*Cyber-Surveillance in Everyday Life: An International Workshop * May 12-15, 2011 *University of Toronto*
DRAFT: DO NOT CITE

Vernacular Resistance To Data Collection And Analysis: A Political Philosophy Of Obfuscation [Toronto Research Paper Condensed Draft]

Finn Brunton	**Helen Nissenbaum**
Media, Culture & Communication	**Media Culture & Communication**
New York University	**New York University**
finnbr@gmail.com	**hfn1@nyu.edu**

Abstract: *Computer-enabled data collection, aggregation, and mining dramatically change the nature of contemporary surveillance. Refusal is not a practical option, as data collection is an inherent condition of many essential societal transactions. We present one vernacular response to this regime of everyday surveillance, a tactic we call obfuscation. With a variety of possible motivations, actors engage in obfuscation by producing misleading, false, or ambiguous data with the intention of confusing an adversary or simply adding to the time or cost of separating bad data from good. Our paper develops a political philosophy of obfuscation, linking contemporary and historical cases to develop a descriptive account of obfuscation that is able to capture key commonalities in systems from radar chaff to BitTorrent.*

1. Introduction: Obfuscation as a Counterstrategy to Data-Centric Surveillance

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. Simply refusing to contribute to these profiles and collections is not a practical option: being profiled is the condition of many essential transactions, from connecting with friends in online social networks to shopping and traveling and engaging with many public and private institutions.

In this paper, we develop a political philosophy of obfuscation. Linking contemporary and historical examples, we provide a descriptive account of obfuscation that captures 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 systems for avoiding the privacy harms of location-based services. 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 (supermarkets, or companies like Acxiom), monitoring for intellectual property violations (the RIAA and MPAA), targeted advertising (sites like Google and Facebook), or police actions by repressive governments.

Additionally, 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 of circles of power and influence, or vice versa? Our political philosophy of obfuscation also addresses normative questions of legitimacy, asking whether 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? They can be deemed legitimate assertions of autonomy, or become problematic instances of economic free ridership (relying on others to be less conscientious in muddying their tracks and therefore better targets); they can be justified in resisting the obligation to acquiesce to monitoring, or be destructive acts, poisoning the wells of collective data. Obfuscation, as a tactic both personal and political, offers a platform for studying legitimate and problematic aspects of surveillance and resistance in an age of ubiquitous data capture.

2. Data Surveillance: Asymmetries of Power and Knowledge

To frame both our case studies and the political arguments concerning obfuscation, we need to describe the two asymmetries that characterize systems of personal digital data collection and analysis today. First, the 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. Second, equally important, is an epistemic asymmetry: 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 can be done to it.

Your data is not accumulated in neutral circumstances, whether the collection involves surveillance at the level of infrastructure with which you must participate, or forms which have to be filled out 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 a major power imbalance, between individual consumers and major corporations, or citizens and governments. Obviously there is nothing inherently wrong with gathering data on individuals -- it is in the combination of data gathering with authority and its interests where the problem begins.

It continues once our data has been collected. We don't know whether the company that gathers it will repackage and resell it, whether it will become part of the schedule of assets after a bankruptcy, or whether a private party like ChoicePoint will be collating it with public records and reassembling it in a very different context from our original provision of it. 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, leaving us unable 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.[1] The inconsequential and even benign can quickly become the problematic and sinister.

Furthermore, we don't know what future techniques and databases will enable. Opportunities for the correlation of information tend increase with time. Institutions very rarely voluntarily destroy materials with as much potential as a rich database, and 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

[1] See, for example, the description of these problems in Solove 2008 and Reiman 1995.

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.

The lack of capacity to assess consequences in full 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 them to symmetrical analysis: such organizations might operate under the veil of national security or proprietary trade secrets, and we likely 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.[2]

But are we therefore obliged to turn to methods like obfuscation? What about more traditional and formal means of redress for threats to individual privacy?

3. Standard Means of Redress and Their Shortcomings for Data Surveillance

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, they have not proven responsive to others, and all have particular short-term flaws, which could compound into a future that worries us.

The steady rhetorical drumbeat in the discussion around data privacy is that refusal is a personal responsibility. If you're so offended by the way these companies collect and deploy your data, just don't 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 just two exceptions), but the social and personal cost of refusal is already substantial, and growing. We pay by loss of utility, efficiency, connection with others in the system, capacity to fulfill 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 the more conventional means of regulation -- or resistance.

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, like DoubleClick and ChoicePoint, or has relied on the data generated

[2] As one among many possible examples of our ignorance of the future uses to which our data may be put —whether it's 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 For analysis of this situation from a specifically legal perspective, see Hildebrandt 2008 and Zarsky 2005.

and provided by its customers to improve its operations, like 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 EU's data protection requirements and directives. While our laws will likely be the eventual site of conversation in which we answer, as a society, hard questions about the harvesting and stockpiling of personal information, it operates slowly, and whatever momentum propels them in the direction of protecting privacy in the public interest it is amply counterweighted by opposing forces of vested corporate and other institutional, including governmental, interests. In the meantime and in the near term, enormous quantities of personal data are already in circulation, packaged, sold, and 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 anonymizing datasets, 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 organizations and institutions mediate much data flow is another matter and fraught with politics.

Tools offered to 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 quite privacy-aware and technologically sophisticated. Additionally, there are utility costs: Tor can be slow, for example, and blocked by many large websites. All privacy-protecting technologies entail trade-offs, and those required by robust approaches like 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, only that the wait for relief from these sources is likely to be long. The status quo offers too much gain from the power and epistemic asymmetries that define and entrench our predicament, and all these approaches still leave a gap. From our specific problem of the gathering and analysis of individual data we turn to an array of historical and contemporary examples of obfuscation so we can see it as a general anti-surveillance strategy with many different forms, media, and motives. These examples illustrate some of the ways obfuscation has worked, and highlight systematic features that will be relevant to its evaluation, before we return to its application and related concerns for our particular moment and crisis.

3. 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: relying on temporal limitations; requiring the "network effect" of cooperation or collaboration by groups of obfuscators; selectively interfering with data; and rendering data ambiguous and doubtful for the long term. What follows is a handful of examples from our much larger collection of obfuscation cases.

3.1 **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 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 planes begin to multiply, their 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 terms it the "decoy strategy," and we can indeed consider obfuscation as the multiplication of plausible data decoys[3]). 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.

[3] Cohen, nd.

The *"Craigslist robber"* offers a minor but illustrative example of time-based obfuscation as a practice turned to criminal ends. At 11 AM on Tuesday, the 30th of September 2008, a man dressed like an exterminator in a blue shirt, goggles and a dust mask, and carrying a spray pump, approached an armored car parked outside a bank in Monroe, Washington, incapacitated the guard with pepper spray, and took a substantial amount of money. When the police arrived, they found thirteen men in the area wearing blue shirts, goggles and dust masks—a uniform they were wearing on the instructions of a Craigslist ad which promised a good wage for maintenance work, which was to start at 11:15 AM at the bank's address. This is one of the few real-world examples of the recurrent trope of obfuscation in movies and television: the many identically dressed actors or objects confusing their pursuers as to the valuable one.[4] Obviously it will only take a few minutes to determine that none of the day laborers is the bank robber—but a few minutes is all he needs.

3.2 Cooperative Obfuscation

The cases described so far can be performed by a single actor (perhaps with some unwitting assistants), but other forms of obfuscation require the cooperation of others. They have the "network effect" of becoming more valuable as more people join.

Loyalty card swapping pools provide a superb real-world example. Quite quickly after the widespread introduction of "loyalty cards" for regular shoppers at grocery stores 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 D.C.-area supermarket chain to print out the barcodes of others, pasting them onto their cards.[5] 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.[6] Cardexchange.org is devoted to exchanging cards by mail, presenting itself as a direct analogue to the physical meetups. These sites also act as clearinghouses 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, the more unreliable the data gets.

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

[4] Netter, 2008.

[5] See Rob Carlson's site: http://epistolary.org/rob/bonuscard/, accessed 25 October 2010.

[6] The Ultimate Shopper project: http://www.cockeyed.com/pranks/safeway/ultimate_shopper.html (accessed 19 October 2010).

from others, agreeing to pass messages sight unseen. In return for running a Tor relay, as the FAQ says, "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're on Tor and not running a relay, then someone surveilling you will know you wrote the message passed to their relay. But if you are letting your computer operate as a relay, the message might be yours or just one among many that you're passing on for other people. Did it start with you or not? The information is now ambiguous, and messages you've written are safe in a flock of other messages you pass along.[7]

3.3 Selective Obfuscation

All of the examples thus far have been about general methods of covering one's tracks. But what if you want this 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.[8] It takes the form of a Firefox plugin 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, 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 authorized, when they browse your Facebook page (using the FaceCloak plugin.) Facebook never gains access to it. 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 unauthorized 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.

Tarek Abdelzaher and others on *privacy-preserving participatory sensing* shows us how this idea could work in a very different domain, on an applied and mathematically sophisticated scale.[9] Where a project like FaceCloak obfuscates the data for all but an authorized 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

[7] 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's been passed and what hasn't.

[8] Luo et al., 2009.

[9] Abdelzaher et al, 2010.

location-based tracking. However, Abdelzaher et al. demonstrate how to perturb the data, letting each vehicle continuously lie about its location and speed while maintaining an accurate picture of the aggregate.

3.4 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, and, as with many of the previous cases of obfuscation, opting-out of web search is not a viable choice for the vast majority of users. (At least since 2006, search companies have acknowledged the problem of the collection of query logs, and have offered ways to address people's concerns, though they continue to collect and analyze these logs.) TrackMeNot, therefore, 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 behaviors. This user 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 searcher please stand up? The activity of individuals is masked by that of many ghosts, making the a pattern harder to discern.

CacheCloak, meanwhile, has an approach to obfuscation suited to its domain of location-based services (LBSs).[10] LBSs take advantage of the locative technology in mobile devices to create various services. If you want the value of an LBS, 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.

[10] Meyerowitz and Choudhury, 2009.

4. The Science of Obfuscation

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 swath 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. Given this diversity, can a science of obfuscation exist? Can we create variables and parameters that will enable us to quantify its value and optimize its utility? Can we be sure obfuscation is working?

With encryption, for example, algorithms have 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. We are tempted, for this reason, to characterize obfuscation as a relatively weak practice. Yet, when strong solutions, such as avoidance, disappearance, hiding (e.g. through encryption) are not available and flat out refusal is not permitted, obfuscation may emerge as a plausible alternative, perhaps the only alternative. It simply has to be good enough, a provisional, ad-hoc means to overcome the challenge that happens to be in its way. In our view, this contingency does not mean we throw up our hands to the challenge of a science. Although proof might not be achievable, it would nevertheless still be valuable to be able to assess how to optimize the value of various obfuscation moves, even if only conditionally. Creating such a model is a challenge, to be sure. If there is to be a science of obfuscation it will need to identify key variables and create a systematic way of looking at the relationships between them. The set of variables will undoubtedly be hybrids of the social and the mathematical, including – goals (i.e. time-based, ambiguating, selective), method (i.e. whether group or individual, whether plausible data or obvious noise, whether hiding or protest), adversarial intent and resources (i.e. time, opportunity cost), ratios (i.e. of noise to signal), cost (i.e. to obfuscator, to target), and more.

5. The Politics of Obfuscation

In "A Tack in the Shoe," Marx writes: "Criteria are needed which would permit us to speak of 'good' and 'bad,' or appropriate and inappropriate efforts to neutralize the collection of personal data." Given that obfuscation constitutes a counter-logic to data gathering and profile generation, an intervention to thwart it directly, we might conclude that obfuscation has no ethical or political valence of its own, only to the ends that it serves. If the surveillance in question is morally defensible, thwarting it by any means may be morally problematic, and, *mutatis mutandis*, obfuscation may be justified by unjust data practices**.** Prior to any analysis of ends, however, other moral and political considerations prompted by the very nature of obfuscation—wastefulness, dishonesty, free-riding, and more—deserve to be critically addressed.

Dishonesty

Implicit in obfuscation is an element of dishonesty—it is meant to mislead. Some people might balk at valorizing any practice that systematizes lying. These critics might prefer encryption or silence to producing streams of lies. Excepting 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. We must ask whether the general benefits of lying in a given instance outweigh costs, and whether valued ends are served better by the lie than truthful alternatives.

Free riding

Many forms of obfuscation rely on the compliance of others. As the maxim of the wild has it, no need to be faster than the predator so long as one is faster than other prey. Obfuscation can be seen as two forms of free riding: taking advantage of the willingness of others to allow their data to be aggregated and processed, or enjoying the benefits of services while denying recompense to the targets of one's obfuscation—continuing to enjoy benefits without contributing to the cost by yielding one's own data into the pool. Loyalty card-swapping might also be understood in this light as participants enjoy the bounty of special offers while not contributing to the information pool that presumably enables these economies. Obfuscation, as a good-enough method, often leaves itself open to this critique, as many of its approaches rely on raising the cost of data gathering and analysis just enough to deter the surveillant, which relies on the cost generally being low.

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.[11]

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.

[11] See, among many other stories, this summary of the salmonella outbreak from Businessweek: http://www.businessweek.com/ap/financialnews/D9EC5QUG0.htm.

FaceCloak demonstrates this problem: "[F]or some profile information (e.g., 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."[12] The risk is not only in the present, but 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 malign, this bug becomes a feature. A database interlarded 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, such as overwhelming communications infrastructure with unnecessary requests. 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. The overarching question that drives this paper is about obfuscation aimed at thwarting data surveillance, aggregation, analysis, and profiling, so we confine our evaluation to this arena, using cases we have introduced. One consideration that is relevant across the board is ends; legitimate ends are necessary, though, clearly, not always sufficient. Once we learn, for example, that the Craigslist robber used obfuscation to rob banks, it hardly seems relevant to inquire further whether the lies or free riding were justifiable. Establishing this point is no slam-dunk, but it opens the way to further questions. The question of ends can also be an issue of proportionality, rather than disapproval. The obfuscator running TrackMeNot may not disapprove of the ultimate purpose of Google's query logs—the company makes its revenue from advertising, and it is reasonable for them to automatically serve keyword-specific ads against a given query—but considers the degree of surveillance too extreme. The suggestion of massive data mining to deliver precisely targeted behavioral ads may be too much than the user feels is fair or proportionate, so they turn TMN on. (In this line, obfuscation could actually be helpful to businesses, sending strong signals when they need to bring their practices back into line with consumer expectations and beliefs—not a demand for total reinvention, but matter of degree.)

In cases such as TrackMeNot, CacheCloak, Tor relays, and loyalty card swapping, the 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. We must also consider broader contexts: when protection by law, technology, and corporate best practice fails, protection by obfuscation presents itself as the only resort. Under duress and with little assurance that those extracting information can be trusted, the obligation to speak the truth is certainly lessened. Contrast this with highly controlled environments, such as a

[12]" Luo et al., 2009, p. 6.

courtroom, where a myriad other constraints circumscribe the actions of all parties; we may still speak under duress but epistemic asymmetries are mitigated because of the strictures of context.

While deception may be justified by asymmetries and the absence of alternatives, other critiques remain. 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.

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 porn a *use*. This 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 scrutinized 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.

The use/waste conundrum gains traction from another assumption about the information flows between individuals and the agencies and service providers that monitor and profile them. Individuals using FaceCloak or CacheCloak, even if not “wasting” resources according to widely held social standards, may still draw the ire of Facebook or location-based services. As businesses see it, the users in question are “wasting” their resources because they are depriving them 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. Do we have reason to believe that the services in question are morally entitled to this positive externality? 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.

We hold off responding to these charges until we have discussed the problem of data "pollution" and the propagation of error and inaccuracy. These may be the trickiest of all, and get to the heart of obfuscation. The intention behind inserting noise into the data stream is precisely to taint the resulting body. But there are various ways it can be tainted and some may be more problematic than others. One misspelled 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 quite fair.[13] Further, while some obfuscation practices might confuse efforts to profile

[13] See the analysis in Gonzalez Fuster, 2009, which provides a cogent explanation of and argument for the process of making data fit for an intended, "primary" use and unfit for further "secondary"—and unconsensual—uses.

individuals accurately, they may not render aggregate analysis useless, for example, as in the case of Abdelhazer's work on perturbing individual data while retaining a reliable total picture. But what if none of these mitigations are possible? Where does this leave the ethics and politics of obfuscation?

Those coerced into providing information into the data pool with insufficient assurance over how it will be used, where it will travel, how it will be secured, are being asked to write a blank check with little reason to trust the check’s recipients. Under coercion, obfuscation is not a luxury but an action of last resort. When pushed to the corner, in cases where the issues of extra load on resources, free riding, and data tainting cannot be denied, where the obfuscator acts earnestly to resist the machinations of monitoring and analysis, obfuscation must be evaluated as an act of reasonable and morally sound disobedience.

6. Conclusions

Obfuscation, as we have presented it here, is at once richer and less rigorous than academically well-established methods of digital privacy protection, like encryption. It is far more ad-hoc and contextual, without the quantifiable protection of a cryptographic methods—a "weapon of the weak" to take a term from James Scott. 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 defense 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 and exigencies it creates. Now, we hope the community of privacy researchers and activists will help 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 optimize different obfuscation tactics for different scenarios, and find weak points in the overall strategy?
- Does our description of obfuscation as 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 like 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 to best 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, whatever claims profit-seeking CEOs make about "human nature" and its transformations. Obfuscation is often cheap, simple, crude, clever rather than intelligent, and lacks the polish or freedom from moral compromises that characterizes 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 figure out if it can be made stronger, a more effective bulwark for those in need.

Acknowledgments

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 benefitted 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, Jonathan Smith, and Mireille Hildebrandt for their astute comments, feedback, suggestions and advice.

References

Tarek Abdelzaher et al., 2010. "Privacy-Preserving Reconstruction of Multidimensional Data Maps in Vehicular Participatory Sensing," WSN '2010: 7th European Conference on Wireless Sensor Networks.

Fred Cohen, nd. "The Use of Deception Techniques: Honeypots and Decoys," Fred Cohen & Associates, at http://all.net/journal/deception/Deception_Techniques_.pdf, accessed 12 December 2010.

Gloria Gonzalez Fuster, 2009. "Inaccuracy as a privacy-enhancing tool," Ethics and Information Technology, issue 1, vol. 12, pp. 87 - 95.

Mireille Hildebrandt, 2008. "Profiling and the Rule of Law," Identity in the Information Society (IDIS), 1.1. (2008): 55-70

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

Joseph Meyerowitz and Romit Roy Choudhury, 2009. "Hiding Stars with Fireworks: Location Privacy through Camouflage," MobiCom’09, September 20–25, 2009, Beijing, China.

Sarah Netter, 2008. "Wash. Man Pulls Off Robbery Using Craigslist, Pepper Spray." ABC News, October 1, 2008.

Jeffrey Reiman, 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, no 1 (March 1995), pp. 27- 44.

Daniel Solove, 2008. "Data Mining and the Security-Liberty Debate," University of Chicago Law Review, vol. 74, p. 343.

Tal Zarsky, 2005. "Online Privacy, Tailoring and Persuasion," in: Katherine J. Strandburg and Daniela Stan Raicu (editors). Privacy and Identity: The Promise and Perils of a Technological Age. New York: Kluwer Publishing.