

Key Dialectics In Cloud Services

Brandt Dainow

Department of Computer Science, Maynooth University

Kildare, Co. Kildare

Ireland

brandt.dainow@nuim.ie

This paper explores dialectics within debates over key ethical issues pertaining to cloud services. These issues concern privacy, responsibility for the actions of systems, and the development of monopoly service providers. Between them these concerns largely dictate the shape and capabilities of current and future cloud-based services. We shall show how the current state of affairs is dominated by a sense of lack of agency in terms of doing things differently from the current reflexive practice, an assumption that no alternatives to current practice are possible. This paper will attempt to organise the key concerns with cloud services by organising them into three dialectical axes:

- The nature of the relationship between personal privacy and service provision.
- The degree to which people who build or operate cloud-based services are ethically responsible for the actions or effects of those services.
- The nature of the marketplace for those services.

Since this paper considers cloud services in the broadest sense, it is appropriate to commence with the definition of cloud computing used within this paper. The *US National Institute of Standards and Technology Special Publication 800-145* defines the essential characteristics of cloud computing as being:

- The ability to provide services whenever desired without human intervention.
- Being available to a wide range of client devices via networking technology.
- The “virtualisation” of computing resources, such that digital operations are not linked to specific servers or locations.
- Scalability – the capability of the systems to scale up or down in response to changes in demand (a necessary corollary of virtualisation).
- Often, but not necessarily, Software as a Service. [29]

Clearly this definition applies to many, if not most, internet systems and digital services, not merely to the virtualisation of server functions previously found in the traditional client-server network. Under this view, Facebook and Google search are both cloud services. I think this is both valid and important - confining discussion of cloud computing to data processing or file storage functions limits discussion to a few contingent uses of a wider system and obscures the essential factors we need to consider.

PRIVACY VERSUS SECURITY

Our first axis is the necessity versus the contingency of reductions to privacy under new digital services. That is to say, there is one body of opinion which holds that the erosion of personal privacy is a necessary and unavoidable consequence of, or precondition for, the delivery of digital services. These positions tend to be a reflexive response within the development community, rarely stated formally, and is a minority view in the literature, as a result of which detailed arguments as to why privacy *must* be reduced to enable cloud services are scarce. However, Professor Bergkamp’s paper, *The Privacy Fallacy* [5] marshals all the arguments in this camp.

Professor Bergkamp argues there should be no privacy protection of any form because preservation of personal privacy is harmful to society in many ways. He provides five main arguments; there is no need for data privacy, data protection reduces individual freedom, personal privacy is contrary to economic growth, EU data legislation is unenforceable and the EU’s data protection regimes put it out of step with the rest of the planet. I will now explore each of these in more depth:

Bergkamp argues there is no need for data protection or digital privacy because no one wants it and it serves no purpose. He states there is no evidence anyone has ever been harmed by privacy violations or personalization of services based on personal data. He does not provide any evidence for this and it is contradictory to the reported activity of many data protection authorities. For example, in 2014 the Data Commissioner of Ireland received 2,264 data breach notifications, investigated 960 complaints and launched 162 prosecutions. Half (53%) of complaints involved disclosing personal data inappropriately, such as disclosure of personal financial data to relatives or the listing of email addresses and passwords on public websites [34]. The *Verizon 2014 Data Breach Investigations Report* [46] covers 63,000 data violations across 93 countries in 2014. It highlights financial theft and the cost of dealing with a breach, such as cancelling credit cards, as the main harms to the individual. Other research exists to show harm from less obvious privacy violations. *RT@Iwantprivacy: Widespread violation of privacy settings in the Twitter social network* details harm from privacy violations in Twitter when people reuse private tweets in public [28]. *Privacy Violations Using Microtargeted Ads* [21] examines harm from privacy violations in Facebook. Privacy violations has also been shown to harm the companies themselves. *How Privacy Flaws Affect Consumer Perception* [2] shows how privacy breaches reduce the chance people will buy from a company, while *Is There a Cost to Privacy Breaches?* [1] shows how privacy violations reduce a company’s share price. Studies also exist to show harm from personalization of advertising and news. The research findings of Sweeney’s *Discrimination in Online Ad Delivery* [42] show how racial stereotyping in ad personalization harms afro-americans in many ways, including job prospects and access to financial services. *Bursting Your Filter Bubble* [38] shows harm from news personalization, while the famous Facebook news manipulation study, *Experimental Evidence of Massive-scale Emotional Contagion through Social Networks* [22] shows how personalizing news feeds to contain more negative contents can depress people. Bergkamp’s proposition that there has never been any harm from privacy violations or personalization appears to be contradicted by such evidence.

Bergkamp also argues there is no need for data protection because no one wants it. He argues that people don't realise that data protection prevents personalization, but that when they do, they always prefer personalization over data protection. He does not cite any evidence for this. By contrast, Culnan's 1993 study of personalization in shopping, *How Did They Get My Name?* [14] shows that when offered the choice, the people he surveyed preferred privacy over personalisation. More recently, the *2013 Comres Big Brother Watch Survey* [11] polled 10,000 people in nine EU countries to find 75% were concerned about privacy and wanted data protection regulations, while 45% believed they were being harmed by corporate data practices.

Bergkamp also argues there is no need to regulate sale of personal data because companies never sell it. However, there is, in fact, a huge industry in the sale and aggregation of personal data, as the 2014 Federal Trade Commission's investigation into data brokers found [7,17].

Bergkamp argues that personalization results in cheaper prices. However, he does not cite any empirical evidence for this or reasons why it should be so. He cites as evidence a statement made by Fred Cate, Professor of Law at Indiana University, that personalization results in cheaper prices, but this was a statement made to a Congressional committee, not a research finding. Prof. Cate's own list of publications does not include any research into personalization, his speciality is data protection law. Later in the paper Bergkamp states that data protection costs money, and that it is so burdensome and expensive that businesses can only survive by ignoring their legal obligations. One may imagine he believes this is the cause of higher prices to consumers, though he does not say so. However, research like Sweeney's *Discrimination in Online Ad Delivery* [42] shows how personalization actually increases costs to afro-american consumers in the USA, while Turow's *The Daily You: How the New Advertising Industry is Defining Your Identity and Your World* [44] shows how personalization can reduce or increase prices, depending on whether you are the consumer companies want or not.

Bergkamp also argues that privacy protection increases identity theft because data protection makes it harder to tell if someone really is who they claim to be. He does not cite any evidence for this and it seems counter-intuitive. Given that privacy protection reduces access to the personal data necessary for identity theft, such protection could be presumed to make it harder to commit, so one could argue the exact opposite of Bergkamp in the absence of any research. Bergkamp's position here allies with his arguments elsewhere in his paper that we all need to know as much as possible about each other in order to protect ourselves from one another, and that privacy directly prevents this. He states that one problem with privacy protection is that it allows an individual to control what they disclose to the world. He does not explicitly say this is a bad thing, but it is clearly implied from his useage. Here it is worth noting research showing the reverse, that lack of privacy restricts human freedom. For example, knowledge one is being watched on the internet has been shown to have a chilling effect on what people say [4] and what they search for [25], even when engaging in legal and socially acceptable activity.

Bergkamp claims there is a vast amount of money to be made acquiring and selling personal data, despite his earlier claim that there are no businesses selling it. He provides no evidence for this economic activity, but the claim is supported elsewhere. For example, in 2013 the OECD estimated the personal data of each Facebook user to range from \$US40/year to \$US400 [33]. Bergkamp claims that this data market alone is sufficient reason to remove privacy protections. However, the mere presence of economic activity does not, in and of itself, mean we should encourage it. There is a vast amount of money to be made in drug smuggling, but no one uses that as an argument for encouraging it.

Bergkamp also states that the EU's data legislation is unenforceable. He says the very concept of personal privacy is too vague to support regulation and that the regulations cannot properly specify what constitutes personal data. Furthermore, he says, each privacy incident must be judged on its own merits. He does not explain how judging a case on its own merits is a problem. Each and every infraction of the law is judged individually, so arguing that this is also the case for privacy issues does not, in and of itself, constitute a sign of poor legislation. Furthermore, it is difficult to imagine what the alternative would be if a regulator or judge was not allowed to consider the specific details of each case they were trying to adjudicate.

As stated earlier, Bergkamp believes that data protection is so onerous that no business can do it properly and survive financially. He claims the only outcome is that data regulations are never enforced. Clearly the many cases of prosecution for privacy violations are not accounted for in this argument. Bergkamp also states that EU data protection legislation is founded on a misunderstanding of how business works, but does not provide any further details regarding the nature of the misunderstanding or what the reality truly is, so this statement is impossible to assess.

Bergkamp's paper also states privacy protection damages society because it involves the government paternalistically interfering in people's relations with each other in a misguided attempt to stop people hurting each other. Such an argument can also be said of laws against violence and theft, so the logical consequence of such a position is that we should move to a state of complete anarchy. However, Bergkamp does not address this implication. Instead he goes on to state that government's should never restrict any information under any circumstances. Again no reasons are provided to justify this proposition. Such a broad statement can also be used as an argument in favour of making child pornography freely available, so some additional clarification would seem appropriate.

Finally, Bergkamp claims that EU data protection legislation is out of step with the rest of the world. He does not provide any evidence to support this, but he clearly thinks this is a bad thing and grounds for abandoning data protection. This is a questionable claim. The EU's data protection regime was intentionally built to accord with pre-existing OECD guidelines, which were first developed in 1980 [32].

Bergkamp never states it is technically impossible to maintain privacy while extending cloud services. His arguments are merely that we should not. My position is that there is no necessary and unavoidable relationship between privacy consequences and functionality. One does not *have* to reduce privacy in order to extend services. Rather, it is always a question of choice, either in how the system is constructed or in the type of business model under which it operates, and there are *always* alternatives. It may be that some of those alternatives are more expensive than the privacy-reducing models, or that alternatives are more technically challenging. However, that, in and of itself, is not an argument for the necessity of privacy-reducing models, but rather an argument underpinning a particular business model or software approach.

Currently those who are building cloud-based services most commonly work on the basis that privacy is exchanged for digital services. However, there is also a growing body of those seeking to develop alternatives, in terms of governance or business model or in terms of code. The most notable is the Privacy by Design movement. However, most of the Privacy by Design material is so vague as to be little more than statements of intent. For example, IBM claim to have moved to Privacy by Design by doing nothing more than implementing awareness training and building an internal system for reporting data breaches [35]. Here Langheinrich's paper, *Principles of Privacy-Aware Ubiquitous Systems* [24] stands out as the exception, being a concrete statement of specific technical

design principles which genuinely do embed privacy considerations into the technical architecture. Langheinrich's paper shows it is possible to build robust systems which have privacy protection embedded within the design and operation of the system.

It is notable that Langheinrich has practical experience in the design of privacy systems, being one of the authors of the W3C's technical standard, *Platform for Privacy Preferences*, or PPP [13]. The PPP standard enables browsers to hold the user's preferences for what data they will allow a website to gather. The server component of PPP allows the web server to list its own data-gathering practices. PPP then enables the browser to compare the web server's practices with the user's preferences. The system provides for warnings to the user and for compact and rapid communication between client and server of data practices. The system was supported in Microsoft Internet Explorer 6 when it first emerged, but lack of support by website owners means PPP is largely unused today.

The first of Lanheinrich's principles is the Principle of Openness, or "Notice." This simply states that no device or service should gather data about someone without telling them. Here he makes reference to PPP as providing a digital vocabulary which could be used to programmatically describe what data is being gathered, for what purpose and by whom. This is paired with the second principle, the Principle of Consent, which encodes the legal necessity for informed consent. A system must allow for someone to opt out of being tracked or recorded, and do so without denying service on a "take it or leave it" basis. Thus, for example, buildings would need to disable tracking for some people and not simply refuse them entry.

The third principle is termed "Anonymity and Pseudonymity." This states that people must have the option to remain anonymous. The issue here is that some services are only possible if they know a user's identity and history. Here Lanheinrich introduces pseudonymity. Under this system a person may have a unique identifier of some form, such as a cookie or RFID chip, which anchors the data systems and forms the index key to their personal data history. However, this identifier contains no personally identifiable information and is discardable at any time. Furthermore, such a system permits people to have multiple pseudonymous ID's and so prevent aggregation of disparate activities by data brokers. It is noteworthy that EU data regulations have recently been updated to add the category of pseudonymous identity between personal and anonymous data [48].

Lanheinrich's fourth principle of "Proximity and Locality" limits the scope of data collection. Looking to a future in which people have many devices capable of recording their surroundings, the principle of proximity states that these devices can only operate in the proximity of their owner. This prevents people leaving devices to record data unseen, then returning for them later. Of wider application is the principle of locality; devices should not transmit data any further than absolutely necessary to fulfil their functions. For example, Samsung's voice activated TV's transmit all conversations they hear to Samsung's central servers. Voice commands are interpreted there and the appropriate command then sent back to the TV. All conversation recordings are stored permanently for later analysis [49]. Under Lanheinrich's principles the TV would have been designed so that it did not need to involve cloud services. Voice recognition chips have been around for 20 years and could easily have been used instead.

The fifth principle is the "Need for Security," in which Lanheinrich advocates various levels of security depending on the nature of the data. More importantly, he illustrates how the previous principles themselves enhance security. If data is not being transmitted many security problems simply vanish. Similarly, if data is not linked to an identifiable individual, but only to a pseudonymous ID, unauthorised access has less potential for harm.

Lanheinrich's final principles are the principles of "Collection and Use Limitation." These state that data collectors should only collect data for a specific purpose and not store it, as Samsung TV does, in case they want to use it in the future. Secondly, they should only collect the data they need in order to fulfil their task and nothing more. Finally, they should only keep data as long as it is necessary for the purpose. While these appear primarily legislative principles, they can be embodied in technical design through the use of the earlier principles. For example, if data is housed in the user's devices in accordance with the principle of locality, then the user can impose usage and storage limitations themselves.

Lanheinrich's principles, if implemented, would solve many privacy concerns, enhance security and actually make many applications of ubiquitous and cloud services easier to construct. What they show is that it is perfectly possible to design cloud services in a manner which enhances both security and privacy at the same time, while permitting all the personalization necessary. They place control of personal data firmly in the hands of the user without compromising technical operations in any way. In fact, their reduced dependence on permanent access to centralised services makes them more robust and reduces the burden of traffic on the internet. These principles are easy to understand and yet produce powerful architectures. They offer a practical and detailed response to the reflexive position that personalized cloud services must reduce privacy. In doing so they provide concrete evidence that it would be possible to move cloud service evolution into a path which fulfils all its potential, yet enhances privacy and security at the same time. Lanheinrich's design principles demonstrate that the reduction of privacy in cloud services is a choice, not a necessity.

ETHICAL RESPONSIBILITY

Our second axis is concerned with the degree to which people who design, build or operate cloud services are ethically responsible for the consequences of the actions of those systems. This question does not arise with regard to all cloud services, but only with the rising generation of autonomous services which process personal data in order to deliver personalised services, such as personalised search results, product recommendations and news feeds. In the near future we will see the rise of more intelligent and more life-critical personalised services, most notably with bio-implantation and other medical services [19]. The question is primarily one of who is responsible when such autonomous services make decisions which result in harm, but where these decisions are not the result of faulty design or incorrect data.

The competing positions are that, on the one hand, programmers and operators are not ethically responsible for the actions of autonomous systems, versus a view that they are. It is difficult to argue that the person holding a hammer is not ethically responsible for the consequences of whatever happens when the hammer hits something because the hammer is totally under the control of the user. However, with large industrially-produced complex automated systems, especially those that include some form of AI functionality, arguments emerge in favour of the position that those who build the systems are not ethically responsible for the decisions those systems make. This argument will no doubt be exacerbated the more powerful and the more intelligent and autonomous these systems become. This issue is discussed most frequently with regard to autonomous military systems, whose lethality makes the question of ethical responsibility both stark and urgent. However, the question is just as pertinent for any form of autonomous system, including those cloud-based personalization systems already in operation.

Andreas Matthias' paper, *The Responsibility Gap* [26] offers a fairly straightforward account of the philosophical logic behind the position that programmers are *not* responsible for the

actions of their autonomous systems, while Robert Sparrow's *Killer Robots* [39] presents the same view via an examination of the practicalities of creating and deploying autonomous systems. Both take the position that no one at all is ethically responsible for the actions of autonomous agents.

Sparrow's argument is based on his particular understandings of the terms 'autonomy' and 'responsibility.' My view is that he defines these terms in such a way as to make any contrary conclusion impossible. Early in the paper, he defines autonomy as being free from external causation:

"Where an agent acts autonomously, then, it is not possible to hold anyone else responsible for its actions. In so far as the agent's actions were its own and stemmed from its own ends, others can not be held responsible for them. Conversely, if we hold anyone else responsible for the actions of an agent, we must hold that, in relation to those acts at least, they were not autonomous." [39:65–66]

Sparrow does not defend this definition of autonomy. However, once it has been defined this way, it becomes a matter of logical necessity that there is no ethical responsibility by the programmers or controllers. It is also worth noting that Sparrow uses autonomy in an absolute sense, as if the agent were free from all influence except their prior experience. In particular, he does not recognise the environment, the capabilities of the device or its internal structures as having any impact on decision-making. He argues the programmer cannot be responsible because the essence of an autonomous system is that it will make unpredictable decisions. He argues the controller of the system is not responsible because they could not anticipate what it would do any better than the programmer. In both cases, he ignores the fact the system is designed to perform a particular role in a particular environment. A software agent is not free to do just anything, it can only recognise inputs of a type it has been designed for, and has a relatively limited range of actions it can take, and can only operate in a specific type of environment. A share-dealing system cannot walk the dog or assess your exercise regime. The type of decisions an autonomous system may make, and the range of options available to it, are not only predictable, they are the basis upon which it was designed and built - they define it. An autonomous system may make its own decisions, even alter its own programming, but its range of actions and the forms of harm it may commit are knowable in advance in virtue of the type of system it is.

Sparrow does not mention Strawson in his paper, but his conception of moral responsibility has close parallels to Strawson's influential work. Strawson's position is that no one is morally responsible for anything because no one is free from external influence [41], though the details of why are beyond the scope of this paper. Though Sparrow does not say so explicitly, his use of responsibility is clearly that one can only be responsible for specific actions. In Sparrow's view, the design of the system and the decision to use it do not carry any ethical responsibility because neither gives one the ability to predict the specifics of an individual act the system may take.

Sparrow also argues it is not possible to hold the system itself responsible because responsibility necessarily requires punishability which requires suffering. Under his definitions, something can only be morally responsible if it can be punished and something can only be punished if it can suffer. Since software systems cannot be made to suffer, they cannot be punished and so cannot be held responsible for their actions. Note here that we have switched from talk of "being responsible" to talk of "being held responsible." This demonstrates that Sparrow has conflated the moral state of being responsible with the social status of being eligible for punishment.

Matthias's *The Responsibility Gap* also argues that no one is responsible for the decisions of autonomous software

systems. His position also links responsibility to individual acts, holding that one can only be responsible if one can know the internal state of the system *and* has control of each act it takes, at least to the degree where one could prevent it. Under this analysis a programmer has no responsibility for the actions of a system once the owner takes control. The owner is not responsible because they cannot know the internal state of the system. Matthias spends some time examining different types of AI learning, showing how each makes their internal state unknowable in different ways, but the differences do not affect his final conclusion.

The narrow understanding of responsibility seen in Matthias and Sparrow is the foundation on which their arguments rest. In contrast, Miller's *Collective Responsibility and Information and Communication Technology* [30] confronts this issue by arguing there are different types of responsibility. In addition to the responsibility for individual acts which Matthias and Sparrow focus on, Miller points out we also recognise one can have "structural" responsibility by creating the conditions which made the act possible or by ordering others to take actions which eventually led to the act. Under Miller's analysis both programmers and controllers of autonomous systems take structural responsibility for every act taken by these systems. Miller then goes deeper, investigating the concept of collective responsibility. He argues that to the degree that individuals contribute something to the shape and operation of an autonomous system, so they share in responsibility for its actions. Here he acknowledges the existence of corporate responsibility, but argues that it does not provide a moral shield for the individual workers, whose individual contributions to a system's operation convey a share in collective responsibility for its actions.

Miller's approach is a step towards recognition that software agents exist within the wider context of human activity. This broader perspective is fully achieved in *Software Agents, Anticipatory Ethics, and Accountability* [20] by Johnson et. al. Reiterating the perspective that technology is socially situated, this paper argues that the concept of any digital service as an autonomous agent is merely metaphorical; that no such system can be autonomous in the sense we apply autonomy to humans in moral debates. As such, the use of the metaphor is justifiable only by its utility. Johnson et. al. criticise the concept of any software as an autonomous agent on the grounds it generates just these ethical problems. Instead, Johnson et. al. argue we should recognise autonomous systems as elements within a larger socio-technical system, made by people and used by people for human purposes. Under this view autonomous systems are not independent entities hermetically sealed from their environments, but systems which can only be understood by reference to the context of their use. Johnson et. al. make implied use of the different forms of responsibility seen in Miller, but do not elucidate the differences. Instead they focus on the arbitrariness of delimiting technical artefacts. They deny that autonomous software agents are different in kind from any other form of automated or semi-automated device, being merely more complicated. Autonomous systems are thus merely, like hammers, extensions of human will and intent. Under this arrangement, ethical responsibility for their actions is not in any way changed by the mere fact of their complexity.

OPEN VERSUS CLOSED

Our final dialectic concerns the form and marketplace of cloud services. Here the dominating dialectic is that of open versus closed systems and open versus closed organisational contexts for such systems.

The scene for this debate is best set Eben Moglen in his presentation, *Freedom in the Cloud*, delivered to the Internet Society in 2010 [31]. Moglen argues that the internet was originally designed as a non-hierarchical peer-to-peer network. However, under the influence of the architectural model of client-

server networking, the services which evolved used a smart-server-dumb-client model, in which both algorithms and data were centralised. Moglen maintains that cloud architecture works on this thin client - fat server model and does not represent a new computing architecture, merely the virtualisation of some server operations within this traditional model. These servers maintain activity logs. These logs can be mined for behavioural data. Marketing companies learned they could mine these logs to understand, predict and influence user behaviour in order to sell advertising. Moglen contends that as the perceived value of this information grew, it spurred the development of a secondary internet infrastructure of tracking services designed to add to the growing database of what we now call "user profiles."

Thus Moglen describes how an architecture which concentrates processing power and data at centralised locations promotes a concentration of both technical proficiency and economic power, while also promoting a top-down hierarchical organisational model and, in a global internet, the development of a limited number of very large monopoly service providers. This has, he argues, produced an extreme power dichotomy between those who own the services and those who use them. The business model which has come to dominate the internet is that of delivering services in exchange for spying on the users all the time. Moglen describes this state of affairs as undesirable for two reasons. Firstly, the price is too high and the services are not worth the loss of privacy. Secondly, the lack of alternative models for access to the same services makes this unfair arrangement unavoidable. He argues that we need an alternative architecture in which the data about us stored on centralised servers is instead housed in devices we own and carry with us. We can then control who accesses this data and how. He argues that this is possible with current technology.

There is an additional element of concern within Moglen's model which he hints at but does not explore. The combination of architecture and business model he describes has produced "walled gardens." These are silos of private technology and proprietary data formats which are not compatible with, or accessible by, other systems or organisations. The patent system combines with a capitalist marketplace to financially reward such behaviour. If I am the sole owner of a system everyone wants to use, I can make money. If I create a system which I give away, I do not benefit. What I therefore need to do is lock everyone into my technology, and then I will "lock in" the market [12].

The effect of this is to lock data and services into a single monopoly provider. The provider becomes the gatekeeper over the knowledge of what they do and how they do it. Users cannot migrate to a competitor without significant effort and loss. For example, if you close your account with Amazon, they will remove all the books from your Kindle [50]. You cannot therefore switch to an alternative, such as Adobe Digital Editions, without re-purchasing your entire digital library. Different legal regimes permit different levels of access inside these walled gardens, but in no case does a society have full knowledge or any substantive control. Such a system has no interest in open standards, interoperability, or a free flow of information. This lack of interoperability and open standards was why the internet and HTML were not developed by commercial enterprises. Early pre-cursors of the web tried the same walled garden approach, including America Online, CompuServe and Lotus Notes. It was only when Tim Berners-Lee gave HTML away that we broke free of this limiting system and gained the web. Berners-Lee gave it away because he saw things in exactly this way and believed that if he patented or sold HTML, it would become just another walled garden [6].

However, as companies have developed services which sit atop these communally-owned standards, so they have developed further proprietary systems. The final result is that companies have built a new layer of walled gardens and data

silos on top of the open platform which is the internet [16]. The scale of the internet user base combines with a shared service delivery infrastructure to enable the rise of extremely large global monopolies, such as Google, Amazon and Facebook. The result is that cloud services are portioned out amongst a limited number of very large hierarchical organisations, each of which hides its use of data from public scrutiny and uses its monopoly position and ownership of personal data as a competitive advantage [8,15,27]. The net effect is that people are locked to service providers like serfs to their lord. However, unlike in the Middle Ages, there is no competing lord to flee to if you are unhappy with your lot. This power is a concern to many. Some argue, for example, that Amazon's potential to control what books are available makes it a political institution as well as an economic one [10], while Google is has consistently been one of the biggest spenders on lobbying in Washington since 2012 [18].

Opposing this state of affairs are a disparate range of alternatives, such as Moglen and his concept of a personal server. Each alternative tends to focus on one aspect of this system, such as technical architecture or business model. Technically, the existence of the internet is based on open standards, such as TCP and IP [40], so alternatives have always been available on a technical level. Here we have the open source activists, such as the Free Software Foundation and the IETF. In addition, we have less obvious alternative architectures based on peer-to-peer (as opposed to client-server) models, such as the BOINC platform for community computing [3] and the BitTorrent protocol [36]. Standards like XML [47] and RDF [37] provide a means of breaking open walled gardens through data exchange, while people such as Chris Marsden in the UK or Robert McChesney in the USA have developed the rationale for breaking down these proprietary data silos.

McChesney argues that the development of monopolies and cartels has so dominated the internet that there has been little economic benefit for the rest of society. He argues that there is so little competition at the point of delivery that service providers constitute a cartel which should be forced into competition with not-for-profit public alternatives. He calls for the monopolistic corporations dominating important services, like Facebook and Google, to be broken into smaller competing units and subject to much more stringent and detailed state control. McChesney's argument is that the size of these corporations is so great they pose a threat to democracy itself through their power to lobby politicians, dominate online debate and skew economic development [27].

Concern over monopoly domination is addressed in a different manner by Brown and Marsden in a number of publications. Instead of seeking a solution by changing the economic structure, they focus on the proprietary data structures which form the foundation of such domination. In addition to rights such as the right to have one's records deleted, they argue for the right to move such data to an alternative provider of the same service [9]. They cite similar historical examples in which Microsoft, IBM and Intel have been forced into making their systems interoperable with competitors, mainly through antitrust approaches in the USA and EU [8]. They argue that state intervention to break up these monopolies is not practical in a world dominated by competing national legislative regimes. Instead, they argue that merely providing users the ability to switch to alternatives would be sufficient. They believe that this would stimulate the development of service providers offering a range of alternative models [10].

The approach of treating monopoly service providers as public utilities is gaining ground in government circles. Recently the UK House of Lords called for the internet to be treated like a public utility rather than a market place of optional luxuries [43]. International bodies, such as the EU and UNESCO, have started calling for wider civic involvement in determining how services are provided [23,45] and for the

development of alternative service provision models. For example, the outgoing EU Vice President, Neelie Kroes, stated in November 2014:

"Why should we have to give up our privacy for a "free" service if we prefer to pay for that same service with cash and keep our privacy?" [23]

CONCLUSIONS - AGENCY

While these three dialectics focus on different issues, the poles of each axis rest on competing perspectives on the possibility of agency. Those who accept things as they are now do not see a possibility for agency, while their opponents do. On the first axis we have those who hold that preservation of privacy and delivery of service are necessarily in opposition. Here they are holding that there is no possibility of agency in the relationship between privacy and service design. To the contrary, we have seen how Langheinrich's design principles show multiple opportunities to intervene in the ways which deliver services while also maintaining privacy. In our second axis of ethical responsibility, the position of there being no ethical connection between the creator of an autonomous

system and that system's effects is also a position of there being no agency. Here lack of agency pertains not to the nature of the system, but to the consequences of the system's actions. Under this view, once an autonomous system is activated, human agency ceases. However, as we have seen, preserving a lack of responsibility requires limiting the conception of where agency lies. Responsibility has to be defined in a very constricted manner which focuses on the making of each individual decision and denies the influence of any context. Instead, services are treated as independent of any human agency, in terms of their design, their environment, their purpose, how they are used and who benefits. By contrast, once autonomous services are contextualised within a field of human practice, human agency becomes apparent throughout the construction and operation of such systems and human ethical responsibility becomes self-evident. Finally, in our third axis of service architecture and business model, we see a historical lack of agency in the development of the broader internet culture. Here the client-server structure was accepted reflexively by developers and users, along with the most obvious reflections of this in organisational and economic models.

In all three debates we see one pole in each dialectic disempowering itself, primarily because it simply fails to recognise that there is a choice and that agency, the power to act differently, exists. The conclusion which emerges from this is that a key step to improving the current ethical status of cloud services is inculcating in programmers and leaders that they possess agency, bringing them to recognise there are alternatives and that they have the power to explore them.

REFERENCES

1. Alessandro Acquisti, Allan Friedman, and Rahul Telang. 2006. Is there a cost to privacy breaches? An event study. *Proceedings of the Twenty-Seventh International Conference on Information Systems (citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.73.2942&rep=rep1&type=pdf)*.
2. Sadia Afroz, Aylin Caliskan Islam, Jordan Santell, Aaron Chapin, and Rachel Greenstadt. 2013. How Privacy Flaws Affect Consumer Perception. *IEEE*, 10–17. <http://doi.org/10.1109/STAST.2013.13>
3. D.P. Anderson and G. Fedak. 2006. The Computational and Storage Potential of Volunteer Computing. *CCGRID '06 Proceedings of the Sixth IEEE International Symposium on Cluster Computing and the Grid*, IEEE Computer Society, 73–80. <http://doi.org/10.1109/CCGRID.2006.101>
4. Katy Glenn Bass. 2013. *Chilling Effects: NSA Surveillance Drives U.S. Writers to Self-Censor*. PEN American Center. Retrieved from http://www.pen.org/sites/default/files/Chilling%20Effects_PEN%20American.pdf
5. Lucas Bergkamp. 2002. The Privacy Fallacy. *Computer Law & Security Report* 18, 1, 31–47.
6. Tim Berners-Lee. 1999. *Weaving the Web*. HarperCollins, New York, N.Y.
7. Nathan Brooks. 2005. *Data Brokers: Background and Industry Overview*. Congressional Research Service, The Library of Congress, Washington, D.C.
8. Ian Brown and Christopher T. Marsden. 2013. *Regulating code: good governance and better regulation in the information age*. The MIT Press, Cambridge, Mass.
9. Ian Brown and Christopher T. Marsden. 2013. Regulating Code: Towards Prosumer Law? *SSRN Electronic Journal*. <http://doi.org/10.2139/ssrn.2224263>
10. Ian Brown and Christopher T. Marsden. 2013. Interoperability as a standard-based ICT competition remedy. *8th International Conference on Standardization and Innovation in Information Technology 2013*, IEEE, 1–8. <http://doi.org/10.1109/SIIT.2013.6774570>
11. Comres. 2013. *Big Brother Watch Online Survey*. Comres, London.
12. Eric P. Crampton and Donald J. Boudreaux. 2003. Does Cyberspace Need Antitrust? In *Who Rules the Net?*, Adam D Thierer and Clyde Wayne Jr. Crews (eds.). Cato Institute, Washington, D.C.
13. Lorrie Cranor, Marc Langheinrich, Massimo Marchiori, Martin Presler-Marshall, and Joseph Reagle. 2002. *The Platform for Privacy Preferences 1.0 (P3P1.0) Specification*. W3C. Retrieved from <http://www.w3.org/TR/P3P/>
14. Mary J. Culnan. 1993. “How Did They Get My Name?”: An Exploratory Investigation of Consumer Attitudes toward Secondary Information Use. *MIS Quarterly* 17, 3, pp. 341–363.
15. James Curran, Natalie Fenton, and Des Freedman. 2012. *Misunderstanding the Internet*. Routledge, London.
16. Tony Dyhouse. 2010. Addressing the silo mentality. *Infosecurity* 7, 2, 43. [http://doi.org/10.1016/S1754-4548\(10\)70043-7](http://doi.org/10.1016/S1754-4548(10)70043-7)
17. Federal Trade Commission. 2014. *Data Brokers: A Call for Transparency*. Federal Trade Commission.
18. Tom Hamburger and Matea Gold. 2014. Google, once disdainful of lobbying, now a master of Washington influence. *The Washington Post*. Retrieved June 23, 2015 from http://www.washingtonpost.com/politics/how-google-is-transforming-power-and-politicsgoogle-once-disdainful-of-lobbying-now-a-master-of-washington-influence/2014/04/12/51648b92-b4d3-11e3-8cb6-284052554d74_story.html
19. Veikko Ikonen, Minni Kanerva, Panu Kouri, Bernd Stahl, and Kutoma Wakunuma. 2010. *D.I.2. Emerging Technologies Report*. ETICA Project.
20. Deborah G. Johnson. 2011. Software Agents, Anticipatory Ethics, and Accountability. In *The Growing Gap Between Emerging Technologies and Legal-Ethical Oversight*, Gary E. Marchant, Braden R. Allenby and Joseph R. Herkert (eds.). Springer Netherlands, Dordrecht, 61–76. Retrieved February 17, 2015 from http://www.springerlink.com/index/10.1007/978-94-007-1356-7_5
21. Aleksandra Korolova. 2011. Privacy Violations Using Microtargeted Ads: A Case Study. *Journal of Privacy and Confidentiality* 3, 1.
22. A. D. I. Kramer, J. E. Guillory, and J. T. Hancock. 2014. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences* 111, 24, 8788–8790. <http://doi.org/10.1073/pnas.1320040111>
23. Neelie Kroes and Carl-Christian Buhr. 2014. Human society in a digital world. *Digital Minds for a New Europe*. Retrieved January 27, 2015 from https://ec.europa.eu/commission_2010-2014/kroes/en/content/human-society-digital-world-neelie-kroes-and-carl-christian-buhr
24. Marc Langheinrich. 2001. Privacy by Design - Principles of Privacy-Aware Ubiquitous Systems. *Proceedings of the Third International Conference on Ubiquitous Computing*, Springer-Verlag, 273–291. Retrieved from <http://www.vs.inf.ethz.ch/publ/papers/privacy-principles.pdf>
25. Alex Marthews and Catherine Tucker. 2014. Government Surveillance and Internet Search Behavior. *SSRN Electronic Journal*. <http://doi.org/10.2139/ssrn.2412564>
26. Andreas Matthias. 2004. The responsibility gap: Ascribing responsibility for the actions of learning automata. *Ethics and Information Technology* 6, 3, 175–183. <http://doi.org/10.1007/s10676-004-3422-1>
27. Robert W. McChesney. 2014. Be Realistic, Demand the Impossible: Three Radically Democratic Internet Policies. *Critical Studies in Media Communication* 31, 2, 92–99. <http://doi.org/10.1080/15295036.2014.913806>
28. Brendan Meeder, Jennifer Tam, Patrick Gage Kelley, and Lorrie Faith Cranor. 2010. RT@IWantPrivacy: Widespread violation of privacy

- settings in the Twitter social network. *Proceedings of the Web*, 1–12.
29. Peter Mell and Timothy Grance. 2011. *The NIST Definition of Cloud Computing*. National Institute for Standards & Technology.
 30. Seumas Miller. 2008. Collective Responsibility and Information and Communication Technology. In *Information Technology and Moral Philosophy*, Jeroen van den Hoven and John Weckert (eds.). Cambridge University Press, Cambridge; New York, 226–250.
 31. Eben Moglen. 2010. Freedom in the Cloud. Retrieved from <http://www.softwarefreedom.org/events/2010/isoc-ny/FreedomInTheCloud-transcript.html>
 32. OECD. 2013. *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*. OECD. Retrieved from <http://www.oecd.org/sti/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm>
 33. OECD. 2013. *Exploring the Economics of Personal Data*. Retrieved June 23, 2015 from http://www.oecd-ilibrary.org/science-and-technology/exploring-the-economics-of-personal-data_5k486qtxldmq-en
 34. Office of the Data Protection Commissioner. 2015. *Annual Report of the Data Protection Commissioner of Ireland 2014*. Data Protection Commission of Ireland.
 35. Office of the Information & Privacy Commissioner of Ontario and IBM. 2011. *Privacy by Design: From Theory to Practice*. Office of the Information & Privacy Commissioner of Ontario, Ontario.
 36. Johan Pouwelse, Paweł Garbacki, Dick Epema, and Henk Sips. 2005. The Bittorrent P2P File-Sharing System: Measurements and Analysis. In *Peer-to-Peer Systems IV*, Miguel Castro and Robbert van Renesse (eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 205–216. Retrieved June 23, 2015 from http://link.springer.com/10.1007/11558989_19
 37. RDF Working Group. 2015. RDF - Semantic Web Standards. *RDF - Semantic Web Standards*. Retrieved June 23, 2015 from <http://www.w3.org/RDF/>
 38. Paul Resnick, R. Kelly Garrett, Travis Kriplean, Sean A. Munson, and Natalie Jomini Stroud. 2013. Bursting your (filter) bubble: strategies for promoting diverse exposure. *Proceedings of the 2013 conference on computer supported cooperative work companion*, ACM Press, 95. <http://doi.org/10.1145/2441955.2441981>
 39. Robert Sparrow. 2007. Killer Robots. *Journal of Applied Philosophy* 24, 1, 62–77. <http://doi.org/10.1111/j.1468-5930.2007.00346.x>
 40. W. Richard Stevens and Gary R. Wright. 1994. *TCP/IP Illustrated Volume 1: The Protocols*. Addison-Wesley, Reading, Mass.
 41. Galen Strawson. 1994. The Impossibility of Moral Responsibility. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 75, 1/2, 5–24.
 42. Latanya Sweeney. 2013. Discrimination in Online Ad Delivery. *SSRN Electronic Journal*. <http://doi.org/10.2139/ssrn.2208240>
 43. The Select Committee on Digital Skills. 2015. *Make or Break: The Digital Future*. The Authority of the House of Lords, UK. Retrieved from <http://www.publications.parliament.uk/pa/ld201415/ldselect/lddigital/111/111.pdf>
 44. Joseph Turow. 2011. *The Daily You: How the New Advertising Industry is Defining Your Identity and Your World*. Yale University Press, New Haven.
 45. UNESCO Secretariat. 2014. *Internet Universality*. United Nations Educational, Scientific and Cultural Organization (UNESCO). Retrieved from http://www.unesco.org/new/fileadmin/MULTIMEDIA/HQ/CI/CI/pdf/news/internet_universality_en.pdf
 46. Verizon. 2014. *2014 Data Breach Investigations Report*. Verizon Enterprise Solutions.
 47. XML Working Group. 2015. XML Core Working Group Public Page. *XML Core Working Group Public Page*. Retrieved from <http://www.w3.org/XML/Core/#Publications>
 48. 2015. *Future-proofing Privacy*. Hogan Lovells, London.
 49. Privacy | Samsung UK. Retrieved June 18, 2015 from <http://www.samsung.com/uk/info/privacy-SmartTV.html>
 50. Amazon.com Help: Kindle Terms of Use. Retrieved June 22, 2015 from <http://www.amazon.com/gp/help/customer/display.html?nodeId=200506200>