7

Political and ethical perspectives on data obfuscation

*Finn Brunton and Helen Nissenbaum**

Asymmetries of data gathering and means of redress: the warrant for obfuscation

Our chapter, like all the others gathered in this volume, is written in light of the fact that computer-enabled data collection, aggregation and mining dramatically change the nature of contemporary surveillance. Innocuous traces of everyday life submitted to sophisticated analytics tools developed for commerce and governance can become the keys for stitching disparate databases together into unprecedented new wholes. This data is often gathered under conditions of profound power imbalance. What can we do when faced with these demands, which are often trivial but whose implications are profound, and which we may not be in a position to refuse?

Being profiled is the condition of many essential transactions, from connecting with friends in online social networks to shopping, travelling and engaging with institutions both public and private. Nor, as we shall discuss below, can we rely on law, technology or the scruples of the data gatherers. What we propose is an alternative strategy of informational self-defence, a method that acts as informational resistance, disobedience, protest or even covert sabotage – a form of redress in the absence of any other protection and defence, and one which disproportionately aids the weak against the strong. We call this method *obfuscation* and, in this chapter, we will argue for the political and ethical philosophy it expresses and embodies.

Obfuscation is the production of misleading, ambiguous and plausible but confusing information as an act of concealment or evasion. It is a term we use to capture key commonalities in systems ranging from chaff, which fills radar's sweep with potential targets; to the circulating exchanges of supermarket loyalty cards that muddle the record of purchases; to peer-to-peer file sharing systems such as BitTorrent, protecting their users from legal action by producing records of many IP addresses, only a few of which may be engaged in file sharing. Through these and other cases we can begin to clarify obfuscation among the other forms of resistance to surveillance, whether that surveillance takes the form of consumer data aggregation (for supermarkets, or by

companies such as Acxiom), monitoring for intellectual property violations (at the behest of the Recording Industry Association of America (RIAA) and the Motion Picture Association of America (MPAA)), targeted advertising (by sites such as Google and Facebook) or police actions by repressive governments (which we will see addressed by obfuscation tactics within platforms for secure private conversation such as Tor).

We distinguish and evaluate different modes of obfuscation as well as motivations and power topologies of key actors: are obfuscation tactics typically the response of the weak against the strong, adopted by those outside circles of power and influence, or vice versa? Our political analysis of obfuscation also addresses normative questions of legitimacy, asking whether such 'smokescreens' to avoid monitoring are morally defensible – ever, never or sometimes? Under what conditions in the political landscape of surveillance are obfuscation's deceptive tactics acceptable? These can be deemed legitimate assertions of autonomy or problematic instances of economic free ridership (relying on others to be less conscientious in muddying their tracks and therefore better targets); they can be hailed as resistance to inappropriate monitoring or damned as the poisoning of wells of collective data. Obfuscation, as a tactic both personal and political, offers a platform for studying legitimate and problematic aspects of both surveillance and its opposition in an age of ubiquitous data capture.

In the context of this volume, we do not need to go out of our way to describe the problematic state of contemporary data gathering and analysis, but we do need to highlight the specific problems of asymmetry these practices, as a matter of fact, often involve. The most mundane points of contact with contemporary life implicate the involuntary production of data on our part: passing security cameras, withdrawing cash, making credit card purchases, making phone calls, using transit (with electronic ticketing systems such as MetroCards, FasTrak tags, Oyster, Octopus, Suica or E-ZPass) – to say nothing of using the internet, where every click and page may be logged and analysed, explicitly providing data to the organisations on whose systems we interact, as well as their associates. This data can be repackaged and sold, collected, sorted and acquired by a variety of means, and reused for purposes of which we, the monitored, know nothing, much less endorse (Gonzalez Fuster 2009). The unreliability of the businesses and public-private partnerships in the information industry gives data mobility still more sinister dimensions, as materials are stolen, leaked, sold improperly or turned to very problematic ends by governments – ChoicePoint's sale of 145,000 records to identity thieves being one particularly egregious example.[1] The nature of these businesses, acquiring new data sets to add to their existing collections, points to a final area of concern. Multiple databases consolidated and cross-referenced, with incidental details linking previously disconnected bodies of information, produce a far more significant whole than any one part would suggest: identities, tendencies, groups and patterns with both historically revelatory and predictive power.[2]

The asymmetry problems to which we alluded above are, first, an asymmetry of power: rarely do we get to choose whether or not we are monitored, what happens to information about us and what happens to us because of this information. We have little or no say when monitoring takes place in inappropriate contexts and is shared inappropriately with inappropriate others. The second asymmetry, equally important, is epistemic: we are often not fully aware of the monitoring, and do not know what will become of the information produced by that monitoring, nor where it will go and what will be done with it.

Your data is not accumulated in neutral circumstances, whether surveillance occurs at the level of infrastructure with which you must participate, through forms that must be completed to receive essential resources, or onerous terms of service to which you must consent before you can use an online product that has become vital to doing business. The context is often one of major power imbalance, between individual consumers and major corporations, or citizens and governments. Obviously there is nothing inherently wrong with gathering and aggregating data on individuals – it is the lifeblood of the work of the epidemiologist, for example, and the starting point for many benefits of the networked society. It is in the combination of data gathering with authority and its arbitrary interests where problems may begin.

These problems continue once our data has been collected: we do not know whether whoever gathers it will repackage and resell it, whether it will become part of a schedule of assets after a bankruptcy or whether it will be collated by a private party such as ChoicePoint with public records for reassembly and used in a different context from the original point of provision. Data mining and related disciplines are complex and intellectually demanding; they often require resources of expertise, software and hardware that people outside large institutions do not possess. We do not have access to the other databases, nor the techniques and the training in mathematics and computer science, to comprehend what can be done with seemingly trivial details from our lives and activities, and how they can provide more powerful, total and revealing analyses than we could have anticipated (Reiman 1995; Solove 2008). The inconsequential and even benign can quickly become the problematic and sinister.

Furthermore, we do not know what future techniques and databases will enable. Opportunities for the correlation of information tend to increase with time. Institutions very rarely voluntarily destroy materials with as much potential as a rich database and, as Templeton (2009) points out, the mechanisms to extract value from databases are only going to get better. Materials from very different contexts, created in conditions of many different norms – telephone call logs, geolocative data, purchase records whether in person or online, demographic and personally identifying information, products of the

data generating machines that are social networking sites – can be combined, correlated and cross-referenced with less and less effort.

Finally, the gravity of the potential consequences of this mass of data is not easily perceived in the many small moments when we are faced with a decision about whether or not to comply, and give up information. The cost to any one individual at any one moment in time is generally very low, becoming apparent only in aggregate and over a longer period – at which point the moment to make a decision is already past. The disproportionate cost, at the moment when you want to join some friends on a social network, get health insurance or purchase airline tickets – or when you are obliged to provide some seemingly insignificant information while facing an asymmetry of power – does not become clear until it scales to the community and longer timescales, and this issue frames the politics of data gathering and analysis.

The lack of capacity to assess consequences that matter to us is deeply troubling. We do not know all that 'they' know about us, how 'they' come to know it or even who all the significant players might be. We cannot easily subject these players to symmetrical analysis: such organisations might operate under the veil of national security or proprietary trade secrets, and we probably would not have the methods or the training to do anything with their data if we could get our hands on it. As people whose data is being collected, what we know of the situation is problematic, and what we do not know is substantial.[3]

In theory, the ways out of our predicament of inescapable, ubiquitous, asymmetric collection and scrutiny of data are numerous and diverse, the palette of options familiar to anyone following the privacy debates: user opt-out, law, corporate best practice and technology. Each offers a prognosis for particular challenges, and each has shortcomings in relation to the asymmetries of data analysis. While useful for certain types of threats, each has not proven responsive to others, and all have particular short-term flaws, which could compound into a future that worries us. The first of these established – even reflexive – approaches is the most common counterargument to the two asymmetries, the opt-out argument, which puts the responsibility on the shoulders of individuals whose data are being gathered. The other three are classic long-term, slow-incentive structures for creating social change; their gradual pace, and investment in existing interests, makes them problematic for short-term protection and sets the stage for self-directed and individually introduced strategies such as obfuscation.

The steady rhetorical drumbeat in the discussion around data privacy is that refusal is a personal responsibility. If you are so offended by the way these companies collect and deploy your data, simply do not use their services – *opt out*. No one is forcing you. To which we reply: yes and no. Many of these systems are not mandatory yet (government systems and various forms of insurance being only two exceptions), but the social and personal cost of refusal is

already substantial and, indeed, growing. We pay by loss of utility, efficiency, connection with others in the system, capacity to fulfil work demands, and even merely being able to engage in many everyday transactions. To rely entirely on personal choice is to leave all but the most dedicated and privacy-obsessed at the mercy of more conventional means of regulation – or resistance.[4]

Why not rely on *corporate best practice?* Private sector efforts are hampered by the fact that companies, for good reasons and bad, are the major strategic beneficiaries of data mining. Whether the company is in the business of gathering, bundling and selling individual data, such as DoubleClick and ChoicePoint, or has relied on the data generated and provided by its customers to improve its operations, such as Amazon and WalMart, or is based on user data-driven advertising revenue, or subcontracts the analysis of consumer data for purposes of spotting credit, insurance or rental risks, it is not in their interest to support general restraints on access to information.

Law and regulation, historically, have been central bulwarks of personal privacy, from the Fourth Amendment of the US Constitution to the European Union's data protection requirements and directives. While our laws will probably be the eventual site of the conversation in which we answer, as a society, hard questions about the harvesting and stockpiling of personal information, they operate slowly; and whatever momentum propels them in the direction of protecting privacy in the public interest is amply counterweighted by opposing forces of vested corporate and other institutional power, including governmental interests. In the meantime, and in the short term, enormous quantities of personal data are already in circulation, packaged, sold, provided freely and growing by the day.

Finally, there is great interest among the technical, particularly research, community in *engineering systems* that 'preserve' and 'enhance' privacy, be it in data mining, surfing or searching the web, or transmitting confidential information. Detecting data provenance, properly anonymising data sets, generating contextual awareness and providing secure, confidential communication: mechanisms supporting these goals pose technical challenges, particularly when embedded in the real world or when working against the grain of features native to infrastructural systems such as the web. Furthermore, no matter how convincing the technical developments and standards, adoption by key societal actors whose organisations and institutions mediate much data flow is another matter and fraught with politics.

Tools offered to the individual directly, such as Tor and other proxy servers, are praiseworthy and valuable but the fact remains that they are not widely understood or deployed outside the relatively small circles of those who are already privacy-aware and technologically sophisticated. Additionally, there are utility costs: Tor can be slow, for example, and is blocked by many large websites. All privacy-protecting technologies entail trade-offs, and those required by robust approaches such as Tor have thus far kept their adoption relatively small.

We are not questioning the ability of law, the private sector and technology to provide relief to individuals from unfettered monitoring, gathering, mining and profiling. The benefits of the status quo to those on the other side of the power and epistemic asymmetries that define and entrench our predicament, however, will not be easily ceded and, even if ultimately they are, the wait for meaningful relief is likely to be long. Turning to obfuscation, therefore, is a way to take matters into our own hands in the interim. Before discussing how it addresses the specific problem of data gathering and analysis, we introduce obfuscation through an array of historical and contemporary examples so that we can see it as a general strategy, with many different forms, media and motives.

Obfuscation in practice: cases and examples

Obfuscation in its broadest and most general form offers a strategy for mitigating the impact of the cycle of monitoring, aggregation, analysis and profiling, adding noise to an existing collection of data in order to make the collection more ambiguous, confusing, harder to use and, therefore, less valuable. (We chose 'obfuscation' for this purpose because of its connotations of confusion, ambiguity and unintelligibility, seeking to distinguish it from other strategies involving concealment or erasure, such as cryptography.) Obfuscation, like data gathering, is a manifold strategy carried out for a variety of purposes, with a variety of methods and perpetrators. Obfuscators may band together and enlist others, or produce misleading information on their own; they might selectively respond to requests for information, or respond so excessively that their contribution skews the outcome. They may engage in obfuscation out of a simple desire to defend themselves against perceived dangers of aggregation, in resentment of the obvious asymmetry of power and knowledge, to conceal legitimate activities or wrongdoing or even in malice, to render the system of data collection as a whole worthless. This diversity of purposes, methods and perpetrators is reflected in the wide range of forms taken by obfuscation tactics.

These forms, across a range of media and circumstances, can be loosely clustered around four themes: time-based obfuscation, which relies on temporal limitations; cooperative obfuscation, requiring the 'network effect' of cooperation or collaboration by groups of obfuscators; selective obfuscation, interfering with data to conceal specific details while leaving others available; and ambiguating obfuscation, which renders data ambiguous and doubtful for future use.

Time-based obfuscation

Whereas some forms of obfuscation try to inject doubt into the data permanently, time-based obfuscation, in many ways the simplest form of the practice, adds need for an onerous amount of processing in a situation where time is of

the essence. *Chaff* offers a canonical example: the radar operator of the Second World War tracks a plane over Hamburg, guiding searchlights and anti-aircraft guns in relation to a phosphor dot whose position is updated with each sweep of the antenna. Abruptly the plane begins to multiply, dots quickly swamping the display. The plane is in there somewhere, impossible to locate for the presence of all the 'false echoes'. The plane has released chaff, strips of black paper backed with aluminum foil and cut to half the target radar's wavelength, floating down through the air, thrown out by the pound and filling the system with signals. Chaff has exactly met the conditions of data the radar is configured to look for and given it more planes, scattered all across the sky, than it can handle. Knowing discovery to be inevitable, chaff uses the time and bandwidth constraints of the discovery system against it by creating too many potential targets (in this regard, Fred Cohen (Cohen 2006: 646) terms it the 'decoy strategy', and we can indeed consider obfuscation as the multiplication of plausible data decoys).That the chaff only works briefly, as it flutters to the ground, and is not a permanent solution, is irrelevant under the circumstances; it only needs to work well enough for the time it will take the plane to get through.

Another contemporary example is the practice of *quote stuffing* in high-frequency trading (HFT). (To be clear, quote stuffing is still only a theoretical obfuscation project, a plausible explanation for recent bursts of anomalous activity on the stock market.) The rarefied world of HFT is built on algorithms that perform large volumes of trades far faster than humans, taking advantage of exceedingly minute spans of time and differences in price that would not be worth the attention of human traders, if it were even physically possible for them to act on the change in price before the advantage was gone. Analysts of market behaviour began to notice unusual patterns of HFT activity over the summer months of 2010 – bursts of quote requests for a particular stock, sometimes thousands a second. Such activity seemed to have no economic rationale, but one of the most interesting and plausible theories (Nanex 2010) is that these bursts are an obfuscation tactic in action: 'If you could generate a large number of quotes that your competitors have to process, but you can ignore since you generated them, you gain valuable processing time'. Unimportant information, in the form of quotes, is used to crowd the field of salient activity, so the generator of the unimportant data can accurately assess what is happening while making it more difficult for competitors to do so in time. The volume of trades creates a cloud of fog that only the obfuscator can see through. In the sub-split-second world of HFT, the act of having to observe and process this hiss of activity is enough to make all the difference.

Finally, two examples of time-based obfuscation in thoroughly concrete contexts. The affair of the 'Craigslist robber' offers a minor but illustrative example of obfuscation as a practice turned to criminal ends. At 11 am on Tuesday30 September 2008, a man dressed as an exterminator in a blue shirt, goggles and a dust mask, and carrying a spray pump, approached an

armoured car parked outside a bank in Monroe, Washington, incapacitated the guard with pepper spray, and took the money. When the police arrived, they found 13 men in the area wearing blue shirts, goggles and dust masks – a uniform they were wearing on the instructions of a Craigslist advertisement which promised a good wage for maintenance work, which was to start at 11:15 am at the bank's address. This incident is one of the few real-world examples of a recurrent trope of obfuscation in movies and television: the many identically dressed actors or objects confusing their pursuers as to the valuable one. Obviously it will only take a few minutes to determine that none of the day labourers is the bank robber – but a few minutes is all the thief needs.

Much of the pleasure and challenge of poker lies in learning to read people and deduce from their expressions, gestures and body language whether they are bluffing, or pretending to hold a weaker hand in hopes of drawing a call. Central to the work of studying opponents is the 'tell', some unconscious habit or tic an opponent will display in response to a strong or weak hand: sweating, a worried glance, leaning forward. Tells play such a crucial role in the informational economy of poker that players will use *false tells*, creating mannerisms which may appear to be part of a larger pattern.[5] According to common poker strategy, the use of a false tell is best reserved for a crucial moment in a tournament, lest the other players figure out that it is inaccurate and turn it against the teller in turn. A patient analysis of multiple games could separate the true from the false tells, but in the time-bound context of a high-stakes game the moment of deception can be highly effective.[6]

Cooperative obfuscation

All of the cases described so far can be performed by a single actor (perhaps with some unwitting assistants), but other forms of obfuscation require the explicit cooperation of others. These obfuscatory cases have a 'network effect' of becoming more valuable as more people join. A powerful legend exemplifies this idea: the often retold, factually inaccurate story that the king and population of Denmark wore the Yellow Star to make it impossible for the occupying Germans to distinguish and deport the Jews. While the Yellow Star was not used in Denmark for fear of arousing more anti-German feeling, '[t]here were documented cases of non-Jews wearing yellow stars to protest Nazi anti-Semitism in Belgium, France, the Netherlands, Poland, and even Germany itself'.[7] The legend is a perfect story of cooperative obfuscation: a small group of non-Jews wearing the Yellow Star is an act of protest; a whole population, into which individual Jews can blend, is an act of obfuscation.

Loyalty card swapping pools provide another superb real-world example. Quite quickly after the widespread introduction of 'loyalty cards', offering discounts to regular shoppers at grocery store chains, came card-swapping networks, where people shared cards – initially in ad hoc physical meetings,

and increasingly in large populations and over wide geographical regions enabled by mailing lists and online social networks – to obfuscate their data. Rob's Giant Bonus Card Swap Meet, for instance, started from the idea that a barcode sharing system could enable customers of the DC-area supermarket chain to print out the barcodes of others, pasting them onto their cards (Carlson (25 October 2010)). A similar notion was adopted by the Ultimate Shopper project, mailing stickers of a Safeway loyalty card barcode and creating 'an army of clones' accruing shopping data (Cockerham (19 October 2010)). Cardexchange.org is devoted to exchanging cards by mail, presenting itself as a direct analogue to the physical meet-ups. These sites also act as clearing houses for discussion, gathering notes, blog posts, news articles and essays on loyalty cards, debating the ethical implications of various approaches and sharing theories and concerns. This is obfuscation as a group activity: the more who are willing to share their cards, the farther the cards travel and the more unreliable the data becomes.

Another form of collective obfuscation appears in the argument for participation in Tor. Tor is a system designed to enable anonymous use of the internet, through a combination of encryption and passing the message through many different independent 'nodes'. If you request a web page while working through Tor, your request will not come from your IP address, but from an 'exit node' on the Tor system, along with the requests of many other Tor users. Data enters the Tor system and passes into a labyrinth of relays, computers on the Tor network that offer some of their bandwidth for handling Tor traffic from others, agreeing to pass messages sight unseen. In return for running a Tor relay, as the FAQ (2012) notes, 'you do get better anonymity against some attacks. The simplest example is an attacker who owns a small number of Tor relays. He will see a connection from you, but he won't be able to know whether the connection originated at your computer or was relayed from somebody else'. If you are on Tor and not running a relay, then an adversary will know you wrote the message you sent to him. But if you are allowing your computer to operate as a relay, the message might be yours or simply one among many that you are passing on for other people. Did it start with you or not? The information is now ambiguous, and messages you have written are safe in a flock of other messages you pass along.[8]

Selective obfuscation

All of the examples thus far have been about general methods of covering one's tracks. But what if you want your data to be useful without diminishing your privacy, or to interfere with some methods of data analysis but not others? This is the project of selective obfuscation. FaceCloak, for example, provides the initial steps towards an elegant and selective obfuscation-based solution to the problem of Facebook profiles. It takes the form of a Firefox plug-in that acts as a mediating layer between a user's personal information and the

social networking site. When you create a Facebook profile and fill in your personal information, including details such as where you live, went to school, your likes and dislikes and so on, FaceCloak offers you a choice: display this information openly, or keep it private? If you let it be displayed openly, it is passed to Facebook's servers like any other normal data, under their privacy policy. If you want to keep that data private, however, FaceCloak sends it to encrypted storage on a separate server only to be decrypted and displayed for friends you have authorised, when they browse your Facebook page (using the FaceCloak plug-in.) Facebook never gains access to the data. Furthermore, by generating fake information for the data that Facebook requires of its profiles, FaceCloak obfuscates its method – the fact that the real data lies elsewhere – from both Facebook and unauthorised viewers. As it passes your real data to the private server, FaceCloak generates a gender, with appropriate name and age, and passes those to Facebook. Under the cover of this generated, plausible non-person, you can connect and exchange with your friends, obfuscating the data for all others.

The theoretical goal for selective obfuscation has been outlined from a policy perspective as obfuscating the data for certain users or the reconstruction of individual acts. In Gloria Gonzalez Fuster's recommendations for EU data processing selective obfuscation is understood as limiting the data to primary processing: structuring the data such that it can be evaluated for its intended purpose, to which the data's subjects consent, but not for unanticipated analyses (Gonzalez Fuster 2009). In this scenario, data gathered for, say, a public health study would be suited to the process used for that study, difficult to use for other public health data mining and impossible to reprocess for any other purpose.

The work of Nam Pham and others (2010) on privacy-preserving participatory sensing shows us how this idea could work in practice, on an applied and mathematically sophisticated scale. Where a project such as FaceCloak obfuscates the data for all but an authorised few, private participatory sensing obfuscates it beyond a certain degree of specificity – the data works generally, but not for identifying or tracking anyone in particular. Vehicular sensors, for instance, which can be used to create a shared pool of data from which to construct maps of traffic or pollution, raise obvious concerns over location-based tracking. However, Pham and his colleagues demonstrate how to perturb the data, letting each vehicle continuously lie about its location and speed while maintaining an accurate picture of the aggregate.

Ambiguating obfuscation

Time-based obfuscation can be quickly seen through; cooperative obfuscation relies on the power of groups to muddy the tracks; selective obfuscation wishes to be clear for some and not others. Ambiguating obfuscation seeks to render an individual's data permanently dubious and untrustworthy as a subject

of analysis. For example, consider the Firefox extension TrackMeNot, developed in 2006. Developed by Daniel Howe, Helen Nissenbaum and Vincent Toubiana, TrackMeNot was designed to foil the profiling of users through their searches. Our search queries end up acting as lists of locations, names, interests and problems, from which not only our identities can be determined but a pattern of our interests revealed regardless of whether our IP addresses are included. As with many of the previous cases of obfuscation, opting-out of a web search is not a viable choice for the vast majority of users. (At least since 2006, search companies have been responsive, although only partially, to users' concerns over the logging and storage of search queries.) TrackMeNot automatically generates queries from a seed list of terms which evolve over time, so that different users develop different seed lists. TrackMeNot submits queries in a manner that tries to mimic user search behaviours. These users may have searched for 'good wi-fi cafe chelsea' but they have also searched for 'savannah kennels', 'exercise delays dementia' and 'telescoping halogen light' – will the real searchers please stand up? The activity of individuals is masked by that of many ghost queries, making a pattern harder to discern.

Similarly, BitTorrent Hydra fights the surveillance efforts of anti-file sharing interests, by mixing genuine requests for bits of a file with dummy requests. The BitTorrent protocol breaks a file up into many small pieces, so that you can share those pieces, sending and receiving them simultaneously with other users. Rather than downloading an entire file from another user, as with the Napster model, you assemble the file's pieces from anyone else who has them, and anyone who needs a piece you have can get it from you (Schulze and Mochalski 2009). To help users of BitTorrent assemble the files they want, the system uses 'torrent trackers', which log IP addresses that are sending and receiving files – if you are looking for these pieces of file *x*, users *a* to *n*, at the following addresses, have the pieces you need. Intellectual property groups, looking for violators, starting running their own trackers to gather the addresses so they could find major uploaders and downloaders of potentially copyrighted material. To protect BitTorrent users, Hydra obfuscates by adding random IP addresses to the tracker, addresses that have been used for BitTorrent connections at some point. This step means that, periodically, as you request pieces of the file you want, you will be directed to another user who does not actually have what you are looking for. It is a small inefficiency for the BitTorrent system as a whole, but it makes address gathering on the part of anti-piracy organisations much less useful. The tracker can no longer be sure that any one address was actually engaged in sharing any particular file. Hydra does not avert data collection, but contaminates the results, making any specific case problematic and doubtful.

CacheCloak, meanwhile, has an approach to obfuscation suited to its domain of location-based services (LBSs). LBSs take advantage of the locative technology in mobile devices to create various services. If you want the value of an LBS – say, to be part of the network that your friends are on so you can

meet if you are nearby – then you will have to sacrifice some privacy and get used to the service provider knowing where you are. 'Where other methods try to obscure the user's path by hiding parts of it', write the creators of CacheCloak, 'we obscure the user's location by surrounding it with other users' paths' – the propagation of ambiguous data. In the standard model, your phone sends your location to the service and gets the information you requested in return. In the CacheCloak model, your phone predicts your possible paths and then fetches the results for several likely routes. As you move, you receive the benefits of locative awareness – access to what you are looking for, in the form of data cached in advance of potential requests – and an adversary is left with many possible paths, unable to distinguish the beginning from the end of a route, where you came from and where you mean to go, still less where you are now. The salient data, the data we wish to keep to ourselves is buried inside a space of other, equally likely data.

Finally, the technique of botnet-resistant coding operates on similar lines to quote stuffing. A botnet is a collection of malware-infected personal computers controlled by a remote attacker, using system resources or snooping for data. One of the more prolific of these botnets, known as Zeus, sits on the network looking for the patterns of data that suggest banking information; when found it sends the information – passwords, account details and so on – back to its controllers, who will use it to make bank withdrawals or commit other forms of identity theft. The defensive solution proposed is an obfuscation move: large quantities of completely plausible but incorrect information would be injected into the transactions between the user's computer and the bank. Banks would know how to filter the false information, because they have generated it, but not the botnet. Faced with this source of confusion, attackers either move on to easier targets or waste resources trying to find the accurate needle in a bank's haystack.

The politics and ethics of obfuscation: a 'weapon of the weak'?

The examples we have compiled show something of the broad range of obfuscation practices, from foiling statistical analysis and escaping visual sensing to thwarting competitors in the stock market. Some methods take advantage of human biases and others the constraints and loopholes of automated systems. Obfuscation is deployed for short-term misdirection, for legal deniability, to encourage an adversary to construct a flawed model of the world and to change the cost-benefit ratio that justifies data collection. The swathe of types, of methods, motives, means and perpetrators are not surprising considering that obfuscation is a reactive strategy and, as such, a function of as many types of actions and practices as it is designed to defeat.

Despite this diversity, we would like to think that obfuscation will become a subject of interest for scientific study, to identify key variables and parameters,

to understand the relationships among them and, ultimately, to quantify its value and optimise its utility. With encryption, for example, algorithms possess standard metrics based on objective measures such as key length, machine power and length of time to inform community evaluations of their strength. By contrast, the success of obfuscation is a function of the goals and motives of both those who obfuscate and those to whom obfuscation is directed, the targets. It simply has to be 'good enough', a provisional, ad hoc means to overcome the challenge that happens to be in its way.

Our task here, however, is not a scientific analysis of obfuscation but an ethical one. There are ways in which obfuscation practices can be unethical, but there are also mitigating conditions that we must consider and details we must resolve – and, along with those ethical particulars, there is a general political analysis to be made before we can claim a full understanding of obfuscation's moral and political dimensions. We discuss each, in turn, below.

Ethics of obfuscation

In 'A Tack in the Shoe' (2003), Gary Marx writes: 'Criteria are needed which would permit us to speak of "good" and "bad", or appropriate and inappropriate efforts to neutralise the collection of personal data'. If we accept that obfuscation works – that, even if weak, it can be a successful and consequential strategy – we must still ask whether it can be defended against charges that it is unethical. Although we are interested in the moral standing of particular uses of obfuscation, our central concern here is with the strategy of information obfuscation itself, whether structurally or inherently unethical. Thus, we address several of the most compelling issues that critics have raised.

Dishonesty

Implicit in obfuscation is an element of dishonesty – it is meant to mislead. Some people might balk at valorising any practice that systematises lying or deception. (Some obfuscation approaches, such as that of CacheCloak, work around this problem by remaining ambiguous instead of providing untrue information – but such an approach depends on an informational relationship where queries can be left vague.) These critics might prefer encryption (that is, hiding, a form of refusal) or silence to producing streams of lies. Whether lying, in general, can be morally justified is an exploration that clearly would take us too far afield from our subject, but that general discussion yields insights that are useful here. Except for the Kantian who holds that lying is always absolutely wrong (famously, prescribing a truthful answer even to the murderer seeking one's friend's whereabouts), in many analyses there are conditions in which the proscription of lying may be relaxed (Bok 1999). We must ask whether the general benefits of lying in a given instance outweigh harms, and whether valued ends are served better by the lie than

truthful alternatives. There are many special circumstances in which lies may be excused; for example, if one is acting under duress or lying to one party to keep a promise to another.

Free riding

Obfuscation may involve two different forms of free riding, both of which take advantage of the compliance of others in the obfuscator's situation. Imperfect as it may be, obfuscation may raise the cost of data gathering and analysis just enough to deter the surveillant or divert him to other data subjects. These may overlap or coexist, but their distinct ethical values are clear. The first takes advantage of the willingness of others to submit to data collection, aggregation and analysis – no need to be faster than the predator so long as one is faster than other prey. It allows others to be victimised while one remains safe oneself, a safety that is the product, however indirectly, of the victimisation of others. The second involves enjoying the benefits provided by the data collector, without paying the price of one's data. (Loyalty card-swapping pools are an instance, as participants enjoy the bounty of special offers while escaping the information pool that presumably supports them.)

Waste, pollution and system damage

A common critique of obfuscation is that it wastes or pollutes informational resources – whether bandwidth and storage, or the common pools of data available for useful projects.

In considering such accusations, we note that 'waste' is a charged word, implying that resources are used improperly, based presumably on an agreed-upon standard. This standard could be challenged; what is wasteful according to one standard might be legitimate use according to another. However, noise introduced into an environment is not only wasteful but may taint the environment itself. On a small scale, obfuscation may be insignificant – what can be the harm of marginal inaccuracy in a large database? On a large scale, however, it could render results questionable or even worthless. To take a recent case, the shopping logs of supermarket loyalty cards were used by the Centers for Disease Control and Prevention to identify a common purchase among a scattered group of people with salmonella, trace that purchase to the source and institute a recall and investigation, a socially valuable project which the widespread adoption of loyalty card swapping pools would have made much slower or even, theoretically, impossible (Mercer 2010).

Data aggregation and mining is used not only to extract social utility but to guide decisions about individuals. If introducing noise into a system interferes with profiling, for example, it might harm the prospects of individuals, innocent bystanders, so to speak. FaceCloak demonstrates this problem: '[F]or some profile information (eg an address or a phone number), it is ethically

questionable to replace it with fake information that turns out to be the real information for somebody else' (Mercer 2010: 6). The risk is not only in the present, but also holds for future uses not yet foreseen, the nightmare of the regularly incorrect United States No-Fly List writ large, or the mistakes of police profiling software compounded by a large pool of alternate, inaccurate names, addresses, activities, search terms, purchases and locations. As a possible counterargument, however, if we believe that these databases and the uses to which they are put are malignant, this bug becomes a feature. A database laden with ambiguously incorrect material becomes highly problematic to act on at all.

Finally, waste includes the potential of damage, possibly fatal damage, to the systems affected by obfuscation. Consider quote stuffing in high-frequency trading, a move which, if broadly adopted, could actually overwhelm the physical infrastructure on which the stock exchanges rely with hundreds of thousands of useless quotes consuming the bandwidth. Any critique of obfuscation based in the threat of destruction must be specific as to the system under threat and to what degree it would be harmed.

Assessing the ethical arguments

The merits of each charge against obfuscation are not easily assessed in the abstract without filling in pertinent details – and these details make all the difference. The overarching question that drives this chapter is about obfuscation aimed at thwarting data monitoring, aggregation, analysis and profiling, so we confine our evaluation to this arena, drawing on the cases we introduced above. One consideration that is relevant across the board is ends; legitimate ends are necessary, although, clearly, not always sufficient. Once we learn, for example, that the Craigslist robber used obfuscation to rob banks or that quote stuffing could bring down the Stock Exchange, it hardly seems relevant to inquire further whether the lies or free riding were justifiable.

The judgment of ends can also take in questions about proportionality and not only whether an action in question is flatly right or wrong. The obfuscator running TrackMeNot may not disapprove of the ultimate purpose of Google's query logs but may consider the degree of surveillance too extreme. The company makes its revenue from advertising, and it is reasonable for it to serve keyword-specific ads automatically against a given query – but if the data mining begins to seem too personal, too precise, or is extended into a previously off-limits private domain and the user feels it is no longer fair or proportionate, he or she will begin using TMN. Astute businesses will be helped by paying attention to customers giving voice to their concerns through soft acts of protest such as these, which signal a need to bring a practice into line with consumer expectations and beliefs. These are not demands for total reinvention but the reassertion of more equitable standing.

Dishonesty

In cases such as TrackMeNot, CacheCloak, Tor relays and loyalty card swapping, the ethical arguments can become quite complex. To justify the falsehoods inherent in obfuscation, the ends must be unproblematic, and other aspects of the case taken into consideration – whether achieving the ends by means other than lying is viable and what claim the targets of falsehood may have to 'real' information. If individuals feel they have little chance of protection through law, technology and corporate best practice, as we discussed above, under duress and with little assurance that those extracting information can be trusted, the obligation to speak the truth is certainly lessened. Contrast this scenario with highly controlled environments, such as a courtroom, where a myriad of other constraints circumscribe the actions of all parties; we may still speak under duress but epistemic asymmetries are mitigated because of these other strictures of context.

Free riding

While deception may be justified by asymmetries of knowledge and power and the absence of alternatives, other critiques remain. The problem of free riding on the contributions of others casts obfuscation efforts in an unseemly light. The obfuscator is presented as not so much the rebel as the sneak, with an interest, however indirect, in the ignorance and foolishness of others: that they fail to 'game the system' as the obfuscator does. (A house's safety from theft, one might say, comes not only from a locked door but from other houses being left unlocked.) Against this charge we can bring in mitigating circumstances and specific details, as we did with dishonesty, but we can also draw on a broader argument which we make below, based in a Rawlsian analysis – free riding has a different ethical valence if it is available to all and disproportionately aids the weak against the strong. As long as the free rider is not actively attempting to keep others from enjoying the same benefit (as though hobbling others in the herd to make them more likely to be caught by predators), the ethical price of their actions is paid by supererogation. Obfuscators cannot be expected to imperil themselves solely because others are in peril; they cannot be morally obligated to starve simply because others are starving.

The second form of free riding – drawing on benefits provided by data collectors without paying the price of personal data – has a different moral pattern. Individuals using FaceCloak or CacheCloak, for example, may draw the ire of Facebook or location-based services because they are depriving these services of the positive externalities of personal information flows, which normally would enrich either their own data stockpiles or those of others to whom this data is sold or exchanged. It is not clear to us that companies are entitled to these externalities. At the very least, these relationships need to be

examined from a broad societal perspective and the flow of costs and benefits (direct and indirect) explicitly recognised. If and only if it can be established that extracting the benefits offered by these services inflicts general, unacceptable costs, and not simply costs to companies, are there grounds to judge such free riding unethical.

Waste

Wastefulness is a charge that may be levelled against systems such as TrackMeNot that 'waste' bandwidth by increasing network traffic and 'waste' server capacity by burdening it with search queries that are not, in reality, of interest to users. A cost-benefit or utilitarian assessment directs us to consider the practical question of how severe the resource usage is. Does the noise significantly or even perceptibly undermine performance? In the case of search queries, which are short text strings, the impact is vanishingly small compared with the internet's everyday uses at this point, such as video distribution, online gaming and music streaming.[9]

Additionally, it is not sufficient to hang the full weight of the evaluation on degree of usage – it is necessary to confront normative assumptions explicitly. There is irony in deeming video streaming a *use* of network but a TrackMeNot initiated search query a *waste* of network, or a TrackMeNot initiated query a *waste* of server resource but a user generated search for pornography a *use*. This claim makes sense, however, once we acknowledge that the difference between waste and use is normative; waste is use of a type that runs counter to a normative standard of desired, approved or acceptable use. The rhetoric of *waste*, however, begs to be scrutinised because, while it may be dressed up as an objective, definable concept, in many cases it is speakers who inject and project their perspectives or interests into defining a particular activity as wasteful.

Pollution and system damage

Data 'pollution' and the propagation of error and inaccuracy may be the trickiest issues of all, and reach to the heart of obfuscation. The intention behind inserting noise into the data stream is precisely to taint the resulting body. Yet there are various ways it can be tainted and some may be more problematic than others. One misspelt name does not a ruined database make; at what point does inaccurate, confusing and ambiguous data render a given project or business effectively worthless? Obfuscation that does not interfere with a system's primary functioning but affects only secondary uses of information might be fair.[10] Further, while some obfuscation practices might confuse efforts to profile individuals accurately, they may not render aggregate analysis useless, for example, as in the case of the work of Pham et al (2010) on perturbing individual data while retaining a reliable total picture.

Yet what if there is no getting around the noise? Where does this reality leave the ethical status of obfuscation? Is it acceptable to coerce people into providing data into the pool for the sake of another party, or even for the common good? And if they are coerced with no assurance as to how the information will be used, where it will travel and how it will be secured, are they not being asked to write a blank cheque with little reason to trust the cheque's recipients? These are akin to many ethical questions confronting individuals, both in relation to other individuals and to society and, as with those questions, there may be no general answers that do not call for further elaboration of the surrounding context. When pushed into a corner, in cases where dishonesty, free riding, resource consumption and data tainting cannot be denied, obfuscation nonetheless may pass the moral test. But establishing this status requires exploration of the specific and general obligations that the obfuscator may owe, whether securing freedom from the machinations of monitoring and analysis is justified and whether the obfuscator, having considered alternatives, is acting in earnest assertion of these freedoms. Explaining the calculus of those freedoms, and what liberties obfuscation defends, is our goal in the remainder of this chapter.

Politics of obfuscation

Reflecting on properties of obfuscation that are potentially morally problematic in the previous section, we found that none by itself implies that data obfuscation is inherently unethical. This finding is relevant to the inquiry of this section, in which we ask about the politics of obfuscation, namely what approach might a just society adopt toward data obfuscation, whether to ban or condone it, and by what lights. Inspired by Rawls's two principles, the first directs us to assess whether data obfuscation violates or erodes basic rights and liberties. If the reasoning above is sound, it seems there are no grounds to assert this categorically. Instead, the details of particular instances or types of instances will matter – for example, whether untruths or dissipation of resources abridge rights of those against whom obfuscation is practised, such as autonomy, property or security and, if they do, whether countervailing claims exist of equal or greater weight and legitimacy (of those who obfuscate), such as autonomy, fair treatment freedoms of speech and political association (that is, various freedoms associated with privacy protection).

Data obfuscation provides a particularly interesting case for Rawls's second 'maximin' principle. Setting aside instances of obfuscation, such as the Craigslist robber, which do not meet the requirements of the first principle, controversial cases may include some in which there are unresolved conflicting rights and liberties, and others in which respective claims are in conflict. The types of cases described above include those in which, say, individuals seek cover through obfuscation for legitimate conditions or behaviours, thus denying positive externalities to data gatherers or that seek secrecy at a cost to

the purity of a data pool. In paradigmatic instances, there are clear power differentials: individuals are reaching for obfuscatory tactics to avoid surveillance, profiling and manipulation, in general, to remain out of reach of a corporate or government actor.

Although obfuscation can be used by the more powerful against the less powerful, there are usually more direct ways for the more powerful to impose their will on the less powerful. Because obfuscation is not a strong strategy, it is only very rarely adopted by powerful actors – and then usually to evade notice by other powerful actors, as in the case of shell companies created to deter journalists and regulators, or the phenomenon in the Guatemalan secret police of multiple 'red herring' evidence plants and false testimonies to suggest that any final determination of what took place in a crime will be impossible (Goldman 2007). There is less need for stronger actors to resort to obfuscation because they have better methods available if they want to hide something – such as secret classifications, censorship and the threat of state violence.

For those who are generally less well off, less politically powerful, not in a position to refuse terms of engagement, technically unsophisticated, without the background in computing to use protections such as encryption, for those who need discounts at the supermarket, free email and cheap mobile phones, obfuscation can be a salve. It can avail some measure of resistance, obscurity and dignity. In this way, obfuscation fits into the domain that James C Scott describes as 'weapons of the weak', the domain of dissimulation, slow-downs, petty theft, gossiping, foot-dragging and other forms of resistance on the part of deeply disempowered actors (in the case of Scott's analysis, agrarian peasants) on the wrong side of severe power asymmetries. These are people without the possibility of armed revolt, without law or legislature on their side – what remains to them is 'passive noncompliance, subtle sabotage, evasion, and deception', terms that nicely capture the dimensions of obfuscation (Scott 1987: 31). As Anatole France put it: 'The law, in its majestic equality, forbids the rich as well as the poor to sleep under bridges and steal bread'. For those whose circumstances and necessity oblige them to give up their data – those who most need the shelter of the bridge, however ad hoc and unsatisfying it may be compared with a proper house – obfuscation provides a means of redress and, as such, is politically justified.

Although these political asymmetries are due in part to traditional sources of power differentials, such as influence, money, social class, education, race and so on, epistemic asymmetries, as discussed above, are also enormously consequential in contemporary, data driven societies. We may reach for obfuscation to shake off unwanted coercive influences, but we may do so simply because we are in the dark; we know that information about us is not disappearing but we know not where it is going nor how it has been or will be used. We are reaching for it to avoid or neutralise a lurking but ill-understood threat. In pushing against not so much the exercise of power and coercion but the threat of it, we are acting against what Philip Pettit might call domination, which he defines as the capacity to interfere in another's choices on an

arbitrary basis (Pettit 1997). From the perspective of the individual on the other side of the epistemic asymmetry, the capacity of those who create and act on profiles of us that they have generated by gathering, aggregating and mining data may seem quite arbitrary.

Rawls's maximin principle demands that a just society opts for 'the alternative the worst outcome of which is superior to the worst outcomes of the others' (Rawls 1971: 153). Because data obfuscation offers a means to the less well off to assert their will against the more well off and powerful, banning data obfuscation either directly or indirectly by supporting measures coercing individuals to provide sound information, in our view, would violate the maximin principle. Where the obfuscator acts earnestly to resist the machinations of monitoring and analysis, obfuscation thus enables acts of reasonable and morally sound disobedience.

Among the toughest challenges to obfuscation are those that point to free riding and database pollution. The obfuscator is faulted for being unwilling to pay the cost for a benefit to him or herself, or for obstructing potential benefits to society at large by being unwilling to pitch in. Although these charges are worth taking seriously, so also is a caution that Jeremy Waldron issues in his discussion of a post-9/11 world in which citizens are expected to accept a rebalancing of security and liberty in favour of the former. Whenever there is talk of achieving a balance among social goods requiring that one be traded off against another, among other objections to such trade offs, one is that all too often we fail to take into consideration that costs and benefits are unevenly distributed (Waldron 2003). It may simply not be the case that *we* collectively give up a certain measure of freedom in return for *our* collective greater safety but that the loss of liberty is concentrated on a small sub-set of our society, who take a massively disproportionate loss for the possible benefit to us as a whole (from which 'they', who lose so much more of their liberty, are now excluded) or for those of us in a different sub-set. According to Waldron, we, collectively, may accept this unfair trade off because, in aggregate, we do not feel the sting very much.

In cases of data obfuscation where we might be inclined to cite free riding or data pollution, Waldron's caution must not be ignored. In these cases, obfuscation might be legitimate acts of resistance by some, carrying the burdens of dataveillance disproportionately, for the sake of others, or for the sake of us all. Obfuscation may be an appropriate response, because it is disproportionately advantageous to the more vulnerable actor against the less vulnerable. Compared with the price of refusal and the difficulties of complete concealment, obfuscation is a relatively simple and intuitive way for the individual to resist, allowing both compliance and protest at the same time.

Conclusions

Obfuscation, as we have presented it here, is at once richer and less rigorous than academically well established methods of digital privacy protection,

such as encryption. It is far more ad hoc and contextual, without the quantifiable protection of cryptographic methods. It is often haphazard and piecemeal, creating only a temporary window of liberty or a certain amount of reasonable doubt. It is for precisely those reasons that we think it is a valuable and rewarding subject for study. Politically, as long as the ends are sound and we take care to avoid certain methods, obfuscation can be a force for good in our contemporary culture of data. These moves are a valuable resource in the defence of our privacy and freedom of action. We have provided an outline of the family, a number of examples, the parameters for quantification and improvement, and a view of the political and ethical problems it creates, as well as arguments in its favour. Now, we hope the community of privacy researchers and activists will help to expand this idea. We face a number of further questions, beginning with one scientific, one moral and one technical:

- Is it possible to create a meaningfully quantified science of obfuscation? Can we optimise different obfuscation tactics for different scenarios, and find weak points in the overall strategy?
- Does our description of obfuscation as a viable and reasonable method of last-ditch privacy protection lead to the same political problems created by other systems of privacy preserving technology and possibilities such as opt out – that is, putting the responsibility back on the private user and side-stepping the need to create a mature civil society around managing data?
- Are there methods for counter-profiling – figuring out how the profilers work to fine-tune our data strategies and how best to stymie them – that could be incorporated into the project of refining obfuscation?

Under duress, in the face of asymmetry, innovative methods for drawing the contextual lines of information flow will emerge; people will create models of informational security and freedom from invasive analysis, irrespective of claims profit-seeking CEOs make about 'human nature' and the transformations of privacy. Obfuscation is often cheap, simple, crude and clever, rather than intelligent and lacks the polish or freedom from moral compromises that characterises more total privacy solutions. Nonetheless it offers the possibility of cover from the scrutiny of third parties and data miners for those without other alternatives. It is the possibility of refuge when other means fail, and we are obliged both to document it and to examine whether it can be made stronger: a more effective bulwark for those in need.

Notes

* This project was researched and written with funding from AFSOR: MURI (ONR BAA 10-002), NSF:PORTIA (ITR-0331542) and NSF-CT-M (CNS-0831124) grants. We are grateful for their support. This work benefited enormously from

the invaluable help and insights of members of the Privacy Research Group at NYU and audiences at Computers, Privacy and Data Protection 2011 and the European Association for the Study of Science and Technology 2010, where developing versions of this work were presented. We would also like to thank Solon Barocas, Ian Kerr and Mireille Hildebrandt for their astute comments, feedback and advice. We are indebted to Luke Stark for providing outstanding research assistance and editorial work.

1 The sale is well documented by the account in CSOonline, http://www.csoonline.com/article/220340/the-five-most-shocking-things-about-the-choicepoint-data-security-breach (accessed 30 October 2012), and the reactions by the FTC and ChoicePoint have been collected in the Privacy Rights Clearinghouse 'Chronology of Data Breaches' (see under 15 February 2005): http://www.privacyrights.org/ar/CPResponse.htm (accessed 30 October 2012). This incident led to the thought-provoking 'Model Regime of Privacy Protection' proposed by Daniel Solove and Chris Jay Hoofnagle; see Solove and Hoofnagle 2005.

2 In making this argument we are drawing on our descriptions of this problem with reference to the received notion of privacy in Nissenbaum (1998, 1999).

3 As one among many possible examples of our ignorance of the future uses to which our data may be put — whether it is records sold by an unscrupulous employee or left in a cab on a USB drive — see the business of scraping social network sites for their data, which can be bundled, sold and used without our ever being aware or giving consent to this use: http://www.readwriteweb.com/archives/bulk_social_data_80legs.php (accessed 30 October 2012). For analysis of this situation from a specifically legal perspective, see Hildebrandt (2008) and Zarsky (2005).

4 Any real opt-out policy would also have to offer the granularity of the process of aggregation and analysis itself, allowing you to make choices that lie between the extremes of refusal and compliance. An opt-out of consequence would enable the receipt of certain benefits in return for a degree of use; data that could be gathered or deployed only in certain contexts or for certain purposes, for a set period of time etc. This does not presently exist, and implementing it relies heavily on the diligence and good behaviour of private corporations. See Barocas and Nissenbaum (2009) for an instance of this problem of consenting to data use after the fact.

5 An anecdotal account of false tells from poker player Phil Hellmuth, from Navarro (2006), can be found online at http://southerngaming.com/?p=62 (accessed 30 October 2012).

6 It is interesting to imagine a poker strategy based around more extensive use of obfuscation — a player generating a constant stream of mannerisms and typical tells, so that anything involuntary is difficult to parse out — but it would probably be so irritating as to get a player ejected!

7 To be clear, that the specific case of the Danes and the Yellow Star is fictional in no way detracts from their heroic wartime history of helping Jews hide and escape.

8 As the FAQ points out, as a practical matter this may not make a difference to a truly empowered adversary with complete oversight of the traffic moving onto and off of your relay — a person who has agents on all sides of you and knows what has been passed and what has not.

9 Some of the quantitative analysis for network and server usage, respectively, will differ for the different 'uses', but the point of the normative argument stands.

10 Again, see the analysis in Gonzalez Fuster (2009), which provides a cogent explanation of an argument for the process of making data fit for an intended, 'primary' use and unfit for further 'secondary' – and non-consensual – uses.

References

Albrecht, K. and Mcintyre, L. (2006) *The Spychips Threat: Why Christians Should Resist RFID and Electronic Surveillance*, Nashville: Nelson Current.

Alexander, J. and Smith, J. (2010) 'Disinformation: A Taxonomy', University of Pennsylvania Department of Computer and Information Science Technical Report No MS-CIS-10-13.

Barocas, S. and Nissenbaum, H. (2009) 'On Notice: The Trouble with Notice and Consent', Proceedings of the Engaging Data Forum: The First International Forum on the Application and Management of Personal Electronic Information, Cambridge, MA, October 2009.

Bok, S. (1999) *Lying: Moral Choice in Public and Private Life (Updated Edition)*, New York: Vintage.

Carlson, R. (2010) 'Rob's Giant BonusCard Swap Meet', available at http://epistolary.org/rob/bonuscard/ (accessed 25 October 2010).

Cohen, F. (2006) 'The Use of Deception Techniques: Honeypots and Decoys', in H. Bidgoli (ed.) *Handbook of Information Security*, Volume 3, New York: Wiley and Sons.

Cockerham, R. (2002) 'The Ultimate Shopper', available at http://www.cockeyed.com/pranks/safeway/ultimate_shopper.html (accessed 19 October 2010).

Duhigg, C. (2009) 'What Does Your Credit-Card Company Know About You?', *The New York Times*, May 12.

Goldman, F. (2007) *The Art of Political Murder: Who Killed the Bishop?*, New York: Grove.

Gonzalez Fuster, G. (2009) 'Inaccuracy as a privacy-enhancing tool', *Ethics and Information Technology*, 12: 87–95.

Hildebrandt, M. (2008) 'Profiling and the Rule of Law', *Identity in the Information Society* (IDIS), 1: 55–70.

Howe, D. and Nissenbaum, H. (2009) 'TrackMeNot: Resisting Surveillance in Web Search', in I. Kerr, C. Lucock and V. Steeves (eds) *Lessons from the Identity Trail: Anonymity, Privacy, and Identity in a Networked Society*, Oxford: Oxford University Press.

Jackson, J. (2003) 'Cards Games: Should buyers beware of how supermarkets use "loyalty cards" to collect personal data?', *Baltimore City Paper*, 1 October.

Lessig, L. (2008) 'Prosecuting Online File Sharing Turns a Generation Criminal', *US News & World Report*, 22 December.

Lieber, R. (2009) 'American Express Kept a (Very) Watchful Eye on Charges', *The New York Times*, 30 January.

Lund, J. and Deak, I. (1990) 'The Legend of King Christian: An Exchange', *The New York Review of Books*, 29 March.

Luo, W., Xie, Q. and Hengartner, U. (2009) 'FaceCloak: An Architecture for User Privacy on Social Networking Sites', Proceedings of 2009 IEEE International Conference on Privacy, Security, Risk and Trust (PASSAT-09), Vancouver, BC, August 2009: 26–33.

Marx, G. (2003) 'A Tack in the Shoe: Neutralizing and Resisting the New Surveillance', *Journal of Social Issues*, 59.

Mercer, D. (2010) 'CDC uses shopper-card data to trace salmonella', *Bloomberg BusinessWeek*, 10 March.

Meyerowitz, J. and Choudhury, R. R. (September 2009) 'Hiding Stars with Fireworks: Location Privacy Through Camouflage', MobiCom'09, Beijing.

Nanex LLC (2010) 'Analysis of the "Flash Crash": Part 4, Quote Stuffing, A Manipulative Device', 18 June 2010, available at http://www.nanex.net/20100506/FlashCrashAnalysis_Part4-1.html (accessed 26 November 2010).

Navarro, J. (2006) *Phil Hellmuth Presents Read 'Em and Reap: A Career FBI Agent's Guide to Decoding Poker Tells*, New York City: Harper.

Netter, S. (2008) 'Wash. Man Pulls Off Robbery Using Craigslist, Pepper Spray', *ABC News*, 1 October.

Nielsen, A. (1952) *What's new in food marketing and marketing research: an address to Grocery Manufacturers of America at Hotel Waldorf-Astoria, New York, N.Y., November 12, 1951*, Chicago: A. C. Nielsen Co.

Nissenbaum, H. (1998) 'Toward an Approach to Privacy in Public: The Challenges of Information Technology', *Ethics and Behavior*, 7: 207–219; reprinted in R. A. Spinello and H. T. Tavani (eds) (2001) *Readings in CyberEthics*, Sudbury: Jones and Bartlett.

— (1999) 'The Meaning of Anonymity in an Information Age', *The Information Society*, 15: 141–44; reprinted in R. A. Spinello and H. T. Tavani (eds) (2001) *Readings in CyberEthics*, Sudbury: Jones and Bartlett.

Pettit, P. (1997) *Republicanism: A Theory of Freedom and Government*, Oxford: Oxford University Press.

Pfaffenberger, B. (1992) 'Technological Dramas', *Science, Technology & Human Values*, 17: 282–312.

Pham, N., Ganti, R. K., Uddin, Y. S., Nath, S. and Abdelzaher, T. (2010) 'Privacy-Preserving Reconstruction of Multidimensional Data Maps in Vehicular Participatory Sensing', WSN 2010: 7th European Conference on Wireless Sensor Networks.

Postman, N. (1990) 'Informing Ourselves to Death', Speech given at the German Informatics Society, Stuttgart, 11 October 1990, available at http://w2.eff.org/Net_culture/Criticisms/informing_ourselves_to_death.paper (accessed 24 November 2010).

Ratcliff, R. A. (2006) *Delusions of Intelligence: Enigma, Ultra and the End of Secure Ciphers*, Cambridge: Cambridge University Press.

Rawls, J. (1971) *A Theory of Justice*, Cambridge, MA: Belknap.

Reiman, J. (1995) 'Driving to the Panopticon: A Philosophical Exploration of the Risks to Privacy Posed by the Highway Technology of the Future', *Santa Clara Computer and High Technology Law Review*, 11: 27–44.

Rothschild, F. and Greko, P. (2010) 'Botnet Resistant Coding: Protecting Your Users from Script Kiddies', paper presented at The Next HOPE, New York, 16 July 2010, available at http://thenexthope.org/talks-list/ (accessed 15 October 2010).

Scott, J. C. (1987) *Weapons of the Weak: Everyday Forms of Peasant Resistance*, New Haven, CT: Yale.

Schulze, H. and Mochalski, K. (2009) *Internet Study 2008/2009*, Leipzig: IPOQUE, available at http://www.ipoque.com/resources/internet-studies/internet-study-2008_2009 (accessed 5 September 2010).

Soghoian, C. (2009) 'Manipulation and abuse of the consumer credit reporting agencies', *First Monday*, 14.

Solove, D. (2008) 'Data Mining and the Security-Liberty Debate', *University of Chicago Law Review*, 74: 343.

Solove, D. and Hoofnagle, C. (2005) 'A Model Regime of Privacy Protection (Version 2.0) (5 April 2005)', GWU Legal Studies Research Paper No 132, available at http://ssrn.com/abstract=699701 (accessed 13 November 2010).

Stead, W. W. and Lin, H. S. (eds) (2009) *Computational Technology for Effective Health Care: Immediate Steps and Strategic Directions*, Committee on Engaging the Computer Science Research Community in Health Care Informatics, National Research Council of the National Academies, Washington, DC: The National Academies Press.

Subramani, M. (2004) 'How Do Suppliers Benefit From Information Technology Use In Supply Chain Relationships?', *MIS Quarterly*, 28: 45–73.

Templeton, B. (2009) 'The Evils of Cloud Computing: Data Portability and Single Sign On', 2009 BIL Conference, Long Beach, California, available at http://www.vimeo.com/3946928 and http://www.acceleratingfuture.com/people-blog/2009/the-evils-of-cloud-computing/ (accessed 5 October 2010).

Waldron, J. (2003) 'Security and Liberty: The Image of Balance', *The Journal of Political Philosophy*, 11: 191–210.

Wohl, R. (1996) *A Passion for Wings: Aviation and the Western Imagination, 1908–1918*, New Haven: Yale.

— (2005) *The Spectacle of Flight: Aviation and the Western Imagination, 1920–1950*, New Haven, CT: Yale.

Zarsky, T. (2005) 'Online Privacy, Tailoring and Persuasion', in K. J. Strandburg and D. Stan Raicu (eds) *Privacy and Identity: The Promise and Perils of a Technological Age*, New York City: Kluwer Publishing.