusive, as a significant difference is found in only one of the two measures used to test this hypothesis.

3.4. Discussion and conclusions

These findings indicate that a SAI-driven recommendations may have advantages over the behavioural targeting model. Even a relatively

simplistic implementation of an SAI system resulted in recommendations that were rated more highly than those offered by the recommendation system of a leading e-commerce platform.

The picture is less straightforward when one considers the second hypothesis, that an individual's beliefs about the process behind a recommendation affects their rating of it. This appears to be false when one compares the scores given to faithfully-represented interest-based recommendations, against scores given to recommendations that are misrepresented as behaviour-based. However, when one compares scores between behaviourally-targeted recommendations that were faithfully represented as such, and those misrepresented as if they were interest-based, there are significant differences. In other words, recommendations based on past behaviour get higher ratings when they appear to be based on self-authored interests, while those that are actually based on self-authored interests get the same rating regardless of their apparent source. It can therefore be concluded that when consumers believe a recommendation is based on their previous behaviour, they tend to like it less. More research is needed to further test, explore and explain this phenomenon.

Despite this complication, our findings indicate that consumer responses to product recommendations are indeed affected by two different factors; the content of the recommendation and the consumer's perception of the process behind it. Consumers are more likely to want to buy a recommended product if the recommendation is presented as deriving from a self-authored interest profile, compared to recommendations deriving from behavioural profiles. By showing misrepresented recommendations in this way, one is able to distinguish the relative importance of these two factors.

3.4.1 Further research

Several avenues for further research remain. There are many additional factors that could have been considered. For instance, I did not attempt to

uncover in a statistical way different consumer types, that may be associated with different consumption styles (e.g. spendthrift, frugal, relaxed, or controlling) more or less suited to the SAI approach. Neither did I consider how these recommendation systems might fare if restricted to different product types, which could have an effect on consumer choices (Senecal & Nantel, 2004). The stage of the consumer's purchasing process – i.e. whether they are 'browsing' or 'searching' - might also make a difference to the observed effects (Schlosser, White, & Lloyd, 2006).

It is also not clear whether a hybrid model, which mixes user control and self-authored interests with more traditional forms of behavioural profile, would elicit similar or different responses (Burke, 2005). What precise aspect of the SAI model appeals to consumers, and what combinations of service design might accentuate or diminish it, remains to be seen.

It may be that consumers are divided, with some preferring SAI and others preferring behavioural targeting (Sørensen et al., 2015). Having the option of both forms of targeting might be a socially optimal outcome, maximising the consumer surplus (H. L. M. Kox, Straathof, & Zwart, 2014).

3.4.2 Implications for industry and policy

These findings raise some important considerations for providers of new profiling services, the digital marketing and advertising industries, and those that are adjacent such as online publishers and providers of advertising-subsidised web services. These considerations are timely considering the current backlash against the existing model and increased interest in alternatives.

First and foremost, the results suggest that the self-authored interest model is worth exploring as an alternative to behavioural targeting. Giving individual consumers transparency and control over their profiles and marketing channels need not be to the detriment of marketing objectives. On the contrary, it could increase consumers' positive responses to the

marketing messages they receive. This is both due to the difference in the content of the recommendations as well as the attitudes consumers have towards the process. Explicitly volunteering information that they feel comfortable sharing for marketing purposes can improve consumers' responses to product recommendations. This suggests SAI profiling is one way to increase the response rate of digital marketing.

On the basis of this study, any concerns that a lack of self-knowledge on the part of consumers will lead to worse recommendations appear to be unfounded. It suggests that consumers know themselves at least well enough to supply a set of interests that generate some appropriate recommendations.

These findings suggest a potential perverse incentive for the designers of personalisation services. Marketers could improve the response rates of behavioural targeting by simply giving individuals the *illusion* of control over their profiles, and thus avoid the negative attitudes towards behavioural profiling without changing their actual practices. But aside from regulatory and ethical risks associated with this strategy, this study suggests that SAI profiles can actually provide more relevant recommendations than behavioural ones anyway. The difference between the pure SAI and pure behavioural conditions suggests that while subjective beliefs about processes make a difference, the advantage of SAI-driven recommendations is primarily due to their actual content, rather than their presentation.

Many questions remain about the design of SAI services. Should consumers who use these service be paid for their data?¹⁹⁰ It is unclear how payment for data might affect consumer's propensity to purchase. Another challenge is how SAI services can convince consumers that they represent a genuinely different approach and can be trusted, since trust is a precondition for consumers to share their data (McKnight, Choudhury, & Kacmar, 2002). If they are to overcome consumer scepticism, SAI services may need to explore different trust models and legal structures to ensure they have the

¹⁹⁰ This question is discussed from an ethical standpoint in (Binns, 2015)

trust of their users.

Part 4: Personal data empowerment

The quantitative study presented above provides some support for the open processing model outlined in the introduction. It suggests that there may be economic incentives for the marketing industry to allow consumers to control their own profiles. This kind of 'open' profiling system is one manifestation of openness for privacy. But even if it aligns with marketer's interests, does it genuinely empower consumers? This section addresses this question. First I present some results of a qualitative study of these systems. Then I address three types of objections to their claims to empower consumers.

To explore possible consumer responses to these platforms, alongside the quantitative study described above, I also gathered qualitative data about attitudes towards these platforms. This took the form of an asynchronous forum discussion hosted in the context of an online course on digital marketing, featuring 274 respondents. Over the course of the discussion, several objections to the idea of these platforms emerged.

As one respondent put it, the incumbent and alternative forms of marketing both had their disadvantages:

'It feels like [being between] a rock and a hard place. Both types of platforms harvest information and regardless of it being personalised, they both still have access to a rich amount of information.'

Some were sceptical about the promises of greater control made by the SAI services on offer. The notion of 'empowerment' through data sharing was questioned:

'representatives from these platforms talk about empowerment ... but I'm not entirely sure that them having full access to all my data is doing that?'

Some participants worried that these services were just another version of the incumbent advertising industry's attempts to integrate elements of consumer control into their existing platforms (such as the 'ad preference' dashboards which allow consumers to edit their profiles). They worried that whilst giving them some control over what messages they receive, such tools ultimately would give more information and power to companies.

'It really is none of their business what my preferences are. Essentially you are ... helping marketing companies even more by providing additional data'.

For those already highly sceptical of the marketing industry, the idea of self-authored profiles seemed like more of the same. There was a strong suspicion from some participants that while some form of payment would be favorable to the current situation, they would inevitably get a raw deal. Amongst those concerned about profligate spending, being paid for one's data seemed to be a false economy.

'Ultimately, I don't want to be convinced by marketing companies

to spend more on products I don't need. Receiving £8 for the privilege does not offset this cost.'

For others, the very possibility of monetary reward brings into question whether they want their data to be shared at all.

'It is interesting how being offered money for the purchasing information we are no doubt giving away for free, changes your perception of its worth, value or influence. Suddenly I do not want this information to be available'¹⁹¹

These remarks suggest that despite any advantages these tools might have for marketers, consumers may be sceptical. In offering to change the way personal data is monetised and exchanged, SAI profiles appear to introduce a range of new considerations which go beyond privacy and convenience, to encompass a broader set of concepts like autonomy, fairness, the role of the market, exploitation and consumerism.

This section outlines three possible political and ethical critiques of open profiling platforms. Versions of these critiques are implied in recent academic literature and commentary which expresses scepticism about the potential for these platforms to empower their users. My aim here is to clarify and develop these critiques. While I do not believe any of them to be decisive arguments against the notion of platforms for personal data empowerment, they do contain important considerations which any defense of that notion needs to respond to.

Scepticism about these platforms can be roughly divided into three distinct claims. The first is that while they purport to be a genuine alternative to prevailing big data systems and business models, they are in fact guilty of uncritically accepting a more fundamental, and more pernicious, kind of market logic. These platforms may present themselves in the language of user control but, so the argument goes, they are not a genuine alternative to the existing big data paradigm. Rather, they are just another way of turning individuals into willing participants in prevailing systems of classification.

A second, related objection is that there is something ethically problematic about markets for personal data. Since many of the services mentioned in the previous section explicitly promote themselves on the basis that they will enable people to sell their own data through such markets, this is taken as a reason to resist them.

A third objection, which is hinted at in the discussions cited above, is whether these supposedly user-centric personal data platforms really support the autonomy and agency of their users. The concern is that while they purport to help their users control their data and support informational self-determination, they might in fact subtly undermine them. A platform with ties to commercial partners might induce its users to make decisions that

191 Such behaviour seems consistent with research in behavioural economics illustrating how monetary incentives can backfire ((Thaler, 2007), (Gneezy & Rustinci, 2000)). It is possible that in offering cash, SAI profile services may actually risk some consumers to rejecting their proposition altogether.

favor its own profit motive rather than the user's own interests (an instance of the principal-agent problem (Grossman et al, 1983)). And even a platform that genuinely does attempt to put its users' interests first might be objectionable on the grounds of covert technological paternalism.

In the following sections, I outline and respond to these three lines of critique in detail. My aim is to give due weight to these criticisms whilst ultimately providing affirmation of the potential for these platforms to be genuinely empowering.

4.1 Open profiling and the logic of big data

The notion that individuals could actively participate in the construction and management of their digital profiles is presented by these platforms as an empowering feature. But critics argue that while purporting to be a revolutionary alternative to the status quo, this approach merely reinforces it. It is, they allege, just the latest in an ever-growing list of supposed 'market solutions' ((Neyland, 2013) (Milyaeva & Neyland, 2015)). Personal data empowerment initiatives are said to reflect a 'neoliberal promise of a responsible citizenry', and act as instruments of 'calculative power'.¹⁹²

These critiques are couched in the terms of the emerging field of critical data studies (CDS), which examines the unique theoretical, ethical and epistemological challenges of big data (Rob Kitchin & Lauriault, 2014). CDS attempts to inject this critical perspective wherever data is 'naively taken to denote objective and transparent informational entities' (Iliadis & Russo, 2015). These platforms, which make strong claims about the empowering potential of personal data, are therefore a prime target for such analysis.

In a more alarmist version of this analysis, Jacob Silverman portrays the idea of personal data stores as a particularly pernicious manifestation of big data-driven neoliberalism;

'this is to give into the logic of Big Data... Rather than trying to dismantle or reform the system...they wish to universalize it... This model would make all of human life part of one vast, automated dataveillance system... No social or behavioral act would be immune from the long arms of neoliberal capitalism. Because everything would be tracked, everything you do would be part of some economic exchange, benefiting a powerful corporation far more than you. This isn't emancipation through technology. It's the subordination of life, culture, and society to the cruel demands of the market.' (Silverman 2015)

The world that these critics portray arising as a result of such personal data platforms is a frightening and plausible one. But attempts to enhance personal data empowerment through tools that give individuals control over their data needn't necessarily lead us into such a dystopia.

¹⁹² In a recent conference panel on personal data (Draper, 2015), (Lehtiniemi, 2015). 'Calculative power' draws from (Callon & Muniesa, 2005). A similar point is made in (Milyaeva & Neyland, 2015). 'Neoliberalisation' here is used in the sense defined by (Brenner & Theodore, 2002).

In fact, a closer look at these tools somewhat dispenses with the idea that their makers implicitly accept the logic of big data. In fact, in asking individuals to question, evaluate and shape the data that defines them, these platforms actually invite their users to engage in critical reflection on some of the epistemic, ontological and normative aspects of data. To explore this further, I draw on the discourses through which these enterprises market their proposition to consumers, to understand how personal data is presented as both valuable and potentially empowering.

The marketing efforts of these platforms frequently appeal to consumers' anxiety and scepticism about the status of data currently used to profile them in big data systems. One argues that 'when it comes to the story of you, this joined up mass of data threads can't even be described as an unauthorized biography'.¹⁹³ The perceived epistemological and ontological deficiencies of big data are thus used to market their alternative form of profiling.

Furthermore, they do not naïvely portray their alternative as inherently more accurate or constitutive of the user's 'true' digital identity. Instead, the performative, fragmented and negotiated nature of marketing profiles is fully embraced. The user can 'step up to the digital mirror to see who you have created and curate, tweak and opt in or out at will'. Another platform allows the creation of multiple profiles to represent different aspects of ones life, with the slogan of 'different people, different you'.¹⁹⁴ These appeals to the ontological and epistemological complexities of online identity are actually quite reminiscent of the critiques of big data to be found in CDS research.

This alternative paradigm of profiling attempts to give consumers a greater influence over their own classification. In one sense, consumers have always played a role in shaping the categories they are placed in. Ian Hacking used the phrase 'interactive kinds' to describe the ways that people and other social entities are 'made up' through processes of categorisation (Hacking, 1990). Similarly, in a personalised digital environment, an individual's behaviour shapes their profile, and their profile can shape their future behaviour in turn. What differs in the case of user-owned and controlled profiles is that individuals are invited to intervene in a far more explicit and deliberate way in their classification, devising their own categories, with the potential to apply different categories to their multiple identities.

In this sense, they also open up the possibility of resistance to big data logic and marketer's systems of classification. That these tools could be a genuine threat to the marketing status quo is revealed quite tellingly in several remarks made by participants in the online discussion who took the position of the digital marketing industry. They were quick to point to the dangers of allowing individual consumers to modify their own data. They felt this would inevitably 'cloud the accuracy of the profile' compared to profiles based on 'genuine' data. Reflecting the notion that big data systems are 'impartial' systems of quantification (Porter, 1995), they claimed that 'the statistical approach may be more accurate'. Sources of bias would arise, they reasoned, because self-authored profiles might feed delusional or dishonest

193 <https://www.citizenme.com/public/wp/?p=175>

194 <https://angel.co/spoorr-me>

tendencies: 'individuals can reflect [that] they live in a mansion (for example) when they cannot afford it? That wouldn't be realistic at all'.

This reaction demonstrates how slippery the notion of 'genuine data' and 'accuracy' can be in the context of marketing profiles. A discussion about apparently objective notions of data quality quickly resolved to a set of normative considerations such as honesty, prudence, and responsibility. This dynamic reveals how the very idea of 'accurate' profiles belies the contested purposes and competing interests underlying their use. Of course, a dossier on an individual composed only of that information which is most useful to a marketer may look different to one composed by the individual themselves. But do marketing profiles exist solely as an instrument for marketers to predict and shape consumer behaviour? Or might they legitimately represent what an individual wants to project about themselves in a given context? The data associated with a profile could be considered 'accurate' in relation to the individual's devised purpose, which may or may not cohere with the marketer's interests in representing the 'real' consumer behind it.¹⁹⁵

In inviting consumers to participate in the construction of their profiles, these platforms render explicit previously opaque processes by which consumers are segmented. In referring to the ontological uncertainties of big data, they attempt to engage consumers in the very subject matter of critical data studies. In this respect, they target users who are capable of exercising a degree of reflexive engagement with discourses about data. They also illustrate the need for what Noortje Marres has termed 'experimental ontology', which directs attention to efforts to purposefully incorporate normative considerations into technological objects (Marres, 2013).

Far from naïvely adopting the precepts of big data, these data empowerment platforms actually embrace elements of critical data studies; the theoretical, ethical and epistemological controversies of big data are appealed to and turned into a marketing strategy to attract disgruntled users. In this sense, these platforms represent both a vindication and a challenge for the critical data studies paradigm out of which this criticism emerges. They vindicate CDS by showing how critical perspectives on the big data paradigm can motivate alternative systems. But they also challenge the notion that CDS scholarship is uniquely placed to uncover the naivety of industry's data hubris; in this case, CDS critiques are already well-understood and used by these emerging platforms to market their alternative.

My aim here has been to highlight how using these tools does not necessarily mean naively accepting the premises and logic of big data. Being able to control and modify profiles about oneself that are exposed to marketers is an empowering ability. This is especially important when marketing profiles are becoming increasingly synonymous with the digital profiles associated with other areas of our online lives. A major concern about the current digital advertising industry is its ability to collate

195 The general notion that supplying false or inaccurate data is always a bad thing is further challenged by recent work on pro-social deception and obfuscation (Brunton and Nissenbaum 2015), (Murray-Rust et al 2014).

information about consumers from multiple disparate sources. If this continues, then the versions of ourselves that we project through our social, civic and professional networks will increasingly be used to populate our consumer profiles. In that scenario, the ability to control one's consumer profile, and to separate it from other kinds of profiles one might have, is an important bulwark against the encroachment of the market. In this sense, far from turning us into willing participants in prevailing systems of big data and neoliberal capitalism, these tools could help us resist the encroachment of the market in the non-consumption aspects our lives.

In defending the empowering potential of open profiling, I do not wish to gloss over the negative implications of the possible commodification of privacy. My defense of these tools is due to their potential to allow individuals to define their own profiles, rather than due to their potential to create new personal data markets. I agree with critics that the extent to which such markets would fundamentally challenge the current economic model of personal data monetisation may be limited (Sevignani, 2013), and that many consumers may balk at the idea of voluntarily sharing their data with marketers (Sørensen et al., 2015). But the notion of empowerment that open profiling enables is separate from, and need not necessarily lead to the creation of such markets.

Unfortunately, the critiques above tend to blend these two aspects – the ability to control ones profile and opportunity to monetise it – together. I have defended the former against the charge of neoliberalism, but not the latter. The ethical status of this kind of direct personal data market therefore remains to be examined.

4.2 The ethics of personal data markets

Having separated the question of control from that of monetisation, we can now focus on the latter. Would the world be a better place if we could sell our own personal data?

Many of these platforms do indeed stake their claims to empowerment on the promise of giving users their fair share of the marketing revenue generated by their profiles. This idea of a property rights approach to privacy has been periodically suggested by scholars, policy-makers and commentators,¹⁹⁶ and more recently by popular technology writers such as Jaron Lanier in *Who Owns the Future?*.

In response to Lanier's proposal, Evgeny Morozov notes that 'to some, the very idea that our every decision is a piece of data to be monetized might seem appalling — and rightly so' ((Lanier, 2013), (Morozov 2013)). The approach allegedly encourages the data subject to become, in Foucault's words, an 'entrepreneur of the self', 'always eager to cash in on some personal trivia' (ibid). The argument here is that monetising personal data only serves neoliberal capitalism, with some inevitable corrupting effects on society and the self.

¹⁹⁶ See: (Spiekermann et al 2015); (P. M. Schwartz, 2004); (Samuelson, 2000); (Lemley, 2000); (Prins, 2006); (Bergelson, 2003); (Litman, 2000); (Murphy, 2012); (Payne & Trumbach, 2009).

Couching the argument in these terms may be a useful starting point, but simply calling an approach neoliberal, in a pejorative sense, is not a sufficient objection to it.¹⁹⁷ The idea that some new good or service ought to be produced and consumed according to market principles isn't necessarily neoliberal, let alone necessarily pernicious. Neoliberalism can be understood as a philosophy according to which society ought to be run according to market logic, a form of market fundamentalism (Davies 2014). But the mere belief that certain things ought to be left to the market does not a neoliberal make. When a new good becomes available – whether that be the contraceptive pill, the hoverboard, or personal data – we need to ask whether its production should be organised according to market principles, and non-neoliberals might answer in the affirmative without being sell-outs. It is therefore not immediately clear why creating a market for personal data should be seen as neoliberal in a problematic sense. Believing that neoliberalism is bad is not a sufficient reason to dismiss personal data markets, just as it is not a reason to dismiss markets for other kinds of new goods.

We must therefore look for independent considerations for or against personal data markets. The remainder of this section assesses whether they are problematic and if so, when and why.¹⁹⁸ This question could be approached from a variety of perspectives. One would be to conduct an economic analysis of a property rights regime in personal data, assessing the likelihood and severity of market failures. This approach has already been the subject of much analysis in law and economics, and as we shall see, it may not be the best way to evaluate whether personal data markets are desirable. Another approach is to ask whether, aside from the classical types of market failure (see 1.1.3), there might be independent moral reasons to limit such markets. This is the approach I pursue here.

Moral and political philosophers have sought to understand why certain market exchanges may raise special and unique ethical challenges.¹⁹⁹ These 'moral market limitation theorists' point to certain examples, including: the sale of human organs for transplants; votes in a democracy; sex and reproductive labour; indentured servitude or slavery; toxic waste disposal; and awards or professional positions normally based on merit.

The problem with markets for these kinds of goods is not necessarily that their exchange leaves buyers or sellers materially worse off. If that were their only failing, such markets might be 'fixed' and allowed to continue, by introducing new regulatory and contractual mechanisms, or in some cases through the redistribution of wealth, to ensure no one is left worse off.

But such markets seem morally problematic regardless of any subsequent

197 For discussion of such simple pejorative use of the term 'neoliberalism', and an overview of attempts to provide a more substantive descriptive account, see (Davies 2014).

198 Unlike some critiques of the equation of privacy and property which focus on their conceptual differences (e.g. (May, 1980)), the argument that follows will be couched in ethical terms

199 See e.g. (Walzer, 1983); (E. Anderson, 1990); (Satz, 2010); (Sandel, 2012); (Sandel, 2013)

redistribution, for reasons that are more readily understood through an ethical rather than economic lens. In other words, there are ethical and political reasons to impose certain limitations to what kinds of things can be bought and sold. For those whose political outlook lies anywhere between extreme libertarianism and extreme socialism, it is important to understand where these limitations might lie and why.

One reason is that there may be something intrinsic to the nature of certain goods that makes their sale morally repugnant. This might be due to the social meaning of the good becoming corrupted (Walzer, 1983). For instance, we might say that friendship becomes meaningless if it is bought; voluntariness may be part of the social meaning of friendship. A slightly different account claims that it is not necessarily the intrinsic nature of goods that makes their market exchange morally problematic, but the context in which the exchange takes place and the effect it has on the relationships between the buyer and the seller (and perhaps also others who are not part of the transaction) (Satz, 2010).

Markets in these goods inevitably put certain people in an unequal social standing even if they increase each party's individual material position. A market for votes, for example, is not wrong because it fails to allocate benefits optimally – it's a win-win for the vote seller and buyer – but rather because it undermines a necessary condition for democracy; that each citizen has equal standing.

We might consider personal data, in some limited respects, to be in this kind category. This is not to say that selling one's data is necessarily as morally problematic as selling one's democratic vote, only that it may be morally problematic for similar reasons. To understand why, we need to look at the context and purposes for which personal data is bought. And in the case of these platforms, we know that personal data is bought by marketers to help them work out how to influence consumers to buy their clients' products.

In this system, marketers are incentivised to seek out and exploit consumers' behavioural biases (a phenomenon Ryan Calo calls 'digital market manipulation' (Calo, 2013b)). Calo notes:

Today's firms fastidiously study consumers and, increasingly, personalize every aspect of their experience. They can also reach consumers any time and anywhere, rather than waiting for the consumer to approach the marketplace. These and related trends mean that firms can not only take advantage of a general understanding of cognitive limitations, but can uncover and even trigger consumer frailty at an individual level. (ibid. p1)

The ability to collect vast amounts of personal data is the catalyst for this kind of activity. If Calo's thesis is correct, personal data contributes to subtle, small, but very real encroachments on individual autonomy. This may sound like an extreme interpretation of the intentions of digital marketers, but consider the following quote from an industry report on the future of marketing, published in 2013:

“In the future advertising will be tasked with planting seeds of desire, expectations, aspirations that intrigue and pull the consumer

along the path to thinking that it was his or her idea, giving a sense of ownership and full decision-making power.” (Billey, 2013)

In these circumstances, selling one's personal data could amount to selling a portion of one's autonomous decision making power. Note that this isn't an objection to sacrificing autonomy in return for money; otherwise it would rule out most forms of employment. What is at issue is a consumer's autonomous decision-making power. Allowing personal data to flow in this way could contribute to unequal standing between consumers and marketers.

It could also create similar problems between consumers, including those who decide not to sell their data at all, who might suffer in more significant ways. As Scott Peppet argues, providing economic incentives for disclosing personal data may lead to a transition from a 'sorting' to a 'signalling' economy (Peppet, 2011). In this scenario, marketers, insurance firms, employers and other actors pay people to provide them with data on their characteristics, rather than attempting to guess them based on statistical models. In such a situation, 'those who refuse to disclose their information will be assumed to be withholding negative information and therefore stigmatised and penalised' (ibid, p. 1156). The scenario thus sets those who choose not to sell their data against those who do.²⁰⁰

These arguments do not imply that the direct sale of personal data should necessarily be banned. Regulation comes in many forms, and exactly how it should be implemented in this context is a complex, delicate and urgent question. The point here is that there may be moral or political reasons to intervene in personal data markets to ensure they do not have a dis-empowering effect – either by undermining individual autonomy or creating unequal relations between consumers.

This points to a deficiency of the economic welfarist approach to privacy and data protection. If we restrict our analysis to the negative externalities imposed on consumers by data trading, the problem is simply that data brokers impose harms on individual consumers without adequately compensating them. In this context, enfranchising those individuals to become sellers of their own data seems a sensible and obvious solution. But, as I have argued, the problem goes deeper than this. If consumers become enfranchised producer-vendors in the current market for personal data, they may still end up sacrificing a degree of autonomous decision-making power, further entrench their unequal standing with marketers, and potentially create new divides between each other. The fact that consumers get paid for their data in the process is not particularly empowering, on balance.

4.3 Personal Data Empowerment and the Ideal Observer

The last two sections have assessed two critiques of the notion of personal data empowerment through open profiling. I will conclude by considering a final challenge, and outlining a principle by which its development can be

200 As Richard Posner notes, the incentives to disclose would likely lead to full disclosure by every rational actor, due to a “pooling equilibrium ... in which privacy is ‘voluntarily’ surrendered, making the legal protection of privacy futile” (R. Posner, 1998)

assessed.

The discussion thus far has focused on personal profiling services. As outlined in the introduction, these are one part of a broader range of personal information management services (PIMS) (Ctrl-Shift, 2014). PIMS aim to empower individuals by equipping them with the tools they need to benefit from their own data on their own terms.

PIMS should not just be seen as an answer to privacy concerns, but also as a means to equip individuals with the capacity to better interact with the complexities of the modern age, particularly in consumer markets. In this context, they provide a new layer between consumers and traditional suppliers of goods and services, acting on behalf of the former to engage with the latter. They aim to bring value to individuals by gathering data, analysing and marshalling it in order to cut through the complexity of confusing, difficult and boring consumer-related tasks.

If a basic function of PIMS is to support, augment or otherwise enhance individuals choices in complex markets, how can this be done in a way that empowers rather than undermines their agency? There is a danger that the decisions and actions that arise from PIMS cease to be attributable to the individuals they operate on behalf of. Users may begin to feel the decisions and actions the service undertakes are not implementations of their own will but rather impositions from an outside force. If individual agency becomes disassociated from the service it is enacted through, then in what sense is the service genuinely empowering? On this view, rather than restoring individual agency, PIMS might undermine it.

A promising answer to this problem of agency lies in the concept of an 'ideal observer', originating in Enlightenment thought and developed further in contemporary moral philosophy.²⁰¹ The general idea is that it is possible to consider not just how individuals actually do choose and act; but how they would choose and act given greater levels of knowledge and the capacity for rational deliberation. Despite being hypothetical, such choices might still be objectively and inextricably linked to the individual as a function of their existing character, circumstances, values and preferences, as opposed to choices made by an outside agent with different values. In this sense, the notion of the ideal observer can help define a non-paternalistic account of what is in the 'best interests' of an individual.

For Bernard Williams, claims about what an individual should do are equivalent to the result of informed deliberation starting from their current motivations; a logical deliberative route must be traceable from an individual's current motivations to their putative 'ideal' choices (Williams & Quinton, 1973). This way of thinking allows us to accept that an individual is capable of misjudging what is in their own interest right now, whilst maintaining that an accurate assessment of their genuine best interests must proceed from their existing values, personality and motivations. The ideal observer is not an external hector laying down universal moral laws, but

²⁰¹ The notion can be traced back to Adam Smith's notion of an 'impartial spectator' (Smith, 1759) [1981] I, i,v, 26), and was developed further in 20th century analytic philosophy, e.g. (Firth, 1952), (Williams & Quinton, 1973), (Railton, 1986).

instead recognises the nuance and detail of a particular situation and the unique attributes of the individual within it. Ignore this, and discussions about an individual's 'best interests' risks treating them as a mere channel between the input of 'the utility network which the projects of others have in part determined' and 'an output of optimific decision' (ibid, p. 260).

This an important consideration for designing PIMS which seek to empower individuals to serve their best interests without undermining their autonomy. On the one hand, individuals' existing choices cannot always be taken at face value.²⁰² But this doesn't mean that PIMS need to be paternalistic, ignoring their users' existing motivations and making decisions for them, or 'nudging' them towards 'better' behaviours ((Sunstein & Thaler, 2008); (Acquisti, 2009)). Instead, PIMS could take the individual's existing motivations, values and preferences as the starting point for decision support and implementation ((Grist, 2010); (Binns, 2013b)).

Whilst the rather lofty goal of embodying an abstract philosophical ideal is unlikely to feature in entrepreneurs' business plans, it nevertheless offers a framework for analysing market-driven empowerment through PIMS. The ideal observer is just that – an ideal – which is impossible to fully embody, but possible to strive towards. Personal information management services have the potential to bring individuals far closer to this ideal than they ever could have on their own. In that sense, they have the capacity to be genuinely empowering.

202 This is because consumers have 'bounded rationality' (Simon, 1972), and are subject to cognitive biases (Kahneman, 2011). In the context of privacy, this is manifested in various 'paradoxes' of consumer behaviour ((D Solove, 2013); (Brandimarte, Acquisti, & Loewenstein, 2012); (Taddicken, 2014); (Xu, Luo, Carroll, & Rosson, 2011); (Wilson, Hall, Az, & Valacich, 2012); (Mainier & O'Brien, 2010); (Acquisti, 2009)).

Part 5: Meta-regulating privacy and the open corporation

I have now covered two different aspects of Openness for Privacy – *open data* and *open processing*. The *open data* approach shows some potential, but is applied at an abstract and general level, not tailored to data processing activities as they apply to different individuals in idiosyncratic ways. By contrast, the notion of *open processing* allows the individual to scrutinise and interact with a system as it processes their personal data. Some key challenges for this approach were raised in the introduction: why would industry adopt it? What are the benefits to individuals?

The section above has provided some answers. The user study suggests that the marketing industry can benefit by giving up some control over consumer profiles, because this may not only increase response rates to targeted messages, but also potentially avoid the negative attitudes some consumers currently have towards behavioural profiling. I then explored whether these platforms could be genuinely empowering for individuals. The overall conclusion is that they can, and that they therefore have the potential to be a mutually beneficial arrangement between consumers and businesses.

It is important that these kinds of arrangements are understood and encouraged by regulators, since they demonstrate that certain technology platforms and business models might serve regulatory goals by design. But they are not a silver bullet; not all organisations can feasibly adopt the open processing approach. They may be a promising development in the context of marketing profiles, but this still leaves familiar data protection and privacy concerns in other contexts untouched. They do not enlighten regulators on how to regulate the risks of personal data processing more generally.

To this end, this final paper explores one way that openness might be integrated into the regulatory practice of data protection. The first and second papers were both about 'opening up' the use of personal data in various ways; similarly, this final paper is about 'opening up' the manner in which personal data is regulated.

It does so in the context of the EU's proposed General Data Protection Regulation (GDPR).²⁰³ There are many aspects of the GDPR proposal that one might consider in relation to the concept of openness, including the frequent references to public consultation,²⁰⁴ information to be provided to data subjects,²⁰⁵ and transparency of privacy policies.²⁰⁶ Some of these aspects have already been discussed by legal scholars.²⁰⁷ Despite their clear connections with the general theme of Openness for Privacy, the final paper does not focus on these aspects; partly because they have already been well-studied, and partly because they are more straightforward manifestations of OfP that have much in common with the previous papers. Instead, it focuses

203 (European Commission, 2010a)

204 E.g. Articles 44(1g), 34,

205 E.g. Recital 48, Articles 51a, 75a

206 Recital 32

207 E.g. (Hildebrandt, 2012)

on part of the proposed GDPR which is less obviously connected to openness; the provisions on Privacy Impact Assessments (PIAs), primarily found in Article 33.

PIAs have been proposed by the Commission partly as a replacement of the system of notification covered in the first paper. As the initial communication noted:

'While [the notification] obligation produces administrative and financial burdens, it did not in all cases contribute to improving the protection of personal data. Therefore such indiscriminate general notification obligation should be abolished and replaced by effective procedures and mechanism which focus instead on those processing operations which are likely to present specific risks... In such cases, a data protection impact assessment should be carried out' (Recital 70)²⁰⁸

The final paper is therefore linked to the first paper in the sense that PIAs aim to replace the notification obligation which underpinned the data used in the first paper.

The final paper also has some general congruences with the second paper. The PIA system proposed in the GDPR foresees an important role for stakeholders in engaging, challenging, revising and otherwise exerting influence on a data controller's policies for handling personal data. In this sense, it is similar to the notion of open processing, which emphasises the ability to scrutinise and intervene in the use of personal data. In this case, rather than being part of a consumer-facing service, engagement by stakeholders is the direct aim of an obligation imposed by the regulator.

Further connections between PIAs and the general theme of Openness for Privacy will be revisited in the conclusion.

208 (European Commission, 2010a)

5.1. Introduction

Abstract

Privacy Impact Assessments (PIAs) are a tool for organisations to manage privacy risks. PIAs have emerged in various jurisdictions since the 1980s,²⁰⁹ initially as a purely voluntary tool. They are now likely to become a mandatory requirement under the European General Data Protection Regulation (GDPR).²¹⁰ This article addresses PIAs from the perspective of regulatory theory. The transition of PIAs from a voluntary tool to a mandatory requirement raises questions about their purpose and role, as well as implications for the direction of data protection in Europe more generally.

Previous analyses have tended to assess PIAs in relation to a limited set of regulatory categories, namely self-regulation, command-and-control regulation, or some form of 'co-regulation'.²¹¹ Drawing from regulatory theory, this article suggests a more nuanced account of the mandatory PIA regime proposed in the GDPR. It argues that this regime can be understood as a form of 'meta-regulation'.²¹² The final section draws on a framework for assessing the prospects of meta-regulation, in order to assess the prospects for a meta-regulatory approach to PIAs.

“It is obvious that technology evolves faster than legislation. The various parties gathered today have recognised this and decided that this Privacy Impact Assessment Framework was the most effective and efficient way to protect the privacy of European citizens without stifling innovation”

- *Neelie Kroes*²¹³

Decades-old regulatory frameworks introduced to deal with the ethical quandaries of the digital era have begun to appear ever more antiquated. Ubiquitous data collection, data mining and profiling of individuals raises concerns about privacy, autonomy, and discrimination.²¹⁴ Policy-makers, anxious to balance the interests of industry and citizens, are eager to find nuanced regulatory mechanisms capable of dealing with the complexities of modern technology.

This is the environment into which Privacy Impact Assessments (PIA) have emerged. A PIA is a process of assessing the possible privacy implications of new uses of personal data.²¹⁵ Proponents of PIAs argue that they could be

209 ((Clarke, 2009),(Tancock et al., 2010c))

210 (Wadhwa & Wright, 2013)

211 (Wright et al., 2014)

212 (Parker, 2002)

213 Vice-President of the European Commission for the Digital Agenda, at the Privacy and Data Protection Impact Assessment Framework Signing Ceremony, Brussels, 6th April 2011. Transcript retrieved from [http://europa.eu/rapid/press-release_SPEECH-11-236_en.htm]

214 For an overview on the threat to privacy, autonomy and discrimination see (Hildebrandt & Gutwirth, 2008) p2; in particular chapters by (Canhoto & Backhouse, 2008), (S. van der Hof & Prins, 2008); Schreurs. Transparency – Hildebrandt. (Barocas & Selbst, 2014);

215 (David Wright & Hert, 2012)

a promising solution to address privacy and data protection concerns.²¹⁶ PIAs are designed to help organisations implement 'privacy by design', by incorporating privacy considerations into their activities and projects from the early stages, thus reducing the risk of privacy violations and any associated regulatory action or reputational damage.²¹⁷ As the following sections describe, PIAs have evolved from a tool used by some organisations voluntarily, into an internationally recognised and increasingly mandated practice. PIAs have been lauded as 'the most comprehensive tool yet available for policy-makers to evaluate new personal data information technologies before they are introduced', capable of imagining the 'unknown unknowns'.²¹⁸

Given the ascendant enthusiasm for PIAs and the perceived risks of data processing, it is not surprising that their use is increasingly urged by regulators.²¹⁹ This has culminated in the inclusion of new provisions in the European Union's proposed General Data Protection Regulation (GDPR).²²⁰ Article 33 of the GDPR requires organisations to conduct impact assessments in a variety of contexts which are likely to 'present specific risks to the rights and freedoms of data subjects by virtue of their nature, their scope or their purposes'.²²¹

A great deal has been written by academics and practitioners about the ideal form of PIAs.²²² However, debates about the merits and purpose of PIAs have generally not drawn significantly from the large body of regulatory theory. This article aims to fill this gap. In particular, it is argued that in making PIAs a regulatory requirement, the Commission have transformed them from a tool of self-regulation into one of 'meta-regulation'.²²³ This approach has the potential to address some of the key challenges identified by the Commission in their motivation for data protection reform.

5.2. Privacy Impact Assessments: Background

This section provides some background on PIAs relevant to the discussion in subsequent sections.

216 See (David Wright et al., 2012), (Stewart, 2012), (Information Commissioner's Office, 2014), (Tancock, Pearson, & Charlesworth, 2010b), (Wadhwa & Wright, 2013), (Adam Warren & Charlesworth, 2012), (Microsoft, 2013), (Tancock et al., 2010a).

217 (Cavoukian, 2006)

218 (G. Marx, 2012)

219 See, for instance ' <http://www.tbs-sct.gc.ca/pol/doc-eng.aspx?id=12451> PIAs are mandatory for public bodies in the US, Canada and elsewhere. Report 'Data Handling Procedures in Government' 2008 made PIAs mandatory in central government departments. RFID mandatory in EU. EDPS fully supports making PIAs mandatory under certain threshold conditions.

220 (European Commission, 2012)

221 Unfortunately, the draft text has not been entirely consistent in its use of terms; for discussion of this, see section X below.

222 For an introduction, see (David Wright & Hert, 2008)

223 (Parker, 2002)

5.2.1 Origin of PIAs

PIAs could be seen as an evolution of provisions set out in early data protection regimes in which organisations were required to register, notify and check with national authorities to ensure compliance prior to processing.²²⁴ However, advocates of PIAs have argued that they go beyond mere compliance checking, citing a need for a wide-ranging, contextually sensitive and 'holistic' approach.²²⁵ According to a history of PIAs produced by Roger Clarke, the concept of a PIA is derived from instruments in other policy areas like environmental law, which allow for this broader perspective.²²⁶ They became established as a concept in policy circles in the late 1990's, primarily in English-speaking common law countries, particularly Canada, Australia and New Zealand.²²⁷

By 2000, PIAs were regarded by a subset of Privacy Commissioners, consultants, data protection professionals and academics as an 'essential tool for data protection'.²²⁸ There was also growing interest from the private sector at this time, partly due to the perceived challenges facing multinational organisations in cross-border compliance.²²⁹ By the mid 2000's, a number of national privacy and data protection authorities, government departments, and regulators had begun producing guidance on conducting PIAs.²³⁰

These guidance documents indicate that policy-makers had converged on a

224 The European Data Protection Directive of 1995 institutionalised 'prior checking' and notification with a national authority (Article 20), but similar provisions are contained in data protection acts pre-dating the 1980 OECD Guidelines, e.g. Sweden (1973), Austria, Denmark, France and Norway (1978). Prior checking has been characterised as a 'forerunner' to PIAs in (Le Grand & Barrau, 2012)

225 For example, PIAs should allow for 'contextual information' (Wright & Wadhwa, 2012, p2), be 'holistic in nature' (Tancock et al., 2010a, p1) and 'wide-ranging' (A Warren, Bayley, & Bennett, 2008).

226 Clarke explains that 'technology assessments' were conducted by the US Congress and in some European contexts, while 'impact statements' and 'impact assessments' are associated with environmental regulations (Clarke, 2009). PIAs were arguably in use as far back as the early 1970's, with the first documented practice resembling a PIA in 1973 (according to (L Hoffman, 1973) as noted in (Clarke, 2009)). The phrase itself is not used until 1994, but the similar phrase 'privacy impact statement' appears in an official document in 1984. According to (Clarke, 2009), this appears in (Information and Privacy Commissioner of Ontario, 1994). See (Clarke, 2009) "It would ... appear that the concept, although not yet the term, was in use in some quarters as early as the first half of the 1970s... the first literature reference to the term 'privacy impact statement' located by this author is ... a 1984 document of the Canadian Justice Committee" (Ibid p. 126).

227 (Clarke, 2009) p. 126

228 (Flaherty, 2000)

229 (Karol, 2001)

230 The PIA guidance produced by national authorities in Australia, the UK and Canada are regarded by some as the 'most comprehensive and practical guidance documents available in any jurisdiction' (Clarke, 2011). European countries with PIA guidance include Slovenia, Spain, Germany and the UK. Other privacy / data protection authorities with guidance documents include Hong Kong, New Zealand, Canada, and the US (multiple government departments).

set of core features of PIAs.²³¹ They stress that a PIA is not just a tool or a report, but a process. The subject matter of a PIA is a 'project' (which encompasses any 'system, database, program, application, service or scheme, and enhancement of any of these, or an initiative, proposal or a review, or even draft legislation'²³²), and its impact on privacy.²³³ They are distinguished from activities with a narrower scope, such as privacy audits, because they begin *before* rather than *after* implementation.²³⁴ Nor are they synonymous with legal compliance assessments, because they deal with 'qualitative matters of legitimacy, participation and proportionality' rather than just compliance with specific rules.²³⁵ They are for 'organisations of all sizes' and can be performed in-house or by external consultants.²³⁶ While these guidance documents often include accompanying templates, organisations are cautioned against seeing the PIA process in terms of a one-size-fits-all solution.²³⁷

More recently, the Privacy Impact Assessment Framework Consortium (PIAF), an EU-funded PIA advocacy group, has defined a PIA as:

'a methodology for assessing the impacts on privacy of a project, policy, programme, service, product or other initiative and, in consultation with stakeholders, for taking remedial actions as necessary in order to avoid or minimise negative impacts'²³⁸

According to their proponents, a primary aim of PIAs is to manage risks associated with threats or vulnerabilities arising from processing personal data.²³⁹ They should therefore be integrated with risk management processes (but are not synonymous with them). As well as impacts on individuals, they

231 The following summary of PIA guidance documents draws heavily from (David Wright & De Hert, 2012) and (Clarke, 2011)

232 (Information Commissioner's Office, 2007)

233 The fact that 'privacy' is a highly contested concept is often noted in discussions of PIAs (e.g. according to Gary Marx, "Privacy is a general term and there are endless arguments about what it applies to and if it is the best term to capture contemporary concerns." (G. Marx, 2012) p. vii). Nonetheless, much of the ensuing debate appears to assume that there is sufficient agreement on its meaning to anchor the idea of PIAs.

234 The ICO defines privacy audits as "the detailed analysis of systems that are already in place against a prevailing legal, management or technology standard" (Information Commissioner's Office, 2014)

235 (De Hert, 2012)

236 The ICO state that they have "published our updated privacy impact assessments code of practice to help organisations of all sizes ensure that the privacy risks associated with a project are identified and addressed at an early stage during a project's development" (Society for Computers and Law, 2014)

237 (Information Commissioner's Office, 2007)

238 According to the project website, 'PIAF (A Privacy Impact Assessment Framework for data protection and privacy rights) is a European Commission co-funded project that aims to encourage the EU and its Member States to adopt a progressive privacy impact assessment policy as a means of addressing needs and challenges related to privacy and to the processing of personal data.' - from [www.piafproject.eu/About%20PIAF.html], Accessed on 6th September 2015.

239 'Determining the risk(s) resulting from various vulnerabilities and threats requires some analysis and assessment, which is what a PIA can and should do.' - p13 (David Wright & De Hert, 2012)

should consider direct and indirect impacts on the organisation (e.g. fines and penalties, opportunity costs or damage to brand reputation). The process of a PIA should be documented in a report, which should include a description of the assessment, as well as terms of reference, deliverables, and responsibilities. The purported benefits for organisations who conduct PIAs include greater transparency and trust, confidence, valuable stakeholder input, understanding and respect, avoiding liabilities and crises down the line, and identifying cost effective solutions.²⁴⁰

At the heart of the PIA philosophy is the idea that privacy problems do not happen completely at random, detached from the actions and policies of organisations. In this sense, 'privacy is not like the weather', as Gary Marx writes in a foreword to a book on PIAs, which epitomises the approach.²⁴¹ Even if specific privacy problems cannot be predicted, the probability of their occurrence and the severity of their effects can be reduced through systematic consideration and planning early on in a project. The implication is that privacy problems occur in semi-predictable ways, as a result of a finite set of causes which can be isolated and addressed.

Marx claims that 'various types of privacy problem do not occur randomly but tend to cluster' at 'different stages of a project',²⁴² and range from what tools are used, who the data subjects are, how data is collected, processed, analysed, and interpreted, how the data is used in action, and what happens to the data when the project is over. Marx argues that 'problems occur at at least one of these stages',²⁴³ and that the costs and challenges of applying limitations or controls are greater the later the stage of the project life-cycle, hence the focus on intervening at the earlier stages.²⁴⁴ He characterises PIAs as a tool for 'thoughtful realists', that can bring 'slices of insight and amelioration ... through transparency and commitment to democratic values'.²⁴⁵

5.2.2 Adoption and implementation of PIAs

Support for PIAs amongst regulatory authorities has gradually led to their use by both public and private organisations, although the drivers of adoption have differed between sectors. PIAs have become mandatory for many public bodies in certain jurisdictions, while private organisations have so far undertaken PIAs at their own discretion.²⁴⁶ The precise extent of PIA activity in either sector is not clear, because much of it goes unreported and a relatively small number of PIAs are published openly.²⁴⁷ However, some

240 p15 (David Wright & De Hert, 2012)

241 (G. Marx, 2012)

242 Ibid. p. xi

243 Ibid. p. xii

244 This echoes a common argument for PIAs, e.g. (Pritchett, 2010):

“Prevention is better than a painful cure, and PIAs provide a structured way to take a timely look at the preventative steps which may be necessary.”

245 Ibid p. xiv

246 For instance, PIAs are mandatory across Europe for projects involving RFID, and in Canada and the US for e-government projects.

247 Various public bodies in the US and Canada maintain comprehensive records of publicly available PIAs. For discussion of the dearth of publicly available PIA reports, see (David Wright, 2014).

reports indicate significant engagement by large organisations in the private as well as public sector, and the level of discussion amongst privacy and data protection professionals suggests a growing awareness of PIAs across sectors.²⁴⁸

Indeed, some companies consider themselves to be at the cutting edge of PIA practice. A book on PIAs features contributions from Siemens, Nokia, and Vodafone describing their extensive PIA frameworks and policies.²⁴⁹ Various sector-specific PIA initiatives are underway or have been proposed,²⁵⁰ and an ISO standard involving PIAs is in progress.²⁵¹ Numerous specialist consultancies offer their own PIA methodologies.²⁵² These are often accompanied by digital tools to aid the PIA process, including automated decision support systems and template-based PIA management and reporting software.²⁵³

European interest in PIAs was arguably precipitated by the 2007 PIA handbook, published by the UK Information Commissioner's Office, and the European Commission's recommendation of PIAs for RFID (Radio Frequency Identification) projects in 2009.²⁵⁴ In both cases, PIAs were encouraged as best practice, and a means of demonstrating compliance with data protection laws.²⁵⁵ By 2011, Neelie Kroes, Vice-President of the European Commission for the Digital Agenda, announced that the use of

248 TrustE report found half of the large organisations they interviewed regularly conducted PIAs. See also: PIAF, SAPIENT projects. (David Wright, Wadhwa, Lagazio, Raab, & Charikane, 2012) found that private sector organisations had completed a higher number of PIAs on average compared to public bodies. Professional bodies also show interest; the International Association of Privacy Professionals has published 39 articles mentioning PIAs over the last year (through a search on www.privacyassociation.org from May 2014 – 2015).

249 See chapters 11-13 in (David Wright & Hert, 2012). Although, “some of this information could be seen as well-reasoned PR for the organisations concerned” as Pritchett 2012 notes

250 These include PIA initiatives for smart metering, financial services, and a web data specific PIA framework (Beaumont, 2014).

251 See ISO 27001 (<http://www.iso.org/iso/home/standards/management-standards/iso27001.htm>)

252 See for example, methodologies from Truste [<http://www.slideshare.net/trusteprivacyseals/tips-tools-for-conducting-effective-pias-truste-webinar>], and Trilateral Research [<http://trilateralresearch.com/services/#impact-assessment>], accessed on the 5th September 2015.

253 The earliest PIA tools took a 'decision tree' approach, according to (Tancock, Pearson, & Charlesworth, 2010a), who describe two new approaches – PRAIS and HPPA (Hewlet Packard). A more recent tool is available from IAPP / AvePoint; an IAPP spokesman claims that 'the need for an increase in the automation of privacy impact assessments is a global one.' (Pfiefler, 2014). However, some criticism has been levelled at automated PIA tools, e.g. those used by US IRS (FEDWeek, 2013)

254 The idea that the handbook catalysed PIA use in Europe is suggested in (Tancock et al., 2010c). The RFID PIA framework, developed by industry, was endorsed by the Article 29 Working Party in February 2011 (Spiekermann, 2011).

255 'Conducting a PIA is not a requirement of the Act, but undertaking one will help to ensure that a new project is compliant' - (Information Commissioner's Office, 2014) p.3

PIAs was 'potentially also the start of a new policy approach, in fact a new commitment to involving all stakeholders in the process of solving privacy problems'.²⁵⁶ Since the ICO's handbook, national authorities have published PIA guidelines in several other EU member-states, including Spain, Finland, Germany, and Slovenia. In recent years, several large-scale research projects examining the use and potential of PIAs have been funded by EU grants.²⁵⁷

Most of this activity in Europe had been pursued on the basis that PIAs are recommended but not required.²⁵⁸ However, in 2012, the Commission proposed the new GDPR, which would make impact assessments mandatory in certain 'high risk' contexts.²⁵⁹ The Parliament stated that impact assessments 'are the essential core of any sustainable data protection framework'.²⁶⁰ Since then, data protection authorities and law firms have begun recommending that organisations start conducting impact assessments now (following existing PIA guidance), in order to pre-emptively comply with the GDPR.²⁶¹

PIAs may soon go from being a purely voluntary tool to a mandatory requirement in Europe – at least, for certain kinds of data processing – in less than a decade. Why has the Commission taken this approach? Is it likely to succeed, or is this transformation of PIAs misguided? To answer these questions, it will help to consider what kind of regulatory instrument a PIA might be, and the merits and problems associated with it.

5.3. Regulatory theory of PIAs

In the existing literature, PIAs have been mainly been discussed in relation to three broad, traditional categories of regulatory approach: legal regulation, self-regulation and co-regulation.

Unfortunately, there is no generally accepted framework or standardised terminology in the field of regulatory theory which might give these terms a single precise definition.²⁶² Regulatory theorists typically conceive of different regulatory approaches as existing on a spectrum between 'pure' legal regulation and 'pure' self-regulation.²⁶³ Legal regulation is often

256 See footnote 210

257 See [www.sapientproject.eu], [www.piafproject.eu].

258 Apart from the aforementioned cases where member states impose mandatory requirements on public bodies to conduct PIAs

259 This change has been widely regarded as replacement to the old system of notification, which was regarded as onerous and indiscriminate – see (Pederson, 2005)

260 Recital (71a)

261 “Businesses that follow the new code of practice on PIAs that the Information Commissioner's Office (ICO) has published will be better prepared for complying with EU new data protection laws when they are introduced.” (Pinsent Masons, 2014)

262 See e.g. (Richards, 2000) which catalogues the many different terms used by different authors to refer to the same kind of regulatory instrument.

263 E.g. (Bartle & Vass, 2007) p889, claim: “at one end of the spectrum is a pure form of self-regulation, the perception being no role for the state beyond normal criminal and civil law”. Similarly, for (Saurwein, 2011): “regulation takes place on a continuum between pure state regulation,

understood as regulation by the state in the form of legal rules backed by criminal or civil sanctions; an approach sometimes referred to (usually pejoratively) as 'command and control'.²⁶⁴ Self-regulation, by contrast, might be characterised as rules that private actors impose on themselves in the absence of state intervention or coercion.²⁶⁵ Between these two extremes lie various terms denoting different configurations of state and private activity, including 'co-regulation', which is generally used to denote some form of collaboration between state and private actors in at least some aspect of the regulatory process.²⁶⁶

All of these terms are fraught with ambiguity, due in part to a lack of consensus on answers to fundamental questions, such as what regulation itself even is.²⁶⁷ However, such terminological and theoretical ambiguities have not prevented the use of these terms in discussions about PIAs in the academic, policy and professional literature.

5.3.1 PIAs as self-regulation

PIAs have traditionally been pursued as a self-regulatory instrument. This is understandable given that, as the history above shows, PIAs emerged primarily as voluntary tools. Apart from certain public bodies in certain jurisdictions, organisations who undertake PIAs today do so without being required by the regulator. Indeed, proponents have often suggested that the benefits of PIAs – such as the mitigation of reputational and other risks – ought to be sufficient to motivate many organisations to undertake them of their own accord.²⁶⁸ The perceived merit of this self-regulatory approach is that organisations are free to develop PIA processes which best suit their specific circumstances. This means PIA practice should develop flexibly and organically to suit the needs of data controllers. They are therefore described as an example of 'reflexive best practice', in contrast to the 'sledgehammer' approach taken in other areas of data protection policy.²⁶⁹

on the one hand, and pure self-regulation, on the other, and can generally be understood as a combination of state/public and societal/private contributions, which are closely interlinked". See also (Lehmkuhl, 2008); (Gunningham & Rees, 1997); (Sinclair, 1997) for articulations of the supposed 'spectrum' of regulatory approaches.

264 E.g. (Black, 2001) defines it as 'regulation by the state, which is often assumed to take a particular form, that is the use of legal rules backed by criminal sanctions: 'command and control' (CAC) regulation' (p.105).

265 Although there are a multitude of different varieties of self-regulation. For a thorough discussion, see (Black, 2001) and (Bartle & Vass, 2007).

266 Linda Senden claims that co-regulation 'can also be said to situate itself somewhere between legislation on the one hand and 'pure' self-regulation on the other' and is defined in the European context as 'the existence of some form of relationship between binding legislation and voluntary agreements in a particular area' (Senden, 2005)

267 As Bettina Lange notes, 'legal regulation has been analysed from various theoretical perspectives, such as welfare economics (Ogus, 1994), Marxism ((Jessop, 2001) 83–92), Foucauldian 'governmentality analysis' (Dean, 1999), discourse analysis (Black, 2002) and systems theory ((B. Lange, 1998): 449–471; (Paterson, 2000): 7)' (Bettina Lange, 2003). See also (Orbach, 2012)

268 (David Wright & Hert, 2012) p15

269 (I. Brown & Marsden, 2013) p. 167.

5.3.2 Ensuring implementation through mandatory PIAs

The main reason for making PIAs mandatory seems straightforward. Unless PIAs are made a mandatory requirement, their use will be confined to those organisations who are already motivated to comply, leaving the risks arising from the processing operations of other organisations to continue unmitigated. If PIAs are indeed as effective at mitigating risks as their proponents claim, regulators ought to ensure they are adopted wherever appropriate. This appears to be the primary argument made by David Wright, in an article in favour of mandatory PIAs published in 2011, prior to the first proposal for the GDPR.²⁷⁰ The perceived need for a mandatory requirement could also be the result of a general scepticism about the market's ability to self-regulate. Many forms of self-regulation are regarded as having severe limitations, particularly in the face of events like the global financial crisis.²⁷¹

5.3.3 Mandatory PIAs as legal regulation: would they suffer the drawbacks of 'command and control' regimes?

However, the trend towards making PIAs mandatory clearly changes their status as a self-regulatory instrument. In the regimes established for public bodies in the US and Canada, and in the regime that will likely be established for both public and private organisations in the EU under the GDPR, PIAs are a legal requirement backed by punitive sanctions. This would appear to place them in the category of traditional legal regulation; a requirement created through state legislation, enforced by regulators through fines, giving little discretion to regulatees. This opens them up to a set of common critiques associated with so-called command-and-control approaches. This section provides an overview of these critiques (although, as we shall see in the following sections, there are alternative accounts of the kind of regulatory approach that a mandatory PIA regime could involve which might avoid these problems).

As Wright rightly acknowledges, making PIAs mandatory raises some potential problems. Primary among these are bureaucratisation. Wright notes that:

“PIAs are only valuable if they have, and are perceived to have, the potential to alter proposed initiatives in order to mitigate privacy risks. Where they are conducted in a mechanical fashion for the purposes of satisfying a legislative or bureaucratic requirement, they are often regarded as exercises in legitimisation rather than risk assessment.”²⁷²

Concerns about bureaucracy have also already been expressed by some

270 (D Wright, 2011)

271 See e.g. (Black, 2012)

272 (D Wright, 2011) p. 11, citing (A Warren et al., 2008)

member states,²⁷³ scholars,²⁷⁴ data protection professionals and industry in response to the provisions of the GDPR.²⁷⁵ They suggest that mandatory PIAs may entrench a command-and-control approach, becoming a 'box ticking' exercise, undermining their intended purpose.²⁷⁶ It is generally recognised that rule-based, coercive and punitive methods applied solely by regulators tend to lead to 'ritualism' (following rules without understanding why they are there), and 'creative compliance' (following the letter of the rules in such a way as to undermine their overall purpose, as in elaborate tax avoidance schemes).²⁷⁷

Those who have developed PIAs as voluntary tools stress that there can be no comprehensive check-list of necessary and sufficient procedures for organisation to follow in undertaking them.²⁷⁸ Rather, organisations should tailor their PIA process to suit their particular circumstances. But if PIAs are mandatory (with the threat of penalties for failure to undertake them adequately), there may inevitably be pressure to create such a check-list, as we have already seen in discussions between the Council and the Parliament. As Bert-Jaap Koops and others have argued, this could result in PIAs focused more on demonstrating compliance with specific procedures rather than on flexible, substantive, and holistic risk assessment and mitigation.²⁷⁹ One only has to look to other contexts in which conducting a PIA was made mandatory, as in US and Canadian public sector organisations, to see what has been described as an overly prescriptive, formulaic system, 'devoid of any content of significance to privacy

273 The Council noted that 'FR, RO, SK and UK warned against the considerable administrative burdens flowing from the proposed obligation.' (Proposal for a regulation of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation) - Chapter IV, footnote 251).

274 E.g. (B. Koops, 2014) p7

275 For instance, one industry blog notes that 'the draft Regulation proposes ... an administrative regime of onerous documentation maintenance to demonstrate compliance with the Regulation.' (Nabarro.com, 2012), while 'the resource and cost burden of employing a data protection officer and carrying out mandatory impact assessments cannot be borne by most small businesses.' These were seen as 'very burdensome provisions.' (Prospect, 2014)

276 (Hosein & Davies, 2013) p. 1613, (Adam Warren & Charlesworth, 2012) p. 16

277 (Haines, 2011), drawing from (Weber & Winckelmann, 1964)

278 E.g. for Roger Clarke: 'a PIA is not a mere checklist ticked through by junior staff or lawyers' (Clarke, 2009) p. 125

279 'I fear that, as long as data protection is not in the hearts and minds of data controllers – and the law so far has done a poor job in reaching those hearts and minds ... - mandatory data protection impact assessments will function as paper checklists that controllers duly fill in, tick off, and file away to duly show to auditors or supervisory authorities if they ever ask for it. Procedure followed, problem solved' (B. Koops, 2014). See also Tancock et al: 'requiring a PIA to be conducted for every project is likely to be counter-productive because it tends to encourage merely formal checklist-filling rather than intellectual engagement with the issues' (Tancock, Pearson, & Charlesworth, 2010c).

protection, beyond the narrowly circumscribed legal requirements'.²⁸⁰

These fears about mandatory PIAs reflect broader critiques of data protection as a command-and-control regime, susceptible to the generic problems of traditional legal regulation.²⁸¹ Critics of the EU regime established by the 1995 Directive argue that the broad and ambiguous definition of some of the core categories, the inflexible nature of the rules, and the onerous systems of notification with national authorities, led to too many organisations and activities being caught up in the net of prescriptive obligations despite posing only minimal risk.²⁸² The current regime has been characterised as an 'inflexible' system, requiring forms and procedures 'more so than the average law'.²⁸³ Critics might therefore see the pursuit of PIAs through traditional legal regulation as just the latest iteration of this problematic approach.

5.3.4 PIAs as 'co-regulation'

Advocates of mandatory PIAs have a different perspective. They tend not to classify the mandatory PIA regime as traditional legal regulation. Instead, they see it as 'co-regulation'.²⁸⁴ In fact, they argue that far from entrenching a 'command-and-control' approach, a mandatory PIA regime would allow regulators to avoid the problems associated with that approach; Wright argues that mandatory PIAs are a 'co-regulatory instrument that may obviate the need for "hard" law'.²⁸⁵ The Commission also appears to see the introduction of mandatory PIAs in this light, part of a move away from the prescriptive and burdensome rules of the existing Directive.²⁸⁶ In contrast to those who see the choice as being between either self-regulation or command-and-control, advocates of mandatory PIAs see co-regulation as a genuine third option.²⁸⁷

280 See (Clarke, 2011)

281 E.g. (B.-J. Koops, 2011), (B. Koops, 2014), (Bergkamp, 2002), (PM Schwartz, 2013), (Charlesworth, 2006), (Tancock et al., 2010a).

However, others have singled out data protection as an area of EU law which is comparably free from the command and control approach (Bignami, 2011)

282 (Bergkamp, 2002) p32

283 (Bergkamp, 2002) p38. This appears to be the author's assertion and comes with no supporting evidence, but it is a common criticism.

284 See (Wright et al., 2014) on the 'co-regulatory privacy impact assessment model'; also (D Wright, 2011), (Wadhwa & Wright, 2013). See also Roger Clarke: 'without 'legislative backing', the co-regulatory model is indistinguishable from self-regulation, since 'voluntary guidelines are not an adequate mechanism... Legal stiffening is needed... to discourage non-compliance' (Clarke, 1998).

285 (David Wright, 2013)

286 'While [the notification] obligation produces administrative and financial burdens, it did not in all cases contribute to improving the protection of personal data'. It should therefore be 'replaced by effective procedures and mechanism which focus instead on those processing operations which are likely to present specific risks', including 'data protection impact assessments' (Recital 70). See also (Pederson, 2005). See (Costa, 2012) for more on the theory of a risk-based approach.

287 For example, Andrew Charlesworth argues that "there has also been a politically-inspired tendency to polarise the data protection debate in terms of a stark choice between command and control regulation or self-

This is understandable in so far as the supposed dichotomy between legal regulation and self-regulation is increasingly seen as too simplistic.²⁸⁸ As Coglianesi and Mendelson argue, 'the dichotomy between free markets and command-and-control regulation fails to capture the full range of options that lie between the polar extremes of absolute discretion and total control'.²⁸⁹ But in what sense are PIAs 'co-regulatory' if the rules are imposed top-down, and punitively enforced by the regulator?

As previously noted, there is unfortunately little clarity over the meaning of terms like 'co-regulation'.²⁹⁰ It is a term that obscures a lot of important detail. For instance, in some cases, industry bodies create and enforce their own rules, which then later gain statutory backing.²⁹¹ In others, regulators attempt to steer competitive market forces towards the pursuit of regulatory goals, such as with transparency schemes like trust marks or mandatory labelling.²⁹² In yet others, regulators set targets and punitive fines but allow regulatees significant discretion in devising their own compliance strategies.²⁹³

Regulatory theorists have introduced a variety of terms to describe such measures (which fall within the umbrella category of 'co-regulation'). Bartle & Vass, for instance, describe three sub-categories, namely:

‘devolved’ (notably established forms of professional self-regulation), ‘delegated’ (a clear act of delegation of regulation by a public regulatory authority to a self-regulatory body), and ‘cooperative’ (co-operation between regulator and regulated on the development of statutory backed regulation).’ (Bartle & Vass 2007, p. 901)

But mandatory PIAs – at least in the forms that they have been proposed thus far – don't appear to quite match any of these sub-categories of co-regulation. It is therefore unclear what advocates of mandatory PIAs have in mind when they describe them as 'co-regulatory'. Perhaps they simply use the term loosely to mean that they fall somewhere between regulation and self-regulation. But even under this interpretation, it is unclear how they really differ from traditional legal regulation (notwithstanding their origins as a voluntary instrument).

regulation” (Charlesworth, 2006) p48.

288 See e.g. (Lichtenstein, 2001). According to (Osuji, 2015) 'the ‘soft’ and ‘hard’ law dichotomy is increasingly acknowledged as too simplistic considering, firstly, the blurred boundaries between the two at national (McBarnet, 2007) and international ((Zerk, 2006) 69-72) levels'

289 (Cary Coglianesi & Mendelson, 2010)

290 For discussion of the varying definitions see e.g. (Black 2001), (Senden, 2005), (Saurwein, 2011); in the context of technology policy (including privacy), see (Marsden, 2011), (Hirsch, 2010).

291 For e.g. (Eijlander, 2005): “An essential aspect of co-regulation is the cooperation between the public and the private actors in the process of creating new rules”. See also (Senden, 2005), (Saurwein, 2011). The UK government's Better Regulation Task Force (BRTF), for example, sees 'coregulation' as involving codes of practice which ‘have a statutory backing or other significant government involvement’ (Bartle & Vass, 2007) p. 20.

292 (Bartle & Vass, 2007) p. 4

293 (Benbear, 2007)

This terminological ambiguity reflects an underlying substantive point; advocates of mandatory PIAs cannot avoid the problems associated with traditional regulation (as they claim to do) by simply labelling their approach co-regulation. At worst, this form of co-regulation could inherit the problems of both command-and-control and self-regulation, without gaining the benefits of either. An account is needed to explain how a mandatory PIA regime, if imposed as proposed in the GDPR, could carve out an effective co-regulatory approach.

To this end, we may find clarification by assessing the Commission's motivations for incorporating PIAs, and by analysing the provisions that appear in the proposed GDPR.

5.4. Analysis of mandatory PIAs in the GDPR

5.4.1 Commission reports prior to the 2012 proposal

The Commission's rationale for making PIAs a mandatory requirement can be gleaned from several documents published in the lead-up to the proposal of the GDPR in 2012.

The first is a study commissioned between 2008-2009, which sought to 'identify the challenges for the protection of personal data produced by current social and technical phenomena', such as 'ubiquitous personal data collection' and 'profiling'.²⁹⁴ According to the study, these new phenomena 'threaten to make the application of the [data protection] principles yet more difficult'.²⁹⁵ It claims that many organisations do not 'pay appropriate attention to the privacy implications of new information systems before they are commissioned'.²⁹⁶

In response to these challenges, the report mentions PIAs as a potential solution, noting the existence of mandatory PIA schemes in other jurisdictions.²⁹⁷ It notes that for some established data protection measures, including privacy impact assessments, 'there have so far been insufficient incentives for their use by data controllers'.²⁹⁸ Hinting at the regulatory approach that might eventually be adopted through the GDPR, it argues that PIAs would need to be pursued with 'the right combination of law and self- or co-regulatory rules and mechanisms'.²⁹⁹ In advocating a mixture of different rules and mechanisms to support PIAs, the report indicates sensitivity towards the dangers of adopting traditional legal or self-regulatory approaches.

In a communication published in 2010, the Commission went on to discuss incorporating PIAs in the new regulation.³⁰⁰ It describes plans to 'explore

294 See (European Commission, 2010b) p. 9

295 Ibid, p. 15

296 Ibid, p. 50

297 Ibid, p. 51

298 Ibid. p. 46

299 Ibid, p. 56

300 (European Commission, 2010a)

ways of ensuring that data controllers put in place effective policies and mechanisms to ensure compliance with data protection rules', one of which is the use of what they call 'Data Protection Impact Assessments' (any potential substantive differences between this term and PIA are not discussed by the Commission).³⁰¹ The communication recommends only making them mandatory in 'specific cases, for instance, when sensitive data are being processed, or when the type of processing otherwise involves specific risks, in particular when using specific technologies, mechanisms or procedures'.³⁰²

5.4.2 The proposed GDPR

This rationale filtered through into subsequent draft proposals for the GDPR, where the requirement for PIAs was laid out in detail. The requirement has changed somewhat over the course of revisions between the Commission's original proposal text (published January 2012), the amended version from the Parliament's LIBE committee (confirmed in March 2014), and the version published by the Council to outline their 'general approach' (adopted June 2015). The initial Commission proposal sets out the main provisions relating to PIAs in Article 33, where the overall rationale of 'enhancing the data controller's responsibility' is cited. What follows discusses some of the main elements of the PIA proposal, particularly those which are relevant to the question of which regulatory approach they embody.

5.4.2.1 When are PIAs required?

Article 33 sets out that PIAs are only mandatory in certain circumstances, where there are likely to be 'high risks'. The LIBE committee made some amendments to the Commission's proposal, to further determine the situations in which an assessment should be conducted (Article 33(2)), and the elements to be assessed (Article 33(3)). In particular, PIAs would specifically be required for profiling and sensitive personal data.³⁰³ These amendments aimed to enhance legal certainty by clearly stipulating 'which specific risks pertain, in an exhaustive manner'.³⁰⁴ Generally, PIAs would be required in situations of uncertainty, where processing operations 'are of a new kind' (Recital 70).

To help data controllers ascertain whether a processing operation is likely to present high risks, supervisory authorities will maintain a list of processing operations which are likely or unlikely to present such risks. The 'European Data Protection Board' (a successor to the Article 29 Working Party, established under the GDPR) will ensure these lists are consistent between national supervisory authorities (Article 57c (1)).

When a PIA of a proposed project reveals that high risks exist, data

301 Ibid. p. 11 NB: At this point, the Commission began to use the term Data Protection Impact Assessments instead of PIA. In what follows I will continue to use the term PIA, however, see Appendix E for a discussion of the significance of this change in terminology.

302 Ibid. p. 12

303 LIBE committee amendments 259 and 260 respectively.

304 Ibid., amendment 258.

controllers must consult with the supervisory authority (Article 34). The PIA must be included in the communication between controller and supervisory authority. The supervisory authority may then come to the opinion that the intended processing would not be compliant, and use their powers (defined in Article 53) to temporarily or indefinitely ban the processing.

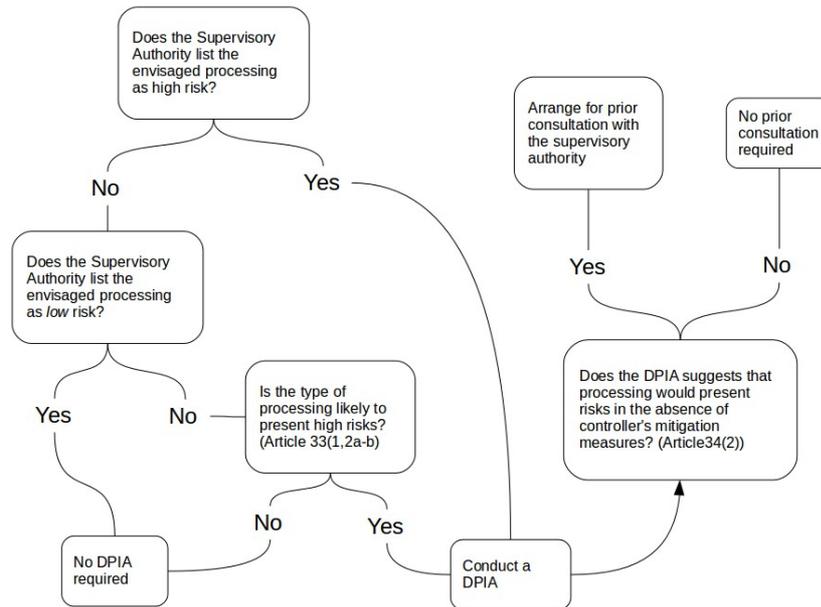


Figure 1. DPIA Triage Process

In combination with the test for determining whether a PIA is required, controllers therefore face a multi-step procedure based on the processing type. This aims to help both data controllers and regulators efficiently prioritise their attention on projects which require greater scrutiny (see figure 1). This rests partly on the detailed lists drawn up by supervisory authorities, but where these do not provide a clear guide, the controller must decide for themselves on the basis of the guidance in Article 33. Further deliberation from the supervisory authority may then occur to determine if the processing will ultimately proceed. The GDPR thus establishes a triage system based on a mixture of regulator and regulatee deliberation.

5.4.2.2 Scope and content of a PIA

The Commission proposal outlines in general terms the aspects that a PIA should encompass. These include a 'general description of the envisaged processing operations', an 'assessment of the risks to the rights and freedoms of data subjects', and 'the measures envisaged to address the risks, safeguards, security measures and mechanisms' (Article 33.3). Initially, the Commission wanted to be able to specify the standards and procedures for carrying out, verifying and auditing DPIAs (Article 33(6-7)), but this was opposed by both Parliament and the Council.

The Parliament's LIBE committee's amendments further defined the scope and content of a PIA to explicitly include consideration of the 'risk of discrimination being embedded in or reinforced by the operation'; 'existing guidelines'; the use of 'modern technologies and methods that can improve

citizens' privacy' (amendment 261); and appropriate means to gain a data subject's consent (amendment 113, (Article 7(1a))).

5.4.2.3 Stakeholder consultation

Initially, there was a strong role for data subjects to play in the formulation of PIAs. The Commission's original proposal would require the PIA process to involve a consultation with data subjects (Article 33(4)). But in the parliament's text, the requirement to consult data subjects was dropped. It argued that this 'represents a disproportionate burden on data controllers' (amendment 262).

The requirement was brought back in a weaker form, in the Council's 'general approach'. There is a reinstated requirement that 'the controller shall seek the views of data subjects or their representatives on the intended processing, without prejudice to the protection of commercial or public interests or the security of the processing operations' (Article 33(4)).

5.4.2.4 Fines and ongoing compliance

Article 79 details the administrative sanctions that may be imposed on data controllers by supervisory authorities. There are three tiers of fines. Failure to undertake a PIA in violation of article 33 could incur the highest tier of fines, up to €1000000 or 2% annual turnover (Article 79a(3de)). However, if an organisation does undertake a PIA, this will mitigate the severity of sanctions they might receive for other violations.

The Commission's original proposal did not include any provisions to ensure that a controller adhered to the measures outlined in their PIA. This could lead to the PIA being a one-off exercise which gets ignored later on; as the LIBE committee noted, 'impact assessments can only be of help if controllers make sure that they comply with the promises originally laid down in them' (amendment 48, recital 74a).

To deal with this danger, they introduced bi-annual 'compliance reviews', to ensure that PIAs would be an ongoing commitment (amendment 130). With this addition, data controllers not only have to comply with the rules outlined in the GDPR, but also ensure they comply with their own self-imposed policies, procedures and safeguards outlined in their PIA. Conducting and monitoring compliance with PIAs would be a key duty of an organisation's data protection officer ('DPO'), which certain organisations will be required to employ (Article 37(1f)).

5.4.3 Summary of the GDPR rationale and provisions

We can infer from the Commission's reports that mandatory PIAs were primarily introduced to deal with new challenges arising from the perceived risks of new technological developments. In particular, for complex areas such as big data and profiling, where data controllers may not be paying sufficient attention to risks before implementing programs, and where there is a lack of clarity on the application of the principles.

The gap between what the data protection principles say and how that ought to apply in a specific complex scenario is particularly large in these cases.

PIAs are seen as a way to address the complexity and need for forward planning. Making them voluntary would not be enough, because as the 2010 report noted, incentives were not sufficient to ensure optimal adoption of PIAs; so they were also made a mandatory requirement for particularly risky cases.

The Commission appears to be using PIAs as a means to allow some form of enforcement even in those complex situations where the basic data protection principles fail to provide firm instructions. Even where data controllers must make their own interpretations of the principles and choose their own mechanisms for risk mitigation, they can still be made accountable for this activity and sanctioned if they fail to do so, or get it wrong. The LIBE committee's amendments even describe the measures laid out in PIAs as a kind of 'promise' which the data controller is expected to uphold (amendment 48, recital 74a).

This is the sense in which mandatory PIAs differ from traditional prescriptive legal regulation; they are a combination of rules prescribed by the regulator, and policies that the regulatees must devise for themselves and impose upon themselves (with input from stakeholders). The proposal therefore has elements of legal regulation, but with a heavy emphasis on controllers coming up with their own measures.

5.5. Meta-regulation as a model of mandatory PIAs

As noted above, the term 'co-regulatory' doesn't really distinguish this particular approach from the many other approaches that fall between traditional legal regulation and 'pure' self-regulation. What is needed is an account of how this particular kind of measure – enforced risk-assessment, and compliance with self-imposed, stakeholder-influenced policies – is supposed to work. Are there other examples of it? In what circumstances are such approaches successful?

I propose that the GDPR's PIA regime can be categorised as an instance of 'meta-regulation', a concept developed by Christine Parker and others.³⁰⁵ It is worth noting that meta-regulation is not a term that the Commission have used to describe their approach in the proposed regulation – it is a term used primarily by academics rather than policymakers themselves – but as I will argue, it is an analytically superior descriptor than any of the alternatives.³⁰⁶

5.5.1 Introducing meta-regulation

Meta-regulation can take many forms. Parker defines it very broadly as 'any form of regulation (whether by tools of state law or other mechanisms) that regulates another form of regulation'.³⁰⁷ However, the primary interest for Parker (and many others who use the term) is a particular form of meta-regulation, namely the 'legal meta-regulation of internal corporate self-regulation'.³⁰⁸ In other words, meta-regulation as a means for the state to make corporations responsible for their own efforts to self-regulate. In what follows, I will use the term with this more specific meaning in mind.

For Parker, one main advantage of meta-regulation over other forms of regulation is that it latches onto companies' inherent capacity to manage themselves, but without letting them off the hook if their self-regulation efforts fall short of regulator (and stakeholder) expectations.³⁰⁹ Meta-regulation differs from self-regulation because its targets don't have the option of not setting up their own rules; individual regulatees are forced to actively self-regulate. Rather than imposing particular rules or technologies on organisations from above, meta-regulation leverages their existing management structures and internal bureaucratic processes in the pursuit of regulatory goals.³¹⁰ Companies may be forced to evaluate and report on their

³⁰⁵ See (Parker, 2002); (Parker, 2007). Parker identifies 'similar uses of 'meta-regulation' or cognate terms', citing (Braithwaite, 2003); (C. Coglianese & Lazer, 2003), p. 691 (government as 'meta-manager'); (Grabosky, 1995) p. 527, 543 ('meta-monitoring'). Similar (albeit non-identical) concepts include 'management based regulation' (Benbear, 2007), and 'enforced self-regulation' (Fairman & Yapp, 2005)

³⁰⁶ In most of the case studies of meta-regulation, it is not explicitly recognised as such by those who design it, but applied retrospectively by academic regulatory theorists. See e.g. (Akinbami 2012), (Dorbeck-Jung & Shelley-Egan 2013).

³⁰⁷ (Christine Parker, 2007) p7

³⁰⁸ Ibid, p. 14

³⁰⁹ Ibid, p. 8

³¹⁰ Ibid, p. 13

own self-regulation strategies so that regulatory agencies can determine whether the ultimate substantive objectives are being met.³¹¹

Meta-regulation must go hand-in-hand with what Parker calls the 'triple loop' of evaluation, which allows regulators and external stakeholders to play an essential evaluative role. The regulator must connect 'the private capacity and practice of corporate self-regulation to public dialogue and justice', by requiring 'companies to gather and disclose information on which corporate self-regulation and its impacts can be judged (by regulators and stakeholders)'.³¹² In Parker's view, this introduces a democratic dynamic: 'in a democracy stakeholders need access to corporate reports of their self-evaluation of their own self-regulation, including how they have identified, prevented and corrected problems'.³¹³

Parker identifies three approaches to fostering meta-regulation; 'building compliance leadership', 'process regulation', and 'education and advice'. The appropriate approach depends on the stage of development of the particular industry and particular regulatees within it. The most relevant for our purposes is the second - 'process regulation' - whereby government teaches regulatees to self-regulate by forcing them to go through *processes* that serve regulatory goals.³¹⁴ This doesn't mean government prescribing the details of the process or mandating precise outcomes, but it does require organisations to take a systematic approach to identifying, controlling and minimising risks.

Examples include occupational safety or food regulations which require firms to engage in their own processes of hazard identification, risk assessment and risk control.³¹⁵ This allows a tailored approach rather than one-size-fits-all regulatory strategy, and gives organisations a chance to integrate regulatory goals into their other business goals and operating procedures. Under such an approach:

'liability is attached to whether the process is in place rather than to its outcomes. The rationale is that by adopting the process the outcomes will generally improve. It also means the regulator can be directly involved in supervising the standard of self-regulation implemented.'³¹⁶

5.5.2 PIAs as meta-regulation

My aim in this section is to demonstrate that the GDPR's mandatory PIA system has many of the hallmarks of meta-regulation, particularly of the kind that Parker outlines. The constitutive features of meta-regulation are manifested in various ways in Article 33 and elsewhere (although, as previously mentioned, I am not claiming that the policymakers who

311 (Gilad, 2010), p. 489

312 (Parker 2007), p. 289

313 Ibid, p. 291

314 Indeed, (Gilad, 2010) relates meta-regulation to other approaches (including systems-based regulation, enforced self-regulation, management-based regulation, principles-based regulation), and argues for 'process -oriented regulation' as an umbrella term.

315 (Parker 2007) p. 27

316 (Gilad 2010) p276

formulated the GDPR were consciously pursuing a meta-regulatory approach as formally described in the academic literature). Table 1 summarizes the constitutive features of meta-regulation, drawn from Parker 2002, alongside the ways that those features are manifested in the GDPR's PIA regime.

Constitutive feature of meta-regulation	Manifestation in the GDPR PIA regime
Requires organisations to take responsibility for their self-regulation efforts	PIAs intend to 'enhance the data controller's responsibility' (Article 33(1))
Requires organisations to undertake risk-assessment processes	A PIA should encompass an evaluation of the risks to the rights and freedoms of individuals (Article 33(3))
Requires organisations to identify risk mitigation strategies	A PIA should involve a description of 'the measures envisaged to address the risk' (Article 33(3))
Does not prescribe specific measures or technologies	No particular measures are prescribed – the controller must identify 'appropriate' measures by themselves
Holds organisations accountable for adhering to their own policies	Controllers expected to 'make sure that they comply with the promises originally laid down' in their PIAs (amendment 48, recital 74a)
Attempts to leverage a corporations' existing management procedures	The GDPR attempts to embed PIAs in management procedures partly through DPO's (Article 37)
Ensures stakeholders can democratically engage in evaluating organisations' measures and policies	Controller must seek input from data subjects or their representatives when conducting a PIA (Article 33(4))
Liability to sanctions is related to failure to undertake the process, rather than focusing on the outcome	Highest penalties are reserved for not undertaking a PIA as required (Article 79a(3de)). Sanctions may be reduced if a PIA has been undertaken (Article 79(2c))

Table 6. Features supporting classification of PIAs as meta-regulation

On the basis of these parallels, I argue that 'meta-regulation' is an apt description of the GDPR's PIA regime, in so far as that regime's intended workings can be surmised from the GDPR texts released by the Parliament, Council and Commission. Of course, if and when the regime is actually put into practice by supervisory authorities, it may end up working somewhat differently, such that the term meta-regulation becomes a less accurate description. But at present, the term is an apt description of the kind of regulatory approach suggested by my close reading of the GDPR.

It is also worth reflecting further at this point on why meta-regulation is a more appropriate label for the regime envisioned in the GDPR than a competing term like co-regulation. Co-regulation is a general term for many different regulatory forms, as described above. While it has often been broken down into more specific sub-categories (such as Bartle and Vass's 'devolved', 'delegated' and 'cooperative' variants introduced earlier), none

of these capture with sufficient specificity what is outlined in the GDPR on PIAs. Meta-regulation, as Parker, Gilad, and others describe it, has a set of particular constitutive features summarized above.

This is consistent with classifying meta-regulation as a *subset* of co-regulatory strategies (although the argument I make here doesn't require or imply such a classification). So the GDPR's PIA regime may be considered a form of co-regulation, whilst also being referred to more specifically as a form of meta-regulation. The advantage of the latter, more specific classification, is that it allows us to evaluate the regime using the resources that have been developed around it. This significant body of theory, analytical frameworks and empirical findings applies to the PIA regime in a way that more general work on co-regulation would not. The following section provides an overview of this work.

5.6 Evaluating meta-regulation

Numerous empirical studies have been undertaken of meta-regulation in a variety of sectors.³¹⁷ Sectors in which meta-regulation initiatives have been studied include anti-pollution; safety of food, toxics and hazardous chemicals; occupational risk prevention; professional ethics; aerospace and financial services.³¹⁸ Generally, while results can be varied, meta-regulation appears to make an overall positive contribution to regulatory goals.³¹⁹

Sharon Gilad has conducted a meta-study of these various empirical case studies.³²⁰ On the basis of this meta-study, Gilad introduces an evaluative framework to identify the conditions under which meta-regulation has most chance of success, and where it is likely to fail. Gilad's framework builds on Parker's version of meta-regulation. The framework proposes that there are several key factors in the success or failure of meta-regulation, which include: the extent to which the effort and expertise of regulatees is leveraged to support regulatory goals; the capacity for independent scrutiny; the degree of stability, trust and support in the regulatory context; the appropriate targeting of regulatory 'tiers'; and the transformation of organisational practices. I outline these key elements of Gilad's framework below, before considering each of them in relation to the PIA regime proposed in the GDPR in the following section.

Leveraging regulatees: Gilad's framework states that meta-regulation is most effective when it leverages regulatees' ability to learn and discover effective measures to achieve regulatory goals. Rather than expecting this to be done by the regulator, who is ill-positioned to uncover the optimal solution for the regulatee's idiosyncratic context, meta-regulation shifts 'the primary responsibility for identifying risks, setting standards, and

317 Case studies where meta-regulation has proved successful (or at least, better than alternatives) include (Coglianese & Lazer, 2003), (Bennear, 2007), (Parker et al., 2010), (Akinbami, 2012), (Dorbeck-Jung & Shelley-Egan, 2013). Mixed results were found in (Hutter, 2001), (Gunningham & Sinclair, 2009), (Black, 2012).

318 See (Bennear, 2007), (C. Coglianese & Lazer, 2003), (Parker, Gordon, & Mark, 2010), (Ford, 2008)

319 (Gilad, 2010), (Haines, 2009)

320 Outlined in (Gilad, 2010)

monitoring compliance to regulated organizations'.³²¹ However, there is a danger that the regulatee's apparently superior knowledge results in ineffective solutions. This brings us to the next factor in Gilad's framework: the need for independent scrutiny.

Independent scrutiny: Gilad notes empirical studies that suggest meta-regulation is less effective when the regulator is unable to evaluate the efficacy of organisations' risk-management plans. This can be due to uncertainty about how to conceptualise risk, and a lack of knowledge about which mitigation strategies are appropriate.³²² This could lead to organisations – intentionally or due to ignorance – pursuing strategies that would result in failure, and the regulators would not be in a position to identify their flaws. Meta-regulation therefore 'depends on regulators' capacity to independently assess and challenge the validity of the information that regulatees generate about their performance'.³²³

Gilad also notes that in addition to well-informed regulator scrutiny, effective meta-regulation requires scrutiny from stakeholders.³²⁴ This helps ensure their values and expertise feed into the regulatory process. These processes need legislative backing; as Parker notes, attempts to make regulatee's activity open to challenge and revision by stakeholders are likely to 'fail badly unless they ... identify and give rights to stakeholders to participate in or contest corporate decisions'.³²⁵

Stability, trust and external support: A third factor identified by Gilad is the need for 'regulators and regulatees [to] enjoy mutual trust and external political and public support, which would provide them with latitude for short-term experimentation in pursuit of long-term improvements'.³²⁶ Regulatees need to feel confident that regulators taking a meta-regulation approach will not just 'shift to them the blame for future failure'.³²⁷

Regulatory tiers: Gilad introduces the concept of 'regulatory tiers' to explain another important factor in the success of meta-regulation.³²⁸ First-tier operations involve 'detailed rules' and 'outcome-oriented standards', i.e. the traditional focus of prescriptive legal regulation. Second-tier operations concern 'the governance structures and controls that regulatees should have in place in order to audit their compliance with first-tier regulatory requirements'. Third-tier operations focus on the 'evaluation, design, and readjustment of ... first-tier production and second-tier controls'. Based on a meta-analysis of case studies, Gilad concludes that meta-regulation works best as part of a hybrid system. Meta-regulation should be aimed at third or second-tier operations, working in combination with more 'prescriptive and outcome-oriented regulation' aimed at the first tier.³²⁹

321 Ibid. p 497

322 Ibid. p 496, citing ((Vaughan, 1996); (Parker & Nielsen, 2011)). A similar phenomenon has also been observed in financial services (Black, 2010).

323 Ibid. p 496, citing (Ford, 2010)

324 Ibid. p 500

325 (Parker 2007), p. 48

326 (Gilad 2010) p. 503

327 Ibid, p. 497, citing (Black, 2008)

328 (Gilad 2010) p. 489

329 Gilad observes: 'In comparison with detailed rules, it is harder to

Shaping organisation's compliance: A final factor in Gilad's framework is the extent to which meta-regulation succeeds in transforming organisations' compliance and capacity to self-regulate. While prescriptive regulation may be suited to 'managing non-compliance by a few bad apples', it is unsuitable 'where non-compliance is persistent and widespread'.³³⁰ In the latter case, what are required are 'profound transformation of industries' resistance to regulation and the constitution of self-regulatory capacity within organizations'.³³¹ To successfully facilitate such transformation, meta-regulation needs to be capable of shaping organisations' self-interest and normative commitment.

Meta-regulation needs to affect every level of the organisation to be effective; commitment to compliance must be 'communicated and internalized beyond the upper echelons of organizations – all the way down to front-level employees across the organization'. Otherwise, the norms of sub-groups within an organisation might continue to 'constitute systematic non-compliance as normal and rational'.³³²

Beyond identifying these key factors – leveraging regulatees, independent scrutiny, stability, regulatory tiers, and shaping organisations' compliance – Gilad also addresses some general problems and limitations raised by others which can hamper attempts to instill meta-regulation.

First, the greater flexibility that meta-regulation affords regulatees doesn't *guarantee* that they will use this flexibility to 'invest in enhanced solutions to regulatory problems'. As Parker acknowledges, if meta-regulation were to simply allow businesses to come up with own rules, it would fail to 'make business accountable for anything – there is nothing to be accountable for, no-one to be accountable to'.³³³ Similarly, Black and Baldwin warn of the danger that 'the firm's internal controls will be directed at ensuring the firm achieves the objectives it sets for itself: namely profits and market share',³³⁴ while Edelman et al argue that 'organizations create symbolic structures as visible efforts to comply with law, but their normative value does not depend on effectiveness so they do not guarantee substantive change'.³³⁵ There is also a potential 'paradox of compliance', where firms complying with meta-regulation will end up engaging in more risky behaviour, believing they have already 'covered' themselves.³³⁶

establish the breach of any form of flexible regulation because broad standards are open to multiple interpretations. Thus, where [meta-regulation] replaces first- and second-tier prescriptive regulation, it could weaken regulatory capacity to deter and to use enforcement against ill-intentioned organizations (Baldwin 1995; Black et al. 2007; Black 2008). Yet as explained in the previous section, in practice [meta-regulation] is likely to complement lower tiers of regulation." Ibid, p 497

330 Ibid. p 498

331 Ibid.

332 (Gilad 2010), p. 499, citing (Vaughan 1996), (Hutter, 2001); (MacLean, 2002)

333 (Parker 2007) summarising (Heydebrand, 2003), p. 326.

334 (Baldwin & Black, 2008) p19

335 (Edelman, Petterson, Chambliss, & Erlanger, 1991)

336 (Laufer, 1999)

Gilad acknowledges these problems, but notes that they may be lessened if meta-regulation can introduce new incentives for compliance. The empirical studies she cites suggest that meta-regulation may provide additional incentives for both 'highly performing organizations' and those who lag behind in their compliance capacity. High performers 'may value the status of industry leaders and the credit that they could gain from that in their interaction with regulators, colleagues, and the public', while for laggards, 'the dissemination of good practice reduces costs of interpretation, and thereby can facilitate cooperation'.

5.7. The prospects for PIAs as meta-regulation

Having explored how the concept of meta-regulation provides an apt description of the PIA regime outlined in the GDPR, and reviewed the theoretical and empirical literature on meta-regulation, we are now in a position to assess the regime's prospects in light of this. Is meta-regulation a promising choice of regulatory style in the contexts in which PIAs are required under the GDPR?

5.7.1 Leveraging regulatees

The GDPR's PIA regime does appear to be designed to leverage regulatees capacity. It attempts to allow room for data controllers to apply their own expertise to a problem. Rather than prescribing the exact measures and safeguards that must be implemented, it requires controllers to identify their own solutions to mitigate risks that are appropriate to their context (Article 33, 3(d-e)).

While member state supervisory authorities may reasonably claim superior understanding of the data protection principles, they may not have a superior understanding of the latest personal data processing techniques, nor the most appropriate privacy-enhancing technologies. For instance, in the case of organisations operating at the cutting edge of data science (an area which may well involve the potentially 'high risk' processing operations covered by Article 33), regulatees are likely to consistently have greater expertise than the regulator.

5.7.2 Independent scrutiny

The success of a meta-regulatory approach to PIAs will significantly depend on the capacity of supervisory authorities to independently scrutinise data controller's proposed mitigation strategies. The proposed GDPR does attempt to ensure such scrutiny happens. First, any PIA must be made available for scrutiny by the supervisory authority on request (Article 33.3(b)). Second, when a PIA identifies 'high risk' processing operations, controllers would have to submit their PIA report to the supervisory authority so that their mitigation strategies can be scrutinised (Article 34.2). Third, they would also have to submit their bi-annual compliance reviews so that their ongoing compliance can be assessed (Article 33a.1). But an obligation on data controllers to allow their programs to be scrutinised does

not guarantee that supervisory authorities will do so competently.

This suggests an important role for the European Data Protection Board, in developing expertise on potential risk identification and mitigation strategies. This could be communicated to supervisory authorities to aid in their consultation processes.

Stakeholder scrutiny is a key part of successful meta-regulation. This is manifested in the GDPR in the requirement to seek input from data subjects or their representatives when conducting a PIA (Article 33.4). It remains to be seen how effective this measure might be. It may strongly depend on the processes by which data subjects or their representatives are identified and consulted; whether they are truly open or simply a tick-box exercise. And if conflicts arise, it's not clear how they'd be easily resolved; stakeholder engagement may be more likely to produce 'dissent, deadlocks, and stultification rather than action'.³³⁷

However, Gilad's overview of empirical studies of meta-regulation suggest the opposite. Stakeholder-regulatee deliberation may actually work *better* where such conflict is greater and therefore stakeholder's motivation to participate is higher.³³⁸ Ideally, regulators need to support stakeholders in holding regulatees to account (in what Ayres and Braithwaite call a 'tripartite' arrangement³³⁹). Such a process is unlikely to lack willing participants. There is a large, diverse, knowledgeable and vocal privacy advocacy community willing to engage on behalf of data subjects.³⁴⁰ In recent years, such groups have expended great efforts in lobbying European regulators over the form of the GDPR and in encouraging their national supervisory authorities to take action against certain companies. In the years that follow, their effort could be re-directed, through a stakeholder-oriented PIA system, into improving the specific activities of data controllers.

5.7.3 Stability, trust and external support

It is uncertain how much stability, trust and external support exists in the context of data protection and PIAs. In terms of its political and regulatory agenda, the EU may provide a relatively stable environment. The substantial time and effort involved in creating and implementing the new data protection regulation means that it is likely to stay in place, unchanged, for a long time.

However, the level of trust between regulators, regulatees, and stakeholders, and the general level of external political and public support, may be less than ideal. For instance, recent relationships between regulators and large technology companies (often based outside the EU), have been adversarial and frayed.³⁴¹ Attempting to develop a co-operative process may prove difficult given this recent history, but equally, it could offer a much-needed

337 (Gilad 2010) p. 178

338 (Gunningham & Sinclair, 2009)

339 (Ayres & Braithwaite, 1992)

340 (Bennett 2010)

341 See, for instance, recent disagreements between European supervisory authorities and Facebook (Schrems, 2014), or Google (P. M. Schwartz, 2013).

fresh start.

Another important element of trust identified by Gilad is that meta-regulation should not simply shift the blame for future failure onto the regulatee. While they do aim to 'increase controller's responsibility', PIAs are not designed in such a way. If a data controller effectively undertakes an PIA as required, and has faithfully implemented the measures outlined in it, this will be taken into account as a mitigating factor if they later face the prospect of an administrative sanction (Article 79(2b(e))). This ought to create conditions in which responsibility is more fairly apportioned and trusting relationships between regulators and data controllers can be built.

5.7.4 Regulatory tiers

Gilad's concept of regulatory tiers is an apt description of the relationship between PIAs and other provisions of the GDPR. Many of the GDPR's provisions, such as the core principles, concern specific rules and outcomes – i.e. 'first tier' operations.³⁴² The provisions on PIAs, by contrast, aim to ensure controllers implement processes for the governance, monitoring, evaluation, design, and readjustment of those first-tier operations.³⁴³ In this respect, PIAs can be seen as second and/or third-tier operations, constituting a complementary layer that sits above the established first-tier data protection rules which are epitomised by a more traditional prescriptive approach.³⁴⁴ The case studies of meta-regulation assessed by Gilad suggest that this is a common and effective arrangement.

5.7.5 Shaping organisations' compliance

As with the introduction of mandatory risk assessments in other industries, by requiring data controllers to identify risks and potential mitigation measures mandatory PIAs may help convince organisations of the long-term net gains to be had from investing in compliance.³⁴⁵

Gilad also talks of the benefits of certain regulatees acting as industry leaders. There is some evidence that some data controllers do indeed value their status as leaders in the industry and are keen to share their knowledge and risk mitigation strategies. This can be seen in efforts by major industry players to publicise and disseminate their best practices, and initiatives within the privacy profession to elevate the status of industry 'thought leaders'.³⁴⁶ This bodes well for a meta-regulatory approach which

³⁴² See GDPR Chapter 2, Articles 5-10

³⁴³ Of course, in attempting to ensure organisations conduct PIAs, the provisions do impose specific rules and processes on controllers – but this is done with the aim of instilling third and second tier operations.

³⁴⁴ There are other aspects of the GDPR which might be classified as second and third-tier operations; for instance the requirement to employ a Data Protection Officer with duties to monitor and implement compliance measures (Article 35.1).

³⁴⁵ (C. Coglianese & Lazer, 2003); (Bennear, 2007)

³⁴⁶ See footnote 247 above on telecoms companies' PIA knowledge-sharing, and the International Association of Privacy Professionals' promotion of 'thought leaders' and 'privacy innovation awards', sponsored by large companies keen to demonstrate their privacy credentials (IAPP, 2015)

emphasises learning and knowledge transfer between regulatees.

Regarding organizational norms, the introduction of data protection officers (DPOs), trained externally by professional bodies with a strong understanding of and normative commitment to privacy, could have an important effect here on the organisation's norms. The GDPR foresees a clear role for DPO's in implementing successful PIAs (Article 37.1(f)). DPOs will have to ensure the results of PIAs reach both 'up' and 'down' the corporate hierarchy, from the C-suite to the shop floor.

5.8 Conclusion

In their early incarnations, PIAs appeared to be part of a self-regulatory approach to data protection. In making PIAs a mandatory requirement under the GDPR, European regulators took them in a different direction. It is understandable that some might fear this change of direction would result in the regulatory pendulum swinging too far the other way; an entrenchment of the command-and-control approach and a re-encroachment of the state.

However, advocates of mandatory PIAs have suggested they can be a cornerstone of a new 'co-regulatory' approach. As we have seen, this term provides little clarity and fails to explain how mandatory PIAs can avoid the typical problems associated with either traditional command-and-control regimes or self-regulation. I have therefore suggested that meta-regulation provides a more accurate description with which they can be better understood and an evaluative framework within which they can be assessed.

Meta-regulation is offered here both as a descriptive account of the mandatory PIA regime laid out in the GDPR, and also as a normative ideal to which policy-makers can aspire. Seeing mandatory PIA regimes as a form of meta-regulation allows us to make sense of the Commission's proposals, as well as outlining their potential benefits and suggesting the kinds of challenges that they might face.

Meta-regulation aims to make organisations responsible and accountable for their efforts to self-regulate, and create a triple-loop evaluative process in which stakeholders can exert influence. By following this approach, mandatory PIAs could allow both the flexibility associated with self-regulation, *and* the benefits of external pressure associated with legal regulation.

Applying Gilad's framework for assessing meta-regulation to the GDPR's proposed provisions on PIAs brings to the fore some potential benefits and likely challenges. The GDPR's PIA regime is strong on several aspects of Gilad's framework, including: the prospects for leveraging the effort and expertise of regulatees; the stability of the regulatory regime; the involvement of stakeholders; the appropriate use of regulatory 'tiers' in a hybrid model; the capacity to engender better compliance norms and introduce new compliance incentives. In each of these respects, the PIA regime proposed in the GDPR appears to accord with successful implementations of meta-regulation in other sectors.

However, the regime is likely to face challenges in other respects. The

capacity for sufficient independent scrutiny by supervisory authorities is uncertain. Trust between regulators and regulatees, and the level of political and public support may be shaky. Transforming compliance cultures within organisations is a fundamentally complex and unpredictable process. Finally, while stakeholders have rights to participate in the PIA process, the GDPR does not guarantee that controllers will facilitate meaningful input from them.

On balance, I tentatively conclude that the GDPR's PIA regime has the potential to create a successful meta-regulatory regime. With this approach, PIAs are not just another hoop for data controllers to jump through, nor yet another way for organisations to cover their backs and avoid liability. In theory, they can add an additional layer which brings responsibility for considering and deliberating on risky and complex data protection issues into the open. They attempt to make the grey areas, which organisations have so far been left to deal with behind closed doors, permeable to external assessment and influence. The success of a meta-regulatory approach is by no means guaranteed, but it is an ideal that regulators would do well to strive towards.

5.9 Epilogue

Having analysed privacy impact assessments and classified them as a form of meta-regulation, I am now in a better position to articulate their link to the concept of Openness for Privacy, and thus connect the final paper to the overarching theme. To this end, I will first outline the wider context of Parker's version of meta-regulation introduced above, to better understand how it maps on to conceptions of openness. I will then explain how this relates to the OfP approach.

The concept of meta-regulation was originally developed by Christine Parker as part of a broader notion of the 'open corporation' (Parker, 2002). The open corporation serves as a guiding concept for an ideal form of regulatory regime towards which policy-makers ought to strive. Writing at a time when scholars of corporate law were sceptical about both the notion of corporate social responsibility (CSR) and the efficacy of command-and-control regimes, Parker aimed to articulate a middle way.³⁴⁷

In Parker's view, earlier attempts to foster corporate citizenship failed to acknowledge the extent to which corporations are normatively closed to external stakeholders – attempts to imbue corporations with social values simply 'bounce off the corporate veil'.³⁴⁸ Drawing from business ethics research, Parker describes how organisational factors fragment and destroy the potential for corporate integrity and democratic responsiveness.³⁴⁹

Parker argued instead that the internal management structures of corporations could be leveraged by regulators to improve corporate citizenship in meaningful ways. The key to achieving this ambitious goal is for regulators to force corporations to be more 'open'.³⁵⁰ Openness, in this case, doesn't just mean engaging in CSR-related activity. Nor is it simply about the organisation communicating their practices and policies in a transparent way. It is about making the corporation's internal processes open to the influence of external stakeholder's values.³⁵¹

This should not be interpreted as a call to refashion corporations in the model of representative democracy. Parker readily acknowledges that requiring decisions to be based on 'collective consent and universally satisfactory resolution of differences' would be an 'unrealistic ideal'.³⁵² Instead, Parker draws on Philip Pettit's 'more practically feasible' notion of *contestatory* democracy, in which 'decisions are legitimate if they are open to contestation in forums and through procedures that are acceptable to all concerned'.³⁵³ Meta-regulation aims to instantiate a version of this concept in the context of a corporation and its stakeholders.

Parker's notion of the *open corporation* helps to define the third aspect of OfP explored in this thesis. This differs from the aspects of OfP explored in

347 (Spender, 2002)

348 (Parker, 2002), p. 28

349 Ibid, p. 31

350 Ibid, p. 2

351 Ibid, p. 2

352 Ibid, p. 38

353 Ibid, p. 38, citing (Pettit, 1997) p. 183-200

the first two papers, both of which operate primarily through the market. By contrast, meta-regulation and the open corporation operate at a regulatory, organisational and societal level.

Approaches to managing the challenges of personal data often result in either of the two failures that Parker identifies. On the one hand, pursuing ever stricter and more prescriptive regulation of personal data is likely to lead to ritualism and creative compliance.³⁵⁴ On the other hand, placing faith in corporate self-regulation and corporate social responsibility to achieve the regulatory goals of privacy, data protection and personal data empowerment may be naïve.³⁵⁵ The open corporation represents a possible third way, based on regulators forcing corporations to make their personal data policies and risk-mitigation strategies open to contestation by external stakeholders.

Mandatory Privacy Impact Assessments are a partial implementation of this approach. The way they have been designed in the GDPR may have its flaws. There may be missed opportunities to instill a truly 'open' approach (for instance, it remains to be seen whether data controllers will even publish their PIAs openly, or whether they will be allowed to be kept secret). But despite these potential imperfections, they take the regulation of personal data in a positive direction of openness via the *open corporation*.

354 (Haines, 2011), drawing from (Weber & Winckelmann, 1964)

355 (Pollach, 2011)

Part 6: Conclusion

The previous sections have explored three aspects of the Openness for Privacy concept in detail. This concluding section has several aims. The first is to provide a summary of the contributions made by each of the three papers presented above. The second is to reflect on the broad approach of Openness for Privacy, laying out some of its general advantages as well as considering potential limitations, challenges and refinements. Finally, I return to the question of the relationship between openness and privacy first raised in the introduction. Drawing on liberal conceptions of both openness and privacy, I offer an account of how the two concepts serve analogous purposes and stem from similar motivations. In so doing, I suggest that there is a deeper philosophical connection between the two principles underlying the Openness for Privacy approach.

6.1 Summary of contributions

6.1.1 Open Data for Data Protection

The first paper analysed a large source of open data comprised of semi-structured notifications from hundreds of thousands of organisations in the UK, to investigate the reasons for data collection, the types of personal data collected and from whom, and the types of recipients who have access to the data. It analysed three specific sectors in detail; health, finance, and data brokerage.

Over the 18 month period, there was growth in the number data controllers, but a steady average number of purposes per controller. A power law was observed in the distribution of types of purposes, data subjects, data classes, and recipients, in accordance with previous studies. Also in line with previous studies, there was evidence of practices that, while not necessarily illegal, certainly conflict with common public expectations. Unlike previous studies of US financial institutions, there was a lack of differentiation between the practices of UK retail banks.

In terms of the number and variety of organisations it contains, the dataset studied in section 2 is an order of magnitude larger than any previous comparable study. It therefore demonstrates the potential for large-scale empirical investigation of organisations' privacy related policies and activities. However, it did reveal some of the limitations of standardised categories and metrics, which can result in errors and omit nuance. But despite these errors, the study demonstrates the potential of a more data-driven approach to policy studies in this area. The fact that this research has already received citations demonstrates the latent demand for more empirical data to inform policy and technical work in this area.

6.1.2 Open Processing

The second paper focused on a new service and business model which gives individuals control over the contents and use of their digital profiles. It sought to answer two questions. First, is user-controlled profiling a viable option for businesses (specifically, marketing and advertising), or will it

negatively impact their revenues in the long term? Second, can it also be a genuinely empowering option for individuals, or are the two mutually exclusive?

The findings indicate that consumer responses to product recommendations are affected by two different factors; the *content* of the recommendation and the consumer's *perception* of the process behind it. Consumers are more likely to want to buy a recommended product if the recommendation is presented as deriving from a profile based on their voluntarily revealed interests, compared to recommendations deriving from predominant behavioural profiling approaches. This suggests that the marketing industry can benefit by giving up some control over consumer profiles, because this may not only increase response rates to targeted messages, but also potentially avoid the negative attitudes some consumers currently have towards behavioural profiling.

The rest of this section went on to explore whether these platforms could be genuinely empowering for individuals, responding to three common objections. First, I argued that far from naïvely adopting the precepts of big data, these platforms actually embrace the theoretical, ethical and epistemological critiques of big data, which are used as a marketing strategy to attract disgruntled users.

I then questioned whether giving individuals the ability to *sell* their own data to marketers might present its own set of ethical concerns. Drawing from theoretical work on the moral limits of markets, I proposed that there may indeed be ethical reasons to intervene in personal data markets to ensure they do not have a dis-empowering effect – either by undermining individual autonomy or creating unequal relations between consumers. Such reasons for regulation are not easily captured through the lens of economic welfarism.

Finally, I considered the problem that these user-centric personal data platforms might become a form of covert technological paternalism. I outlined a framework under which platform designers might avoid this problem, by appealing to the notion of an 'ideal observer'. Personal information management services have the potential to bring individuals far closer to this ideal than they ever could have on their own. In that sense, they have the capacity to be genuinely empowering.

6.1.3 Meta-regulating privacy and the open corporation

The third paper explored Privacy Impact Assessments (PIAs), an aspect of the EU's proposed General Data Protection Regulation (GDPR). The concept of meta-regulation was offered both as a descriptive account of the mandatory PIA regime laid out in the GDPR, and also as a normative ideal to which policy-makers can aspire. I argued that Parker's concepts of meta-regulation and the open corporation demonstrate a possible new direction for data protection, which would be based on regulators forcing corporations to make their personal data policies and risk-mitigation strategies open to contestation by external stakeholders.

My analysis suggests the GDPR's PIA regime is in accordance with

successful implementations of meta-regulation in other sectors in at least some respects. These include: the prospects for leveraging the effort and expertise of regulatees; the stability of the regulatory regime; the involvement of stakeholders; the appropriate use of regulatory 'tiers' in a hybrid model; the capacity to engender better compliance norms and introduce new compliance incentives. However, the regime is not perfect; likely shortcomings include the lack of capacity for regulator scrutiny, low levels of trust, compliance cultures that may be resistant to change, and the uncertainties around what would constitute meaningful stakeholder engagement.

Despite these potential challenges, the meta-regulatory approach to the regulation of personal data embodied in the GDPR signals a potentially positive move towards greater openness.

6.1.4 Summary table

The contributions of this thesis can be divided into *theoretical*, *empirical* and *policy implications*. Each paper makes at least one such contribution. These contributions are also summarised in the table below (separated into theoretical, empirical or policy contribution):

	Type of contribution		
Paper	Theoretical	Empirical	Policy, Industry and Design Implications
1. Open Data for Data Protection		Overview of UK organisations personal data practices	Design of transparency systems
2.1 Personal Profiling (quantitative user study)		Consumer response to targeting – objective and subjective	Marketers may increase response rates by giving consumers more transparency and control
2.2. Personal Profiling (critical data study)	Theory of personal data empowerment, ethics of data markets		New business models may achieve meaningful personal data empowerment.
3. PIAs as Meta-Regulation	Regulatory theory of data protection		Regulators, civil society and forward-thinking businesses should consider a meta-regulatory approach to data protection.

Table 7. Summary of contributions

6.2 Evaluating the Openness-for-Privacy approach

The three papers differ in terms of their methodology, disciplinary boundaries, and their relevance to different stakeholders, but they are unified by the concept of Openness for Privacy. This is an approach which motivates the particular research questions addressed in each paper.

Like similarly abstract principles (such as Privacy by Design), OfP could motivate a wide range of different research projects; the papers collected here are just a selection. In each case there are opportunities for further research, and other potential manifestations of OfP outlined in the introduction (section 1.4.4) also merit further exploration. However, the papers included in sections 2-5 provide enough material for some general reflections on the OfP approach.

I initially introduced OfP as an overarching theme to link the papers together. While it may be a convenient narrative device for the purposes of this thesis, the question remains as to whether it is also a promising approach to addressing the challenges of personal data. This section aims to re-evaluate OfP in light of the theoretical, empirical and policy contributions from each paper, as well as raising some considerations on the potential and limitations of the approach.

6.2.2 The promise of OfP

There are several features of OfP that are worth re-iterating at this point. As discussed in the introduction, there are many ways to apply the notion of 'openness' to challenges of personal data. The preceding chapters suggest that openness can be manifested in different relationships, operations and stages of the personal data 'life-cycle'. The level on which openness is traditionally pursued in this context can be seen in legal requirements

imposed on organisations to publish or respond to requests for general information on how they use personal data. As we have seen, this model often fails to provide the kind of information resources required for effective notice-and-choice systems.

The first paper suggests that an Open Data approach might be more effective, both as a basis for traditional notice-and-choice models, but also for other possible third party-supported uses. At the very least, effective privacy decisions and informed policy choices require some form of open approach to managing the mass of relevant information about organisations' personal data practices.

The second paper examines openness as a feature of data processing activities, with a particular focus on profiling. The platforms referred to here enable users to understand and control how their profiles are constructed and targeted against. Rather than dealing with relatively coarse-grained, public information about organisations' general activities, this kind of openness is attached to particular instances of processing relating to particular data subjects. Furthermore, unlike the open data approach explored in the first paper, not only is the processing here transparent, it is also *manipulable*. The content and meaning of profiles is to some extent open for revision. We can distinguish the two kinds of openness in regards to personal data processing as 'read-only' and 'read/write' respectively.³⁵⁶ The former only aims to make an organisation's practices transparent, while the latter (open processing) also aims to give external actors the power to change them.

Finally, the third paper explores an 'open' approach to data protection regulation, based on Parker's notions of the *open corporation* and *meta-regulation*. Meta-regulation, applied to the challenges of personal data and privacy, can be seen as a manifestation of Openness for Privacy. It calls for organisations to open up their use of personal data to external scrutiny and modification by regulators and relevant stakeholders. In this sense, it is a strong complement to the service components described in the second paper, in that both attempt to allow transparency and manipulation by actors outside the organisation. The differences with meta-regulation are that it attempts to achieve these aims across whole industries; it is facilitated and enforced by regulators (rather than voluntarily by firms); and the external influence is typically enacted through representative stakeholder groups engaging in governance processes, rather than by individuals themselves using digital tools.

We might further differentiate these manifestations of OfP by the types of actors and interests they might appeal to. The second paper, for instance, relates to the idea of OfP as a *business strategy*. Industry is beginning to realise the potential value of empowering individuals with their own data for their own purposes; there is a potentially large market for PIMS (Ctrl-Shift, 2014). To this extent, the idea of open processing may prove compelling. Of course, it does not cohere with every business model. For instance,

³⁵⁶ I use the terms 'read-only' vs 'read-write', normally used in the technical context of filesystem permissions, here to draw a more abstract distinction between types of socio-technical systems, as in e.g. (Berners-Lee & O'Hara, 2013), (Lessig, 2006)

traditional data broker revenues are based on the scarcity and integrity of the data they hold. Allowing individuals to control and edit their data is antithetical to this model. Meanwhile, the first and third papers suggest various ways that OfP might help civil society and privacy advocacy groups. Meta-regulation of open corporations might allow such groups to better channel their efforts into substantive changes.

These examples demonstrate how OfP can operate in very different ways, and embed different commercial, legal, technical, political and cultural assumptions. One might be optimistic about one but sceptical about another. These heterogeneous manifestations of OfP might therefore pull in different directions. This needn't render the OfP approach itself a failure, since the original purpose was simply to expand our understanding of the myriad ways we could appeal to openness from within privacy and data protection discussions. If some of the resulting policy implications or business strategies are in conflict with one another, at least the conceptual horizons of the debate have been broadened.

However, I don't think the various aspects of OfP outlined here are necessarily contradictory. On the contrary, pursuing OfP at different levels may provide for a cohesive overall approach. If privacy challenges are heterogeneous and occur in multiple ways, solutions to them may need reflect this, as the various manifestations of OfP do. However, particular aspects of OfP may be more favorable within certain paradigms; for instance, advocates of self-regulation may prefer the market-driven forms of OfP discussed in the second paper.

To illustrate how the various aspects of OfP might work together, consider an ideal case of personal data processing which fully embodies the OfP approach. We begin at the design stage, where an organisation is considering a project involving a new use of personal data. Since the organisation adheres to Parker's ideal of the open corporation, they immediately begin a process to engage stakeholders (or their representatives) in outlining their interests, identifying the risks, and influencing the design of the project, its safeguards and mitigation strategies. The processing operations would also be designed to be 'open', giving data subjects the ability to scrutinise and modify the contents of any profile the system creates of them. Information about the system, its purpose and use of data, would be published in a machine-readable format that can be aggregated into a common database for use by regulators, researchers and third parties. Finally, regulators and stakeholders would be able to evaluate and contest the ongoing use of the system, monitoring its development and assessing the organisations efforts to stick to the policies and risk mitigation strategies agreed during the PIA.

6.2.3 Limitations and challenges of OfP

There are various challenges facing this ideal form OfP. In what follows, I consider some of the most significant.

One limitation of OfP is the fact that, for better or worse, most software is proprietary and subject to commercial secrecy. This means there will be inevitable limitations on the extent to which the details of data processing can be scrutinised by individuals. Furthermore, transparency may be

practically impossible so long as the epistemic possibilities of data science outstrip the bounded cognitive capacity of humans to understand them ((Hildebrandt, 2013) p. 239).

OfP may not be compatible with certain other approaches to privacy. There are some notable strategies employed in the name of privacy which would conflict with it. These include, for instance, proposals to enable individuals to control organisations' uses of their data by installing physically secure hardware on all computers that might use it (a 'trusted computing' approach).³⁵⁷ Such an approach is antithetical to FOSS principles, since it depends on certain processes being inaccessible to the user.³⁵⁸ More generally, privacy enhancing technologies and other kinds of PIMS which are 'closed' or proprietary do not sit well with the OfP approach. Another example would be web plug-ins that block online tracking based on a proprietary database of blacklisted trackers (as is the case for the most popular versions of such tools). In such a case the tool might protect privacy, but important decisions about which particular entities to block are not available for scrutiny and revision.³⁵⁹

It is important to recognise that even if openness enables us to better address privacy challenges, there is no guarantee that we will do so. For instance, open data about organisations' privacy practices might go unused, as was largely the case for the register of data controllers. Given the power to edit their profiles, consumers might end up divulging even more harmful information. Similarly, stakeholder representatives participating in PIA consultations might fail in defending the interests of the data subjects they are supposed to represent. In these ways, OfP has the potential to make things worse.

Similarly, a potential shortcoming of the OfP approach is that it may only empower the already empowered, and result in unequal outcomes for different groups.³⁶⁰ It may be that only certain kinds of people have the capacity to engage with the opportunities offered by open processing. In the case of an open corporation that is permeable to external stakeholder influences, certain stakeholder groups will inevitably be better represented than others. In these senses, OfP might inadvertently contribute to differential levels of empowerment. How problematic this is may depend on one's political outlook. A welfarist might argue that if OfP provides a net benefit (in economist's terms, a pareto-superior outcome), then the fact that its benefits may be unequally distributed is not necessarily a problem. One might also prosaically rejoin that just because some people are illiterate, that's no argument against funding public libraries. OfP may lead to more egalitarian outcomes in some cases and less in others.

357 E.g. (Iliev & Smith, 2005).

358 See (R. Anderson, 2004)

359 For instance, cookie-blocking tool Ghostery states that 'We do not publicly expose our library since it represents our view/take on what should or should not be in it.' Statement retrieved from [https://getsatisfaction.com/ghostery/topics/export_the_list_of_trackers_found_on_a_site_to_support_complaints#reply_8363654]

360 This objection is a specific application of a general critique of openness, mentioned in the introduction (see 1.3.3).

6.2.4 Refining OfP

Having outlined the promise of OfP, considered some of its limitations and challenges, I will now suggest some potential refinements which may need to be made to the approach for it to be successful.

The comparison between OfP to Privacy by Design prompts considerations of what the scope and audience of the OfP approach is. Privacy by Design was first presented as an approach to systems engineering, and therefore targeted at engineers of potentially privacy-invasive technologies. As the approach became known in wider circles, including legal teams, policymakers, and others, its scope and audience expanded. It is now seen not only as something for engineers to embed in systems, but more holistically as something to be pursued by organisations at multiple levels. Numerous forms of guidance, expertise and working knowledge to help organisations embed PbD in an ever-widening array of circumstances.

OfP could follow a similar trajectory, and be further developed into a set of practical set of 'best practices' and techniques to embed appropriate forms of openness in personal data systems. The findings from the three papers presented here could be developed into specific guidance on how to embed openness in each of the domains covered. This would need to be based on a deep understanding of the user needs and aims in each case, as well as contextualised to the specific technologies involved. For instance, given the problems identified in section 2, regulators and other third parties may need detailed guidance on more effective means of collecting, structuring and publishing open data about organisations' use of personal data in order for the vision of OfP to be realised.

However, as the significant amount of theoretical discussion above attests, OfP is not simply a set of practical best practices regarding personal data. It is also intended to be a contribution to normative and theoretical debates about privacy. It provides a novel perspective on the relationship between privacy and openness, a synthesis and extension of existing concepts in data protection law, and an articulation of personal data empowerment. These normative and theoretical dimensions of OfP could be further explored. For instance, a systematic comparison of the OfP approach with particular conceptions of privacy might help us to understand its benefits and shortcomings.

One possibility is that an OfP approach might complement the conception of privacy as contextual integrity (Nissenbaum 2009). This theory holds that privacy is best understood in terms of violation of contextually-derived social norms about the proper flow of personal information. Since social norms are set by a range of stakeholders in a particular context – e.g. doctors, patients, researchers and health institutions – and OfP aims to ensure that the full range of stakeholders' contributions, desires and ideas are considered by systems involving personal data, it could be argued that an OfP approach might therefore support contextual integrity by better supporting and surfacing contextually-derived social norms.

This also suggests that OfP will work best when put into particular contexts by relevant stakeholders. Dedicated privacy advocacy may need to become

integrated into the work of specific interest groups who already have a deep understanding of their particular domain. OfP would need to proceed alongside a growing awareness of privacy issues amongst those who already represent interests of different sectors of society, such as trade unions, consumer rights groups, NGOs, civil society organisations and professional bodies. Efforts that were previously directed at influencing privacy regulators, could then be directed at specific practices affecting specific populations.

6.3 Openness and privacy: mutually supportive principles

The narrative which introduced OfP in the introduction began with Popper and ended with a proposal for a new approach to addressing the challenges of personal data. Having fleshed out that approach in the intervening sections, this narrative can be revisited and considered with fresh eyes.

Popper's concerns, formulated in the wake of the second world war and in response to what he regarded as the twin evils of fascism and communism, may seem to have little relevance in the current global political climate. But from a certain perspective, they resonate with the concerns of many privacy advocates today.

Like his theory of knowledge, Popper's liberalism is based partly on the notion of fallibilism (Popper, 1963). Just as scientists tolerate and critically assess each others' claims, and creatively seek out ways to test their best theories about the universe, so societies ought to tolerate dissent and alternative views. Science progresses by being open to critique and refutation, and so society progresses by allowing the contestation of received wisdom, including government policies.

Closed approaches tend to fail, according to Popper, because knowledge and conceptions of value are dispersed and varied amongst the population (an idea which Hayek later expanded upon while advocating for market-driven decentralisation). Closed societies, even if they are based on a substantial proportion of true claims, are incapable of incorporating competing views and new knowledge, and are therefore likely to end up wreaking more harm than good.

Openness is Popper's antidote to closed societies; it aims to instantiate the right conditions for pursuing knowledge and making good decisions. It permits, even encourages, objection and deviation from the knowledge and value claims of the powerful. The key virtue of an open society is the ability of individuals within it to deviate from and challenge dominant knowledge claims. This encourages a society's inherent capacity to develop diverse viewpoints, explore alternative values and contest received wisdom. Institutions ought to be set up to foster such approaches to the creation and evaluation social and political programs, allowing for what Popper calls 'piecemeal social engineering' rather than grand utopian projects.

Moving from political programs to computer programs, modern forms of openness promote a similar approach in our digital milieu. Software, data, and information, according to this view, ought to be made available so that they can be scrutinised, modified and reproduced to suit a wide variety of

ends. Since no one software vendor, data generator or knowledge holder can possibly know how their offerings might best serve any and all relevant stakeholders, there is a prima facie reason for openness. The piecemeal software engineer is allowed to pragmatically experiment and pursue new ways of doing things by trial and error.

The switch in context between the old and new forms of openness – from political philosophies which started world wars to the mundane world of software – might seem morally insensitive. Whatever the inconvenience of proprietary software, it pales in comparison to the brutality and indignity of life under Nazism or Communism (or, one can imagine, the life of a slave in Plato's imaginary Republic). To assert a moral equivalence between the architects and implementors of communism and fascism and the CEOs of current technology giants seems somewhat harsh to the latter.

The new politics of openness may have a less dramatic and violent backdrop than their predecessor, but they are played out in subtle, banal ways that can still be highly consequential. The consequences of the political programs Popper was concerned about were not only felt as a result of direct state violence, but also in the mundane aspects of everyday life, such as waiting in queue for bread. Data and digital systems are now increasingly embedded in everyday life, affecting what we buy, how we are taught, how we interact with the state, what news we consume, where and how we travel, our social lives and career paths; the cumulative effect of all this could be just as significant. The focus of the latest iteration of open politics on data and software is justifiable given their increasingly pervasive effects on wider society.

Openness can be seen as a means to potentially resist or mitigate the pernicious consequences of these hierarchically controlled data systems. I have explored how this new iteration of openness might address a particularly pressing set of problems around personal data. Whatever the prospects of openness in general, I have argued that it has a particularly promising role to play in addressing these problems.

Having traced a trajectory from Popper, to open source, to the new iteration of openness for the digital age, before finally focusing on personal data and privacy, the connection between openness and privacy may seem tenuous and circumstantial. This brings us to a final question about the relationship between these two central concepts which make up what I have called the Openness for Privacy approach. I have so far not attempted to posit any deep underlying connection between openness and privacy, just as advocates of Privacy by Design don't posit any such connection between the concepts of 'privacy' and 'design'; in that approach, design is just a means to the end of privacy. However, I would now like to suggest that there is in fact a more fundamental connection between openness and privacy.

First, it is worth noting that at a political level, the enemies of privacy and the enemies of openness have common interests. As Julie Cohen writes:

“Advocates of strong copyright and advocates of weak privacy share interests in strengthening the commodification of information and in developing infrastructures that render individual activity

transparent to third-party observers”.³⁶¹

As the proverb goes, 'the enemy of my enemy is my friend'.³⁶² In this sense, we might argue that advocates of openness and privacy should be 'friends'. But this would be a highly contingent alliance, not an underlying philosophical connection between the concepts, which is what I am arguing for here.

I contend that a connection between the two concepts can be found in the parallel roles they play in aspects of political thought. These parallels are most obvious in the case of liberal conceptions of privacy and openness, which I now turn to. (By focusing on liberalism alone, I do not mean to imply that the parallels do not also obtain between alternative conceptions of privacy and openness offered by other traditions.)

In its rejection of utopian social engineering, Popper's political philosophy exemplifies some key liberal claims. One is anti-perfectionism (alternatively called 'state neutrality'), the view that the state should not justify its actions and policies by appealing to an objective account of the good.³⁶³ Another is pluralism, the view that the state ought to tolerate multiple value systems.³⁶⁴

Advocates of the new digital openness marshal similar arguments against proprietary software, data hoarding and the copyright system. In each case, openness aims to ensure that the natural and valuable human capacities to explore, question and innovate are able to flourish unencumbered. I argue that privacy, according to various liberal conceptions, plays an analogous role to this conception of openness. This is most clear in accounts of privacy's relation to autonomy.

Autonomy is a central strand in liberal political thought, from 19th century liberals like Wilhelm von Humboldt and John Stuart Mill, to contemporary exponents of liberalism such as Joseph Raz.³⁶⁵ As Ben Colburn writes, there is a common thread running throughout these various liberal conceptions of autonomy, which is the claim that:

“What is distinctive and valuable about human life is our capacity to decide for ourselves what is valuable in life, and to shape our lives in accordance with that decision”.³⁶⁶

Like Popper's conception of openness, this conception of autonomy is also associated with a philosophical commitment to anti-perfectionism and pluralism.

Many liberal defences of privacy are based on the claim that it is necessary for autonomy. This should be distinguished from the definitional claim, that privacy *means* autonomy over the use of one's personal information. Rather, the liberal defence of privacy is that it is a necessary condition for the

361 (J. E. Cohen, 2012) p. 1

362 (Kautilya, 1929) p. 296

363 See (Popper, 1945) chapter 9. For a definition of anti-perfectionism, see (Raz, 1986) p. 108.

364 (Mason, 2015)

365 (Colburn, 2010) p. 2, citing (Humboldt, 1810), (Mill, 1859), and (Raz, 1986).

366 Ibid, p. 3

development of autonomous individuals.³⁶⁷ To return to the aforementioned phrases from German constitutional law, we might say that 'informational self determination' is a necessary basis for the 'free development of the personality'. This is because the inability to decide how data about oneself will be used may lead to 'anticipatory conformity', or a 'chilling effect', where one constrains ones behaviour.³⁶⁸ This could interfere with an individual's ability to engage in the experimentation necessary to decide for themselves what is valuable in life, and to shape their lives in accordance with that decision. Thus privacy is necessary for autonomy.

Julie Cohen has provided a compelling articulation of this argument. She is critical of the traditional liberal account of privacy because of its conception of the self.³⁶⁹ But her defence is similar, albeit updated in light of work on self-hood from cognitive science, sociology and philosophy. On Cohen's account, privacy protects critical independence of thought, self-actualization and reason – which are 'essential tools for identifying and pursuing the material and political conditions for self-fulfillment and more broadly for human flourishing'.³⁷⁰ It 'shelters dynamic, emergent subjectivity' from attempts by 'commercial and government actors to render individuals and communities fixed, transparent, and predictable'.³⁷¹ It is therefore 'an indispensable structural feature of liberal democratic political systems'. Innovation also depends on privacy, to the extent that it emerges from 'processes of play and experimentation' which privacy shelters.³⁷²

On each of these accounts, privacy and openness have strong similarities. Both aim to protect and promote the inherent and valuable human proclivity to explore and experiment, to contest dominant conceptions of knowledge and value. Both are claimed to be necessary to the vitality of a vibrant and innovative society. In this sense, privacy and openness play similar roles.

There may be differences in emphasis between the two principles; privacy deals primarily with the free development of personality, while openness is concerned with a diversity of political ideology (on Popper's account), and a general model of knowledge and innovation (in more recent discourse). But these different ends, while not necessarily substitutable, are certainly related. For instance, self-development is linked to diversity in the political and economic sphere (as Cohen argues), because a citizenry with diverse personalities is more likely to result in an expression of different ideologies and innovations. Similarly, an individual who grows up in a society in which a multitude of ideologies and views are 'on offer' may be more likely to find a conception of the good which allows them to manifest their personality.³⁷³

367 E.g. (van Dijk, 2009) ("The protection of personality however is one of the main principles behind privacy" p. 58), (J. E. Cohen, 2000), (Allen, 1999), (Paul Schwartz, 1999),

368 Terms borrowed from (Finn, Wright, & Friedewald, 2013) p. 12

369 For Cohen, liberal privacy scholarship is rooted in the false notion of an inherently autonomous unfolding self. She argues instead that 'the liberal self' that these accounts posit is not a reality, but 'an aspiration—an idealized model of identity formation' (p. 6).

370 Ibid, p. 6

371 Ibid, p. 2

372 Ibid, p. 2

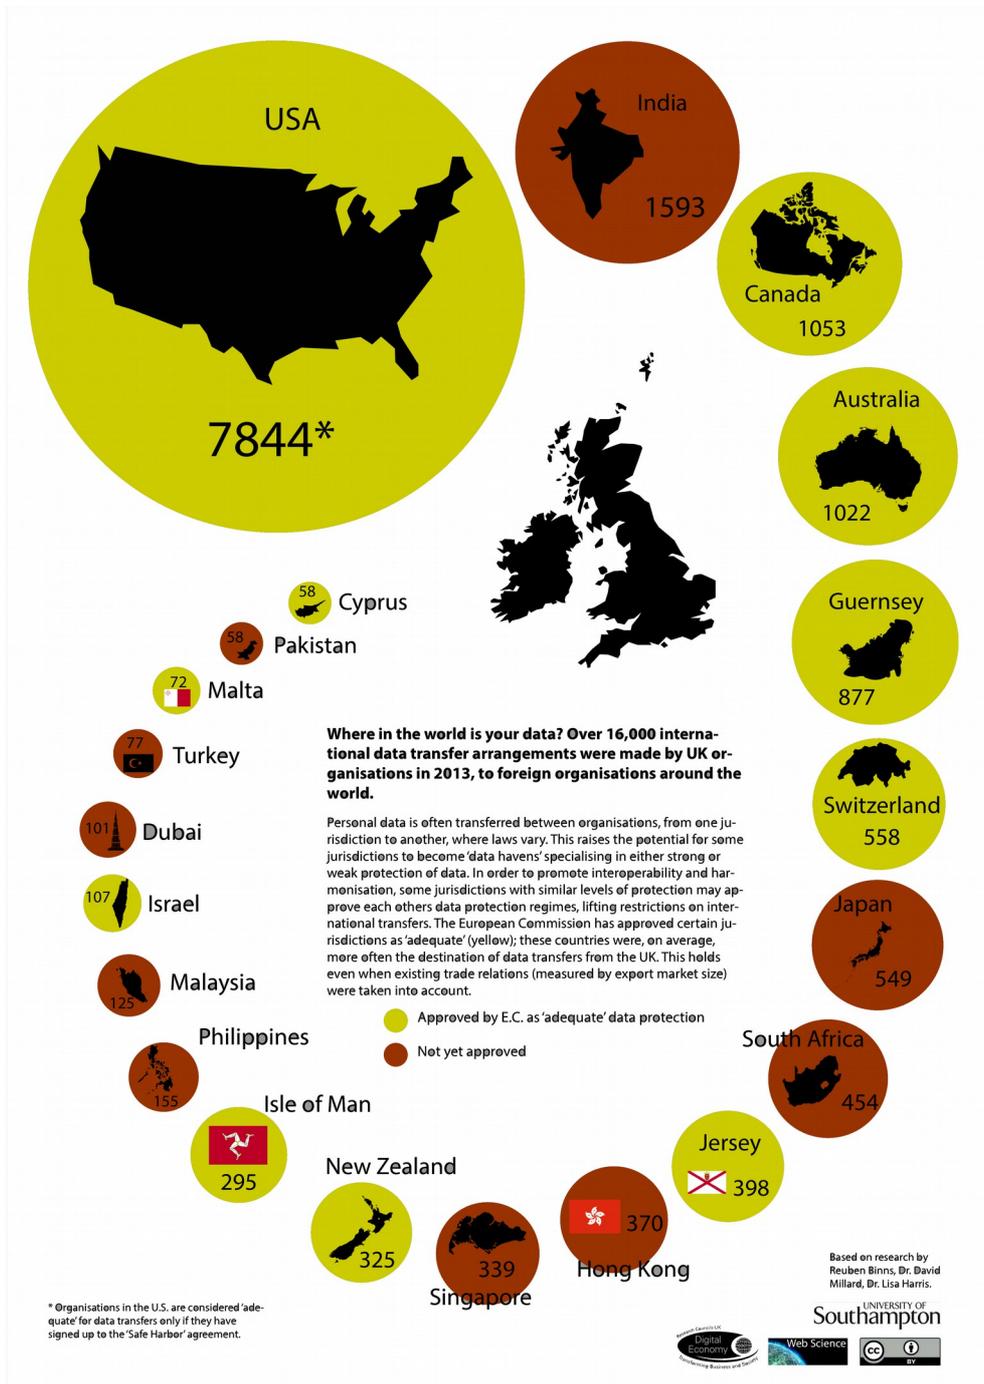
373 This recalls John Stuart Mill's defence of the value of 'experiments in

This suggests that the notion of 'openness for privacy' is not just a conjunction of two otherwise unrelated concepts. Openness and privacy are connected at a deeper philosophical level. As the rest of this thesis suggests, combining the two may be a promising strategy for addressing the challenges of personal data.

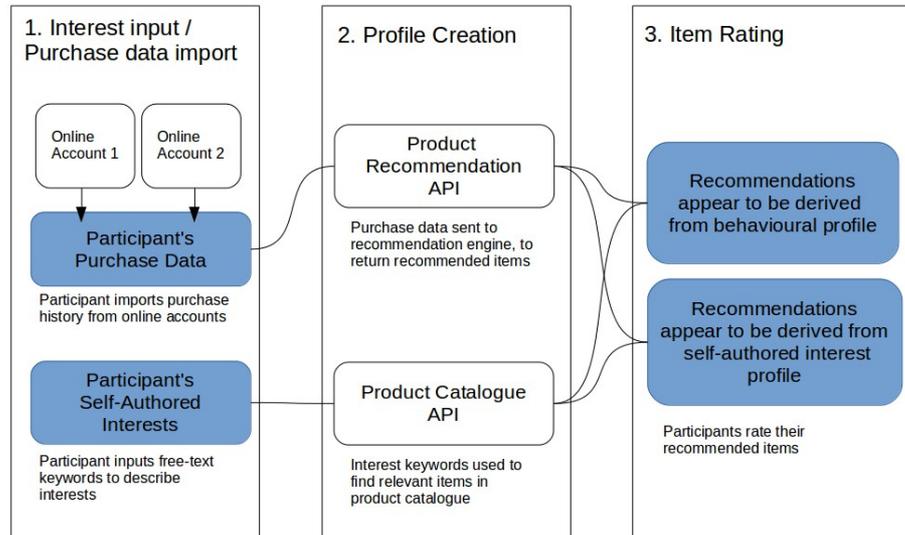
living' ((Mill, 1859) II 23, 38; III 1).

Appendices

A. Visualisation of international data transfers



B. Study design flowchart



Study Design Flowchart

C. Study design considerations

This appendix provides further details on the design of the study described in part 3, as well as some further methodological considerations relating to it.

Procedure

This section describes in further detail the procedure of the experiment.

The experiment was conducted through a website which was purpose built for the study. Upon arriving at the website, participants were given a brief description of the aims of the study and further information to enable them to decide whether to take part. Upon clicking a button to indicate their informed consent and agreement to take part, participants were randomly assigned to one of two experimental conditions (A and B). Participants in both conditions were then presented with a page which asked them to provide details of five recently purchased products, using free text input boxes. In order to aid their recall, at this point they were also provided with a direct link to their purchase histories on Amazon and Ebay (if they had accounts with these services). After the participant submitted descriptions of five recent purchases, the application called the Amazon Product API to attempt to match the user's descriptions to products listed in the catalogue. In cases where the participant made use of their Amazon purchase history data, this matching process tended to be successful, since the vast majority of items they had bought on Amazon were still listed in the catalogue. However, in some cases it was not possible to find a matching item in the Amazon catalogue. If less than 3 products were successfully identified, the participant would be asked to provide descriptions of alternative products until at least 3 were successfully identified in the Amazon catalogue. Next, all participants were presented with a page which asked them to provide five types of products they would be interested in receiving recommendations for. These could be specific products like 'digital camera', general categories like 'poetry', or particular brands (see appendix D).

After completing these steps, the participant had contributed both behavioural data (their purchase histories) and interest data (in the form of free text descriptions). These two types of data were then processed to generate two different kinds of recommendations. The behavioural data was used as an input to a query on the Amazon Product API to find further product recommendations. These recommendations were based on Amazon's product recommendation algorithm which is roughly based on what other Amazon users who had bought the same products subsequently went on to buy. The interest data was also used to derive a set of recommendations. These interest-based recommendations were generated by matching keywords to categories in the Amazon product catalogue, and returning top-selling products in those categories. Participants progressed to the final stage in the experiment. They were presented with a selection of five recommendations. Participants in condition A were presented with explanatory text which described the recommendations as deriving from an analysis of their previous purchases. Participants in condition B were presented with explanatory text which described the recommendations as

deriving from their declared interests.

In condition A, four out of five of the recommendations were indeed derived from an analysis of the participants' previous purchases, as the explanatory text stated. However, one out of five of the recommendations was in fact derived from their declared interests, contrary to the explanatory text (this was presented in a random order amongst the other recommendations). In condition B, this setup was reversed. Four out of five of the recommendations were indeed derived from an analysis of the participant's declared interests, as the explanatory text stated. However, one out of five of the recommendations was in fact derived from their previous purchases, contrary to the explanatory text.

Participants were asked to rate the recommendations on a 5 point scale according to how likely they would be to purchase the recommended item.

Demographic data.

This section explains demographic aspects of this study. There are two broad reasons why demographic data might be collected as part of a user study. The first is when the study is intended to examine the relationship between one or more variables that are themselves demographic in nature. In these cases, the collection of relevant demographic data from participants is directly involved in the study. This was not the case in this study. The second reason why demographic information might be important is for use as ancillary data to account for differences between the study in question and other studies of the same phenomena. For instance, if two studies examining the same phenomena find different results, then demographic differences between samples might suggest there is some additional factor which accounts for the difference. This study is an instance of this second case, where general demographic information may prove useful in future if we were to compare the results with another study and find a discrepancy.

As noted in section 3.2, participants for this study were recruited in two batches. The first batch were recruited by advertising for volunteers through online networks associated with my academic institution, the University of Southampton, and therefore comprised mostly undergraduate and graduate HE students. The second batch were recruited via an online platform for conducting user studies (Prolific Academic). This platform was able to provide additional demographic data for each of the participants who took part in the study. 37 were male and 21 were female. The most common nationalities were United Kingdom (34), United States (13), and India (5). Participants average age was 29. Such general demographic information might be useful if the results of this study are found to be incongruous with those of similar studies performed on different populations.

A further consideration regarding the collection of additional demographic data is whether it could have helped to explain the asymmetry between the effect of subjective attitudes between conditions A and B. To recap; at the outset, one might expect that if participants' attitudes had a negative effect on their appraisal of behavioural profiling, they would also have a correspondingly positive effect on their appraisal of SAI profiling. But the experiment found that when consumers believe a recommendation is based

on their previous behaviour, they tend to like it less, but there is not an equivalent 'boost' in ratings when they believe a recommendation is based on their stated interests. This asymmetry between observations is intriguing, and it would be interesting to examine further. However, I believe that would require a full independent study, with a different design. It could not simply be based on demographic data of the participants in the original study. If some demographic attribute were indeed causally relevant to the observed asymmetry, it would be because a disproportionate number of participants with that attribute were assigned to one or the other condition. This is statistically unlikely, since the groups were randomly assigned and a power analysis was performed to ensure the number of participants recruited would be sufficient for an even distribution between conditions. In any case, the demographic data supplied from the online platform (see above) showed no indication of an uneven distribution of the demographic variables collected (including sex, age, and location) between conditions A and B.

Furthermore, any attempt to use such data to explain the asymmetry between groups of randomly assigned participants would probably be unwarranted (even if an uneven distribution were to be discovered). First, since there is necessarily a limited number of demographic attributes, the putative causally relevant attribute may have been left out and therefore the demographic data would fail to provide any explanation. Second, the more demographic variables one collects, the higher the likelihood of finding differences between the two groups, even if those differences have nothing to do with the observed asymmetry. Attempts to explain the asymmetry through demographic data, without conducting an independent, hypothesis-driven study, would therefore likely be unsound (otherwise known as 'fishing for correlations'). Further research aimed at explaining the discrepancy will therefore require an independent, hypothesis driven study.

D. Study interface

What kind of recommendations would you like to receive?

Rather than using previous purchases to select products to recommend, some services allow consumers to say what kinds of products they are interested in. To simulate this, enter some categories of products you are interested in receiving recommendations for. These could be:

- Specific products, (e.g. 'portable coffee maker' or 'digital camera').
- Categories of products (e.g. 'poetry')
- Names associated with products you like, (e.g. your favourite author, sports team or clothing brand).

Please enter some keywords below:

The following product recommendations have been selected by analysing your previous purchases. Other people who bought those same products went on to buy the following items.

Would you consider buying these items?



Polaroid Bumper Pendant Case (Black) for the Polaroid CUBE HD Action Lifestyle Camera - Includes 90cm Lanyard & Metal Hook

- Very unlikely
- Unlikely
- Maybe
- Likely
- Very likely



36 Cactus Misc 2inch Potted Cactus Collection

- Very unlikely
- Unlikely
- Maybe
- Likely
- Very likely



GoPro HERO

- Very unlikely
- Unlikely
- Maybe
- Likely
- Very likely

E. What's in a name? Privacy Impact Assessments and Data Protection Impact Assessments

Despite the established pedigree of the term 'privacy impact assessment', the alternative term 'data protection impact assessment' (DPIA) has been used in the text of the proposed General Data Protection Regulation since at least 2010. This terminological difference inevitably raises the question of what's in a name. Is a DPIA just a PIA by another name? Or might the different terms reflect different underlying assumptions about the purposes, scope and shape of impact assessments?

On one hand, the difference of terminology might indicate differing ideas about the role and scope of an impact assessment. From this perspective, the Commission's use of *DPIA* instead of *PIA* in the GDPR suggests an attempt to define the scope of impact assessments in terms of data protection (Hosein & Davies, 2013) p25. Whether this constitutes a widening or a narrowing of their application depends on the scope of and relationship between data protection and privacy, a debate which is itself complex and controversial. It should be noted that early proponents and scholars of privacy impact assessments were keen to differentiate them from 'data protection law compliance checks' and 'data protection audits', which they regarded as narrower in scope, focusing on compliance with data protection laws rather than exercising best practice with respect to privacy (Clarke, 2009). Yet others might perceive the rebranding of PIAs as DPIAs as a case of an expansionist tendency in data protection law.

But its also possible that the change of term has less significance. *Data protection impact assessments* (DPIAs) might be the preferred term of European policymakers simply as a way to signify the legal basis of this new requirement as data protection law. Under this interpretation, the term DPIA might be intended to be functionally equivalent to *privacy* impact assessments (PIAs). This implies that policymakers believe that the respective domains of privacy and data protection are similar enough that an assessment of the impacts of technology would be the same in either case.

Despite having invented the term DPIA as a replacement for PIA, not even EU institutions themselves have maintained a consistent vocabulary. A leaked document released in March 2015, containing the Council and Commission's proposed changes, breaks with the Parliament's prior convention by referring in some places to *privacy* impact assessments.³⁷⁴ Meanwhile, the European Data Protection Service (EDPS) refers to 'personal data impact assessments', a unique term whose ambiguity suggests either incoherence or, perhaps, careful diplomacy.³⁷⁵ In any case, it seems that even at this late stage in the negotiation there is little consistency in terminology. In so far as this reflects underlying conflicts in the envisioned role of PIAs, it only sows the seeds for further confusion down the line.

As discussed in the introduction (1.1.2), privacy and data protection are

374 (https://edri.org/files/EP_Council_Comparison.pdf - (23e))

375 See EDPS Opinion on the Reform Package, 7 March 2012

(https://secure.edps.europa.eu/EDPSWEB/webdav/site/mySite/shared/Documents/Consultation/Opinions/2012/12-03-07_EDPS_Reform_package_EN.pdf)

legally distinct, but equivocation between the two concepts is common. The European Court of Justice has allegedly 'treated privacy and data protection as if they are interchangeable' (Lynskey 2014); the mistake is repeated, according to DeHert and Gutwirth by "many scholars [who] hold data protection and privacy to be interchangeable", with data protection perceived as a "late privacy spin-off" (Data Protection in the Case Law of Strasbourg and Luxemburg: Constitutionalisation in Action). Despite being jurisprudentially incorrect, equivocation between the two concepts is prevalent enough that the difference in terminology between PIA and DPIA may not signify any intention by European legislators to mark substantive differences between the two.

And in practice, it appears that most who have commented on this terminological difference treat PIAs and DPIAs as roughly equivalent. For instance, the terms are equated by ((De Hert & Papakonstantinou, 2012) footnote 91), ((Hosein & Davies, 2013), p25), ((Wright & Raab, 2014), p1)). Industry commentators also seem to regard the two as equivalent; privacy impact assessment systems have been marketed as helping compliance with Article 33 of the GDPR,³⁷⁶ while a leading law firm refers to them as "data protection, or privacy, impact assessments".³⁷⁷

Exactly what activities organisations do to comply with Article 33 remains to be seen. There is still room for differing interpretation, so these terminological differences may turn out to be more significant as the GDPR is put into practice.

376 e.g. [<http://www.avepoint.com/community/avepoint-blog/navigate-european-union-general-data-protection-reform-gdpr-avepoint-privacy-impact-assessment-apia-system/>]

377 [<http://www.out-law.com/en/articles/2015/january/data-protection-impact-assessments--when-will-eu-businesses-be-required-to-carry-them-out/>]

F. PIA Triage Process

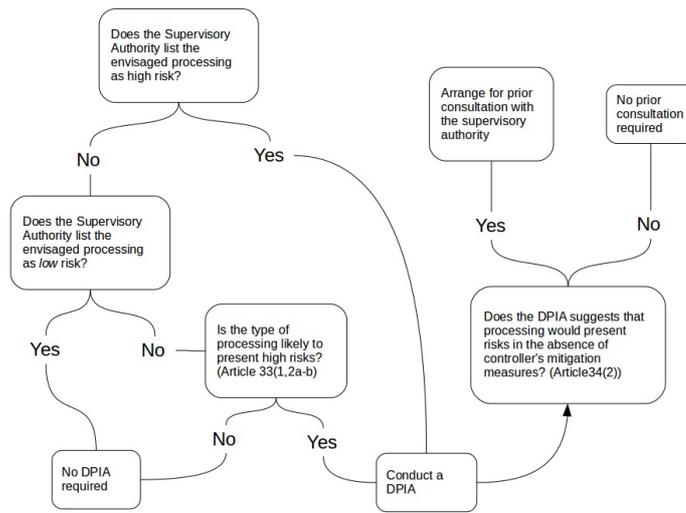


Figure 1. DPIA Triage Process

References

- 35th International Conference of Data Protection and Privacy Commissioners. (2013). Resolution on Openness of Personal Data Practices.
- Abiteboul, S., André, B., & Kaplan, D. (2015). Managing your digital life. *Communications of the ACM*, 32–35.
- Abraham, S. (2015). Privacy Vs Transparency. *New Internationalist Magazine*. Retrieved from <http://newint.org/features/2015/01/01/privacy-transparency/>
- Ackerman, M., & Cranor, L. (1999). Privacy critics: UI components to safeguard users' privacy. *CHI'99 Extended Abstracts on Human Factors*, 3–4.
- Acquisti, A. (2008). Identity Management, Privacy, and Price Discrimination. *IEEE Security & Privacy Magazine*, 6(2), 46–50.
- Acquisti, A. (2009). Nudging Privacy. *Security & Privacy Economics*, (December).
- Acquisti, A., & Grossklags, J. (2007). What Can Behavioral Economics Teach Us About Privacy? In *Digital Privacy: Theory, Technologies and Practices* (pp. 363–377). Taylor & Francis Group.
- Adler, M. D. (2009). Regulatory Theory. *University of Pennsylvania Law School Faculty Scholarship, Paper 301*.
- Ahmed, M. (2015, July 10). Sir Tim Berners-Lee seeks revival of plan to open up NHS data. *Financial Times*.
- Akinbami, F. (2012). Is meta-regulation all it's cracked up to be? The case of UK financial regulation. *Journal of Banking Regulation*, 14(1), 16–32.
- Alder, G. S., Schminke, M., Noel, T. W., & Kuenzi, M. (2008). Employee reactions to internet monitoring: The moderating role of ethical orientation. *Journal of Business Ethics*, 80(3), 481–498.
- Allen, A. L. (1999). Coercing Privacy. *WM. & MARY L. REV.*, 40, 738–40.

- Allen, A. L. (2000). Privacy-as-Data Control : Conceptual, Practical , and Moral Limits of the Paradigm. *Penn Law : Legal Scholarship Repository*.
- Alnahdi, S., & Ali, M. (2014). The effectiveness of online advertising via the behavioural targeting mechanism. *The Business & Management Review*, 5(1), 2014.
- Alwitt, L. F., & Abhaker, P. R. (1994). Identifying who dislikes television advertising: Not by demographics alone. *Journal of Advertising Research*, 34(6), 17–29.
- Amin, A. (1994). *Post-Fordism: A Reader*. Oxford: Blackwell.
- Anciaux, N., Bouganim, L., Pucheral, P., Guo, Y., & Le, L. (2013). MILo-DB : a personal , secure and portable database machine. *Distributed and Parallel Databases*, 1–43.
- Anderson, E. (1990). The ethical limitations of the market. *Economics and Philosophy*, 6(2), 179–205.
- Anderson, H. R. (2011). The Mythical Right to Obscurity: A Pragmatic Defense of No Privacy in Public. *ISJLP*, 7, 543.
- Anderson, R. (2004). Cryptography and competition policy-issues with “trusted computing.” In L. J. Camp & S. Lewis (Eds.), *Economics of information security*. (pp. 35–52). Springer US.
- Archer, M. S. (2007). *Making our way through the world: Human reflexivity and social mobility*. Cambridge University Press.
- Article 29 Data Protection Working Party. (2007). *Opinion 4/2007 on the concept of personal data, WP 136*.
- Article 29 Data Protection Working Party. (2010). *Opinion 3/2010 on the principle of accountability, WP 173*.
- Assange, J., Appelbaum, J., Müller-Maguhn, A., & Zimmermann, J. (2012). *Cypherpunks*. OR Books.
- Association of National Advertisers. (2014). *The Bot Baseline: Fraud in Digital Advertising*. Retrieved from www.ana.net/getfile/21853
- Atorough, P., & Donaldson, B. (2012). The relationship between regulatory focus and online shopping–perceived

risk, affect, and consumers' response to online marketing. *International Journal of Internet Marketing and Advertising*, 7(4), 333–358.

Austin, L. M. (2014). Enough About Me: Why Privacy is About Power, Not Consent (or Harm). In A. Sarat (Ed.), *A World Without Privacy? What Can/Should Law Do* (pp. 1–51). Cambridge University Press.

Aviv, R., Boardman, R., & Jones, W. (2004). SIG: Personal Information Management, 1598–1599.

Ayres, I., & Braithwaite, J. (1992). *Responsive regulation: Transcending the deregulation debate*. Oxford: Oxford University Press.

Baack, S. (2015). Datafication and empowerment: How the open data movement re-articulates notions of democracy, participation, and journalism. *Big Data & Society*, 2(2), 1–11.

Baek, T. H., & Morimoto, M. (2012). Stay away from me: Examining the determinants of consumer avoidance of personalized advertising. *Journal of Advertising*, 41(1), 59–76.

Baldwin, R., & Black, J. (2008). Really responsive regulation. *The Modern Law Review*.

Balebako, R., Leon, P. G., Almuhimedi, H., Kelley, P. G., Mungan, J., Acquisti, A., ... & Sadeh, N. (2011, May). Nudging users towards privacy on mobile devices. In *Proc. CHI 2011 Workshop on Persuasion, Nudge, Influence and Coercion*.

Bamberger, K. A., & Mulligan, D. K. (2013). Privacy in Europe : Initial Data on Governance Choices and Corporate Practices, 81(5), 1529–1664.

Banisar, David. (200) *Privacy and Human Rights 2000: An International Survey of Privacy Laws and Developments*. Electronic Privacy Information Center,.

Banisar, D. (2011). The right to information and privacy: balancing rights and managing conflicts. *World Bank Institute Governance Working Paper*.

Barbrook, R., & Cameron, A. (1996). The Californian ideology. *Science as Culture*, 6(1), 44–72.

- Barlow, J. P. (1996). A Declaration of the Independence of Cyberspace. Retrieved from <https://projects.eff.org/~barlow/Declaration-Final.html>
- Barnard, L. (2014). *Thesis: The cost of creepiness: How online behavioral advertising affects consumer purchase intention*. The University of North Carolina at Chapel Hill.
- Barnes, S. (2006). A privacy paradox: Social networking in the United States. *First Monday*, 11(9).
- Barocas, S., & Nissenbaum, H. (2009). On notice: The trouble with Notice and Consent. In *Proceedings of the Engaging Data Forum: The First International Forum on the Application and Management of Personal Electronic Information*.
- Barocas, S., & Selbst, A. D. (2014). Big Data's Disparate Impact. *Social Science Research Network Working Paper Series*.
- Bartle, I., & Vass, P. (2007). Self-Regulation Within the Regulatory State: Towards a New Regulatory Paradigm? *Public Administration*, 85(4), 885–905.
- Bass, B. G. D., Brian, D., & Eisen, N. (2014). *Why Critics of Transparency Are Wrong*. Washington, DC, Brookings Center for Effective Public Management.
- Bates, J. (2012). This is what modern deregulation looks like: co-optation and contestation in the shaping of the UK's Open Government Data Initiative. *Journal of Community Informatics*, 8(2).
- Beatty, P., Reay, I., Dick, S., & Miller, J. (2007). P3P Adoption on E-Commerce Web sites. *Internet Computing, IEEE 11.2*, (April), 65–71.
- Beaumont, R. (2014). Privacy Impact Assessments and the DPR. *EU Data Protection Law*. Retrieved from <http://www.eudataprotectionlaw.com/privacy-impact-assessments-and-the-dpr/>
- Bell, D. (2007). *Information Society*. Basic Books.
- Bell, E. A., Ohno-Machado, L., & Grando, M. A. (2014). Sharing My Health Data: A Survey of Data Sharing Preferences of Healthy Individuals. *AMIA Annual Symposium Proceedings. Vol. 2014*.

- Benkler, Y. (2002). Coase's Penguin, or, Linux and "The Nature of the Firm." *The Yale Law Journal*, 112(3), 369.
- Benkler, Y. (2006). *The wealth of networks: How social production transforms markets and freedom*. Yale University Press.
- Benkler, Y. (2010). The Idea of Access to Knowledge and the Information Commons: Long-Term Trends and Basic Elements. In *Access to Knowledge in the Age of Intellectual Property*.
- Benbear, S. (2007). Are Management-based Regulations Effective? Evidence from State Pollution Prevention Programs. *Journal of Policy Analysis and Management*, (26), 327–348.
- Bennett, C. J. (1992). *Regulating Privacy: Data Protection and Public Policy in Europe and the United States*.
- Bennett, C. J. (2010). *The privacy advocates*. Mit Press.
- Bennett, C. J., & Raab, C. D. (1997). The adequacy of privacy: The European Union data protection directive and the North American response. *The Information Society*, 13(3), 245–264.
- Bennett, C. J., & Raab, C. D. (2006). *The governance of privacy: Policy instruments in global perspective*.
- Bergelson, V. (2003). It's Personal But Is It Mine? Toward Property Rights in Personal Information. *Rutgers Law School (Newark) Faculty Papers*, (33).
- Bergkamp, L. (2002). EU Data Protection Policy: The Privacy Fallacy. *Computer Law & Security Review*, 18(1), 31–47.
- Berners-Lee, T. (2004). *Semantic Web Road Map*. Retrieved from www.w3.org/DesignIssues/Semantic.html
- Berners-Lee, T. (2006). *Linked Data*. Retrieved from www.w3.org/DesignIssues/LinkedData.html
- Berners-Lee, T., & Fischetti, M. (1999). *Weaving the web: The original design and ultimate destiny of the world wide web by its inventor*. San Francisco: Harper Collins.
- Berners-Lee, T., & O'Hara, K. (2013). The read–write Linked Data Web. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and*

Engineering Sciences, 371(1987), 20120513.

- Berners-Lee, T., Weitzner, D. J., Hall, W., O'Hara, K., Shadbolt, N., & Hendler, J. a. (2006). A Framework for Web Science. *Foundations and Trends in Web Science*, 1(1), 1–130.
- Bettman, J. R., Luce, M. F., & Payne, J. W. (1998). Constructive consumer choice processes. *Journal of consumer research*, 25(3), 187-217.
- Bignami, F. (2011). Cooperative Legalism and the Non-Americanization of European Regulatory Styles: The Case of Data Privacy. *American Journal of Comparative Law*, 59(2), 411–461. doi:10.5131/AJCL.2010.0017
- Billey, L. (2013). Responsiveness and Research, Focusing on Individuals. In *Advertising 2020. The Wharton Future of Advertising Program*. Retrieved from <http://wfoa.wharton.upenn.edu/perspective/lori-billey/>
- Binns, R. (2013a). 5 Stars of Personal Data Access. Retrieved from <http://www.reubenbinns.com/blog/5-stars-of-personal-data-access/>
- Binns, R. (2013b). Nudge Yourself. Retrieved from <http://www.reubenbinns.com/blog/nudge-yourself/>
- Binns, R. (2014a). Personal Data Empowerment and the Ideal Observer. In K. O'Hara, M. C. Nguyen, & P. Haynes (Eds.), *Digital Enlightenment Yearbook 2014 : Social Networks and Social Machines, Surveillance and Empowerment*.
- Binns, R. (2014b). Standardised Privacy Policies: A Post-mortem and Promising Developments. In *W3C Privacy Workshop: Privacy and User-Centric Controls*. Berlin. Retrieved from http://www.reubenbinns.com/blog/wp-content/uploads/2014/12/W3C_Privacy_User.pdf
- Binns, R. (2015). Caveat Venditor : Should We Sell Our Own Data ? In *Websci15: Workshop on the Economics of Surveillance*.
- Binns, R., & Lizar, M. (2012). Opening up the online notice infrastructure. In *W3C Privacy Workshop: Do Not Track and Beyond*. Retrieved from <http://eprints.soton.ac.uk/345931/>

- Binns, R., & Matthews, D. (2014). Community Structure for Efficient Information Flow in “ ToS ; DR ”, a Social Machine for Parsing Legalese. In *Proceedings of the companion publication of the 23rd international conference on World wide web companion* (pp. 881–884). Seoul, South Korea.
- Binns, R., Millard, D., & Harris, L. (2014, June). Data havens, or privacy sans frontières?: a study of international personal data transfers. In *Proceedings of the 2014 ACM conference on Web science* (pp. 273-274). ACM.
- Binns, R., Millard, D., & Harris, L. (2015). The Who, What and Why: An Analysis of Personal Data Transparency Notices in the UK. *Journal of Open Access to Law*, 3(1).
- Birnback, M. (2008). The EU Data Protection Directive: An engine of a global regime. *Computer Law & Security Review*, 24(6), 1–23.
- Black, J. (2001). Decentring Regulation: Understanding the Role of Regulation and Self-Regulation in a'Post-Regulatory'World. *Current legal problems*, 54(1), 103.
- Black, J. (2002). Critical reflections on regulation. *Austl. J. Leg. Phil.*, 27, 1.
- Black, J. (2008). Forms and paradoxes of principles-based regulation. *Capital Markets Law Journal*.
- Black, J. (2010) *The rise, fall and fate of principles based regulation*. LSE Law, Society and Economy working papers, 17-2010. Department of Law, London School of Economics and Political Science, London, UK.
- Black, J. (2012). Paradoxes and Failures: “New Governance” Techniques and the Financial Crisis. *The Modern Law Review*, 75(6), 1037–1063.
- Bleier, A., & Eisenbeiss, M. (2015). The Importance of Trust for Personalized Online Advertising. *Journal of Retailing*.
- Block, W., Whitehead, R., & Kinsella, N. S. (2005). The Duty to Defend Advertising Injuries Caused by Junk Faxes: An Analysis of Privacy, Spam, Detection and Blackmail. *Whittier Law Review*, 27, 925–950.
- Blom, J. (2002). A theory of personalized recommendations. *CHI '02 Extended Abstracts on Human Factors in*

Computing Systems - CHI '02, 540.

- Bloustein, Edward J. "Privacy as an aspect of human dignity: An answer to Dean Prosser." *NYUL Rev.* 39 (1964): 962.
- Bohm, P. (1987). *Social efficiency: a concise introduction to welfare economics*. Macmillan.
- Bonneau, J., & Preibusch, S. (2010). The Privacy Jungle: On the Market for Data Protection in Social Networks. *Economics of Information Security and Privacy*, 121–167.
- Boudreau, K. (2010). Open platform strategies and innovation: Granting access vs. devolving control. *Management Science*, 56(10), 1849-1872.
- Boutang, Y. M. (2011). *Cognitive Capitalism*. Polity Press.
- Boyd, D. (2008). Facebook's Privacy Trainwreck: Exposure, Invasion, and Social Convergence. *Convergence: The International Journal of Research into New Media Technologies*, 14(1), 13–20.
- Boyle, J. (2002). Fencing off ideas: enclosure & the disappearance of the public domain. *Daedalus*, 131(2), 13-25.
- Boyle, J. (2008). *The Public Domain: Enclosing the Commons of the mind*. The Bull Classics.
- Bradwell, P. (2010). *Private Lives: A People's Enquiry into Personal Information*. Retrieved from <http://www.demos.co.uk/publications/privatelives>
- Braithwaite, J. (2003). Meta-risk management and responsive regulation for tax system integrity. *Law and Policy*, 25.
- Bramwell, A. (2014). The Uncanny Valley and Why Big Data Marketers Are Headed Right For It. Retrieved from <https://www.linkedin.com/pulse/uncanny-valley-why-big-data-alex-bramwell>
- Brandeis, L. (1913). What publicity can do. *Harper's Weekly*, 20.
- Brandimarte, L., Acquisti, A., & Loewenstein, G. (2012). Misplaced Confidences: Privacy and the Control Paradox. *Social Psychological and Personality Science*, 4(3), 340–347.

- Brenner, N., & Theodore, N. (2002). Cities and the geographies of “actually existing neoliberalism”. *Antipode*, 34(3), 349-379.
- Brin, D. (1999). *The transparent society: will technology force us to choose between privacy and freedom?* Basic Books.
- Brockdorff, N., & Appleby-Arnold, S. (2013). What consumers think. *EU CONSENT Project, Workpackages*, 7.
- Brown, I. (2013). *Future Identities: Changing identities in the UK—the next 10 years*. Retrieved from https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/275760/13-509-how-will-ideology-affect-identity.pdf
- Brown, I. (2015). The economics of privacy, data protection and surveillance. In M. Latzer & J. M. Bauer (Eds.), *Handbook on the Economics of the Internet*. Cheltenham: Edward Elgar, Forthcoming.
- Brown, I., & Marsden, C. T. (2013). *Regulating code: Good governance and better regulation in the information age*. MIT Press.
- Brown, J. O., Broderick, A. J., & Lee, N. (2007). Word of Mouth Communication within Online Communities: Conceptualizing the Online Social Network. *Journal of Interactive Marketing*, 21(3), 2–20.
- Brown, M., & Muchira, R. (2004). Investigating the relationship between Internet privacy concerns and online purchase behavior. *Journal of Electronic Commerce Research*, 5(1), 62–70.
- Brownsword, R. (2008). *Rights, Regulation, and the technological revolution*. New York: Oxford University Press.
- Brunton, F., & Nissenbaum, H. (2012). Political and ethical perspectives on data obfuscation. In K. de Vries (Ed.), *Privacy, Due Process and the Computational Turn* (pp. 164–188).
- Bühler, S., Dewenter, R., & Haucap, J. (2006). Mobile number portability in Europe. *Telecommunications Policy*, 30(7), 385-399.
- Burke, R. (2005). Hybrid systems for personalized

recommendations. In *Intelligent Techniques for Web Personalization* (pp. 133–152.). Springer Berlin Heidelberg.

- Burton, G. (2014, April 2). ICO says anonymous data “not covered” by Data Protection Act - until it’s de-anonymised. *Computing*. Retrieved from <http://www.computing.co.uk/ctg/news/2337679/ico-says-anonymous-data-not-covered-by-data-protection-act-until-its-de-anonymised>
- Butin, D., Chicote, M., & Métayer, D. Le. (2012). Accountability by Design for Privacy. In *Privacy and Emerging Science and Technologies* (12).
- Byers, S., Cranor, L., & Kormann, D. (2003). Automated analysis of P3P-enabled web sites. *Proceedings of the 5th International Conference on Electronic Commerce*.
- Bygrave, L. (1998). Data protection pursuant to the right to privacy in human rights treaties. *International Journal of Law and Information Technology*, 6(3), 247–284.
- Bygrave, L. (2010). Privacy and data protection in an international perspective. *Scandinavian studies in law*, 56, 165-200.
- Bygrave, L. A., & Bing, J. (Eds.). (2009). *Internet Governance: Infrastructure and Institutions: Infrastructure and Institutions*. OUP Oxford.
- Caetano, a. (2014). Defining personal reflexivity: A critical reading of Archer’s approach. *European Journal of Social Theory*, 18(1), 60–75.
- Callon, M., & Muniesa, F. (2005). Economic Markets as Calculative Collective Devices. *Organization Studies*, 26(8), 1229–1250.
- Calo, R. (2012). Against Notice Skepticism in Privacy (and Elsewhere). *Notre Dame Law Review*, 87(3), 1027.
- Calo, R. (2013a). Code, Nudge, or Notice? *Iowa Law Review*, 9(9), 773–802.
- Calo, R. (2013b). Digital Market Manipulation. *University of Washington School of Law Legal Studies Research Paper*, 2013(27).

- Canhoto, A., & Backhouse, J. (2008). General description of the process of behavioural profiling. In M. Hildebrandt & S. Gutwirth (Eds.), *Profiling the European Citizen* (pp. 47–63). Springer.
- Carrera, S., Fuster, G. G., Guild, E., & Mitsilegas, V. (2015). *Access to Electronic Data by Third-Country Law Enforcement Authorities. Challenges to EU Rule of Law and Fundamental Rights*. Brussel: CEPS.
- Casellas, N., Nieto, J., Meroño, A., & Roig, A. (2006). Ontological Semantics for Data Privacy Compliance: The NEURONA Project, (1), 34–38. Retrieved from <http://www.aaai.org/ocs/index.php/SSS/SSS10/paper/viewFile/1071/1476>
- Caspar, J. (2015). The CJEU Google Spain decision. *Datenschutz Und Datensicherheit-DuD*, 39.9, 589–592.
- Castells, M. (1999). *The Information Age, Volumes 1-3: Economy, Society and Culture*. Oxford: Wiley-Blackwell.
- Cate, F. H. (1995). The EU Data Protection Directive, Information Privacy, and the Public Interest. *Iowa Law Review*.
- Cate, F. H. (2001, September 2). Invasions of Privacy? We'll All Pay the Cost if We Cut Free Flow of Information. *BOSTON GLOBE*.
- Cate, F. H. (2010). The limits of notice and choice. *Security & Privacy, IEEE*, 8(2), 59-62.
- Cavoukian, A. (2006). Privacy by Design: The 7 Foundational Principles. Implementation and Mapping of Fair Information Practices. *Information and Privacy Commissioner of Ontario, Canada*.
- Cavoukian, A., & El Emam, K. (2014). *The unintended consequences of privacy paternalism*. Canadian Electronic Library. Canadian Public Policy Collection.
- Charlesworth, A. (2006). The future of UK data protection regulation. *Information Security Technical Report*, 11(1), 46–54.
- Chauhan, S., & Rathore, S. (2015). Ethics in Behavioural Targeting: Mapping Consumers Perceptions. In IRMA (Ed.), *Business Law and Ethics: Concepts, Methodologies*,

Tools, and Applications (p. 303). AmSci Publications Office.

- Cho, C.-H., & Cheon, H. J. (2004). Why do people avoid advertising on the internet? *Journal of Advertising*, 33(4), 89–97.
- Chopra, S., & Dexter, S. (2010). Free software and the economics of information justice. *Ethics and Information Technology*, 13(3), 173–184.
- Chopra, S., & White, L. (2011). *A Legal Theory for Autonomous Artificial Agents*. University of Michigan Press.
- Clarke, R. (1993). Profiling: A Hidden Challenge to the Regulation of Data Surveillance. *Journal of Law and Information Science*, 4(3).
- Clarke, R. (1998). Serious flaws in the National Privacy Principles. *Privacy Law and Policy Reporter*, 4(9).
- Clarke, R. (2009). Privacy impact assessment: Its origins and development. *Computer Law & Security Review*, 25(2), 123–135. doi:10.1016/j.clsr.2009.02.002
- Clarke, R. (2011). An Evaluation of Privacy Impact Assessment Guidance Documents. *International Data Privacy Law*, 1(2), 111–120.
- Coase, R. (1937). The Nature of the Firm. *Economica*, 4(16), 386–405.
- Coglianesse, C., & Lazer, D. (2003). Management-based regulation: Prescribing private management to achieve public goals. *Law & Society Review*, 37(4), 691-730.
- Coglianesse, C., & Mendelson, E. (2010). Meta-regulation and self-regulation. In Baldwin, Cave and Lodge (Eds.) *The Oxford Handbook of Regulation*.
- Cohen, J. (1988). *Statistical Power Analysis for the Behavioral Sciences* (2nd ed.). Hillsdale: Earlbaum.
- Cohen, J. E. (2000). Examined Lives: Informational Privacy and the Subject as Object. *Stanford Law Review*, 52, 1424–25.
- Cohen, J. E. (2012). Introduction: Imagining the Networked Information Society. In *Configuring the Networked Self*

(pp. 1–24).

- Colburn, B. (2010). *Autonomy and Liberalism*. Routledge.
- Coles-Kemp, L., & Kani-Zabihi, E. (2010, September). On-line privacy and consent: a dialogue, not a monologue. In *Proceedings of the 2010 workshop on New security paradigms* (pp. 95-106). ACM.
- Cooley, T. (1879). *Law of Torts*.
- Costa, L. (2012). Privacy and the precautionary principle. *Computer Law & Security Review*, 28(1), 14–24.
- Cranor, L. (2012). Necessary but not sufficient: Standardized mechanisms for privacy notice and choice. *Journal on Telecommunications & High Technology Law*, 10, 273–307.
- Cranor, L. F., Egelman, S., Sheng, S., McDonald, A., and Chowdhury, A.. (2008) P3P deployment on websites. *Electronic Commerce Research and Applications* 7, no. 3 (2008): 274-293.
- Cranor, L. F., Idouchi, K., Leon, P. G., Sleeper, M., & Ur, B. (2013). Are They Actually Any Different? Comparing Thousands of Financial Institutions' Privacy Practices. *The Twelfth Workshop on the Economics of Information Security (WEIS 2013)*, June 11–12, 2013, Washington, DC.
- Ctrl-Shift. (2014). Personal Information Management Services : An analysis of an emerging market, (June).
- Culnan, M. J. (2000). Protecting Privacy Online: Is Self-Regulation Working? *Journal of Public Policy & Marketing*, 19(1), 20–26.
- Culnan, M. J., & Armstrong, P. K. (1999). Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation. *Organization Science*, 10(1), 104–115.
- Custers, B., & Hof, S. Van Der. (2013). Informed consent in social media use. The gap between user expectations and EU personal data protection law. *SCRIPTed*, 10(4).
- Danezis, G., & Domingo-Ferrer, J. (2015). Privacy and Data Protection by Design-from policy to engineering. *arXiv*

Preprint, (December).

- Davies, W. (2014). Neoliberalism: A bibliographic review. *Theory, Culture & Society*, 0263276414546383.
- De Hert, P. (2012). A Human Rights Perspective on Privacy and Data Protection Impact Assessments. In D. Wright & P. DeHert (Eds.), *Privacy Impact Assessment* (pp. 33–74). Springer Science & Business Media.
- De Vries, K. (2010). Identity, profiling algorithms and a world of ambient intelligence. *Ethics and Information Technology*, 12(1), 71–85.
- Dean, M. (1999). *Governmentality, Power and Rule in Modern Society*. London: Sage Publications.
- Deane, J. K., Meuer, T., & Teets, J. M. (2011). A longitudinal analysis of web surf history to maximise the effectiveness of behavioural targeting techniques. *International Journal of Electronic Marketing and Retailing*, 4(2-3), 117–128.
- DeNardis, L. (2011). *Opening standards: the global politics of interoperability*. MIT Press.
- Denegri-Knott, J. (2006). Consumers behaving badly: Deviation or innovation? Power struggles on the web. *Journal of Consumer Behaviour*, 5(1), 82–94.
- Dennett, D., & Roy, D. (2015). How Digital Transparency Became a Force of Nature. *Scientific American*, 312(3).
- Doctorow, C. (2004). Microsoft research DRM talk. (Transcript), Microsoft Research Group, Redmond, WA 17, USA. Retrieved from <http://www.cs.ucdavis.edu/~rogaway/classes/188/materials/doctorow.pdf>
- Dorbeck-Jung, B., & Shelley-Egan, C. (2013). Meta-regulation and nanotechnologies: the challenge of responsabilisation within the European Commission's code of conduct for responsible nanosciences and. *Nanoethics*, 55–68.
- Draper, N. (2015). The Promise of Small Data: Regulating Individual Choice Through Access to Personal Information. In *Data Power Conference, Sheffield 2015*.
- Drucker, P. (1969). *The Age of Discontinuity; Guidelines to Our Changing Society*. New York: Harper and Row.

- Duff, B. R. L., & Faber, R. J. (2011). Missing the mark: Advertising avoidance and distractor devaluation. *Journal of Advertising*, 40(2), 51–62.
- Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2011). Fairness Through Awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference. Computational Complexity; Computers and Society*.
- Dwork, C., & Mulligan, D. (2013). It's not privacy, and it's not fair. *Stanford Law Review Online*, 4(2000), 35–40.
- Eberle, E. J. (1998). Human Dignity, Privacy, and Personality in German and American Constitutional Law. *Utah Law Review*, 4.
- Edelman, L., Petterson, S., Chambliss, E., & Erlanger, H. (1991). Legal ambiguity and the politics of compliance: Affirmative action officers' dilemma. *Law & Policy*, (13).
- Edwards, J. (2013). BEHOLD: The First Banner Ad Ever — From 1994. *Business Insider*. Retrieved from <http://www.businessinsider.com/ behold-the-first-banner-ad-ever--from-1994-2013-2?IR=T>
- Egelman, S., & Tsai, J. (2006). Studying the impact of privacy information on online purchase decisions. *Workshop on Privacy and HCI: Methodologies for Studying Privacy Issues at CHI, 2006* 1–4.
- Eijlander, P. (2005). Possibilities and Constraints in the Use of Self-regulation and Co-regulation in Legislative Policy: Experiences in the Netherlands-Lessons to Be Learned for the EU? *European Journal of Comparative Law*, 9(1).
- Electronic Privacy Information Center. (2000). *Pretty Poor Privacy: An Assessment of P3P and Internet Privacy*.
- Emarketer. (2014). Digital Ad Spending Worldwide to Hit \$137.53 Billion in 2014. *Emarketer*. Retrieved from <http://www.emarketer.com/article.aspx?R=1010736>
- Eslami, M., Rickman, A., Vaccaro, K., Aleyasen, A., Vuong, A., Karahalios, K., Hamilton, K. and Sandvig, C., 2015, April. I always assumed that I wasn't really that close to [her]: Reasoning about Invisible Algorithms in News Feeds. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (pp. 153-162).

ACM.

- Etzioni, A. (1999). *The Limits of Privacy*. Basic Books.
- European Commission. (2010a). *COMMUNICATION FROM THE COMMISSION TO THE EUROPEAN PARLIAMENT, THE COUNCIL, THE ECONOMIC AND SOCIAL COMMITTEE AND THE COMMITTEE OF THE REGIONS A comprehensive approach on personal data protection in the European Union* /* COM/2010/0609 final */. Retrieved from <http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52010DC0609>
- European Commission. (2010b). Comparative Study on Different Approaches to New Privacy Challenges, In Particular in the Light of Technological Developments.
- European Commission. (2012). Proposal for a Regulation of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation), /* COM/2012/0011 final */.
- European Council and Parliament (1995). Directive 95/46/EC of 24 October 1995 on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of Such Data, 1995.
- Fairman, R., & Yapp, C. (2005). Enforced Self-Regulation, Prescription, and Conceptions of Compliance within Small Businesses: The Impact of Enforcement*. *Law & Policy*, 27(4), 491–519. doi:10.1111/j.1467-9930.2005.00209.x
- Federal Trade Commission. (2013). *FTC Testifies on Data Brokers Before Senate Committee on Commerce, Science and Transportation*. Retrieved from <http://www.ftc.gov/news-events/press-releases/2013/12/ftc-testifies-data-brokers-senate-committee-commerce-science>
- FEDWeek. (2013). IG Calls for Improvements to IRS Privacy Impact Assessment Process. *FEDWeek*. Retrieved from <http://www.fedweek.com/federal-managers-daily-report/ig-calls-for-improvements-to-irs-privacy-impact-assessment-process/>
- Finn, R., Wright, D., & Friedewald, M. (2013). Seven types of

privacy. In S. Gutwirth, R. Leenes, P. de Hert, & Y. Poullet (Eds.), *European data protection: coming of age*. Dordrecht: Springer Netherlands.

- Firth, R. (1952). Ethical Absolutism and the Ideal Observer. *Philosophy and Phenomenological Research*, 12(3), 317–345.
- Flaherty, D. (2000). Privacy impact assessments: an essential tool for data protection. *Privacy Law & Policy Reporter*, 5, 85.
- Flavián, C., & Guinalú, M. (2006). Consumer trust, perceived security, and privacy policy: Three basic elements of loyalty to a website. *Industrial Management and Data Systems*, 106, 601–620.
- Floridi, L. (2014). Open data, data protection, and group privacy. *Philosophy & Technology*, 1–3.
- Ford, C. L. (2008). New Governance, Compliance, and Principles-based Securities Regulation. *American Business Law Journal*, 25, 1–60.
- Ford, C. L. (2010). Principles-based Securities Regulation in the Wake of the Global Financial Crisis. *McGill Law Journal*, (55), 257–310.
- Francis, J. G., & Francis, L. P. (2014). Privacy, Confidentiality, and Justice. *Journal of Social Philosophy*, 45(3), 408–431.
- Fuchs, C. (2012). The political economy of privacy on Facebook. *Television & New Media*. Retrieved from <http://tvn.sagepub.com/content/13/2/139.short>
- Fukuyama, F. (2014). *Political Order and Political Decay: From the Industrial Revolution to the Globalization of Democracy*. NY: Farrar, Straus and Giroux.
- Füllera, J., Mühlbacherb, H., Matzlerb, K., & Jaweckic, G. (2009). Consumer Empowerment Through Internet-Based Co-creation. *Journal of Management Information Systems*, 26(3).
- Fuster, G. G. (2014). *The Emergence of Personal Data Protection as a Fundamental Right of the EU* (p. 274). Springer International Publishing.
- Fuster, G., & Gutwirth, S. (2013). Opening up personal data

protection: A conceptual controversy. *Computer Law & Security Review*, 9(1). doi:10.1016/j.clsr.2013.07.008

- Gandon, F., & Sadeh, N. (2003). Semantic web technologies to reconcile privacy and context awareness. *Institute for Software Research, Paper 848*. Retrieved from <http://www.sciencedirect.com/science/article/pii/S1570826804000022>
- Gandy, O. (2010). Engaging rational discrimination: exploring reasons for placing regulatory constraints on decision support systems. *Ethics and Information Technology*, 12(1), 29–42.
- Geist, M. (2012). Privacy Commissioner Should Disclose the Identities of Privacy Leakers. Retrieved from <http://www.michaelgeist.ca/2012/10/privacy-commish-on-leaky-sites/>
- Gellert, R., & Gutwirth, S. (2012). Beyond accountability, the return to privacy. In D. Guagnin, L. Hempel, C. Ilten, I. Kroener, D. Neyland, & H. Postigo (Eds.), *Managing privacy through accountability* (pp. 261–284).
- Gilad, S. (2010). It runs in the family: Meta-regulation and its siblings. *Regulation & Governance*, 4(4), 485–506. doi:10.1111/j.1748-5991.2010.01090.x
- Grabosky, P. (1995). Using non- governmental resources to foster regulatory compliance. *Governance*, 8.
- Gramsci, A. (1995). Americanism and fordism. In *Prison Notebooks*. University of Minnesota Press.
- Greenleaf, G. (2011). Asia-Pacific data privacy: 2011, year of revolution? *Kyung Hee Law Journal*, *Forthcoming*. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?Abstract_id=1914212
- Greenleaf, G. (2012a). Global data privacy in a networked world. In I. Brown (Ed.), *Research Handbook on Governance of the Internet*. Edward Elgar.
- Greenleaf, G. (2012b). Global data privacy laws: 89 countries, and accelerating. *Privacy Laws & Business International Report*, (98).
- Greenstadt, R., & Smith, M. (2005). Protecting Personal

- Information: Obstacles and Directions. *WEIS*, 1–22.
Retrieved from
<http://www.infoecon.net/workshop/pdf/48.pdf>
- Grimmelmann, J. (2010). Privacy as Product Safety. *Widener Law Journal*, 793–827.
- Grist, M. (2010). *Steer: Mastering our behaviour through instinct, environment and reason*. Retrieved from
http://www.thersa.org/_data/assets/pdf_file/0017/313208/RSA-Social-Brain_WEB-2.pdf
- Grossman, S. J., & Hart, O. D. (1983). An analysis of the principal-agent problem. *Econometrica: Journal of the Econometric Society*, 7-45.
- Gummerus, J., Liljander, V., Weman, E., & Pihlstrom, M. (2012). Customer engagement in a Facebook brand community. *Management Research Review*, 35(9), 857–877.
- Gunningham, N., & Rees, J. (1997). Industry self-regulation: an institutional perspective. *Law & Policy*, 1. Retrieved from <http://onlinelibrary.wiley.com/doi/10.1111/1467-9930.t01-1-00033/abstract>
- Gunningham, N., & Sinclair, D. (2009). On the Limits of Management Based Regulation. *Law and Society Review*, 2(43), 865–900.
- Gurstein, M. (2011). Open Data: Empowering the empowered or effective data use for everyone? *First Monday*, 16(2).
- Gutwirth, S., & DeHert, P. (2006). Privacy, Data Protection and Law Enforcement. Opacity of the Individual and Transparency of Power. In E. Claes, A. Duff, & S. Gutwirth (Eds.), *Privacy and the criminal law*. Intersentia nv.
- Hacking, I. (1990). *The Taming of Chance*. Cambridge University Press.
- Hadija, Z., Barnes, S. B., & Hair, N. (2012). Why we ignore social networking advertising. *Qualitative Market Research*, 15(1), 19–32.
- Haines, F. (2009). Regulatory Failures and Regulatory Solutions: A Characteristic Analysis of the Aftermath of

- Disaster. *Law and Social Inquiry*, (34), 31–60.
- Haines, F. (2011). *The Paradox of Regulation: What Regulation Can Achieve and What It Cannot*. Edward Elgar Publishing.
- Halford, S., Pope, C., & Carr, L. (2010). A manifesto for web science. *Web Science Conf. 2010, April 26-27, 2010, Raleigh, NC, USA.*, 1–6. Retrieved from <http://eprints.soton.ac.uk/271033/>
- Hayes, C. (2008). *Popper, Hayek and the open society*. Routledge.
- Heath, W., Alexander, D., & Booth, P. (2013). Digital enlightenment, Mydex, and restoring control over personal data to the individual. In M. Hildebrandt, K. O'Hara, & M. Waidner (Eds.), *Digital Enlightenment Yearbook 2013: The Value of Personal Data* (pp. 253–269). IOS Press.
- Heimstädt, M., Saunderson, F., & Heath, T. (2014). From Toddler to Teen: Growth of an Open Data Ecosystem. *eJournal of eDemocracy & Open Government*, 6(2).
- Helberger, N. (2013). Form matters: informing consumers effectively. *IVIR*, 1(51).
- Henkin, L. (1974). Privacy and Autonomy. *Columbia Law Review*, 74(8), 1410–1433.
- Heydebrand, W. (2003). Process Rationality as Legal Governance: A Comparative Perspective. *International Sociology*, 18(2), 325–349.
- Hildebrandt, M. (2012). The dawn of a critical transparency right for the profiling era. *Digital Enlightenment Yearbook 2012*, 12(2008), 41–56. Retrieved from http://books.google.com/books?hl=en&lr=&id=D_ZAHbaJXO8C&oi=fnd&pg=PA41&dq=The+Dawn+of+a+Critical+Transparency+Right+for+the+Profiling+Era&ots=d9n6HvAgqf&sig=RCTddE3Onj3iL1pvnB_1lyKdcfw
- Hildebrandt, M. (2013). Profile transparency by design: Re-enabling double contingency. In M. Hildebrandt & K. De Vries (Eds.), *Privacy, Due Process and the Computational Turn*. Routledge.

- Hildebrandt, M., & Gutwirth, S. (2008). *Profiling the European citizen* (Vol. c, pp. 1–390). Springer.
- Hilgers, D., & Ihl, C. (2010). Citizensourcing: Applying the Concept of Open Innovation to the Public Sector. *The International Journal of Public Participation*, 4(1).
- Hiriart, Y., Martimort, D., & Pouyet, J. (2004). On the optimal use of ex ante regulation and ex post liability. *Economics Letters*, 33(March), 1–5. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0165176504000862>
- Hirsch, D. D. (2010). The Law and Policy of Online Privacy. *Seattle UL Rev.*, 34, 439–480.
- Hof, R. (2006). If I Had a Nickel For Every Click. *Bloomberg Business*. Retrieved from <http://www.bloomberg.com/bw/stories/2006-09-24/if-i-had-a-nickel-for-every-click-dot-dot-dot>
- Hof, S. van der, & Prins, C. (2008). Personalisation and its Influence on Identities, Behaviour and Social Values. In S. Gutwirth & M. Hildebrandt (Eds.), *Profiling the European Citizen* (pp. 111–127). Springer.
- Hong, W., Thong, J. Y. L., & Tam, K. Y. (2004). Does Animation Attract Online Users' Attention? The Effects of Flash on Information Search Performance and Perceptions. *Information Systems Research*, 15(1), 60–86. doi:10.1287/isre.1040.0017
- Hoofnagle, C. J. (2009). Beyond Google and evil: How policy makers, journalists and consumers should talk differently about Google and privacy. *First Monday*, 14(4).
- Hosein, G., & Davies, S. (2013). Empirical research of contextual factors affecting the introduction of privacy impact assessment frameworks in the Member States of the European Union. *PIAF Deliverable D2*.
- Humboldt, W. (1810). *Ideen Zu Einem Versuch, Die Grenzen Der Wirksamkeit Des Staats Zu Bestimmen*.
- Hutter, M. (2001). *Regulation and Risk: Occupational Health and Safety on the Railways*. Oxford University Press.
- Hwang, T., & Kamdar, A. (2013). *The Theory of Peak*

Advertising and the Future of the Web (pp. 1–12).

- IAPP. (2015). HP-IAPP Privacy Innovation Awards. *www.iapp.org*. Retrieved from <https://iapp.org/about/annual-awards/hp-iapp-privacy-innovation-awards/>
- Iliadis, A., & Russo, F. (2015). Call for Proposals: Special Theme on “Critical Data Studies.” *Big Data & Society*. Retrieved from <http://bigdatasoc.blogspot.co.uk/2015/06/call-for-proposals-special-theme-on.html>
- Iliev, A., & Smith, S. W. (2005). Protecting client privacy with trusted computing at the server. *IEEE Security & Privacy*, 2, 20–28.
- Information and Privacy Commissioner of Ontario. (1994). *Suggested changes to the municipal freedom of information and protection of privacy act: submission to the standing committee on the legislative assembly*.
- Information Commissioner’s Office. (2007). *Privacy Impact Assessment Handbook*.
- Information Commissioner’s Office. (2014). Conducting privacy impact assessments code of practice. *Ico.org.uk*. Retrieved from https://ico.org.uk/for_organisations/guidance_index/~media/documents/library/Data_Protection/Practical_application/pia-code-of-practice-final-draft.pdf
- Information Commissioner’s Office. (2015a). Conditions for Processing. *www.ico.org.uk*. Retrieved from <https://ico.org.uk/for-organisations/guide-to-data-protection/conditions-for-processing/>
- Information Commissioner’s Office. (2015b). *Summary of feedback on Big data and data protection and ICO response* (pp. 1–10).
- Innes, R. (2004). Enforcement costs, optimal sanctions, and the choice between ex-post liability and ex-ante regulation. *International Review of Law and Economics*, 24(1), 29–48. doi:10.1016/j.irl.2004.03.003
- Involve UK. (2014). Data sharing open policy process. Retrieved from

<http://www.involve.org.uk/blog/2014/05/02/data-sharing-open-policy-process/>

- Iyer, G., Soberman, D., & Villas-Boas, J. (2005). The targeting of advertising. *Marketing Science*.
- Janssen, K., & Hugelier, S. (2013). Open Data: A New Battle in an Old War Between Access and Privacy. In *Digital Enlightenment Yearbook 2013: The Value of Personal Data2*.
- Jarvis, J. (2011). *Public parts: How sharing in the digital age improves the way we work and live*. Simon and Schuster.
- Jensen, M., & Meckling, W. H. (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4), 305–360.
- Jessop, B. (2001). State Theory, Regulation and Autopoiesis: Debates and Controversies. *Capital and Class*, (75), 83–92.
- Jin, C. H., & Villegas, J. (2007). Consumer responses to advertising on the Internet: The effect of individual difference on ambivalence and avoidance. *CyberPsychology and Behaviour*, 10(2), 258–266.
- JISC Legal. (2012). ICO Consults on Proposed Changes to Data Protection Notification. JISC Legal. Retrieved from <http://www.jisclegal.ac.uk/ManageContent/ViewDetail/ID/2638/ICO-Consults-on-Proposed-Changes-to-Data-Protection-Notification.aspx>
- Johnson, J. A. (2014). From open data to information justice. *Ethics and Information Technology*, 16(4), 263–274.
- Jonas, J., & Harper, J. (2010). Open Government: The Privacy Imperative. In Lathrop, D., & Ruma, L. (eds.), *Open government: Collaboration, transparency, and participation in practice*. O'Reilly Media, Inc. (pp. 321–330).
- Jones, W. (2007). Personal information management. *Annual Review of Information Science and Technology*, 41(1), 453–504.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. Macmillan.
- Kaplow, L., & Shavell, S. (2002). *Fairness versus Welfare*.

- Cambridge, Massachusetts: Harvard University Press.
- Karol, T. J. (2001). Cross-Border Privacy Impact Assessments: An Introduction. *Information Systems Control*, 3, 50–52.
- Kassen, M. (2013). A Promising Phenomenon of Open Data: A Case Study of the Chicago Open Data Project. *Government Information Quarterly*, (508).
- Kautilya. (1929). *Arthashastra* (translated.). Wesleyan Mission Press.
- Kay, J. (1994). The um toolkit for cooperative user modelling. *User Modeling and User-Adapted Interaction*, 4(3), 149–196.
- Kay, J., & Kummerfeld, B. (2012). Creating personalized systems that people can scrutinize and control: Drivers, principles and experience. *ACM Transactions on Interactive Intelligent Systems (TiiS)*, 2(4), 24.
- Kelley, K. (2014). Why You Should Embrace Surveillance, not Fight it. *Wired*. Retrieved from <http://www.wired.com/2014/03/going-tracked-heres-way-embrace-surveillance/>
- Kelley, P., Cesca, L., Bresee, J., & Cranor, L. (2010). Standardizing privacy notices: an online study of the nutrition label approach. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.
- Kelty, C. (2008). *Two bits: The cultural significance of free software*. Durham: Duke University Press.
- Kirkham, T., & Winfield, S. (2011). A personal data store for an internet of subjects. In *2011 International Conference on Information Society (i-Society)* (pp. 92–97).
- kitabeta [twitter user]. (2015). .@katecrawford on “the great algorithm panic of 2015” & accountability through lens of the deodand (subtweeting assemblage theory...) #TtW15 [Twitter post]. Retrieved from <https://twitter.com/kitabeta/status/589561477578756096>
- Kitchin, R. (2014). Big Data, new epistemologies and paradigm shifts. *Big Data & Society*, 1(1), 1–12.
- Kitchin, R., & Lauriault, T. (2014). Towards critical data studies: Charting and unpacking data assemblages and

their work. *The Programmable City Working Paper 2*. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2474112

- Kleek, M. Van, Smith, D. A., & Shadbolt, N. R. (2012). A decentralized architecture for consolidating personal information ecosystems : The WebBox. Retrieved from <http://eprints.ecs.soton.ac.uk/23200/1/webbox-pim.pdf>
- Kleek, M. Van, Smith, D., & Packer, H. (2013). Carpé data: supporting serendipitous data integration in personal information management. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.
- Kling, C. C., Kunegis, J., Hartmann, H., Strohmaier, M., & Staab, S. (2015). Voting Behaviour and Power in Online Democracy: A Study of LiquidFeedback in Germany's Pirate Party. *arXiv preprint arXiv:1503.07723*.
- Kolovski, V., Katz, Y., & Hendler, J. (2005). Towards a policy-aware web. *Semantic Web and Policy Workshop at the 4th International Semantic Web Conference*.
- Kolstad, C. D., Ulen, T. S., & Johnson, G. V. (1990). Ex Post Liability for Harm vs. Ex Ante Safety Regulation: Substitutes or Complements? *The American Economic Review*, 80(4), 888–901.
- König, U., & Hansen, M. (2011). Extending Comparison Shopping Sites by Privacy Information on Retailers. In *Privacy and Identity Management for Life* (pp. 171-186). Springer Berlin Heidelberg.
- Koops, B. J. (2013). On decision transparency, or how to enhance data protection after the computational turn. In Hildebrandt, Mireille, and Katja De Vries (eds.) *Privacy, due process and the computational turn: the philosophy of law meets the philosophy of technology*, 196-220.
- Koops, B. J. (2014). The trouble with European data protection law. *International Data Privacy Law*, 4(4), 250-261.
- Koops, B.-J. (2011). The Evolution of Privacy Law and Policy in the Netherlands. *Journal of Comparative Policy Analysis: Research and Practice*, 13(2), 165–179.
- Kost, M., Freytag, J.-C., Kargl, F., & Kung, A. (2011). Privacy Verification Using Ontologies. *2011 Sixth International Conference on Availability, Reliability and Security*, 627–

- Kox, H. (2014). The online advertising and tracking industry: technology, business model, and market structure. From *The Selected Works of Henk LM Kox*, (January). Retrieved from http://works.bepress.com/cgi/viewcontent.cgi?article=1127&context=henk_kox
- Kox, H. L. M., Straathof, B., & Zwart, G. (2014). Targeted advertising, platform competition and privacy. *CPB Discussion Paper*, (280). Retrieved from http://works.bepress.com/henk_kox/65/
- Kramer, A. DI, Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, *111*(24), 8788–8790.
- Krasnova, H., Günther, O., Spiekermann, S., & Koroleva, K. (2009). Privacy concerns and identity in online social networks. *Identity in the Information Society*, *2*(1), 39–63.
- Krikorian, G., & Kapczynski, A. (2010). *Access to Knowledge in the Age of Intellectual Property*. Zone Books.
- L Hoffman. (1973). *Security and privacy in computer systems*. Wiley.
- Lange, B. (1998). Understanding Regulatory Law: Empirical vs Systems-Theoretical Approaches? *Oxford Journal of Legal Studies*, *18*, 449–71.
- Lange, B. (2003). Regulatory spaces and interactions: An introduction. *Social & Legal Studies*, *12*(200312), 411–423.
- Langheinrich, M. (2001). Privacy in Ubiquitous Computing.
- Laufer, W. (1999). Corporate Liability, Risk Shifting, and the Paradox of Compliance. *Vanderbilt Law Review*, (54), 1343–9.
- Le Grand, G., & Barrau, E. (2012). Prior Checking, A Forerunner to Privacy Impact Assessments. In (Wright, D., and De Hert, P., (eds.) *Privacy Impact Assessment* (pp. 97–115).
- Leeper, J. (2000). Choosing the Correct Statistical Test. Retrieved from <http://bama.ua.edu/~jleeper/>.

- Lehmkuhl, D. (2008). Control Modes in the Age of Transnational Governance. *Law & Policy*, 30(3), 336–363.
- Lehtiniemi, T. (2015). The Calculative Power Over Personal Data. In *Data Power Conference, Sheffield 2015*.
- Lemley, M. A. (2000). Private property. *Stanford Law Review*, 1545-1557.
- Leon, P., & Cranor, L. (2010). Token attempt: the misrepresentation of website privacy policies through the misuse of p3p compact policy tokens. *ACM Workshop on Privacy*.
- Leon, P.G., Ur, B., Wang, Y., Sleeper, M., Balebako, R., Shay, R., Bauer, L., Christodorescu, M. and Cranor, L.F., (2013). What matters to users?: factors that affect users' willingness to share information with online advertisers. In *Proceedings of the ninth symposium on usable privacy and security* (p. 7). ACM.
- Lessig, L. (1999). *Code and other Laws of Cyberspace*. Basic Books.
- Lessig, L. (2004). *Free culture: How big media uses technology and the law to lock down culture and control creativity*. Penguin.
- Lessig, L. (2006). The Read/Write Society. In *Lecture delivered on Friday, 15 September 2006, at Humboldt University, Berlin*.
- Li, H., Sarathy, R., & Xu, H. (2010). Understanding situational online information disclosure as a privacy calculus. *Journal of Computer Information Systems*, (2010), 1–29.
- Lichtenstein, C. C. (2001). Hard Law v. Soft Law: Unnecessary Dichotomy? *The International Lawyer*, 35(4), 1433–1441.
- Lim, K., & Benbasat, I. (2000). The effect of multimedia on perceived equivocality and perceived usefulness of information systems. *MIS Quarterly*, 24(3), 449–471.
- Linden, G., Smith, B., & York, J. (2003). Amazon. com recommendations: Item-to-item collaborative filtering. *Internet Computing, IEEE*, (February).
- Lindenberg, F. (2014). How Can Online Research Tools Help Investigative Reporters? *Global Investigative Journalism*

- Network*. Retrieved from <http://gijn.org/2014/09/26/how-can-online-research-tools-help-investigative-reporters/>
- Litman, J. (2000). Information Privacy/Information Property. *Stanford Law Review*, 52(5), 1283.
- Lomas, N. (2013). Handshake Is A Personal Data Marketplace Where Users Get Paid To Sell Their Own Data. *TechCrunch*. Retrieved from <http://techcrunch.com/2013/09/02/handshake/>
- Longo, J. (2011). Open Data: Digital-Era Governance Thoroughbred or New Public Management Trojan Horse?'. *Public Policy and Governance Review*, 2(2), 38–52.
- Lucas, G. R. (2014). NSA management directive# 424: secrecy and privacy in the aftermath of Edward Snowden. *Ethics & International Affairs*, 28(1), 29–38.
- Lundblad, N., & Masiello, B. (2010). Opt-in dystopias. *SCRIPTed*, 7(1).
- Lynskey, O. (2014). Deconstructing data protection : the 'added-value' of a right to data protection in the EU legal order. *International & Comparative Law Quarterly*, 1–20.
- Lyon, D. (2001). *Surveillance Society: Monitoring Everyday Life*. Open University Press.
- MacLean, T. (2002). Reframing Organizational Misconduct: A Study of Deceptive Sales Practices at a Major Life Insurance Company. *Business and Society*, (41), 242–250.
- Mainier, M. J., & O'Brien, M. (2010). Online social networks and the privacy paradox: A research framework. *Issues in Information Systems*, XI(1), 513–517.
- Malheiros, M., Jennett, C., & Patel, S. (2012). Too close for comfort: A study of the effectiveness and acceptability of rich-media personalized advertising. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.
- Mann, S., Nolan, J., & Wellman, B. (2002). Sousveillance: Inventing and Using Wearable Computing Devices for Data Collection in Surveillance Environments. *Surveillance & Society*, 1.3, 331–355.

- Mantelero, A. (2014). The future of consumer data protection in the E.U. Re-thinking the 'notice and consent' paradigm in the new era of predictive analytics. *Computer Law & Security Review*, 30(6), 643–660.
- Marres, N. (2013). Why political ontology must be experimentalized: On eco-show homes as devices of participation. *Social Studies of Science*, 43(3), 417–443.
- Marsden, C. T. (2011). *Internet Co-Regulation: European Law, Regulatory Governance and Legitimacy in Cyberspace*. Cambridge University Press.
- Marx, G. (2012). Foreword: Privacy is not quite like the weather. In D. Wright & P. De Hert (Eds.), *Privacy Impact Assessment* (pp. v–xiv). Springer.
- Marx, K. (1939). *Grundrisse der Kritik der politischen Ökonomie*. Europäische Verlags-Anstalt.
- Mascetti, S., Ricci, A., & Ruggieri, S. (2014). Introduction to special issue on computational methods for enforcing privacy and fairness in the knowledge society. *Artificial Intelligence and Law*, 109–111.
- Mason, E. (2015). Value Pluralism. In *The Stanford Encyclopedia of Philosophy*. Retrieved from <http://plato.stanford.edu/archives/sum2015/entries/value-pluralism/>
- May, L. (1980). Privacy and property. *Philosophy in Context*, 10(6), 40–53.
- McBarnet, D. (2007). Corporate Social Responsibility Beyond Law, Through Law, For Law. In D. McBarnet, A. Voiculescu, & T. Campbell (Eds.), *The New Corporate Accountability: Corporate Social Responsibility and the Law*. Cambridge University Press.
- McCole, P., Ramsey, E., & Williams, J. (2010). Trust considerations on attitudes towards online purchasing: The moderating effect of privacy and security concerns. *Journal of Business Research*, (63), 1018–1024.
- McDonald, A., & Cranor, L. (2008). The Cost of reading privacy policies. *ISJLP*, 0389, 1–22.
- McDonald, A. M., Reeder, R. W., Kelley, P. G., & Cranor, L. F.

- (2009). A comparative study of online privacy policies and formats. *Proceedings of the 5th Symposium on Usable Privacy and Security - SOUPS '09*, 1.
- McEwan, P. (2014). The Uncanny Valley of Interactive Advertising. Retrieved from <http://tribalyell.com/the-uncanny-valley-in-interactive-advertising/>
- McKnight, D. H., Choudhury, V., & Kacmar, C. (2002). The impact of initial consumer trust on intentions to transact with a web site: a trust building model. *The Journal of Strategic Information Systems*, 11(3), 297–323.
- McLaughlin, E. C. (2014). After Eric Garner: What's point of police body cameras? *CNN*. Retrieved from <http://edition.cnn.com/2014/12/04/us/eric-garner-ferguson-body-cameras-debate/index.html>
- McNee, S. M. (2003). Interfaces for eliciting new user preferences in recommender systems. In *User Modeling* (pp. 178–187). Springer Berlin Heidelberg.
- McQuillan, D. (2014). Activism and the Internet of Things. Retrieved from http://www.internetartizans.co.uk/activism_and_internet_of_things_abstract
- Mcquillan, D. (2015). Algorithmic States of Exception. *European Journal of Cultural Studies*, 18.4(5).
- Meijer, A. (2009). Understanding modern transparency. *International Review of Administrative Sciences*, 75(2), 255–269.
- Meijer, A. (2013). Understanding the complex dynamics of transparency. *Public Administration Review*, 73, 429–439.
- Meijer, R., Conradie, P., & Choenni, S. (2014). Reconciling contradictions of open data regarding transparency, privacy, security and trust. *Journal of theoretical and applied electronic commerce research*, 9(3), 32-44.
- Mestdagh, C. N. J.D.V, & Rijgersberg, R. W. (2015). Legisprudence Internet Governance and Global Self Regulation. *Legisprudence*, IV(3), 37–41.
doi:10.1080/17521467.2010.11424719
- Microsoft. (2013). Privacy impact assessments, (February). Retrieved from <http://www.rogerclarke.com/DV/PIA.html>

- Mikians, J., Gyarmati, L., Erramilli, V., & Laoutaris, N. (2012). Detecting price and search discrimination on the internet. *Proceedings of the 11th ACM Workshop on Hot Topics in Networks - HotNets-XI*, 79–84.
- Mill, J. S. (1859). *On Liberty*. (S. Collini, Ed.) (On Liberty.). Cambridge University Press.
- Millard, C., & Church, P. (2007). Tissue Sample and Graffiti: Personal Data and the Article 29 Working Party. *Computers & Law*, 18(3), 27–29.
- Millard, C., & Kuan Hon, W. (2012). Defining “Personal Data” in e-Social Science. *Information, Communication & Society*, 15(1), 66–84.
- Miller, A. (2014). What do we worry about when we worry about price discrimination? The law and ethics of using personal information for pricing. *Journal of Technology Law and Policy*, 41–104.
- Milne, G. R., & Culnan, M. J. (2002). Using the Content of Online Privacy Notices to Inform Public Policy: A Longitudinal Analysis of the 1998-2001 U.S. Web Surveys. *The Information Society*, 18(5), 345–359.
- Milyaeva, S., & Neyland, D. (2015). *On re-devising markets, re-locating value: Online personal data and “empowering” privacy* (Vol. 4).
- Misuraca, G., Mureddu, F., & Osimo, D. (2014). Policy-making 2.0: Unleashing the power of big data for public governance. In *Open Government* (pp. 171–188). Springer New York.
- Moiso, C., & Minerva, R. (2012). Towards a user-centric personal data ecosystem The role of the bank of individuals’ data. In *2012 16th International Conference on Intelligence in Next Generation Networks* (pp. 202–209). IEEE.
- Mont, M. C., Sharma, V., & Pearson, S. (2012). EnCoRe: dynamic consent, policy enforcement and accountable information sharing within and across organisations. *HP Laboratories technical Report#*: HPL-2012-36.
- Monteleone, S. (2011). *Ambient Intelligence and the Right to Privacy: The Challenge of Detection Technologies*.

European University Institute. Retrieved from
[http://papers.ssrn.com/sol3/papers.cfm?
abstract_id=1953939](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1953939)

- Moore, M. A., Huxford, J., & Hopper, K. M. (2014). Whistleblower as news source: A complex relationship examined through a survey of journalists' attitudes. *Journal of Applied Journalism & Media Studies*, 3(3), 355–374.
- Mori, M., & Minato, T. (1970). The Uncanny Valley. *Energy*, 7(4), 33–35.
- Morin, J., & Glassey, O. (2012). ThinkData: a Data Protection and Transparency Awareness Service based on Storytelling. *Legal Knowledge and Information Systems*, 1–4.
- Morozov, E. (2011). The Internet Intellectual. *New Republic*. Retrieved from
<http://www.newrepublic.com/article/books/magazine/96116/the-internet-intellectual>
- Morozov, E. (2013, March 16). Open and Closed. *The New York Times*. Retrieved from
<http://www.nytimes.com/2013/03/17/opinion/sunday/morozov-open-and-closed.html>
- Morozov, E. (2013, May 3rd) Review of Who Owns the Future, by Jaron Lanier. Retrieved from
https://www.washingtonpost.com/opinions/who-owns-the-future-by-jaron-lanier/2013/05/03/400f8fb0-ab6d-11e2-b6fd-ba6f5f26d70e_print.html
- Mortier, R., Greenhalgh, C., McAuley, D., Spence, A., Madhavapeddy, A., Crowcroft, J., & Hand, S. (2010). The personal container, or your life in bits. *Proceedings of Digital Futures*, 10.
- Moses, L. B. (2007). Recurring Dilemmas: The Law's Race To Keep Up With Technological Change. *Journal of Law, Technology and Policy*, 239–285.
- Moura, P. T. (2014). *The Sticky Case of Sticky Data : An Examination of the Rationale, Legality, and Implementation of a Right to Data Portability Under European Competition Law*. (Thesis) London School of

Economics and Political Science.

- Mun, M., Hao, S., Mishra, N., Shilton, K., Burke, J., Estrin, D., ... & Govindan, R. (2010). Personal data vaults: a locus of control for personal data streams. In *ACM Proceedings of the 6th International Conference CoNEXT* (p. 17).
- Murphy, Dominic, (2015). Concepts of Disease and Health. In Edward N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (Spring 2015 Edition). Retrieved from <http://plato.stanford.edu/archives/spr2015/entries/health-disease/>
- Murphy, R. S. (1995). Property rights in personal information: An economic defense of privacy. *Geo. LJ*, 84, 2381.
- Murray, A. (2007). *The regulation of cyberspace: control in the online environment*. Routledge.
- Murray, A. (2012). *Entering into contracts electronically: the real WWW*. Hart Publishing.
- Murray-Rust, D., Kleek, M. Van, Dragan, L., & Shadbolt, N. (2014). Social Palimpsests – Clouding the Lens of the Personal Panopticon. In K. O’Hara, M. C. Nguyen, & P. Haynes (Eds.), *Digital Enlightenment Yearbook 2014 : Social Networks and Social Machines, Surveillance and Empowerment*.
- Nabarro.com. (2012). New EU data protection laws will hit local authorities. Retrieved from <http://www.nabarro.com/insight/alerts/2012/february/new-eu-data-protection-laws-will-hit-local-authorities/>
- Narayanan, A. (2014). Eternal vigilance is a solvable technology problem: A proposal for streamlined privacy alerts. *Freedom to Tinker*. Retrieved from <https://freedom-to-tinker.com/blog/randomwalker/eternal-vigilance-is-a-solvable-technology-problem-a-proposal-for-streamlined-privacy-alerts/>
- Narayanan, A., Huey, J., & Felten, E. W. (2016). A Precautionary Approach to Big Data Privacy. In *Data Protection on the Move* (pp. 357-385). Springer Netherlands.
- Narayanan, A., & Shmatikov, V. (2008). Robust de-anonymization of large sparse datasets. In *Security and Privacy, 2008. SP 2008. IEEE Symposium on* (pp. 111-

125). IEEE.

- Narayanan, A., Toubiana, V., Barocas, S., Nissenbaum, H., & Boneh, D. (2012). A Critical Look at Decentralized Personal Data Architectures. *arXiv Preprint*. Retrieved from <http://arxiv.org/abs/1202.4503>
- National Telecommunications and Information Administration (NTIA). (2013). *Short Form Notice Code of Conduct to Promote Transparency*.
- Neuendorf, K., Xiong, C., Blake, B., & Hudzinski, K. (2014). Need for Presence, Enjoyment, and Attitude toward Vendor: Predicting Purchase Intent in the Online Shopping Environment. In *Midwest Association for Public Opinion Research, Chicago, IL* (pp. 0–22).
- Neyland, D. (2013). Can Markets Solve Problems? *MISTS Working Paper No. 1*, 8(2008).
- Nissenbaum, H. (2009). *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press.
- Notturmo, M. (2014). *Hayek and Popper: On Rationality, Economism, and Democracy*. Routledge.
- O'Hara, K. (2010). Intimacy 2.0: Privacy rights and privacy responsibilities on the World Wide Web. In *Web Science Conf. 2010, April 26-27, 2010, Raleigh, NC, USA*. Retrieved from <http://eprints.soton.ac.uk/268760/>
- O'Hara, K. (2011). Transparent Government, Not Transparent Citizens: a report on privacy and transparency for the Cabinet Office.
- O'Hara, K. (2012). Can Semantic Web Technology Help Implement a Right to Be Forgotten?. *Computers and Law*, 22(6).
- O'Hara, K. (2012b). Transparency, open data and trust in government: shaping the infosphere. In *Proceedings of the 4th Annual ACM Web Science Conference*.
- O'Neill, O. (2002). *A question of trust: The BBC Reith Lectures*. Cambridge University Press.
- O'Reilly, T. (2011). Government as a Platform. *Innovations*, 6(1), 13–41. Retrieved from http://www.mitpressjournals.org/doi/abs/10.1162/inov_a_0

- OECD (1980). OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (1980).
- OECD. (2011). The Evolving Privacy Landscape : 30 Years After the OECD Privacy Guidelines. *OECD Digital Economy Papers*, (176).
- Office of the Privacy Commissioner of Canada. (2013). *Online Privacy Transparency: Annual Report to Parliament. Report on the Personal Information Protection and Electronic Documents Act*.
- Ogus, A. (1994). *Regulation: Legal Form and Economic Theory*. Oxford: Clarendon Press.
- Ohm, P. (2010). Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA law review*, 57, 1701.
- Open Identity Exchange. (2014). An Open Market Solution for Online Identity Assurance. *White Paper*. Retrieved from <http://openidentityexchange.org/wp-content/uploads/2014/03/oix-white-paper-2010-03-02.pdf>
- Orbach, B. (2012). What Is Regulation? *Yale Journal on Regulation Online*, 30(1), 1–10.
- Osman, F. Y., & Rahim, N. Z. A. (2011). Self-disclosure and Social network sites users' awareness. *IEEE International Conference on Research and Innovation in Information Systems (ICRIIS)*.
- Osuji, O. (2015). Corporate Social Responsibility, Juridification and Globalization: 'Inventive Interventionism' for a 'Paradox'. *International Journal of Law in Context*, 44(0), 1–60.
- Özpolat, K., Gao, G., Jank, W., & Viswanathan, S. (2010). The Value of Online Trust Seals: Evidence from Online Retailing. *SSRN Electronic Journal*, 1–32.
- Pagallo, U., & Bassi, E. (2013). Open Data Protection: Challenges, Perspectives and Tools for the Re-use of PSI. In M. Hildebrandt, K. O'Hara, & M. Waidner (Eds.), *Digital Enlightenment Yearbook 2013: The Value of Personal Data* (pp. 179–189).

- Parker, C. (2002). *The Open Corporation: Effective Self-regulation and Democracy*. Cambridge University Press.
- Parker, C. (2007). Meta-regulation: legal accountability for corporate social responsibility. In D. McBarnet, A. Voiculescu, & T. Campbell (Eds.), *The New Corporate Accountability : Corporate Social Responsibility and the Law* (Vol. 29, pp. 1–49). Cambridge University Press.
- Parker, C., Gordon, T., & Mark, S. (2010). Regulating Law Firm Ethics Management: An Empirical Assessment of an Innovation in Regulation of the Legal Profession in New South Wales. *Journal of Law and Society*, (37), 466–500.
- Parker, C., & Nielsen, V. L. (2011). *Explaining Compliance: Business Responses to Regulation*. Edward Elgar Publishing.
- Pasquale, F. (2010). Beyond Innovation and Competition: The Need for Qualified Transparency in Internet Intermediaries. *Northwestern University Law Review*, 104(I), 68–71.
- Pasquale, F. (2014). *The Black Box Society*. Cambridge, MA: Harvard University Press.
- Patel, R. (2015). Why privacy is not enough: Big Data and predictive analytics. In *Seventh Workshop on the Philosophy of Information, University College London*.
- Paterson, J. (2000). *Behind the Mask: Regulating Health and Safety in Britain's Off-shore Oil and Gas Industry*. Dartmouth.
- Patil, S. (2013). Synthesizing Findings of Privacy Studies using Meta-Analysis. *Workshop on Measuring Networked Privacy, CSCW 2013: The 2013 ACM Conference on Computer Supported Cooperative Work*.
- Payne, D., & Trumbach, C. C. (2009). Data mining: proprietary rights, people and proposals. *Business Ethics: A European Review*, 18(3), 241–252.
- Pearson, S., & Charlesworth, A. (2009). Accountability as a way forward for privacy protection in the cloud. In *Cloud computing* (pp. 131-144). Springer Berlin Heidelberg.
- Pederson, A. (2005). Notification – what is the point? *Privacy*

Laws & Business United Kingdom Newsletter, 20(7).

- Peixoto, T. (2013). The Uncertain Relationship Between Open Data and Accountability: A Response to Yu and Robinson's The New Ambiguity of "Open Government." *UCLA Law Review* 2.
- Peppet, S. R. (2014). Regulating the Internet of Things: First Steps toward Managing Discrimination, Privacy, Security and Consent. *Tex. L. Rev.*, 93, 85.
- Peppet, S. R. (2011). Unraveling Privacy: The Personal Prospectus and the Threat of a Full-Disclosure Future. *Northwestern U. L. Rev.*, 105(3), 1153–1204.
- Perez, R. A., & Moreau, L. (2008). Provenance-based auditing of private data use. Retrieved from <http://eprints.soton.ac.uk/266580/>
- Pettit, P. (1997). *Republicanism: A Theory of Freedom and Government*. Oxford University Press.
- Pfiefle, S. (2014, March). IAPP Heads to Singapore with APIA Template in Tow. *The Privacy Advisor*. Retrieved from <https://iapp.org/news/a/iapp-heads-to-singapore-with-apia-template-in-tow>
- Pinsent Masons. (2014). New ICO guidelines on privacy impact assessments pre-empts EU reforms, says expert. *Out-Law.com*. Retrieved from <http://www.out-law.com/en/articles/2014/february/new-ico-guidelines-on-privacy-impact-assessments-pre-empts-eu-reforms-says-expert/>
- Pires, G. D., Stanton, J., & Rita, P. (2006). The internet, consumer empowerment and marketing strategies. *European Journal of Marketing*, 40(9/10), 936–949.
- Poikola, A., Kuikkaniemi, K., & Honko, H. (2015). *MyData - A Nordic Model for Human-Centered Personal Data Management and Processing*. Retrieved from <http://www.lvm.fi/documents/20181/859937/MyData-nordic-model/2e9b4eb0-68d7-463b-9460-821493449a63?version=1.0>
- Pollach, I. (2011). Online privacy as a corporate social responsibility: an empirical study. *Business Ethics: A European Review*, 20(1), 88–102.

- Pollock, R. (2008). *The economics of public sector information*. University of Cambridge, Faculty of Economics.
- Pollock, R. (2012). Making a real commons: Creative Commons should drop the non-commercial and no derivatives licenses. Retrieved from <http://blog.okfn.org/2012/10/04/making-a-real-commons-creative-commons-should-drop-the-non-commercial-and-no-derivatives-licenses/>
- Pollock, R. (2015). Putting Open at the Heart of the Digital Age. Retrieved from <http://blog.okfn.org/2015/06/05/putting-open-at-the-heart-of-the-digital-age/>
- Popper, K. (1945). *The Open Society and its Enemies*. London: Routledge.
- Popper, K. (1963). *Conjectures and refutations* (7th ed.). London: Routledge.
- Porter, T. (1995). *Trust in Numbers*. Princeton University Press.
- Posner, R. (1979). Utilitarianism, Economics and Legal Theory. *The Journal of Legal Studies*, 103–140.
- Posner, R. (1998). Privacy. In P. Newman (Ed.), *The New Palgrave Dictionary of Economics & The Law* (pp. 103–107).
- Posner, R. A. (1973). *Economic Analysis of Law*. Little Brown and Company.
- Prins, C. (2006). When personal data, behavior and virtual identities become a commodity: Would a property rights approach matter? *SCRIPT-Ed*, 3(4), 270–303.
- Pritchett, S. (2010). Using Privacy Impact Assessments. *Privacy & Data Protection Journal*, 10(6), 7–9.
- Prospect. (2014). EU data protection regulations: will they work, and how do we prepare? *Prospect Magazine*. Retrieved from <http://www.prospectmagazine.co.uk/economics-and-finance/eu-data-protection-regulations-will-they-work-and-how-do-we-prepare>
- Prosser, W. (1960). Privacy (A Legal Analysis). *California Law Review*, 48(3), 338–423.

- Rabinowitz, P. (2015). Street/Crime: From Rodney King's Beating to Michael Brown's Shooting. *Cultural Critique*, 90(1), 143–147.
- Railton, P. (1986). Moral realism. *The Philosophical Review*, 95(2), 163–207. Retrieved from <http://www.jstor.org/stable/2185589>
- Rainer, B., & Stefan, K. (2010). Trained to Accept ? A Field Experiment on Consent Dialogs. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 2403–2406).
- Rauhofer, J. (2008). Privacy is dead, get over it! Information privacy and the dream of a risk-free society. *Information & Communications Technology Law*, 17(3), 185–197.
- Raymond, E. (1999). The cathedral and the bazaar. *Knowledge, Technology & Policy*, 12(3), 23-49.
- Raz, J. (1986). *The Morality of Freedom*. Oxford: Clarendon Press.
- Richards, K. (2000). Framing Environmental Policy Choice. *Duke Environmental Law & Policy Forum*, 10, 221–85.
- Romanosky, S., & Acquisti, A. (2009). Privacy costs and personal data protection: Economic and legal perspectives. *Berkeley Technology Law Journal*, 24(3), 1061-1101.
- Romanosky, S., Acquisti, A., Hong, J., Cranor, L. F., & Friedman, B. (2006). Privacy patterns for online interactions. *Proceedings of the 2006 Conference on Pattern Languages of Programs - PLoP '06*, 1.
- Rosen, L. (2005). *Open source licensing* (Vol. 692). Prentice Hall.
- Rushkoff, D. (2003). *Open source democracy: How online communication is changing offline politics*. London. Retrieved from <http://rushkoff.com/books/open-source-democracy>
- Russel, S., & Norvig, P. (2009). *Artificial Intelligence: A Modern Approach* (3rd ed.). Pearson.
- Salmon, F. (2011). The Uncanny Valley of Advertising. *Wired Magazine*. Retrieved from <http://www.wired.com/2011/04/uncanny-valley-of-ads/>

- Samuelson, P. (2000). Privacy as intellectual property? *Stanford Law Review*, 18, 1–42.
- Sandel, M. (2012). *What money can't buy: the moral limits of markets*. Macmillan.
- Sandel, M. (2013). Market Reasoning as moral Reasoning: Why Economists Should Re-engage with Political Philosophy. *Journal of Economic Perspectives*, 27(4), 121–140.
- Sandvig, C., Hamilton, K., Karahalios, K., Langbort, C., Stevenson, D., Davies, T., & Woolley, S. (2014). *Data and Discrimination: Collected Essays*. (V. Eubanks & S. Barocas, Eds.). Open Technology Institute / New America.
- Satz, D. (2010). *Why Some Things Should Not Be For Sale*. Oxford University Press.
- Saurwein, F. (2011). Regulatory Choice for Alternative Modes of Regulation: How Context Matters. *Law & Policy*, 33(3), 334–366.
- Schestowitz, R. (2015) The Disturbing Rise of Openwashing: Today's Case of Apple and Microsoft. Retrieved from <http://techrights.org/2015/06/12/openwashing-apple-and-microsoft/>
- Schlosser, A. E., White, T. B., & Lloyd, S. M. (2006). Converting Web Site Visitors into Buyers: How Web Site Investment Increases Consumer Trusting Beliefs and Online Purchase Intentions. *Journal of Marketing*, 70(2), 133–148.
- Schneider, J., Passant, A., & Breslin, J. G. (2010). A Content Analysis : How Wikipedia Talk Pages Are Used. In *Web Science Conf. 2010, April 26-27, 2010, Raleigh, NC, USA*. (pp. 1–7).
- Schrems, M. (2014). *Kämpf um deine Daten*. Verlag.
- Schroeder, R. (2014). Big Data and the brave new world of social media research. *Big Data & Society*, 1(2), 1–11.
- Schultz, D. E. (2006). IMC is do or die in new pull marketplace. *Marketing News*, 40(13), 77.
- Schwartz, A. (2009). Looking back at P3P: lessons for the

- future. *Center for Democracy & Technology*, (November). Retrieved from https://www.cdt.org/files/pdfs/P3P_Retro_Final_0.pdf
- Schwartz, P. (1999). Privacy and Democracy in Cyberspace. *Vanderbilt Law Review*, 52, 17.
- Schwartz, P. M. The EU-US Privacy Collision: A Turn to Institutions and Procedures' (2013). *Harvard Law Review*, 126, 1966-1975.
- Schwartz, P. M. (2004). Property, Privacy, and Personal Data. *Harvard Law Review*, 117(7), 2056.
- Schwartz, P. M. (2013). EU Privacy and the Cloud: Consent and Jurisdiction Under the Proposed Regulation. *12 BNA Privacy and Security Law Report* 718.
- Schwartz, P. M., & Solove, D. J. (2014). Reconciling Personal Information in the United States and European Union. *California Law Review*, 102(4).
- Searles, D. (2013). *The intention economy: when customers take charge*. Harvard Business Press.
- Sedgwick, P. H. (1999). *The market economy and Christian ethics*. Cambridge University Press.
- Senden, L. (2005). Soft law, self-regulation and co-regulation in European law: where do they meet? *Electronic Journal of Comparative Law*, 9(January), 1–39.
- Senecal, S., & Nantel, J. (2004). The influence of online product recommendations on consumers' online choices. *Journal of Retailing*, 80(2), 159–169.
- Seneviratne, O., & Kagal, L. (2014a). Enabling privacy through transparency. *2014 Twelfth Annual International Conference on Privacy, Security and Trust*, (1), 121–128.
- Sevignani, S. (2013). The commodification of privacy on the Internet. *Science and Public Policy*, 40(6), 733–739.
- Shadbolt, N. (2013). Midata: Towards a Personal Information Revolution. In M. Hildebrandt, K. O'Hara, & M. Waidner (Eds.), *Digital Enlightenment Yearbook 2013: The Value of Personal Data* (pp. 202–224).
- Shadbolt, N., & O'Hara, K. (2013). Linked Data in Government. *IEEE Internet Computing*, 17(4), 72–77.

- Shapiro, C., & Varian, H. R. (1998). Versioning: the smart way to sell information. *Harvard Business Review*, 107(6).
- Sherratt, Y. (2006). *Continental philosophy of social science: hermeneutics, genealogy, critical theory*. Cambridge University Press.
- Silverman, J. (2015). *Terms of service: social media and the price of constant connection*. Harper Collins.
- Simon, H. A. (1972). Theories of bounded rationality. *Decision and organization*, 1(1), 161-176.
- Sinclair, D. (1997). Self-Regulation Versus Command and Control? Beyond False Dichotomies. *Law & Policy*, 19(4), 529–559.
- Slee, T. (2012). Seeing Like a Geek. *Crooked Timber Blog*. Retrieved from <http://crookedtimber.org/2012/06/25/seeing-like-a-geek/>
- van der Sloot, B. (2014). Do data protection rules protect the individual and should they? An assessment of the proposed General Data Protection Regulation. *International Data Privacy Law*, 4(4), 307.
- Slovic, P. (1995). The construction of preference. *American Psychologist*, 50(5), 364–371.
- Smith, A. (1759). *The Theory of Moral Sentiments*
- Smith, A. (1776). *An Inquiry into the Nature and Causes of the Wealth of Nations*.
- Society for Computers and Law. (2014). ICO's Privacy Impact Assessments Code Published. Retrieved from <http://www.scl.org/site.aspx?i=ne35906>
- Solomon, M., Russell-Bennett, R., & Previte, J. (2012). *Consumer Behaviour*. Pearson Higher Education AU.
- Solove, D. (2002). Conceptualizing Privacy. *California Law Review*, 90(4).
- Solove, D. J. (2012). Introduction: Privacy self-management and the consent dilemma. *Harv. L. Rev.*, 126, 1880.
- Solove, D. J. (2006). A Taxonomy of Privacy University of Pennsylvania. *University of Pennsylvania Law Review*, 477.

- Solow, R. M. (1956). A contribution to the theory of economic growth. *The quarterly journal of economics*, 65-94.
- Sørensen, L., Sørensen, J. K., & Khajuria, S. (2015). Privacy for Sale ? – Analysis of Online User Privacy. *CMI Working Paper*, (8).
- Speck, P. S., & Elliott, M. T. (1997). Predictors of advertising avoidance in print and broadcast media. *Journal of Advertising*, 26(3), 61–76.
- Spender, P. (2002). Book Review: The Open Corporation: Effective Self-Regulation and Democracy. *University of New South Wales Law Journal*, 25(2), 1–6.
- Spiekermann, S. (2011). The RFID PIA - Developed By Industry, Agreed By Regulators. In D. Wright & P. De Hert (Eds.), *Privacy Impact Assessment: Engaging Stakeholders in Protecting Privacy* (pp. 1–22). Springer.
- Spiekermann, S., & Pallas, F. (2006). Technology paternalism—wider implications of ubiquitous computing. *Poiesis & praxis*, 4(1), 6-18.
- Spiekermann, S., Acquisti, A., Böhme, R., & Hui, K. L. (2015). The challenges of personal data markets and privacy. *Electronic Markets*, 25(2), 161-167.
- Stacey, K., & Aglionby, J. (2013, November 7). Top UK Spies Accept Need for More Openness. *Financial Times*.
- Stallman, R. M. (2002). What is free software? In *Free Software, Free Society: Selected Essays of Richard Stallman*. Boston, MA: GNU Press.
- Steinke, G. (2002). Data privacy approaches from US and EU perspectives. *Telematics and Informatics*, 19(2), 193–200.
- Sterling, B. (2014). *The Epic Struggle for the Internet of Things*. Streika.
- Stewart, B. (2012). Privacy Impact Assessment: Optimising the Regulator’s Role. *Law, Governance and Technology*, 6, 437–444.
- Strong, C. (2014). Are brands entering an uncanny valley? Retrieved from <http://colinstrong.net/2014/05/30/are-brands-entering-an-uncanny-valley/>

- Su, X., & Khoshgoftaar, T. M. (2009). A Survey of Collaborative Filtering Techniques. *Advances in Artificial Intelligence, 2009* (Section 3), 1–19.
- Summers, C., Smith, R., & Reczek, R. (2014). Learning about the Self through Advertising: The Effect of Behaviorally-Targeted Advertising on Consumer Self-Perceptions and Behavior. *Advances in Consumer Research, 42*, 693–694.
- Sundar, S. S., & Marathe, S. S. (2010). Personalization versus Customization: The Importance of Agency, Privacy, and Power Usage. *Human Communication Research, 36*(3), 298–322.
- Sundholm, M. (2001). *The open method of co-ordination: the Linux of European integration*. (Thesis) College of Europe. Retrieved from <http://eucenter.wisc.edu/OMC/Papers/Archive/sundholm2001.pdf>
- Sunstein, C., & Thaler, R. (2008). *Nudge: Improving decisions about health, wealth and happiness*. Yale University Press.
- Sweeney, L. (2013). Discrimination in Online Ad Delivery. *Queue, 11*(3), 10.
- Swire, P. (2012). Social Networks, Privacy, and Freedom of Association: Data Protection vs. Data Empowerment. *North Carolina Law Review, 90*.
- Taddicken, M. (2014). The “Privacy Paradox” in the Social Web: The Impact of Privacy Concerns, Individual Characteristics, and the Perceived Social Relevance on Different Forms of Self-Disclosure. *Journal of Computer-Mediated Communication, 19*(2), 248–273.
- Tamang, Suzanne, and Gregory T. Donovan (2015) Introduction: Media and Methods for Opening Education. *Journal of Interactive Technology and Pedagogy 2015: 5*
- Tancock, D., Pearson, S., & Charlesworth, A. (2010a). A Privacy Impact Assessment Tool for Cloud Computing. *2010 IEEE Second International Conference on Cloud Computing Technology and Science, 667–676*.
- Tancock, D., Pearson, S., & Charlesworth, A. (2010b). Analysis of Privacy Impact Assessments within Major jurisdictions. In *Eighth Annual International Conference on Privacy*,

Security and Trust (pp. 118–125).

- Tancock, D., Pearson, S., & Charlesworth, A. (2010). The emergence of privacy impact assessments. Retrieved from <http://www.hpl.hp.com/techreports/2010/HPL-2010-63.pdf>, [May. 21, 2010].
- Tang, Z., Hu, Y., & Smith, M. D. (2008). Gaining trust through online privacy protection: Self-regulation, mandatory standards, or caveat emptor. *Journal of Management Information Systems*, 24(4), 153-173.
- Tapscott, D. (2010). Why Transparency and Privacy Should Go Hand in Hand. *Huffington Post*. Retrieved from http://www.huffingtonpost.com/don-tapscott/why-transparency-and-priv_b_643221.html
- Taylor, D. G., Lewin, J. E., & Strutton, D. (2011). Friends, fans, and followers: Do ads work on social networks? *Journal of Advertising Research*, 51(1), 258–275.
- Tene, O. (2013). Privacy Law's Midlife Crisis: A Critical Assessment of the Second Wave of Global Privacy Laws. *Ohio St. LJ*, (JANUARY 2014).
- The White House. (2014). *Big Data: seizing opportunities, preserving values (Report for the President)*. Retrieved from http://www.whitehouse.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf
- Thiel, S., Hermann, F., Heupel, M., & Bourimi, M. (2013). Unlinkability Support in a Decentralised, Multiple-identity Social Network. In *Proceedings of the Open Identity Summit* (pp. 32–42).
- Tkacz, N. (2012). From open source to open government: a critique of open politics. *Ephemera: Theory and Politics in Organization*, 12(4), 386–405.
- Tsai, J. Y., Egelman, S., Cranor, L., & Acquisti, A. (2010). The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Information Systems Research*, 22(2), 254–268.
- Tsiavos, P., Stefaneas, P., & Karounos, T. (2013). The Transposition of European Union Open Data/Public Sector Information Policies in Greece: A Critical Analysis. *Policy & Internet*, 5(4), 402–417.

- Tuffield, M. M., & Shadbolt, N. (2008). Lifelogging : Privacy and empowerment with memories for life. *Identity in the Information Society*, 1(1), 155–172.
- Turow, J. (2008). *Niche Envy*. MIT Press.
- Turow, J., Hennessy, M., & Draper, N. (2015). *The Tradeoff Fallacy: How Marketers are Misrepresenting American Consumers and Opening them up to Exploitation*.
- Turow, J., & McGuigan, L. (2014). Retailing and Social Discrimination: The New Normal? In *Data and Discrimination: Selected Essays*.
- Tzanou, M. (2013). Data protection as a fundamental right next to privacy ? 'Reconstructing' a not so new right. *International Data Privacy Law*, 1–12.
- U.S. White House Office of Management and Budget. (1983). *The President's Annual Report on the Agencies' Implementation of the Privacy Act of 1974* (118 (Dec. 4, 1985)).
- Ueckert, F., Goerz, M., Ataian, M., Tessmann, S., & Prokosch, H.-U. (2003). Empowerment of patients and communication with health care professionals through an electronic health record. *International Journal of Medical Informatics*, 70(2), 99–108.
- UK Cabinet Office. (2012). Midata: 2012 Review and Consultation. Retrieved from <https://www.gov.uk/government/consultations/midata-2012-review-and-consultation>
- UK Cabinet Office. (2015). *Open Policy Making Toolkit*. Retrieved from <https://www.gov.uk/open-policy-making-toolkit>
- UK Information Commissioner's Office. (2007). *Data Protection Technical Guidance: Determining what is personal data*. Retrieved from <https://ico.org.uk/media/for-organisations/documents/1554/determining-what-is-personal-data.pdf>
- Ullrich, P., & Wollinger, G. R. (2011). A surveillance studies perspective on protest policing: the case of video surveillance of demonstrations in Germany., *Interface*,

3(1), 12–38.

- United States Government Accountability Office. (2013). *Information Resellers*. (September).
- Ur, B., Leon, P. G., Cranor, L. F., Shay, R., & Wang, Y. (2012). Smart, Useful, Scary, Creepy: Perceptions of Online Behavioral Advertising. *CMU-CyLab Working Paper*.
- Valvi, A. C., & West, D. C. (2013). E-loyalty is not all about trust, price also matters: Extending expectation–confirmation theory in bookselling websites. *Journal of Electronic Commerce Research*, 14(1), 99–123.
- van der Woert, Nicolai, Robert Schuwer, and Martijn Ouwehand. (2015). Connecting various forms of openness: seeking a stronger value proposition. in *2015 Open and Online Education Trend Report*
- Van Dijck, J., & Nieborg, D. (2009). Wikinomics and its discontents: a critical analysis of Web 2.0 business manifestos. *New Media & Society*, 11(5), 855–874.
- Van Dijk, N. (2009). Property, privacy and personhood in a world of ambient intelligence. *Ethics and Information Technology*, 12(1), 57–69.
- Varian, H. R. (1997). Economic aspects of personal privacy. In *Privacy and Self-regulation in the Information Age*. Washington, DC: US Department of Commerce.
- Vaughan, D. (1996). *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. University of Chicago Press, Chicago, IL.
- Viktor Mayer-Schonberger, K. C. (2013). *Big Data: A Revolution that Will Change How We Live, Work and Think*. Houghton Mifflin Harcourt.
- Vimercati, G., Paraboschi, S., & Pedrini, E. (2009). Primelife policy language. Retrieved from http://spdp.di.unimi.it/papers/w3c_wsacas_2009_02.pdf
- Von Neumann, J. (1953). A certain zero-sum two-person game equivalent to the optimal assignment problem. In *Contributions to the Theory of Games 2* (pp. 5–12).
- Wadhwa, K., & Wright, D. (2013). Introducing a privacy impact assessment policy in the EU Member States.

International Data Privacy Law, 3(1), 13–28.

- Walker, K. (2000). Where Everyone Knows Your Name: A Pragmatic Look at the Costs of Privacy and the Benefits of Information Exchange. *Stanford Technology Law Review*, 4(4).
- Walzer, M. (1983). *Spheres of justice: A defense of pluralism and equality*. Basic Books.
- Warren, A., Bayley, R., Bennett, C., Charlesworth, A., Clarke, R., & Oppenheim, C. (2008). Privacy Impact Assessments: International experience as a basis for UK Guidance. *Computer Law & Security Review*, 24(3), 233-242.
- Warren, A., & Charlesworth, A. (2012). Privacy impact assessment in the UK. In *Privacy Impact Assessment* (pp. 205-224). Springer Netherlands.
- Warren, S., & Brandeis, L. (1890). The right to privacy. *Harvard Law Review*, 4(5), 193–220.
- Wasinger, R., Wallbank, J., Pizzato, L., Kay, J., Kummerfeld, B., Böhmer, M., & Krüger, A. (2013). Scrutable user models and personalised item recommendation in mobile lifestyle applications. In *User Modeling, Adaptation, and Personalization* (pp. 77-88). Springer Berlin Heidelberg.
- Watson, S. M. (2014). The Uncanny Valley of Personalization. *The Atlantic*. Retrieved from <http://m.theatlantic.com/technology/archive/2014/06/data-doppelgangers-and-the-uncanny-valley-of-personalization/372780/>
- Weber, M., & Winckelmann, J. (1964). *Soziologie: Welgeschichtliche Analysen*.
- Webster, F. (2014). *Theories of the information society. Radical thought in Italy: A potential politics*. Routledge.
- WEF. (2011). Personal Data : The Emergence of a New Asset Class.
- WEF. (2012). Rethinking Personal Data : Strengthening Trust.
- WEF. (2013). Unlocking the Value of Personal Data : From Collection to Usage.
- Wellcome Trust. (2013). Summary Report of Qualitative Research into Public Attitudes to Personal Data and

- Linking Personal Data. Retrieved from
http://www.wellcome.ac.uk/stellent/groups/corporatesite/@msh_grants/documents/web_document/wtp053205.pdf
- Seltzer, W. (2014). *Network Security as a Public Good*. Retrieved from
<https://www.w3.org/2014/strint/papers/60.pdf>
- Westin, A. F. (1967). *Privacy and Freedom*. Book (Vol. 97, p. xvi, 487 p.). Atheneum.
- White, C. (2014). Patients still in the dark about care.data, warn doctors' leaders. *OnMedica*. Retrieved from
<http://www.onmedica.com/newsarticle.aspx?id=91f4222a-dc30-4487-96c4-6f59797b822f>
- Williams, B., & Quinton, A. (1973). *Utilitarianism, For and Against*. Cambridge University Press.
- Wilson, D. W., Hall, M., Az, T., & Valacich, J. S. (2012). Unpacking the Privacy Paradox : Irrational Decision-Making within the Privacy Calculus. In *Thirty Third International Conference on Information Systems, Orlando 2012* (pp. 1–11).
- Winn, J. K. (2009). Technical Standards as Data Protection Regulation. In *Reinventing Data Protection?* (pp. 191–206).
- Wirtz, J., & Lwin., M. O. (2009). Regulatory focus theory, trust, and privacy concern. *Journal of Service Research*, 12(2).
- Wittgenstein, L. (1968). *Philosophical Investigations § 65 (1958)*. (G. E. M. Anscombe, Ed.).
- Worthy, B. (2014). Making Transparency Stick: The Complex Dynamics of Open Data. *Available at SSRN 2497659*.
- Wright, D. (2011). Should privacy impact assessments be mandatory? *Communications of the ACM*, 54(8), 1–17.
- Wright, D. (2013). Making Privacy Impact Assessment More Effective. *The Information Society*, 29(5), 307–315.
- Wright, D. (2014). How Good are PIA Reports – And Where are They? *European Business Law Review*, 25(3), 407–426.

- Wright, D., & De Hert, P. (2012). Part I: Setting the Scene. In *Privacy Impact Assessment*. Springer Science & Business Media.
- Wright, D., & Hert, P. De. (2008). *Privacy Impact Assessment*. Springer Science & Business Media.
- Wright, D., & Hert, P. De. (2012). Introduction to Privacy Impact Assessment. *Law, Governance and Technology*, 6, 3–32.
- Wright, D., Kroener, I., Lagazio, M., Finn, R., Gellert, R. B., Gutwirth, S., & Vermeulen, M. (2014) Deliverable 5.3 Final Report: Findings and Recommendations. *SAPIENT Project*. Retrieved from <http://www.sapientproject.eu/D5.2%20-%20Final%20report.pdf>
- Wright, D., & Wadhwa, K. (2012, April). A step-by-step guide to privacy impact assessment. In *Second PIAF Workshop, Sopot, Poland* (Vol. 24).
- Wright, D., Wadhwa, K., Lagazio, M., Raab, C. D., & Charikane, E. (2012). *Managing Risk with Privacy Impact Assessment* (pp. 1–20).
- Wright, Glover, Prakash, P., Abraham, S., & Shah, N. (2009). *Open Government Data Study: India. 2011*.
- Xu, H., Luo, X. (Robert), Carroll, J. M., & Rosson, M. B. (2011). The personalization privacy paradox: An exploratory study of decision making process for location-aware marketing. *Decision Support Systems*, 51(1), 42–52.
- Yan, J., Liu, N., Wang, G., Zhang, W., Jiang, Y., & Chen, Z. (2009). How much can behavioral targeting help online advertising? *Proceedings of the 18th International Conference on World Wide Web - WWW '09*, 261.
- Zanfir, G. (2014). Forgetting About Consent. Why The Focus Should Be On “Suitable Safeguards” in Data Protection Law. In *Reloading Data Protection* (pp. 237–257). Springer Netherlands.
- Zarsky, T. (2014). Understanding Discrimination in the Scored Society. *Washington Law Review*, 4, 1375–1412.

Cloud Storage Networks. In *15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid), 2015*.

Zerk, J. A. (2006). *Multinationals and Corporate Social Responsibility: Limitations and Opportunities in International Law*. Cambridge: Cambridge University Press.

Ziewitz, M., & Brown, I. (2013). A Prehistory of Internet Governance. In *Research Handbook on Governance of the Internet* (pp. 1–36). Edward Elgar Publishing.

Zittrain, J. (2008). *The future of the internet – and how to stop it*. Yale University Press.

Zuckerberg, M. (2010, May 24). From Facebook, answering privacy concerns with new settings. *The Washington Post*.