CIGI

CHATHAM HOUSE
The Royal Institute of International Affairs

Global Commission on Internet Governance

ourinternet.org

PAPER SERIES: NO. 26 — MARCH 2016

A Pragmatic Approach to the Right to Be Forgotten

Kieron O'Hara, Nigel Shadbolt and Wendy Hall

A PRAGMATIC APPROACH TO THE RIGHT TO BE FORGOTTEN

Kieron O'Hara, Nigel Shadbolt and Wendy Hall

CHATHAM HOUSE
The Royal Institute of International Affairs

Copyright © 2016 by Kieron O'Hara, Sir Nigel Shadbolt and Dame Wendy Hall.

Published by the Centre for International Governance Innovation and Chatham House.

This work was carried out with the aid of a grant from the International Development Research Centre (IDRC), Ottawa, Canada.

The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the Centre for International Governance Innovation or its Board of Directors.

The views expressed herein do not necessarily represent those of IDRC or its Board of Governors.

This work is licensed under a Creative Commons Attribution — Non-commercial — No Derivatives License. To view this licence, visit (www.creativecommons.org/licenses/by-nc-nd/3.0/). For re-use or distribution, please include this copyright notice.

Centre for International Governance Innovation, CIGI and the CIGI globe are registered trademarks.

67 Erb Street West
Waterloo, Ontario N2L 6C2
Canada
tel +1 519 885 2444 fax +1 519 885 5450
www.cigionline.org

10 St James's Square
London, England SW1Y 4LE
United Kingdom
tel +44 (0)20 7957 5700 fax +44 (0)20 7957 5710
www.chathamhouse.org

TABLE OF CONTENTS

ABOUT THE GLOBAL COMMISSION ON INTERNET GOVERNANCE

The Global Commission on Internet Governance was established in January 2014 to articulate and advance a strategic vision for the future of Internet governance. The two-year project conducts and supports independent research on Internet-related dimensions of global public policy, culminating in an official commission report that will articulate concrete policy recommendations for the future of Internet governance. These recommendations will address concerns about the stability, interoperability, security and resilience of the Internet ecosystem.

Launched by two independent global think tanks, the Centre for International Governance Innovation (CIGI) and Chatham House, the Global Commission on Internet Governance will help educate the wider public on the most effective ways to promote Internet access, while simultaneously championing the principles of freedom of expression and the free flow of ideas over the Internet.

The Global Commission on Internet Governance will focus on four key themes:

- enhancing governance legitimacy — including regulatory approaches and standards;
- stimulating economic innovation and growth — including critical Internet resources, infrastructure and competition policy;
- ensuring human rights online — including establishing the principle of technological neutrality for human rights, privacy and free expression; and
- avoiding systemic risk — including establishing norms regarding state conduct, cybercrime cooperation and non-proliferation, confidence-building measures and disarmament issues.

The goal of the Global Commission on Internet Governance is two-fold. First, it will encourage globally inclusive public discussions on the future of Internet governance. Second, through its comprehensive policy-oriented report, and the subsequent promotion of this final report, the Global Commission on Internet Governance will communicate its findings with senior stakeholders at key Internet governance events.

www.ourinternet.org

ABOUT THE AUTHORS

Kieron O'Hara is a senior lecturer and principal research fellow in electronics and computer science at the University of Southampton. His interests are in the philosophy, sociology and politics of technology, particularly the World Wide Web; key themes are trust, privacy, transparency and openness. He is the author of several books on technology and politics, the latest of which, *The Devil's Long Tail* (Oxford University Press, 2015), looks at online extremism. He has also written extensively on political philosophy and British politics. He is one of the leads on the UK Anonymisation Network, which disseminates best practices in data anonymization, and writes the Digital Citizen column for *IEEE Internet Computing*.

Sir Nigel Shadbolt is professor of computer science at the University of Oxford and Principal of Jesus College. He is also the chairman and co-founder of the Open Data Institute. Since 2009, Sir Nigel has acted as an information adviser to the UK government, helping transform public access to government information, including the widely acclaimed data.gov.uk site. With more than 500 publications, he researches and publishes on computer science, artificial intelligence, open data and web science. He has also worked in philosophy, psychology and linguistics. Since 2000, he has secured 17 projects as principal investigator with a value of more than £20 million. He is currently principal investigator on a £6.14-million EPSRC-funded program grant researching the theory of social machines — Web-scale problem-solving systems comprising large numbers of humans and computers. In 2013, he was awarded a knighthood for services to science and engineering.

Dame Wendy Hall, DBE, FRS, FREng, is professor of computer science at the University of Southampton. She was dean of the Faculty of Physical Science and Engineering from 2010 to 2014 and head of the School of Electronics and Computer Science from 2002 to 2007. The influence of her work has been significant in many areas, including digital libraries, the development of the Semantic Web, and the emerging research discipline of Web Science. She is now executive director of the Web Science Institute at Southampton. She was president of the Association for Computing Machinery from 2008 to 2010, a member of the Prime Minister's Council for Science and Technology from 2004 to 2010 and a founding member of the Scientific Council of the European Research Council. She is currently a member of the Global Commission on Internet Governance and the World Economic Forum's Global Agenda Council on Artificial Intelligence and Robotics.

ACRONYMS

AEPD	Agencia Española de Protección de Datos
CJEU	Court of Justice of the European Union
DPAs	data protection authorities
DPD	Data Protection Directive
EEA	European Economic Area
EFTA	European Free Trade Association
ICO	Information Commissioner's Office
PDMA	Personal Data Management Architecture
PDSs	Personal Data Stores
PIMS	Personal Information Management Services
URL	uniform resource locator

EXECUTIVE SUMMARY

This paper considers the shape that a "right to be forgotten" is taking in the online world, in the aftermath of the so-called Google Spain decision, in which the Court of Justice of the European Union (CJEU) found (against Google) that European data subjects had the right to request that search engines de-index webpages that feature in searches on their names.

The right to be forgotten is a contested concept. This paper considers various pre-Internet understandings of a right to be forgotten, arguing that, although it is linked to ideas about human and social memory, justice and forgiveness, and to social developments such as information management and the rise of bureaucracy, any rights in that area have been limited and partial, such as the treatment of "spent" convictions in the United Kingdom's Rehabilitation of Offenders Act. Only since the rise of data protection laws has a right to be forgotten become more general and feasible. The Google Spain decision is a plausible interpretation of the European Union's 1995 Data Protection Directive (DPD). The paper explains how the court reasoned, and what measures have been put in place by Google in response.

The judgment, and Google's response, raises a series of questions that are addressed in this paper. In particular, the judgment affects the nature of the balance between free speech and privacy on the Internet. Google's presentation of its search as a neutral reflection of the state of the Web (and for that reason, a valuable resource for Web users) was found wanting by the court, and indeed Google itself has often adjusted its PageRank algorithm to improve its output by excluding, for example, spam, link farms and child pornography. Such methods cannot be transparent, since they would then be gamed by the spammers, and so Google has to present as a corporate "black box." Yet it is a big step to devolve issues of privacy and freedom to an opaque process — even if it is accepted that a private sector actor can legitimately make decisions in this area.

The final section of the paper considers whether individuals might manage their personal data with flexible architectures that could act as points of contact for those wishing to use the data. Many of the issues discussed earlier could, in such a technological ecosystem, be addressed within a system that respected the autonomy of the data subject in providing limited abilities to control self-presentation. However, this remains a thought experiment at this stage — such technologies, though technologically feasible, are not yet the subject of great demand or takeup from consumers, while the state of current regulation means that business models favour sidelining data subjects from decisions made about the use of their data.

INTRODUCTION

In May 2014, the world of privacy regulation, data handling and the World Wide Web changed dramatically as a result of judgment C-131/12 in the CJEU.[1] The so-called Google Spain decision confirmed that EU data protection legislation gives data subjects the right to request search engines to de-index webpages that appear in the search results on their names. The search engine is not obliged to agree to such requests — certain conditions have to be met and tests applied — but it is not free simply to ignore them. The decision drew on the 1995 DPD[2] and the Charter of Fundamental Rights of the European Union,[3] and is consistent with a general direction toward more aggressive protection of privacy rights in Europe, as evidenced by the annulment of the Data Retention Directive, also in 2014 (CJEU 2014). Nevertheless, despite these antecedents, it has been seen as a major step in establishing a right to be forgotten.

The right to be forgotten is primarily a legal concept, therefore much of the discussion in this paper will be to do with the law. This is not a legal opinion, however, and the authors are not lawyers. The right to be forgotten covers moral and political issues, and raises technical and institutional problems. Our issue as engineers of the Web is not only how we respond to the politico-legal debate, but also how to influence it by theorizing about the art of the possible. Any "solution" to the conundrums of privacy, deletion and free expression that, for example, balkanizes the Internet, will arguably produce worse effects than the problems it attempts to solve. This paper is set, broadly, in the current context of data protection. It will not speculate on how the proposed revisions to the EU data protection law will affect the position (Zanfir 2014), nor does it demand particular changes to or interpretations of the law. It will, however, consider the possibility of a technological

1 See http://curia.europa.eu/juris/liste.jsf?num=C-131/12.

2 See http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:31995L0046.

3 See www.europarl.europa.eu/charter/pdf/text_en.pdf.

contribution to what is currently being fixed by a relatively controversial process.

The paper consists of four substantial sections between this introduction and a conclusion. The first considers the nature of the right to be forgotten, and what it could mean, closing with the debate that developed around it as the European Union began to consider revising the DPD. The next section will look at C-131/12, the decision of the CJEU about an appeal made by Google Spain against a judgment of the Spanish data protection authority, the Agencia Española de Protección de Datos (AEPD). This is the most visible assertion of data rights in the European Union in this area. The third section will consider a few of the many issues that this contentious judgment has raised. Fourthly, given this judgment and the controversy it has provoked, a discussion will be presented of the potential of one particular technology to deliver (some of) the aspirations of the right to be forgotten, and a framework of norms in which that potential would be maximized.

THE RIGHT TO BE FORGOTTEN, BEFORE GOOGLE SPAIN

Traditionally, the right to be forgotten has not been understood as a natural right; we have no offline analogue. It does not appear, for example, in the *Declaration of the Rights of Man and of the Citizen* (1789). When, in one's medieval village, one committed a faux pas, the upshot of centuries of folk wisdom was that one would have to live with the consequences. In the splendid story from *One Thousand and One Nights* called "The Historic Fart," Abu Hassan flees from his wedding in shame after emitting "a thunderous fart which echoed from wall to wall and silenced every voice in the room." He travels in the East for 10 years, homesick but too embarrassed to return. When he finally plucks up the courage to go back, hoping that everyone has forgotten, he discovers that far from having been consigned to obscurity, his solecism has become a temporal standard. A child asks his mother when he was born; she replies that he must be 10, because he was born in the year Abu Hassan farted. "And with these words, hope died in his heart forever. He fled the land and was never seen again."

He might well have wished for a mechanism to suppress memories of his embarrassment, but the humour of the story revolves around the ways in which collective memory sometimes seizes upon apparently inconsequential events, over whose interpretation and (accurate or inaccurate) recollection their protagonists have no control.

Psychological Forgetting

Forgetting, of course, *takes place*, and has its uses (Schacter 2001). One might put misdeeds behind one, or live them down. The passage of time helps, as does the creation of a worthier identity. One could even imagine the science fiction experiment of "editing" experience to remove unpleasant memories, as in the film *Eternal Sunshine of the Spotless Mind*. But this doesn't help us understand a right to be forgotten, for two key reasons. First, the locus of forgetting is the rememberer; the right to be forgotten, by contrast, is a right *to be* forgotten, not a right to forget. If Z commits a faux pas in front of X and Y, X may forget, but Y may not (and then may remind X); Z's forgetting the event is neither here nor there. Not only is the forgetting of Z's faux pas a random event, but it is very unlikely to happen simultaneously over all rememberers; the collective memory, taken as the union of the memories of its members, is quite robust against forgetting.

Second, forgetting in this psychological sense is morally neutral. It may be that one's good deeds help one's youthful indiscretions be forgotten by a society, and that shows a mature society. However, it is just as likely that the memory of the indiscretion will hinder the creation of a positive reputation, or that a later bad deed will eclipse the collective memory of all one's past good deeds, as Shakespeare laments in Sonnet 25: "The painful warrior famousèd for worth/ After a thousand victories once foiled/ Is from the book of honour razèd quite." So even if society has mechanisms for forgetting, they will not always serve the purposes of the individual or of society. Clearly, this fact about psychological forgetting distinguishes it from the right to be forgotten, which all agree is morally charged (whether positively or negatively).

Thus, the disanalogy between the right to be forgotten (collective forgetting) and psychological forgetting must be kept in mind. In psychology, the individual forgets; in the digital world, the individual is, or hopes to be, forgotten. In the former case, the individual's memory is wiped, while in the latter, the individual hopes to wipe the "memory" of others. Nevertheless, the mechanisms of psychological forgetting (or the failures of the mechanisms of memory) are still relevant.

From an information processing view,[4] there are three basic operations that make up memory:

- registration (the transformation of input into a form in which it can be stored);
- storage (the holding of information in memory); and
- retrieval (extracting stored information).

Forgetting could be seen as another basic operation (that of clearing up used and out-of-date material), but it is more usually conceptualized in this framework as a failure in one of the three other operations. In an Internet-based analogue, failure of registration is not the issue — the assumption of

4 For a review, see, for example, Gross and McIlveen (1999).

the current debate over the right to be forgotten is that the information is stored somewhere online, and the issue is access to it.

Hence, the psychology of forgetting reminds us that the two relevant concepts are failures of *availability* (i.e., the information is no longer stored) and failures of *accessibility* (i.e., it is stored but cannot be retrieved). These map onto the ideas of deleting information from the Web and removing (some) links to it, making it harder to find, and correspond to, respectively, (a right to) erasure and (a right to) de-indexing or de-linking. Removal of all links is effectively indistinguishable from erasure, while removal of some links reduces the likelihood of retrieval. Clearly, the fewer links removed, the less the likelihood of retrieval is reduced.

Justice, Forgiveness and Bureaucracy

A related concept to forgetting is forgiveness (Margalit 2002). Forgiveness goes beyond forgetting; it requires remembering, while ceasing to judge harshly. Paul Ricoeur (2006, 19) argues that forgiveness is not intended "to extinguish memory: on the contrary, the goal it has of cancelling the debt is incompatible with that of cancelling memory." Horrendous deeds should not be forgotten, but we conduct our affairs in such a way that there is a route for their perpetrators to become useful members of society. Forgiveness, whatever its moral overtones, implies a learning process such that the original crime will not be committed again.

It has traditionally been hard to institutionalize forgiveness; it often seems to rely on individual case-by-case judgment that resists translation into systems. The urge to forgive can manifest itself against the background of a rigid, impartial system; the social justice of a system that is "blind" can throw up examples of individual injustice. Bureaucracies emerge to handle complexity, records are kept and the past becomes harder to shake off. The plot of Charles Dickens's *Bleak House*, for example, revolves around the mysterious past of Lady Dedlock, the truth of which is painstakingly revealed from legal documents hitherto lost or concealed, with tragic consequences.

It may be that an individual can reinvent himself or herself — in American terms, by "going West" to new territory where the memory of the original wrongdoing is less vivid. Improved communications and transport links mean that one is not confined to particular locations. In Victor Hugo's *Les Misérables*, Jean Valjean shakes off his convict past through travel to new places. It is no coincidence that the novels just cited are of the mid-nineteenth century, when urbanization, globalization and the professionalization of bureaucracy were beginning to have important effects on the lives of ordinary people. Collective memory became decoupled from particular locations and geographical communities, and its content and durability far less contingent.

Power and social status are also important in determining which features of one's past or reputation will be acted upon in the present or future. Both Valjean and Lady Dedlock are in positions of power, but are undone by impersonal and unstoppable forces of the law that are devoid of compassion. In satires such as *Moll Flanders* and *Vanity Fair*, perceptions of the flighty pasts of young ladies are subtly altered by marriages, social position and wealth.

Forgiveness suggests that the debt of the past misdeed has indeed been paid, and that the perpetrator needs to move on, "to find faith in the everyday again and mastery over their time" (Augé 2004, 88). This is part of the justification for a right to be forgotten. There are many examples of permanent records that affect the individual's social standing after taking a punishment or suffering online humiliation.[5] In the United Kingdom, for instance, a 14-year-old boy found himself on the national news because he had "sexted" a naked image of himself to a girl who had shared it with others (BBC 2015). His action was logged on a police database as an instance of the crime of making and sharing indecent images of a child (i.e., himself), with potentially disproportionate consequences for him in later life (for example, if he attempts to work with children).

The injustice to the boy was illuminated against the rule-based machinery in which he was caught, rules drafted by politicians concerned with the specific problem of online pedophilia and necessarily insensitive to the details of an everyday situation — ultimately, the same problem faced by Jean Valjean. This illustrates a paradox inherent in the right to be forgotten. If machinery for institutional forgetting is in place, it will be just as insensitive to the individual situation as the machinery for institutional remembering. In such a case, the subject acts upon their own initiative to show that the past information is outdated according to some definition, but without having to make the case to wider society that they have also moved on in the sense of being a different, better or more socially attuned person. Forgiveness morphs back into forgetting, as the focus of the system is on the information, not the person. The right to be forgotten would be a means of an individual's regaining his "faith in the everyday," but it would be his choice to pursue. Offline, forgiveness is a decision of others; a right to be forgotten — like all rights — is a matter for the individual. In a world of mass data collection, forgiveness may simply not scale. To facilitate individuals' moving on, the power to decouple information from its social effects may have to be devolved to individuals (through a right to be forgotten, or other powers of deletion), not to wider society.

5 See Mayer-Schönberger (2009) for several examples.

Forgetting and the Law

In more recent years, targeted forms of institutional forgetting, explicitly associated with a forgiving or a debt-paying process, have been enshrined in legal practice for more or less utilitarian reasons. The rehabilitation of offenders has often been facilitated by reducing access to information about convictions once the sentence has been served. The UK Rehabilitation of Offenders Act (1974) allows offenders to withhold evidence of "spent" convictions in certain contexts, such as applying for a job or conducting civil proceedings; a conviction is considered spent after a specified period of time (which depends on the severity of the original sentence) has elapsed since the sentence was served, as long as the offender has not since reoffended. It is, however, a very weak protection. In Germany, criminals' names can be withheld from news reports once the sentence is served, which led to a high-profile case when two convicted murderers sued Wikipedia for naming them in its account of the crime. The German courts have developed a number of criteria for balancing the interests of offenders in protecting their personality rights and ability to reintegrate into society, and the interests of publishers, historians and journalists in writing publicly about such events (Siry and Schmitz 2012). In the criminal justice setting, the UK Law Commission proposed a requirement that the media take down material that might prejudice a fair trial if a juror were to find it (Law Commission 2013), but the government declined to implement the proposal in full, recognizing the "disquiet" the proposal had generated (Oswald 2014).

Such forgetting is seen as benefiting both the individual and society via the individual's rehabilitation and reintegration. Amitai Etzioni (1999) has argued against this, that disclosure of convictions —for example, of sex offences — is a justifiable invasion of offenders' privacy, given the dangers to communities from their presence within. In the UK Rehabilitation of Offenders Act, a crime that received a sentence of four years or more can never be spent, presumably on the grounds that information about a serious offence must remain in the public domain for reasons of public safety. Similarly, certain classes of responsible people, ranging from those working with children, to those involved in the humane destruction of animals, to financial managers, to (somewhat bizarrely) butlers, must disclose all convictions when applying for jobs, even if the convictions are spent.

Such laws are part of the tapestry of legislation, regulations and rights that might fall under the rubric of a right to be forgotten grounded in the general right to privacy, in the context of the public exposure of an individual's personal life (Ambrose and Ausloos 2013). However, despite the term *droit à l'oubli* that is sometimes applied to them, they cannot collectively be seen as constituting a general right to be forgotten, if only because of their narrow coverage, focusing on convictions for criminal behaviour, and limited to specific contexts such as employment issues. The impetus for the development of a right to be forgotten has come, in recent years, rather more strongly from a different route, via data protection, which is concerned with managing the effect on individuals of information about them that is or has been publicly available.

The Debate over Data Protection Reform

The adoption of the Charter of Fundamental Rights of the European Union in 2009 made clear, for the first time, the status of data protection within the European Union. The European Convention of Human Rights, ratified in 1953, has traditionally provided the European human rights framework, and contains a right to a private life, but no specific mention of data protection. The DPD of 1995 provides for data protection, of course, but in the context of ensuring the free flow of information across borders in the single European market, rather than defending or demarcating particular rights. The charter is the first document to include data protection as a human right.

The debate over the right to be forgotten was transformed in the early part of this decade by a series of muscular speeches by European Commissioner Viviane Reding (2010), in the context of moves to revise the now antiquated DPD. Her speeches, floating the right to be forgotten as a key part of Europe's data protection regime, caused an immense amount of comment. Initial debate focused on how far-reaching the proposal might be — would it mean, for example, a right to erase? Could one get unauthorized (or even authorized) photographs of oneself taken down from others' social media sites? Would it ensnare private citizens in a bureaucratic net? Or, alternatively, did it refer to better enforcement of the very much more minor rights that are enshrined already in the DPD — for example, rights to have data deleted if it is held for longer than it should be, or to object to unauthorized use? Reding (2012)claimed that a right to be forgotten would clarify and strengthen existing rights.

The distinction between memory failures of availability versus failures of accessibility is replicated on the Internet. One paper made the distinction among the following:

- a right to erasure after due process and time;
- a right to a "clean slate" (i.e., regulating the use of data so that it is not used against you after a sufficient period has elapsed); and
- a right to free expression without the danger that your utterances or behaviour will be used against you in future.

The first is a reduction of availability, while the second and third are reductions of accessibility (Koops 2011). Most commentators argued, or assumed, that a right to be forgotten, if it was to extend beyond the current data

protection right to erase false content, must be tantamount to a right to erasure (Bernal 2011; Markou 2014). Meanwhile, web scientists estimated how technically feasible some of the more draconian interpretations might be, usually with negative results (O'Hara 2012).

The lack of a defined context produced something of a vacuum that was filled with commentary (some thought that the use of the term "right to be forgotten" was inflammatory and probably going to be misleading [Markou 2014]). Jeffrey Rosen (2012) called this a "proposal to create a sweeping new privacy right," which "represents the biggest threat to free speech on the Internet in the coming decade." A leading Google lawyer called the right to be forgotten "foggy thinking" (Fleischer 2011). Meanwhile, many scholars argued that some kind of right to be forgotten was already implicit in the network of data protection jurisprudence (Zanfir 2014), although there was little guidance to date about how a data controller might strike the balance between the right to be forgotten and exceptions where that right could be overridden (Ambrose and Ausloos 2013), and some in Europe argued that these rights, if they existed, were limited in scope and no big deal anyway (Ausloos 2012). Mayer-Schönberger argued that all data should have an expiration date, so that forgetting became a default — although it was hard to see how that suggestion would help with issues such as the greater powers of the search engines and social networks (not to mention governments) to set the terms of data collection, and so his idea probably serves the purpose of (first-person) forgetting, more so than the desire *to be forgotten* (Mayer-Schönberger 2009).

The root of this dispute was the philosophical divergence between the United States and the European Union on privacy. In the former, it is taken to facilitate liberty, while in the latter it supports dignity, and conceptions differ according to how privacy should interact with other norms and institutions to produce different desired effects (Post 2001; Whitman 2004). Furthermore, the US First Amendment is one of the most complete protections of free speech, and is prioritized over many other rights. For instance, the right to free speech was recently taken as the basis for calling some restrictions on political campaign finance unconstitutional, for example, in the cases of *Citizens United v Federal Election Commission* 2010 and *McCutcheon et al v Federal Election Commission* 2014 (Mutch 2014). There would seem little doubt that a right to be forgotten, however it was enacted, would fall foul of First Amendment rights — hence Rosen's response.

THE GOOGLE SPAIN DECISION

The Google Spain decision C131-12 (European Commission 2014) was based on a case brought by Google Spain against the AEPD. The AEPD had, from 2007 on, pursued a couple of hundred similar cases in which individuals protested that data about them online, although true, was excessive or outdated (Daley 2011). These are cardinal sins in the data protection world — the DPD specifically requires that data should be "adequate, relevant and not excessive in relation to the purposes for which they are processed;… such purposes must be explicit and legitimate and must be determined at the time of collection of the data; [and] the purposes of processing further to collection shall not be incompatible with the purposes as they were originally specified" (Recital 28).[6] Nevertheless, this was something of a lone crusade for the AEPD, which was not generally supported or copied by other data protection authorities (DPAs).[7]

The cases the AEPD took on often resulted from digitization, of newspaper archives or public gazettes, for example. Minor but embarrassing judgments (a conviction for urinating in a public street, for example) became prominent for certain citizens via Google searches. Sometimes the newspaper archive did not tell the full story. A charge or a conviction would be reported, but the acquittal or the successful appeal would not, so the archive, although it told the truth, could not be said to have told the whole truth, and taken *in toto* could be seriously misleading.

The problems are sometimes less with the content of the webpages, and more with the style of presentation of the search results. For example, given that result ordering is crucial, there are many cases where the charge/conviction features prominently in the first couple of pages of search results, but the acquittal/appeal appears so low down that a searcher would be unlikely to get that far. Sometimes, the problem is not that the webpage's information is misleading, but the extract from the page that accompanies the result gives a false impression.

The objection raised by the person who brought the key case against Google Spain was against information he argued to be outdated and irrelevant to his current professional life. Some time previously, after some issues with his tax authority, his home had been repossessed and auctioned off. The auction was publicized in a newspaper in order to help maximize revenue for the auction. Once the newspaper's archive was digitized, the auction notice resurfaced, and the complainant argued that his privacy was being infringed because the proceedings had been fully resolved for several years, they were irrelevant to his current life and indeed had the potential to harm his professional career. He therefore argued that the newspaper should take down the piece from its

6 In general, article 6 of the DPD provides five data quality principles. Data must be: processed fairly and lawfully; collected for specific and explicit purposes; adequate, relevant and not excessive relative to those purposes; accurate and up to date; and kept in a form where data subjects are identifiable for no longer than required for the purpose.

7 See, for example, Information Commissioner's Office (ICO) (2011), which is somewhat lukewarm toward the evolving proposals.

archive, and that Google Spain should cease to index it in searches on his name. Although the AEPD rejected his case against the newspaper, whose archival function it respected, it found in his favour with respect to the search engine (thereby implicitly endorsing the complainant's assessment of the information), and Google Spain took the case to a resolution in the CJEU.

It is fair to say that many observers thought that the AEPD was not going to succeed in the case, particularly when the advocate-general, the CJEU's special adviser on legal matters, upheld crucial parts of Google Spain's case (European Commission 2013; Lynskey 2013). However, the court chose to reject the advocate-general's non-binding view, and came down in favour of the AEPD's original decision.

The Substance of the Judgment

In its judgment, the CJEU rejected all four key aspects of Google Spain's defence. Its responses to the italicized defences are summarized in the next four paragraphs.

- *Search is not data processing: it involves locating, indexing and even temporarily storing data, but not processing.* The DPD is clear that processing happens when data is "collected," "organized," "stored," "retrieved," "disclosed," etc. (article 2(b)), and the court was clear that this was indeed happening.

- *The European Union has no jurisdiction over the case, as the search engine was run from the United States by Google Inc., while Google Spain, which does fall under its jurisdiction, does no processing.* The CJEU ruled that Google Spain is an EU establishment, as it is based in Spain (this was not in contention). Furthermore, Google Inc.'s processing of the data took place in the context of the activities of Google Spain (on the territory of the member state Spain) that were "intended to promote and sell… advertising space offered by the search engine which serves to make the service offered by that engine profitable." Hence the search engine's data processing, even though it happened in the United States, took place within the context of Google Spain's business (it wouldn't have happened otherwise), which, the court argued, brought the processing within the European Union's jurisdiction.

- *Neither Google Spain nor Google Inc. is a data controller; they are merely passive intermediaries that make no distinction between personal data*[8] *and other kinds of data, have no control over it, and make no decisions relating to its management.* This was the key contention, with which the advocate-general concurred, arguing that to be a controller, "the data processing must appear to him as processing of personal data, that is 'information relating to an identified or identifiable natural person' in some semantically relevant way and not a mere computer code" (European Commission 2013). However, the CJEU rejected the argument because the search engine "determined the purposes and means of processing" *within the context of the activities of Google Spain*. This processing, controlled by Google Inc., was the subject of the case, not the processing performed by third-party webmasters, and it consisted in the creation of "a structured overview of the information" relating to the individual searched for, which could not be created in the absence of the search engine. The processing of personal data by search engines is distinct from and additional to that of the third parties, and also plays a decisive role in its dissemination.

- *The information was already public, and there was no right (and Google had no power) to erase it.* The court agreed that the information did not have to be taken down, assuming it was true. However, it also concluded that Google Spain was performing an extra privacy-relevant function, by bringing links to public information together on a single webpage. In this, the CJEU followed the US Supreme Court, which had recognized the privacy interest in collecting public information, and the privacy protection of what was termed *practical obscurity*. A 1989 judgment argued that FBI rap-sheets need not be released under Freedom of Information requests because "a third party's request for law enforcement records or information about a private citizen can reasonably be expected to invade that citizen's privacy, and that, when the request seeks no 'official information' about a Government agency, but merely records that the Government happens to be storing, the invasion of privacy is 'unwarranted'."[9] In other words, someone wanting to know about the FBI could have access to the information, but not someone wanting to know about the person. The CJEU's argument was roughly parallel (Goodman 2015). The public information upon which the search results would be based was to be unchanged, and the information could be made available through the search engine as long as the searcher's interest was not in the person involved, as evidenced by the search terms she used.

The Upshot of the Judgment

The victory of the AEPD showed that data subjects had the right to apply to Google to remove outdated, inaccurate or excessive information from Google searches within

8 Personal data is defined in the DPD as data from which an individual is identifiable. Different data protection acts implement the DPD across the European Union, and these differ in their interpretation of "identifiable." For instance, the UK Data Protection Act specifically defines "identifiable" as "identifiable by the data controller," which weakens its privacy-protecting provisions relative to other acts.

9 See https://supreme.justia.com/cases/federal/us/489/749/case.html.

Europe, as long as they were searches for information on the data subject him- or herself. So, for example, if one had committed some youthful misdemeanour that was referred to in a webpage, then one could go to Google with a request to de-index that specific uniform resource locator (URL) from searches on one's name. The webpage would remain online, and it could be reached via a different search — for instance, if one searched for examples of the specific misdemeanour, the offending webpage might legally appear in the search results. In the judgment, "forgetting" does not involve deletion, and so a right to be forgotten is distinct from a right to erasure. In that sense, the concept is somewhat closer to the notion of forgiving and moving on discussed earlier. Erasure is already a data protection right "where personal data storage is no longer necessary or is irrelevant for the original purposes of the processing for which the data was collected" (article 32 of the DPD). Furthermore, as this is a right, it is not necessary for the data subject to show that he has been harmed or the information is prejudicial; it is sufficient that he objects. However, it is accepted that archives have special requirements to hold information and to keep full records.

The key parameter to be provided to Google would be the URL of the webpage, not the information itself. If the offending information was present on a series of webpages, Google would only be obliged to de-index the particular pages of the URLs it had notified.

Google can turn down any such request. In that event, the complainant has the right to go to their national DPA (or straight to court), which can override Google's judgment. The judgment suggested a number of grounds for refusing a data subject's request. Although the economic interest of the search company was not deemed sufficient reason to overturn a European citizen's data protection rights, those rights would have to be balanced on a case-by-case basis against rights to freedom of expression and of the media, and also against the interests of the public in having access to the information via a search on the subject's name. The status of the complainant as a public figure would therefore be a contributory factor. Google has no obligation to inform third-party webmasters of its decision to remove a webpage from searches (though it often does), and those third-party webmasters therefore might, as far as the law is concerned, remain ignorant of a decision.

The decision only counts in the jurisdiction of the European Union, and applies to any searches carried out in the context of a business or enterprise established in the European Union, even if the actual servers carrying out the search are outside the European Union. Google Spain is certainly established in the European Union (as is Google Ireland, which sells the advertising), and so the California-based searching falls under the European Union's jurisdiction. The court said nothing about what the limits were to that judgment, but the most probable interpretation is that a search from a non-EU webpage — say, google.ca, which is based in Canada and intended for Canadian users — would be unaffected by the ruling. However, searches within the European Union — for example, on google.co.uk, google.be, google.fr and of course google.es — would be affected across all EU domains. Where Google has agreed to de-index an item in one domain in the European Union, it will follow suit across Europe.

It is finally important to point out that the key part of the CJEU's judgment was the finding that Google was a data controller. This role brings with it responsibilities under EU data protection legislation, and conversely if the court had *not* found that Google was a data controller, it would have been powerless. A data controller is defined as: "… the natural or legal person, public authority, agency or any other body which alone or jointly with others determines the purposes and means of the processing of personal data" (DPD, article 2(d)).

The advocate-general noted that Google makes no distinction between the personal data and the non-personal data that it processed, and that it does not treat personal data as personal (for example, it does not try to identify people from the data it processes). However, the CJEU ruled that those factors are not relevant to the question; Google processes personal data, whether or not it is aware that the data is personal, for purposes of its own, using means determined by itself, and for that reason it is a data controller. *But for that*, it would not fall under the jurisdiction of EU data protection law.

The Implementation of the Judgment

Following the judgment, Google attempted to drive the debate on privacy, data, free speech and a right to be forgotten by setting up a neutral advisory council of philosophers, politicians and entrepreneurs. It reported in 2015 (Advisory Council 2015), shortly after guidance was released by the Article 29 Working Party of EU DPAs (WP29 2014). The two documents set out somewhat divergent views. The WP29 document emphasized the judgment that search engines are data controllers, whose processing of personal data is in addition to the processing done by third-party websites. It played down the potential impact of the ruling, but at the same time made the strong claim that "de-listing should…be effective on all relevant domains, including .com." On the other hand, it suggested that the balance between the rights (and interests) of data subjects, and those of the data controller and the public, was perhaps finer than the CJEU had implied, and set out a series of criteria that would be relevant to making a judgment. The Google Advisory Council report also recommended criteria, but generally supported a weaker interpretation of the judgment. It recommended that publishers of information should be kept in the loop, informed of a de-listing where possible and given an opportunity to challenge a judgment. It also challenged the WP29 interpretation of the ideal geographical scope of

the judgment, concluding that "removal from nationally directed versions of Google's search services within the EU is the appropriate means to implement the Ruling at this stage" (Advisory Council 2015, 20). The council wasn't shy of asking a private corporation to make judgments in this space, because "assessing legal removal requests is already the norm for, and expected behavior of, search engines and other intermediaries in contexts outside data protection" (ibid., 18). Both reports recommended transparency in principles if not in the details of actual judgments.

However, simultaneously, Google constructively worked with DPAs to develop a procedure for dealing with the issues created by the judgment. The agreed procedure with Google is outlined below. It is likely that other search engines established within the European Union will implement something similar if they haven't already, since they will fall under the scope of the ruling.

If someone objects to a webpage appearing in a name-based search for them, they first contact the search engine to ask them to de-index the page from searches based on their name. There is a fairly Byzantine process in Google to do that — Google recommends that they contact the third-party webmaster first — but ultimately they are asked to fill in a form giving the reasons to de-index. Based on the information provided, Google makes a decision. In June 2015, the statistics released by Google showed that in the year following the judgment, there were 272,940 requests to remove 991,074 pages across the European Economic Area (EEA),[10] of which 58.6 percent were rejected by Google.[11]

If the person is unsatisfied by Google's decision, they can contact their national DPA, which makes an assessment and informs the search engine of its preliminary view. Information provided by the UK DPA, the ICO, reveals that the number of complaints at this stage in the process is currently small and manageable: in the first year, there were about 250 (Bourne 2015). In the United Kingdom, Google had 34,346 requests to take down 134,931 pages, of which it rejected 62.4 percent, which equates to 1.17 percent of failed complaints to Google going forward to the DPA.

In the United Kingdom, the ICO bases its decisions to uphold or reject Google's judgment on at least the following criteria (ibid.).[12]

- Does the search result relate to a natural person (an individual), and does it come up against a search on that person's name? Pseudonyms and nicknames will also be considered if the complainant can show that such names are linked to their basic identity.
- Does the individual play a role in public life? The ICO makes judgments here on a case-by-case basis, while recognizing that the public interest in information about public figures is stronger. One important question it will ask is whether the information whose de-indexing is requested could help protect the public against improper professional conduct.
- Does the data relate to an individual's working life? Not all personal data is private, and the less the data reveals about someone's private life, the more likely the ICO is to accept its availability in search results. Again, this judgment will depend on whether the individual in question is a public figure, although even such people have rights to privacy.
- Was the original published in a journalistic context? The law provides protection for journalism that is not available to search engines, so in that context, the ICO will take public rights to know and media rights to freedom of expression into account.
- Does the data relate to a criminal offence? The ICO takes into account public policy with respect to rehabilitation of offenders, and the existence of mechanisms outside Web search to protect the public. It handles these on a case-by-case basis, but is likelier to favour de-indexing for cases that are more minor, and that happened longer ago. The balance between public safety in particular (as many right-to-be-forgotten cases concern previous criminal convictions) and privacy is one that exercises the ICO in its thinking.

Some media outlets deliberately provide extra links to stories that have been de-indexed, for example via a central page linking to all such stories, either as a protest against a threat to their business models, or as a principled stand for free speech. This is perfectly legal, and is far less of a threat to privacy as the searcher would need to know the substance of the story in order to find something relevant to an individual. Such pages, as a matter of fact, provide researchers with interesting material for trying to work out

10 Actually, it is across the European Union and the European Free Trade Association (EFTA). Switzerland is a member of EFTA but not the EEA, while at the time of writing Croatia is a member of the European Union but only a provisional member of the EEA. Both are covered by Google's de-indexing regime. However, the EEA is a useful shorthand.

11 For up-to-date figures released by Google, see www.google.com/transparencyreport/removals/europeprivacy/?hl=en.

12 The following bullets are taken directly from an ICO presentation of its policy toward right-to-be-forgotten cases (Bourne 2015). An anonymous referee for this paper pointed out that, although the ICO sets out its policies in terms of the aspects of the context that it will take into account, its resources are limited, and it may struggle to live up to these ideals if it were presented with a large number of cases.

what kinds of requests are made. On the other hand, if the outlet republishes the content on a new page, then this will also circumvent the judgment (as it would be a different URL), and could lead to the search engine re-indexing the to-be-forgotten page. This, in contrast, is a notable threat to privacy.

ISSUES ARISING FROM THE JUDGMENT

An enumeration of several issues, positive and negative, arising from the judgment, can be found in Kieron O'Hara (2015). This section will briefly review a few of the most pressing and salient issues — in particular, the debate between privacy and free speech; the judgment's implicit view of the status of search results; the jurisdictional issues that European data protection activism has thrown up; the transparency of the de-indexing process; the potential difficulty individuals have with information that is proliferating or being spread; and the barriers to entry that may have been created.

Privacy versus Free Speech

The law is not new. The CJEU's task was to determine what was already implicit within the DPD, and it has argued that it merely interpreted DPD in the context of search. There is no extra right to erasure created, and information de-indexed remains online, findable by going direct to the site, and by following existing hyperlinks. Indeed, it can be found by standard search, as long as the search term is not the name of the data subject (it could be the name of another data subject who has not objected to the page). In this sense, the judgment has driven a wedge between rights of erasure or deletion, and rights to restrict access to information. The right to be forgotten falls under the latter, consistent with earlier critiques that erasure was not consistent with forgetting (Markou 2014), while also disappointing those who wished erasure or deletion rights to go further (Mayer-Schönberger 2014; Bernal 2011).

So, for instance, it could be argued that the financial difficulties of the original complainant should be accessible to, say, future employers or potential business partners. Employers could not be sure of getting that information by searching on his name after the Google Spain judgment (of course, they could not be sure of getting the information before Google Spain either, depending on what had been prominently linked to on the Web). But if they are entitled to that information, they can still go to official bankruptcy records to check. The difference is that in the latter case, there is a targeted search within the accepted scope of the employer's interests, while in the former there is a generalized search for any information, which may turn up relevant or irrelevant material.

One of the judgment's most controversial suggestions is that rights to privacy "override, as a general rule" (paragraph 81) freedom of information and expression rights. This is debatable, but the claim does help counterbalance a major asymmetry between privacy and free speech. In making a free speech argument, no one asks Google to show that it (or anyone) has been harmed by the de-indexing of certain pages; the cry of "censorship" is enough. The CJEU, in rejecting the requirement for the data subject to show harm, levels the playing field between privacy rights and free speech rights. Granted, rights to privacy might have to be balanced against others' rights (for example, the right to free speech), in which case the level of harm might become a factor in the deliberation. But it should not be a necessary condition in a rights-based discourse.

Yet some of the arguments that a right to be forgotten is a major blow to free speech have involved exaggerated claims that trade on the asymmetry. Speaking at an event, one prominent Internet scholar argued that a right to be forgotten was censorship. "It's like saying the book can stay in the library, we just have to set fire to the catalog" (quoted in Roberts 2015). The simile is overdrawn. It is more like saying the book can stay in the library, but we will remove the single catalogue entry that refers directly to X's name, while all the other catalogue entries remain in place (and we also, for good measure, keep the book in its right place on the shelves, so that you can also find it if you know the author's name). That is not to say that such a measure would not also be controversial, but it clearly does not support the analogy. Similarly, Jimmy Wales' argument, in his dissenting comments from the Google Advisory Council report, that publishers' works "are being suppressed" (Advisory Council 2015, 27) is an overstatement of the actual effect on the publishers, if we take "suppression" to mean the prevention of publication.

Not all commentators have gone so far. In his dissenting comments to the Google Advisory Council report, Frank LaRue argued that "we cannot make a difference between the information that exists, on files, official records or news papers, and that is obtained through a search engine" (ibid., 28). This seems like a category mistake — the information obtained through a search engine *is* the information that exists on files, etc. However, that is no reason not to distinguish between means of getting that information, given the privacy interest in dossiers of public information as recognized in the practical obscurity doctrine (Goodman 2015). There is little sign that this doctrine would constrain search engines in the United States, but it seems incorrect to suggest that there is no difference in either functional or privacy terms between 1,000 catalogues of 1,000 documents, where each catalogue contains one document that refers to Person X, and a single list of the 1,000 documents that refer to X. The judgment assumes a significant difference between these two circumstances.

The judgment should not inhibit serious journalism. A researcher in search of information about someone will have to invest more resources in finding public information, because the efficacy of "fishing expeditions" to find unspecific information is reduced. If the researcher or journalist is looking for something of any specificity at all, then they should be able to craft an effective set of search terms. The privacy threat to an individual is flagged by the use of the individual's name as a search term. Yet, as argued above, there is no pre-Internet right to be forgotten, and so erasure is not supported by the judgment. History, in the sense of what information is available on the Internet, is unchanged.

On the other hand, search engines play another important role with respect to journalism, in getting journalistic output before the public. Removal from search results could have a serious effect on the dissemination of journalism, as well as its pursuit. However, there are exemptions for journalism in the DPD, and DPAs will weigh the public interest in having access to the information. As noted above, the ICO in the United Kingdom, for example, will take that issue (and other issues, such as the public interest in knowing about perpetrators of serious crimes) into account.

And as noted, a determined searcher is unlikely to be disadvantaged for too long. There are many ways around the restriction, which means that the immediate effects of the judgment will be relatively minor. The judgment does not go as far as many privacy campaigners had originally demanded (Bernal 2011), and favours impeding the search for information over the more radical measures of policing and restricting misuse, or erasure (Oswald 2014).

Opening the Corporate Black Box

The judgment rejects the claim that search is a neutral "black box" that merely reflects the structure of the Web at a particular time. A search is a construct that mediates between the user and the Web of documents, and its ordering is a key factor in the likelihood of a link to a page being followed. Google, as a giant corporation employing many fine minds, will be able to cope with the further overhead created by a right to be forgotten. It has, after all, mapped the world, its search algorithms are already able to weed out items such as copyright material, link farms and users of the robot.txt exclusion protocol, and at the time of writing it is planning to de-index revenge porn on request (Singhal 2015). Necessarily, much about these algorithms is confidential (otherwise spammers could game them), but that very confidentiality speaks against search engines being trusted, neutral interfaces to the Web.

Google's marketing and market dominance depend on trust in the system, which in turn rest on a myth of completeness; its search is marketed as a non-selective neutral instrumentation of the conversations on the Web. Even some who want a strong right of erasure argue that Google's formal indifference to content should not be interfered with (Markou 2014). But, of course, Google doesn't index the entire Web, and eliminates and ignores many sources of information, and so this myth should be dispelled. Google is not the Web, although it is of course a marvellous tool for navigating the mass of information, possibly indispensable in the age of digital networks. Neither is the Internet or the Web a privileged version of history. Even when an aggregation of pages provides a true narrative, it is not necessarily the whole truth (as with a newspaper archive publishing a conviction for an offence but not the successful appeal).

Google is a partial view of a partial repository of information. For serious engagement with history, or attempts to hold people to account for their actions, or defence of the public against harm, Google, like Wikipedia, is an excellent starting point, but a starting point alone. It is not the whole Web, and the Web is not the whole truth.

Jurisdiction

The Internet and the Web have often been held up as exemplars of a new type of space, independent of the constraints and confines of the nation state, perhaps most famously in John Perry Barlow's *Declaration of the Independence of Cyberspace*. More prosaically, issues to do with regulation and law enforcement across different jurisdictions have often been problematic, and regulators have tended to work at a slower pace than innovators. Data protection law is a classic case where different interpretations of EU and US law, and the right to be forgotten, as well as other privacy issues, have long threatened to drive a wedge between the two jurisdictions (Whitman 2004; Bamberger and Mulligan 2011; Ambrose and Ausloos 2013; Bygrave 2014).

The CJEU's judgment has been implemented by Google only on its EU and EFTA domains, such as .es, .uk, .fr, .de and so on. The main .com site, which is US-facing, does not de-index pages on data protection/right-to-be-forgotten grounds. The rationale for this decision is that Google has a large share of the European search market, most of which goes on the national domains such as google.co.uk. Someone wishing to use google.com in Europe is diverted to the national domain, and it takes a little persistence to get to google.com (or indeed any other non-EU national domain). It is not much of a barrier to the determined (indeed, you can make google.com your home page to circumvent the defaults), but the power of default (plus linguistic preferences) means that most searchers end up using their national domain. This minor (but, in practice, significant) barrier reduces the radicalism of a right to be forgotten, and meets the *desideratum* that it protects Europeans in Europe, where data protection rights are recognized, while not protecting anyone elsewhere. For

most Europeans, their reputations matter most in Europe, and so the level of protection is useful and not insignificant.

This view is not universally held. Following the Google Spain judgment, little has been heard of Google's defence that the European Union should have no jurisdiction over the actions of a US company operating equipment in California, but presumably that feeling has not gone away (a Republican Congress might one day consider the argument). On the other side, the Article 29 Working Party went beyond the CJEU's judgment to demand that it should also apply to the .com domain, as this was (easily) reachable from Europe.

> In order to give full effect to the data subject's rights as defined in the Court's ruling, de-listing decisions must be implemented in such a way that they guarantee the effective and complete protection of data subjects' rights and that EU law cannot be circumvented. In that sense, limiting de-listing to EU domains on the grounds that users tend to access search engines via their national domains cannot be considered a sufficient mean to satisfactorily guarantee the rights of data subjects according to the ruling. In practice, this means that in any case de-listing should also be effective on all relevant domains, including .com. (WP29, 2014)[13]

The argument over jurisdiction continues, and will remain live for some time. At the time of writing, Google is reported (Fioretti 2016) to be about to implement a judgment from the Commission nationale de l'information et des libertés, the French DPA, to extend the right-to-be-forgotten procedure to all domains globally, having initially resisted it (Fleischer 2014). It will only apply, at present, to searches within European territory (so a European search on google.com would be de-indexed, but not one from outside Europe).[14] However, in practical terms, it is hard to see how the European Union could enforce global compliance. Furthermore, the same logic could be applied to EU-based search engines by more repressive governments.

The position of enforcing a right to be forgotten in EU territory — and not elsewhere — is enforceable, largely effective given the percentage of searches done on European domains in Europe (where most Europeans have their main privacy interests), and not over-restrictive. It respects the different intuitions, rules and norms that obtain outside Europe, while simultaneously remaining consistent with the CJEU's reasoning and the imperative for data protection within the European Union. It also appropriately constrains a right to be forgotten.

Transparency

The original judgment gave little guidance as to the criteria for the decision to de-index or not, although since then the Article 29 Working Party has provided non-binding guidance (WP29 2014). Google's Advisory Council has also given its advice on the topic (Advisory Council 2015, 7–14). Google itself has made decisions on hundreds of thousands of requests, of which only a tiny percentage have been referred to DPAs. Teams of lawyers, paralegals and engineers deal with the many "easy" cases, while hard cases are referred to the executive level (Fleisher and Schechner 2015). Google, as noted, releases statistics on its decisions, which have stabilized at an acceptance rate of about 40 percent. It is certainly important that jurisprudence should emerge (Jones 2015).

There is no doubt that the decisions Google has been asked to make (and this is not a power it sought) are important ones involving censorship and information flow. It is not ideal that such decisions be privatized at all,[15] but even given that privatization was the solution, it is essential that decisions be transparent. Google's Global Privacy Counsel has argued that it is "building a rich programme of jurisprudence," but this program is, in the words of an open letter to Google by 80 scholars requesting greater transparency, "built in the dark" (Goodman, Powles et al. 2015).[16] There is, of course, tension between the needs of transparency and privacy, but aggregate statistics — for example, of categories of successful and unsuccessful claimants, or of the types of requests made and granted (crime victims? health information? false accusations? old and minor misdemeanours? political opinions no longer held?) — should be possible to generate without threatening privacy. At the time of writing, Google is considering this request (Collins 2015).

It is also possible that third-party publishers might be more readily involved in the judgment process (especially the media, given the protections for journalism in the DPD, although one would not wish accidentally to inform, say, a revenge porn site that a subject had invoked their right to be forgotten). This would allow input of more relevant information, from the publisher, into the decision-making process. It would also allow publishers to take a case to

13 See also Sabine Leutheusser-Schnarrenberger's dissenting argument to the same effect in Advisory Council (2015, 26).

14 See www.reuters.com/article/us-google-eu-privacy-idUSKCN0VJ29U.

15 There is debate on this. Even on the Google Advisory Council some, such as Leutheusser-Schnarrenberger, argued that "this is a typical relationship between a private user on the one hand … and a private company on the other hand ... [whose] right to decide cannot be taken away" (Advisory Council 2015, 25), while La Rue took the opposite view (ibid., 29).

16 Disclaimer: one of the authors of this paper, Kieron O'Hara, was a signatory to the letter.

the DPA, which is important, given that most DPAs have the dual function of protecting privacy and freedom of information. The risk, however, with this option is that it would also allow publishers to identify and republish de-indexed pages with a new URL, which would take them out of the scope of the judgment, and would then require the individual to make a new approach to the search engine.

Onus on the Individual

The system as it has evolved places the onus of work on privacy-aware individuals, and in this sense is part of a general trend (Van der Sloot 2014). In particular, they have to specify particular URLs to be considered, and the statistics show that the average individual specifies about three or four. Yet these individuals are less interested in making access to particular webpages harder than lowering the likelihood that someone specifically interested in them in particular can easily get hold of outdated or excessive information about them, or information that puts them in a false (usually bad) light. So interconnected is the Web that information is likely to be distributed across several pages, and may feature in a range of contexts. It may also be disseminated maliciously.

The key variable is not the webpage, but the information, yet the individual is not allowed to specify the association or information that is embarrassing, misleading or outdated. If information proliferates, they can only try to keep track of which URLs the information appears on, and contact search engines accordingly. It does not seem to be the case that it is *easy* to reduce access to information, particularly if it is widely distributed (*pace* Jones 2015). Indeed, despite the arguments of the judgment's opponents, there is little evidence about how much individuals have benefited from it. Maybe de-indexed pages simply get posted under alternative URLs routinely, to reappear in search results. Without extensive evaluation, it cannot be known how effective a protection the system provides.

Barriers to Entry to Search

The final point that will be emphasized in this paper is that, although Google can cope perfectly well with the extra burden, this is because it is a well-resourced company. DPAs, in contrast, could not deal with all requests directly. At a point when the European Commission is concerned about competition issues in search (European Commission 2015), it may be a perverse effect to increase the barriers to entry to the search market by insisting on the implementation of a right to be forgotten by search engines other than Google. Having said that, it may also be the case that Google's machinery for dealing with de-indexing requests has been over-engineered and that there would be cheaper, more transparent and less burdensome ways of dealing with them (Powles and Floridi 2014). Ultimately, it will be essential to explore the means to increase transparency, and to make the interactions between search engines and complainants (and DPAs) more routine, in order that a right to be forgotten can be implemented without large-scale resources.

Data Protection in the Digital Age

This section has discussed a number of issues arising from the Google Spain judgment, but their effects can be detected beyond both the individual case, and the relatively narrow class of cases to which the judgment applies. The issues of privacy, free expression, transparency and the asymmetries of power that have been discussed here all play out in a number of ways as our digital technologies record ever more data, and increasingly many of our actions and interactions are symbolized and recorded, becoming visible and shareable in new and unfamiliar ways.[17] Our means of negotiating these difficult and uncertain waters will vary widely, and will include changes in law, social norms, business models and education. In the final section of this chapter, we will consider one possible technological approach that has been advocated in the context of these wider themes of data protection in the digital age, and sketch (lightly) a possible approach to rebalancing power.

PERSONAL DATA MANAGEMENT: EMPOWERING AND MAINTAINING TRUST

Currently, the discussion has been at the level of law. However, it is also possible that technology could play a part in the solution. There are a number of potential technological fixes for (parts of) the problem, including improving accountability for the misuse of information, enriching search with sentiment analysis, and a clearer process for reporting and dealing with disputes. This section will consider one particular technology that may be part of the solution, given the appropriate supporting background of regulation, digital literacy and social norms.

However, our aim is not primarily to argue for the introduction of this technology. This is a thought experiment — the idea is to show that a different relationship, mediated by technology, between data subjects and data consumers is possible, and that many of the issues arising from the right-to-be-forgotten judgment, and from problems with privacy in general, could be addressed in a different world. We will develop the thought experiment to highlight what is lacking in the current regime. In particular, if the world contained a vibrant market for *personal data management*, then more equitable relationships, with fewer information asymmetries, could be sustainable.

17 Two interesting and contrasting critiques of this new tendency are Hildebrandt (2015) and Zuboff (2015).

Personal Data Management Architectures

The Web was designed as a decentralized information and communication tool, but recently this model has been frayed by the economic forces of network effects, technological lock-in and low marginal costs of adoption, which have favoured large corporations able to amass giant user bases for their walled gardens (Zittrain 2008). Data is harvested from users and consolidated in giant databases where analytics produce monetizable insight to the benefit of data gatherers. People are decoupled from their data, unable to manage, curate or police it, and identity management and partitioning are hard, leading to a lack of trust (Coll 2015).

One class of technologies with the potential to rebalance asymmetries and restore trust are architectures that allow the data subjects some measure of control over, or input into the exploitation of, their personal data, including both data they have collected themselves and data collected or inferred about them. Let us call these Personal Data Management Architectures (PDMAs), intending the term to be agnostic over particular architectures, affordances and business models. It includes, but is not restricted to, Personal Data Stores (PDSs) and Personal Information Management Services (PIMSs) (Heath, Alexander and Booth 2013; Nguyen et al. 2013; Ctrl-Shift 2014; Van Kleek and O'Hara 2014; Abiteboul, André and Kaplan 2015). There is some skepticism about the PDS model of information management, often on the grounds of security or usability (Lemley 2000, Narayanan et al. 2012). The technology is certainly not mature, and although there are a number of products available there is still much work to do. Furthermore, regulation and business models do not work to its advantage. This paper does not address these problems directly, but as a thought experiment let us assume that next generation data management is possible, with a mature industry in which security and interface issues have been largely resolved. To reiterate, our aim is not to provide a road map of how to get from here to a PDMA world, but rather to envisage a different relationship between data consumers and data subjects.

The services PDMAs might provide include user-centric consent management tools, preventing external access to data except under approved conditions, negotiating privacy policies, handling credentials and even allowing access to rich sources of data from personal data collection devices (for example, health-care monitors such as the FitBit) for payment, free services or other benefits. It is important to note that such services *do not* depend on the PDMA storing data, and it should not be assumed that they will necessarily provide storage services (although PDSs do, and there is no reason why a PDMA might not store *some* data). They might merely point to data, or handle our interactions around it.

The PDMA could act as a privacy and identity assistant, with an understanding of context (such as interaction history), mapping multiple identities to different activities, and establishing trust credentials from those requesting access to the data. Forced identity consolidation as favoured by the walled gardens would no longer be appropriate (or possible), and data would have portability across at least some contexts. The PDMA would manage interactions so that external parties need not be aware that, for example, the employee of a well-known bank, the player of World of Warcraft, the denizen of a fetish site and the campaigner for immigration rights are all the same person. There is also no implication in this account that anyone would be restricted to a single PDMA. One could partition identity across PDMAs, and use them for different purposes.

PDMA technology is certainly not mature, and may never make a market breakthrough, but in this speculative section let us assume that innovation capable of providing the above-mentioned services is with us. Assuming a mature market of critical mass emerges, the Web, currently centralizing around the major platforms, could be re-decentralized by socially aware PDMAs.

PDMAs and the Right to be Forgotten

PDMAs might help with the de-indexing issues raised in this paper by being the locus for dialogue and interaction with search engines, publishers and DPAs. This arrangement would require the development of new norms and possibly new regulation, but would not require a critical mass of PDMA users to work. All that is assumed in this section is a PDMA ecosystem that would allow privacy-aware individuals to manage their relations with search engines. Nothing precludes PDMAs being used alongside other technologies to interface with search engines.

The following functions or practices, integrated with the PDMA, would help craft a holistic approach to the issues raised by the Google Spain judgment.

Storing details of information or data to which its owner would wish to reduce public access by exercising their data protection rights. This would include URLs of webpages with excessive or outdated information, but might also include a specification of the problematic event(s) or information. Given that information, the PDMA could periodically search for pages that referred to it. Discovery of a prominently placed webpage with the offending information would prompt the PDMA to contact the relevant search engine automatically, or to send an alert to the user.

Associating with this database of URLs the metadata that search engines would require to assess whether the criteria for de-indexing were met — for instance, how old the information in question might be, whether the PDMA-owner was a public figure, and so on.

Cooperating with search engines. When a search came up on a person's name, a search engine could also look for PDMAs owned by people named by the identifying string, and proactively look for offending URLs in the search results, and even look for pages containing the offending information. Of course, the engine would not be obliged to de-index those pages, but could test them against its de-indexing criteria if it had access to the relevant metadata as well through the PDMA. Currently, there is no mechanism to allow search engines to do this.

Hosting dialogue with search engines, third-party publishers and DPAs. Whenever the PDMA's owner invoked a right to be forgotten, they must expect dialogue, explanation and discussion of the importance or otherwise of the information, its context, its prominence in the search results, the motives for publication, the age of the incident reported and the owner's status with respect to the public space. Such a dialogue would of course require careful monitoring of access and management of credentials. If the PDMA hosted this dialogue, there would be a central venue for the debate, and if another search engine found itself with the same right-to-be-forgotten case before it, it could immediately visit the discussion, to see, for example, how the DPA treated the case, and what courses of action other search engines had taken, thereby reducing the costs of enforcement of the right to be forgotten across the search industry.

Informing publishers. This is a risk, of course, but the above dialogue could also lead to a successful request to erase the webpage altogether, if it was sufficiently misleading or false to ring standard data protection alarm bells, if it wasn't covered by exemptions for journalism or archives, and if the jurisdiction of the website's owner was within Europe. The publisher may or may not be given access to the nature of the offending information, depending on how sensitive it was. Even so, at least the publisher would be able to annotate the database of URLs within the PDMA to give his side of the story. Such annotations would be available to search engines, the PDMA's owner and ultimately the DPA (if alerted by another party), to enable a balanced decision to be made about de-indexing, both now and in future cases.

The PDMA, therefore, could handle the database of problematic URLs, the nature of the information to be de-indexed, the metadata, the discussion, the interaction with the search engine and DPA, and the requests for de-indexing — all in a handy place that can be readily accessed during a search on the individual's name. And if a search engine wished to consider problematic pages proactively, then it could include relevant PDMAs in its search whenever it received a search request on a name or identifier. These functions would improve the interaction between data subjects and search engines in a number of ways.

First, it would reduce the effort for an individual to patrol the Internet (PDMAs generally have the aim of reducing data management demands while increasing an individual's power over their data). The onus of complaint would remain on the individual, but searching for content could be automated, and so the effort required would be lower. The PDMA could handle communication with the search engine itself, or it could merely warn a data subject of a problem. It could also structure the complaint, based on the metadata it held about the offending incident or information.

Second, by doing this it would help rebalance the power asymmetries between data users and data subjects, even if only to a small degree. Third, it would lower the barriers to entry to the search market, by providing a guide for new entrants to previous decisions and actions by search engines, publishers and DPAs. Fourth, it would lower the burden on DPAs to collect discussion and argument in one place. Fifth, it would provide a route to introduce third-party publishers into the debate to defend their position. If search engines played an active role in consulting PDMAs and annotating their databases — perhaps a big "if" — then the gains would be larger. The cause of transparency would be served, while much of the uncertainty that currently surrounds this issue — for data subjects, search engines, other data controllers and DPAs alike — could be dispelled. Search engines' cooperation is also the simplest means of genuinely reducing the onus on the individual (rather than merely automating their responsibilities).

Why would search engines collude in redrafting the social contract between data user/gatherers and data subjects? One reason might have to do with one of the other issues discussed above, that of opening up the corporate black box. Much of the search engine myth depends on an assumption of formal indifference. They, in theory, do not care what their users say or do, or what they search for; they are non-judgmental. They want as much data about as many actions as possible, however subjective, to get a full picture of the range of human endeavour, noble or embarrassing, idealistic or cynical, significant or trivial, selfless or prurient. All that matters is that the data is captured.

This is an important picture, but it is an ideal. As noted earlier, formal indifference is an ideal that Google tries to approach, rather than expects to achieve — it weeds out link farms, copyright material, child porn and revenge porn. There are campaigns to suppress more content, such as real-life torture videos (Overton 2015). Yet beyond these special categories of content, data protection legislation provides a series of quality principles (see footnote 8). Augmenting the semantically neutral calculations about the links to a page, a commentary based around data protection principles — is this information outdated? Is it excessive? — is also potentially helpful for searchers. It is arguable that if information has been judged (either

by an internal process in the search engine, or more formally by a DPA) to commit one of the data protection sins, then its value to a searcher is correspondingly less than it otherwise would be. Thus the search engine, by taking this into account, is adding value to its searches, not diminishing them. Which searcher would prefer misleading information to relevant information? The information in the Google Spain decision, after all, had been found misleading by three courts and regulators.

Currently, search engines' business models are usually focused around data processing, surveillance and advertising, but at the heart of the business is the search function, which competes on quality. Formal indifference is not a guarantor of quality; the moment search takes account of malicious content, a distinction is made between the "useful" web of content, and the "parasitic" web of spam. The Google Spain judgment has introduced the data protection framework as a competing quality vector, which may ultimately work to search engines' advantage.

The mechanisms embedded in PDMAs described above would ease the requirements on search engines that took this line, by streamlining debate with aggrieved data subjects and DPAs, giving a voice to third-party publishers, recording the rationale for decisions and avoiding duplication of decision making.

Not all search engines would have to adopt this position; those that did, or those, such as Google, that found themselves legally obliged to, would find valuable resources for the task. It would also not be the case that each search engine would have to come to the same conclusion about whether a particular item should be de-indexed or not. Not only would different national DPAs sometimes differ, but search engines might have different policies about when the quality of search results was compromised.

CONCLUSION

The privacy/free speech issues that Google Spain has raised, together with the potential jurisdictional conflicts, are not intractable, as our speculative thought experiment about new norms for interaction between search engines and individuals, mediated by PDMAs, shows. In particular, if search engines agreed to include consideration of statements about the quality of information on websites collected in PDMAs during searches on names, many of the conflicts based around the use of law to protect privacy, and much of the unfinished business of the present situation, would be ameliorated.

Of course, there would be a question as to why search engines might adopt such a code. One answer could be based on a revision of business models — the task (and cost) of remaining DPD-compliant might be eased by interaction with PDMAs, and there may be other benefits (for example, access to greater quantities of other data) that follow. Another reason might be that search engines' own assessments of the quality of search results they put out could be augmented by the five data protection principles of data quality. Or it may be that the intangible benefits of goodwill and a proper respect for privacy and data protection would bring the tangible business benefits of corporate social responsibility.

Clearly, the use of PDMAs in the maximal sense would reduce the onus on the individual. Individuals are interested in protecting their reputation, and in informational self-determination, not in the identification of specific webpages, and are unlikely to have the resources to police the Web and detect every single threat to their privacy. An ecosystem in which search engines cooperated with individuals using PDMAs would no doubt not be perfect either, but the balance would at least be redressed and the task less Sisyphean.

It would also help open the corporate black box to sunlight and scrutiny. This would help lower the barriers to entry, as the PDMA would be an early port of call for a search engine, which would then be able to access any existing discussion relating to a particular complainant and make an earlier, speedier, more informed and less risky decision without the need to employ a complex evaluation process in all cases. Transparency may be an issue, however, as too much information revealed to the outside world about an interaction could identify someone as an objector to the dissemination of a particular piece of information, which in turn might alert third parties to what that information was, thereby counterproductively revealing what was to have been concealed (known as the Streisand effect). However, it would still be possible for search engines to flag all searches that may have been amended because of the right to be forgotten, as Google does now, and to release accurate and fine-grained statistical information.

The past is over; its interpretation is not. In our digital age, searches are not preambles to the interpretation and understanding of the past, neutral providers of raw materials. Search is itself a vital part of the interpretative process. This important truth must stay in the forefront of our minds as we work to regulate in this space.

It must also be remembered that this kind of forgetting (and certainly anything stronger) is a conscious decision to interrupt the flow of information. This is an active process, and so it is paramount to make sure that it takes place within a framework of accountability. It should also be ensured that records of the past remain accessible to challenge contemporary narratives and current tropes. Given the controversy that surrounds it, the scope and power of any implemented right to be forgotten should surely be, in the first instance at least, limited and constrained. The lack of an offline analogue, the potential clash with free expression, and the potential for the powerful to erase

traces of wrongdoing all point in that direction. In this paper, it is argued that the CJEU's judgment, as currently interpreted and implemented, meets these desiderata, and that the technological resources to cement a new and more equitable relationship between data consumers and subjects within this framework are not out of reach.

Acknowledgements

This research was carried out under the UK Engineering and Physical Sciences Research Council programme grant Social Machines (SOCIAM), grant number EP/J017728/1. The authors would like to thank Anni Rowland-Campbell of Intersticia, Iain Bourne and Julia Parr of the ICO, and two anonymous reviewers for comments on this paper.

WORKS CITED

Abiteboul, Serge, Benjamin André and Daniel Kaplan. 2015. "Managing Your Digital Life." *Communications of the ACM* 58 (5): 32–35.

Advisory Council. 2015. *The Advisory Council to Google on the Right to be Forgotten*. www.google.com/advisorycouncil/.

Ambrose, Meg Leta and Jef Ausloos. 2013. "The Right to be Forgotten Across the Pond." *Journal of Information Policy* 3: 1–23.

Augé, Marc. 2004. *Oblivion*. Minneapolis: University of Minnesota Press.

Ausloos, Jef. 2012. "The Right to be Forgotten — Worth Remembering?" *Computer Law and Security Review* 12 (2): 143–52.

Bamberger, Kenneth A. and Deirdre K. Mulligan. 2011. "Privacy on the Books and on the Ground." *Stanford Law Review* 63 (2): 247–316.

BBC. 2015. "Sexting Boy's Naked Selfie Recorded As Crime By Police." BBC News, September 3. www.bbc.co.uk/news/uk-34136388.

Bernal, Paul Alexander. 2011. "A Right to Delete?" *European Journal of Law and Technology* 2 (2). http://ejlt.org/article/view/75.

Bourne, Iain. 2015. "Where Now for the 'Right to be Forgotten'?" Paper presented at the Second Conference on Trust, Risk, Information and the Law, Winchester, April.

Bygrave, Lee A. 2014. "Data Privacy Law and the Internet: Policy Challenges." In *Emerging Challenges in Privacy Law: Comparative Perspectives*, edited by Normann Witzleb, David Lindsay, Moira Paterson and Sharon Rodrick, 259–89. Cambridge: Cambridge University Press.

CJEU. 2014. *The Court of Justice Declares the Data Retention Directive to be Invalid*. CJEU press release, April 8. http://curia.europa.eu/jcms/jcms/P_125951/.

Coll, Liz. 2015. "Personal Data Empowerment: Time for a Fairer Data Deal?" London: Citizens' Advice Bureau. www.citizensadvice.org.uk/Global/CitizensAdvice/Corporate%20content/Publications/Personal%20data%20empowerment%20report.pdf.

Collins, Katie. 2015. "Google 'Considers' Further 'Right to be Forgotten' Transparency." *Wired*, May 14.

Ctrl-Shift. 2014. "Personal Information Management Services: An Analysis of an Emerging Market." London: Nesta. www.nesta.org.uk/sites/default/files/personal_information_management_services.pdf.

Daley, Suzanne. 2011. "On its Own, Europe Backs Web Privacy Fights." *The New York Times*, August 9.

Etzioni, Amitai. 1999. *The Limits of Privacy*. New York: Basic Books.

European Commission. 2013. *Opinion of Advocate General Jääskinen*. June 25. http://curia.europa.eu/juris/document/document.jsf?docid=138782&doclang=EN.

———. 2014. *Factsheet on the "Right to be Forgotten" Ruling (C131-12)*. http://ec.europa.eu/justice/data-protection/files/factsheets/factsheet_data_protection_en.pdf.

———. 2015. *Antitrust: Commission Sends Statement of Objections to Google on Comparison Shopping Service; Opens Separate Formal Investigation on Android*. European Commission press release, April 15. http://europa.eu/rapid/press-release_IP-15-4780_en.htm.

Fioretti, Julia. 2016. "Google to scrub web search results more widely to soothe EU objections." Reuters, February 10. www.reuters.com/article/us-google-eu-privacy-idUSKCN0VJ29U.

Fleischer, Peter. 2011. "Foggy Thinking About the Right to Oblivion." Blog, March 9. http://peterfleischer.blogspot.co.uk/2011/03/foggy-thinking-about-right-to-oblivion.html.

———. 2014. "Implementing a Global, Not European, Right to be Forgotten." Europe blog, July 30. http://googlepolicyeurope.blogspot.fr/2015/07/implementing-european-not-global-right.html.

Fleisher, Lisa and Sam Schechner. 2015. "How Google's Top Minds Decide What to Forget: as 'Right to be Forgotten' Ruling Turns One Year Old, Google Offers Glimpse Into its Decision-Making Process." *Wall Street Journal*, May 12.

Goodman, Ellen P. 2015. "Practical Obscurity and the Right to be Forgotten: 'Pretty Much' Privacy is Enough." *medium.com* (blog), February 4. https://medium.com/@ellgood/practical-obscurity-and-the-right-to-be-forgotten-pretty-much-privacy-is-enough-c321bdaffa08.

Goodman, Ellen P., Julia Powles et al. 2015. "Open Letter to Google From 80 Internet Scholars: Release RTBF Compliance Data." *medium.com* (blog), May 14. https://medium.com/@ellgood/open-letter-to-google-from-80-internet-scholars-release-rtbf-compliance-data-cbfc6d59f1bd.

Gross, Richard and Rob McIlveen. 1999. *Memory*. London: Hodder & Stoughton.

Heath, William, David Alexander and Phil Booth. 2013. "Digital Enlightenment, Mydex, and Restoring Control Over Personal Data to the Individual." In *Digital Enlightenment Yearbook 2013: The Value of Personal Data*, edited by Mireille Hildebrandt, Kieron O'Hara and Michael Waidner, 253–69. Amsterdam: IOS Press.

Hildebrandt, Mireille. 2015. *Smart Technologies and the End(s) of Law*. Cheltenham, UK: Edward Elgar.

ICO. 2011. "The Information Commissioner's (United Kingdom) Response to 'A Comprehensive Approach on Personal Data Protection in the European Union'." European Commission. January 14. http://ec.europa.eu/justice/news/consulting_public/0006/contributions/public_authorities/ico_infocommoffice_en.pdf.

Jones, Meg Leta. 2015. "Forgetting Made (Too) Easy." *Communications of the ACM* 58 (6): 34-5.

Koops, Bert-Jaap. 2011. "Forgetting Footprints, Shunning Shadows: A Critical Analysis of the 'Right to be Forgotten' in Big Data Practice." Social Science Research Network, December 20. http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1986719.

Law Commission. 2013. *Contempt of Court (1): Juror Misconduct and Internet Publications*. HC 860. London: The Stationery Office.

Lemley, Mark A. 2000. "Private Property: A Comment on Professor Samuelson's Contribution." *Stanford Law Review* 52: 1545–57.

Lynskey, Orla. 2013. "Time to forget the 'Right to be Forgotten'? Advocate General Jääskinen's opinion in C-131/12 Google Spain v AEPD." European Law Blog, July 3. http://europeanlawblog.eu/?p=1818.

Margalit, Avishai. 2002. *The Ethics of Memory*. Cambridge, MA: Harvard University Press.

Markou, Christiana. 2014. "The 'Right to be Forgotten': Ten Reasons Why it Should be Forgotten." In *Reforming European Data Protection Law*, edited by Serge Gutwirth, Ronald Leenes and Paul de Hert, 203–26. Dordrecht: Springer.

Mayer-Schönberger, Viktor. 2009. *Delete: The Virtue of Forgetting in the Digital Age*. Princeton, NJ: Princeton University Press.

———. 2014. "Omission of Search Results is Not a 'Right to be Forgotten' or the End of Google." *The Guardian*, May 13.

Mutch, Robert E. 2014. *Buying the Vote: A History of Campaign Finance Reform*. New York: Oxford University Press.

Narayanan, Arvind, Solon Barocas, Vincent Toubiana, Helen Nissenbaum and Dan Boneh. 2012. "A Critical Look at Decentralized Personal Data Architectures." arXiv. http://arxiv.org/abs/1202.4503.

Nguyen, M.-H. Carolyn, Peter Haynes, Sean MacGuire and Jeffrey Friedberg. 2013. "A User-Centred Approach to the Data Dilemma: Context, Architecture, and Policy." In *Digital Enlightenment Yearbook 2013: The Value of Personal Data*, edited by Mireille Hildebrandt, Kieron O'Hara and Michael Waidner, 227–42. Amsterdam: IOS Press.

O'Hara, Kieron. 2012. "Can Semantic Web Technology Help Implement a Right to be Forgotten?" *Computers and Law* 22 (6).

———. 2015. "The Right to be Forgotten: The Good, the Bad and the Ugly." *IEEE Internet Computing* 19 (4), 73–79.

Oswald, Marion. 2014. "Seek and Ye Shall Not Necessarily Find: The Google Spain Decision, the Surveillant on the Street, and Privacy Vigilantism." In *Digital Enlightenment Yearbook 2014: Social Networks and Social Machines, Surveillance and Empowerment*, edited by Kieron O'Hara, M.-H. Carolyn Nguyen and Peter Haynes, 99–115. Amsterdam: IOS Press.

Overton, Iain. 2015. "What Will It Take to End the Pornography of Videoed Torture?" *The Guardian*, September 7.

Post, Robert C. 2001. "Three Concepts of Privacy." *Georgetown Law Journal* 89.

Powles, Julia and Luciano Floridi. 2014. "A Manifesto for the Future of the 'Right to be Forgotten' Debate." *The Guardian*, July 22.

Reding, Viviane. 2010. "Privacy Matters: Why the EU Needs New Personal Data Protection Rules." Speech presented for the European Commission, November 30. http://europa.eu/rapid/press-release_SPEECH-10-700_en.pdf.

———. 2012. "The EU Data Protection Reform 2012: Making Europe the Standard Setter for Modern Data Protection Rules in the Digital Age." Speech presented at the Digital Life Design Conference, Munich, January 24. http://europa.eu/rapid/press-release_SPEECH-12-26_en.htm.

Ricoeur, Paul. 2006. "Memory — Forgetting — History." In *Meaning and Representation in History*, edited by Jörn Rüsen, 9–19. Oxford: Berghahn Books.

Roberts, Jeff John. 2015. "The Right to be Forgotten From Google? Forget it, Says U.S. Crowd." *Fortune*, March 12. http://fortune.com/2015/03/12/the-right-to-be-forgotten-from-google-forget-it-says-u-s-crowd/.

Rosen, Jeffrey. 2012. "The Right to be Forgotten." *Stanford Law Review* 88.

Schacter, Daniel L. 2001. *The Seven Sins of Memory: How the Mind Forgets and Remembers*. New York: Houghton Mifflin.

Singhal, Amit. 2015. "'Revenge Porn' and Search." Google Public Policy Blog, June 19. http://googlepublicpolicy.blogspot.co.uk/2015/06/revenge-porn-and-search.html.

Siry, Lawrence and Sandra Schmitz. 2012. "A Right to be Forgotten? How Recent Developments in Germany May Affect the Internet Publishers in the US." *European Journal of Law and Technology* 3 (1). http://ejlt.org/article/download/141/222.

Van der Sloot, Bart. 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): 307–25.

Van Kleek, Max and Kieron O'Hara. 2014. "The Future of Social is Personal: the Potential of the Personal Data Store." In *Social Collective Intelligence: Combining the Powers of Humans and Machines to Build a Smarter Society*, edited by Daniele Miorandi, Vincenzo Maltese, Michael Rovatsos, Anton Nijholt and James Stewart, 125–58. Cham: Springer.

Whitman, James Q. 2004. "Two Western Cultures of Privacy: Dignity Versus Liberty." *Yale Law Journal* 113 (6).

WP29. 2014. *Guidelines on the Implementation of the Court of Justice of the European Union Judgment on "Google Spain and Inc v. Agencia Española de Protección de Datos (AEPD) and Mario Costeja González" C-131/12*. Brussels: Article 29 Data Protection Working Party.

Zanfir, Gabriela. 2014. "Tracing the Right to be Forgotten in the Short History of Data Protection Law: The 'New Clothes' of an Old Right." In *Reforming European Data Protection Law*, edited by Serge Gutwirth, Ronald Leenes and Paul de Hert, 227–49. Dordrecht: Springer.

Zittrain, Jonathan. 2008. *The Future of the Internet — And How to Stop It*. New Haven, CT: Yale University Press.

Zuboff, Shoshana. 2015. "Big Other: Surveillance Capitalism and the Prospects of an Information Civilization." *Journal of Information Technology* 30: 75–89.

CIGI PUBLICATIONS
ADVANCING POLICY IDEAS AND DEBATE

Global Commission on Internet Governance

The Global Commission on Internet Governance (GCIG) was established in January 2014 to articulate and advance a strategic vision for the future of Internet governance. The two-year project conducts and supports independent research on Internet-related dimensions of global public policy, culminating in an official commission report that will articulate concrete policy recommendations for the future of Internet governance. These recommendations will address concerns about the stability, interoperability, security and resilience of the Internet ecosystem. Launched by two independent global think tanks, the Centre for International Governance Innovation and Chatham House, the GCIG will help educate the wider public on the most effective ways to promote Internet access, while simultaneously championing the principles of freedom of expression and the free flow of ideas over the Internet.

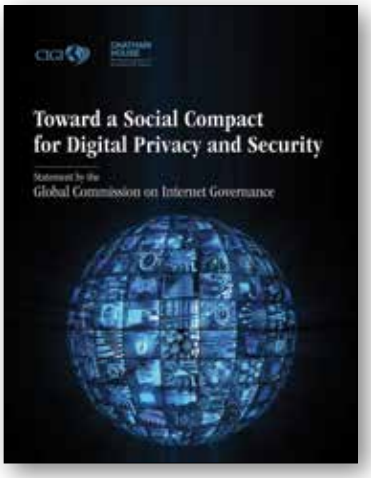

Toward a Social Compact for Digital Privacy and Security

Statement by the Global Commission on Internet Governance

On the occasion of the April 2015 Global Conference on Cyberspace meeting in The Hague, the Global Commission on Internet Governance called on the global community to build a new social compact between citizens and their elected representatives, the judiciary, law enforcement and intelligence agencies, business, civil society and the Internet technical community, with the goal of restoring trust and enhancing confidence in the Internet. It is now essential that governments, collaborating with all other stakeholders, take steps to build confidence that the right to privacy of all people is respected on the Internet. This statement provides the Commission's view of the issues at stake and describes in greater detail the core elements that are essential to achieving a social compact for digital privacy and security.

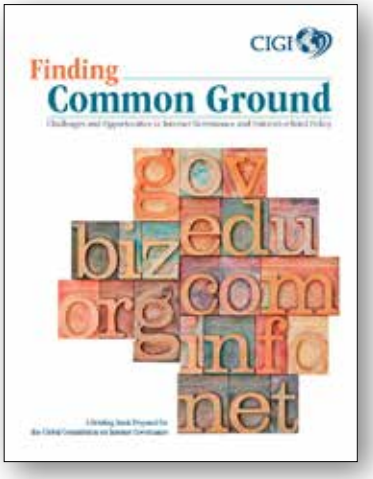

Finding Common Ground

A Briefing Book Prepared for the Global Commission on Internet Governance

This briefing book contextualizes the current debate on the many challenges involved in Internet governance. These include: managing systemic risk — norms of state conduct, cybercrime and surveillance, as well as infrastructure protection and risk management; interconnection and economic development; and ensuring rights online — such as technological neutrality for human rights, privacy, the right to be forgotten and the right to Internet access.

ABOUT CIGI

The Centre for International Governance Innovation is an independent, non-partisan think tank on international governance. Led by experienced practitioners and distinguished academics, CIGI supports research, forms networks, advances policy debate and generates ideas for multilateral governance improvements. Conducting an active agenda of research, events and publications, CIGI's interdisciplinary work includes collaboration with policy, business and academic communities around the world.

CIGI's current research programs focus on three themes: the global economy; global security & politics; and international law.

CIGI was founded in 2001 by Jim Balsillie, then co-CEO of Research In Motion (BlackBerry), and collaborates with and gratefully acknowledges support from a number of strategic partners, in particular the Government of Canada and the Government of Ontario.

Le CIGI a été fondé en 2001 par Jim Balsillie, qui était alors co-chef de la direction de Research In Motion (BlackBerry). Il collabore avec de nombreux partenaires stratégiques et exprime sa reconnaissance du soutien reçu de ceux-ci, notamment de l'appui reçu du gouvernement du Canada et de celui du gouvernement de l'Ontario.

For more information, please visit www.cigionline.org.

ABOUT CHATHAM HOUSE

Chatham House, the Royal Institute of International Affairs, is based in London. Chatham House's mission is to be a world-leading source of independent analysis, informed debate and influential ideas on how to build a prosperous and secure world for all. The institute: engages governments, the private sector, civil society and its members in open debates and confidential discussions about significant developments in international affairs; produces independent and rigorous analysis of critical global, regional and country-specific challenges and opportunities; and offers new ideas to decision-makers and -shapers on how these could best be tackled from the near- to the long-term. For more information, please visit: www.chathamhouse.org.

CIGI MASTHEAD

Executive

President	Rohinton P. Medhora
Director of the International Law Research Program	Oonagh Fitzgerald
Director of the Global Security & Politics Program	Fen Osler Hampson
Director of Human Resources	Susan Hirst
Director of the Global Economy Program	Domenico Lombardi
Vice President of Finance	Mark Menard
Chief of Staff and General Counsel	Aaron Shull

Publications

Managing Editor, Publications	Carol Bonnett
Publications Editor	Jennifer Goyder
Publications Editor	Patricia Holmes
Publications Editor	Nicole Langlois
Publications Editor	Kristen Scott Ndiaye
Publications Editor	Lynn Schellenberg
Graphic Designer	Sara Moore
Graphic Designer	Melodie Wakefield

Communications

For media enquiries, please contact communications@cigionline.org.

CIGI
67 Erb Street West
Waterloo, Ontario N2L 6C2
tel +1 519 885 2444 fax +1 519 885 5450
www.cigionline.org

CHATHAM HOUSE
The Royal Institute of International Affairs
10 St James's Square
London, England SW1Y 4LE, United Kingdom
tel +44 (0)20 7957 5700 fax +44 (0)20 7957 5710
www.chathamhouse.org

FSC www.fsc.org
RECYCLED
Paper made from recycled material
FSC® C023070