

Voter manipulation: Why psychological profiling can
violate autonomy, and how privacy paternalism can
help

MA thesis (Philosophy, Politics and Economics)

18,161 words

Joseph David Clarke

s1790498

j.d.clarke@umail.leidenuniv.nl

Thesis supervisor: Dorota Mokrosinska

Leiden University, December 2016

Contents

Introduction	2
1. The story of the floating voter	2
2. An intuitive account of autonomy	4
2.1. Developing the intuitive account	5
2.2. Assessing Amy's autonomy using the intuitive account	8
2.3. The problem with the intuitive account	9
3. A hierarchical account of autonomy	11
3.1. Structural hierarchicalism	12
3.2. Proceduralist hierarchicalism	13
3.3. A solution to the problem of authenticity	17
4. Autonomy and paternalism	18
4.1. Internalism, content neutrality and soft paternalism	19
4.2. Anti-aggregation legislation	21
4.3. A hard paternalistic intervention	23
5. The problem with Christman's account	25
6. An externalist-substantivist account of autonomy	27
6.1. Externalism	27
6.2. Substantivism	29
6.3. Four conditions for autonomy	31
7. In support of hard privacy paternalism	33
Conclusion	36
Bibliography	37

"Whenever a campaign or other big organization knows much more about you and your habits than you know about them, any voter is open to manipulation." - Chris Calabrese, privacy lobbyist for the American Civil Liberties Union (in an interview with Beckett, 2012)

Introduction

Many modern political campaigns use psychological profiling in order to influence voting decisions. I argue that this practice threatens the autonomy of voters. In doing so, I develop a theoretical account of autonomy. In order to protect voters from psychological profiling, I suggest a form of "privacy paternalism," which prevents people from acquiescing to the trade and aggregation of their personal data.

My thesis involves two separate claims. The first is that psychological profiling is capable of violating autonomy. The second is that, because psychological profiling can violate autonomy, it should be outlawed. The structure of my paper is as follows. The opening three sections deal with the first claim of my thesis. In section one, I present an imaginary case study in which Amy, an undecided voter, is influenced by a psychological profiling team. In section two, I argue that an intuitive account of autonomy fails to accurately determine whether or not Amy's voting decision is autonomous, because it cannot distinguish between persuasive and manipulative influences. In section three, I argue that John Christman's hierarchical account of autonomy is capable of making such a distinction. Using Christman's account, it becomes clear why psychological profiling violates Amy's autonomy.

The remaining sections deal with my second claim; that psychological profiling should be outlawed. In section four, I propose that this can be done via a form of "privacy paternalism," which conflicts with Christman's account of autonomy. In section five, I expose the flaws in Christman's account. In section six, I argue that Christman's account should be modified. In section seven, I argue that, under my alternative account of autonomy, "privacy paternalism" is justified.

1. The story of the floating voter

Amy is an undecided, "floating" voter, taking a keen interest in this year's presidential election. One of the candidates is a populist politician who voices strongly xenophobic opinions. At first, Amy found his views morally abhorrent, because she sincerely identifies as a global egalitarian; she believes that all human beings should be treated as moral equals. However, most of Amy's friends feel the opposite, and they frequently share articles and posts on social media which are supportive of the

politician. Over time, Amy has begun to normalize and internalize these xenophobic opinions, occasionally mimicking them in her own online posts.

Amy watches the presidential debates on cable television, reads online news articles, and voices her opinions about the candidates on social media. Information about all of these activities is collected and stored in databases, a process to which Amy agreed when signing up to each service. She also acquiesced to the potential future sale of that data to third parties.

Part of the populist politician's campaign involves the use of a psychological profiling team. Because Amy exchanged her personal data in return for access to digital services, the team is able to influence her actions. They purchase data about Amy from the digital services that she uses, and from other sources including government censuses, voter registration databases, call logs from mobile phone companies and consumer purchase records. This data is then aggregated into a complex psychological profile, capable of revealing Amy's innermost beliefs and desires. An algorithm identifies Amy as a "persuadable voter" who cares primarily about security, but also, increasingly, about immigration. A psychologist, working for the profiling team, accurately ascertains that Amy has a "neurotic" personality.

On election day, Amy remains undecided. She has strong reservations regarding the morality of voting for the populist. Then she receives a telephone call from someone on that politician's campaign team. He reads from a script which has been specially tailored to influence Amy. Given Amy's "neuroticism", the campaigner's computer prompts him to use emotive, fearful language in their conversation. He strongly emphasizes the fact that some immigrants are planning to perform terrorist attacks within the country. This stimulates Amy's fears about security.

Later that day, Amy decides to vote for the populist candidate. She walks to the polling station and puts a cross in the appropriate box on her ballot paper.

Two years later, Amy has moved to a different area and left her old circle of friends behind. When debating moral principles with her family, she continues to promote her beliefs in global egalitarianism. Her teenage son recalls that she voted for the populist politician, and points out the contradiction between his policies and Amy's deeply held beliefs. That night, Amy reflects on her voting decision and is full of regret – she doesn't feel as if that action authentically represented her values.

While this particular case study is a work of imagination, modern political campaign teams employ psychological profiling in these exact ways in order to influence voting decisions. The technique was pioneered in the United States by Barack Obama in his 2008 and 2012 electoral campaigns, which

“used precision targeting to build rich data-driven profiles of every potential voter in the United States, using this data to precisely tune its messaging at the person level” (Leetaru, 2016). Obama’s team identified “persuadable voters” and then “directed volunteers to scripted conversations ... with the objective of changing minds” (Issenberg, 2012). In 2016 the British “Vote Leave” campaign (Payne, 2016), the Republican nominee Ted Cruz (Hamburger, 2015) and the president-elect Donald Trump (Kranish, 2016) employed the services of Cambridge Analytica. This company “incorporates private profile data from tens of millions of Facebook users” with more than “50,000 data points gathered from voting records, popular websites and consumer information such as magazine subscriptions, car ownership and preferences for food and clothing” (Leetaru, 2016). The chief executive of the company claimed that his behavioral psychologists were able to use this data to “determine the personality of every single adult in the United States of America,” and could thereby “identify millions of voters who are most open to being persuaded to support Trump” (Kranish, 2016).

Although most of the publically available information about this form of psychological profiling comes from the United Kingdom and the United States, it would be a mistake to assume that it is an isolated phenomenon. Cambridge Analytica claims to have “worked on campaigns in 22 countries” (Ibid.). This fact, coupled with the global influence wielded by the United Kingdom, the United States, and their corporations, entails that the use of psychological profiling to influence voting decisions should be of worldwide concern.¹

2. An intuitive account of autonomy

I claim that the use of psychological profiling to influence voting decisions should be outlawed, because it puts the autonomous status of these voting decisions under threat. The strength of my thesis thus depends on a convincing descriptive account of personal autonomy. In this section, I present some common sense, intuitive assumptions surrounding the concept of autonomy. Then, I show that an intuitive account of autonomy fails to accurately determine whether or not Amy’s voting decision was autonomous, because it does not give us a clear answer as to whether the influence of the profiling team upon Amy’s desires was persuasive or manipulative. This entails that a theoretical account of autonomy is required.

¹ This paper focuses exclusively on the relatively new use of psychological profiling to influence voting decisions. However, it must be noted that the practice has been an established part of the advertising industry for the past twenty years, as an effective method for influencing consumer decisions (Gunter, 2016). Arguments similar to those put forward in this paper could also be used to demonstrate that psychological profiling threatens the autonomy of *consumers*, and that targeted advertising should thereby be outlawed.

2.1. Developing the intuitive account

We can attain an elementary understanding of the concept of autonomy by studying its etymology. “Auto” means self, and “nomos” means governance. “Self-governance” loosely refers to other concepts such as self-legislation, self-determination, self-ownership, and personal sovereignty; all of which pertain to the notion of having control over one’s life (Ashley, 2012, p.1). My concern is solely with personal autonomy (the self-governance of an individual) rather than collective autonomy (the self-governance of a group). Personal autonomy can be attributed to a variety of objects: actions, decisions, and agents (an agent being a person that acts). Since I am discussing the status of voting decisions, I will usually speak in terms of whether or not an agent’s *decision to act* is autonomous. Nothing will turn upon this terminological point, however, since the objects of autonomy attribution are often interchangeable. That a *decision to act* is autonomous implies that the *action* based on that decision is autonomous, from which it also follows that the *agent* performing the action is (at least at that particular time) acting autonomously.

It must also be noted that I am not discussing the notion of Kantian autonomy; under which, to be autonomous, an agent must *not* be motivated by their personal beliefs and desires. Many contemporary theorists hold instead that the inner mental world of the agent does play a vital role in self-governance (Christman, 1991; Dworkin, 1988; Frankfurt, 1988). Under this conception of personal autonomy, an agent is autonomous if she is capable of governing herself so that the desires that motivate her choices are those “by which one genuinely wants to be motivated, rather than ... by unconscious prejudices, compulsive obsessions, or the like” (Anderson, 2008, p.8). An important intuition here is that most adults, most of the time, should be described as autonomous; as having the ability to decide for themselves which actions to perform. As Gerald Dworkin (1988, p.31) notes, “any feature that is going to be fundamental in moral thinking must be a feature that persons share.” A satisfactory descriptive account of autonomy must thereby delineate the conditions that must be present in order for us to ascribe autonomy to agents. At the same time, these conditions must not be so stringent that they exclude everyday cases of self-governance.

Since my thesis is that psychological profiling should be outlawed because it threatens autonomy, it presupposes the *importance* of autonomy. It is not controversial to claim that autonomy is valuable. When we ascribe value to something, we either do so intrinsically (*x* is valuable for its own sake) or instrumentally (*x* is valuable because it is a means to some other end that has intrinsic value). Autonomy certainly has instrumental value; if agents are able to exercise their autonomy, then they are able to pursue whatever course of life they think is best for them (Wall, 2003). But it is not the case that self-governance is only one means, among many possible separate means, of achieving

well-being. Rather, it seems that self-governance is a *necessary* element of well-being (Mill, 2003, p.121) and that autonomy therefore also has intrinsic value. We describe slavery as oppressive and abhorrent precisely because it is a life in which individuals are not sovereign over their own actions. Thus, the opportunity to determine and carry out one's life plans seems necessarily constitutive of a good life. John Christman (2009) explains why one has an inherent interest in being able to govern one's own actions, rather than those actions being governed by external forces:

“to be autonomous is to be one's own person, to be directed by considerations, desires, conditions, and characteristics that are not simply imposed externally upon one, but are part of what can somehow be considered one's authentic self. Autonomy in this sense seems an irrefutable value, especially since its opposite — being guided by forces external to the self and which one cannot authentically embrace — seems to mark the height of oppression”.

The intrinsic value of autonomy is also used as justification for one of the fundamental tenets of political liberalism.² *The principle of state neutrality* holds that, since societies are composed of citizens with different values, the state should refrain from “endorsing any particular comprehensive conception of the good life and what is truly valuable” (O’Shea, 2011, p.25). This principle is supposed to prevent undue external interference by the state with the autonomous decisions of its citizens. The value of autonomy is thus thought to protect citizens from both the perfectionist conviction that they ought to extol particular values, and paternalistic laws which enforce such values upon them.³ Therefore, for both moral and political reasons, autonomy is taken to be valuable for its own sake.

Since autonomy is a valuable part of human life, it is important for us to be able to accurately distinguish between autonomous and non-autonomous actions. The layman might equate autonomy with freedom, but it is not the case that an agent is autonomous merely because they have the physical or legal freedom to act upon their desires. Autonomy refers (at least primarily, if not exclusively) to the inner mental world of individual agents, rather than the external state of affairs that agents are situated in. For example, it is possible for an agent to be physically free, and yet non-autonomous; imagine a person with bipolar disorder who makes impulsive choices during a manic episode. Conversely, an agent could lack physical freedom and yet be fully autonomous, such as a political prisoner who refuses to renounce his revolutionary views. Thus, an important facet of autonomy ascription is that an agent has the internal ability to decide for herself which actions to perform. This common sense intuition - that autonomy depends crucially on the internal mental conditions of the agent - is labeled by theorists as “internalism.”

² Political liberalism is perceived to be the foundation of modern Western political thought.

³ The relationship between the concept of autonomy and the justification of government policies is pertinent throughout this paper, especially where my version of “privacy paternalism” is considered.

Tom Beauchamp's (2009, p.65) analysis of autonomy begins with three assumptions regarding the internal mental conditions of autonomous agents - that they act (1) intentionally, (2) with understanding, and (3) voluntarily. Although this paper is not concerned with Beauchamp's own separate theory of autonomy, the three intuitions he identifies help us determine the common sense view of the concept. I concur with Beauchamp (2009, p.62) that "any theory that classifies acts as nonautonomous when these acts are well understood, intentional, and not controlled by others is a conceptually dubious theory of autonomy." Since the aim of this section is merely to lay some preliminary building blocks for my theory of autonomy, I will paint these intuitions in as simple and uncontroversial way as possible here. Throughout the course of my analysis of autonomy, they will be scrutinised in more detail (the condition of voluntariness in particular).

(1) Intentionality. We have the intuition that autonomous actions involve intentionality. That an action is intentional means that the agent meant to carry out the action – she had a plan involving the reason for the action and the hoped effects of the action. Intentional actions can be contrasted with accidental actions. Take Donald Davidson's (1963) famous example of a burglar who is startled when the homeowner happens to switch on her kitchen light. Intuitively, the homeowner's turning on of the light was an autonomous action, but her alerting of the burglar was not. The difference between intentional and non-intentional acts depends on the beliefs and desires of the agent. The homeowner desired to turn on the light, and believed that her action would turn on the light. She did not believe that, or desire for, her action to alert the burglar. Thus, it seems that, to be autonomous, an agent must choose to perform an action on the basis of internal reasons such as beliefs, desires and values.

(2) Understanding. Relatedly, some sort of appropriate understanding of one's action is required for autonomy. The homeowner understood that flicking the switch would cause the light to come on, but she did not understand that flicking the switch would cause the burglar to be alerted. For an action to be described as autonomous, then, an agent must understand the nature and consequences of her actions. It does not appear, however, that an *exhaustive* level of understanding is requisite for autonomy. Many people autonomously drive their cars without fully comprehending what is going on underneath the bonnet. So, intuitively, what is necessary for autonomy is merely some sort of minimal understanding of the pertinent information related to one's action.

(3) Voluntariness. Imagine that, as the homeowner enters the room, the burglar puts a gun to her head and commands her to turn on the light. In this case, her action is not autonomous. This is because her decision to turn on the light is not a free and voluntary choice – it is influenced by the burglar's demand. *Influence* occurs when one's beliefs, desires or actions are affected by others. If

an agent's action has been influenced by someone else, we often describe them as non-autonomous. But this is not always the case. Sometimes, a choice can be influenced, and yet still appear to be voluntary and autonomous. Influence can be understood as comprising of three subcategories – coercion, persuasion and manipulation (Beauchamp, 2009, p.69). Coercion, which often involves a threat of physical violence (like the burglar's above), is a relatively obvious form of influence that I will not pay much attention to here. A more interesting topic involves the blurred lines between persuasive and manipulative influences. Persuasive influences seem compatible with voluntariness, and thereby with autonomy. An agent can be persuaded to change their mind on the basis of a reasoned argument. For example, one can influence a friend's decision to become vegetarian by appealing to their pre-existing desire to protect the environment. Intuitively, if their subsequent decision to become vegetarian is a voluntary choice based on consciously entertained reasons, we describe such a change of heart as autonomous, despite the external influence. On the other hand, manipulation - influence that involves lies or deceit - seems incompatible with voluntariness, and thereby with autonomy. In one of Shakespeare's (1999) famous plays, Othello decided to murder Desdemona for revenge, because Iago deceitfully convinced him that she had been unfaithful (when she had not). Here, Othello's decision to commit murder was not fully voluntary; he did not appear to be in complete control of his actions. Thus, voluntariness (and thereby an absence of manipulation) seems necessary for autonomy.

A common sense, intuitive account of autonomous action thus requires that agents act intentionally, with understanding, and voluntarily. Is this intuitive account of autonomy sufficient for our understanding of the concept? I argue that it is not. Amy seemed to be acting intentionally, with understanding, and voluntarily. Yet she also appeared, intuitively, not to be fully in control of her actions. In light of cases such as these, where the distinction between persuasive and manipulative influence is uncertain, our intuitions about autonomy must be supplemented with other, more theoretical considerations.

2.2. Assessing Amy's autonomy using the intuitive account

Should Amy's decision to vote for the populist politician be deemed autonomous or non-autonomous? Employing Beauchamp's intuitive criteria for autonomy, it seems like Amy does act intentionally, with understanding, and voluntarily. Certainly, Amy's vote was an intentional action. We saw that she internalised the politician's opinions, identifying with them herself. Thus, at the time of her placing the vote, the xenophobic opinions which motivated this action did in an important sense belong to Amy. Furthermore, the action was not accidental; her pencil did not slip in the voting booth. She believed that going to the polling station and putting her cross in the

relevant box would enable her to vote for the populist, and she had a corresponding desire to do so. She performed the action on the basis of those internal reasons.

Did Amy understand the nature and consequences of her actions? Granted, she did not have an exhaustive understanding of the consequences of her vote; none of us can predict the future, especially not in the political realm. But she surely did have the appropriate minimal level of understanding that we intuitively deemed as requisite for autonomous action. She understood that by voting for the populist, she was increasing the chances of his election. Furthermore, she understood that his election would lead to the implementation of some of his xenophobic policies.

Was Amy's vote voluntary? Nobody physically or coercively forced Amy to vote for the populist. Of course, her decision was influenced by the campaign team, but does that influence count as autonomy-conferring persuasion or autonomy-undermining manipulation? It is not unreasonable to argue that Amy was merely *persuaded*. The campaign team did not plant a chip in Amy's brain that radically changed her beliefs and desires. They did not even lie to her in order to influence her decision. Rather, they identified the issues in which Amy had already expressed an interest, and used them to persuade her of their case, much like the example of the vegetarian who appealed to the environmental concerns of her friend to persuade him to stop eating meat. Since we deemed that decision to become vegetarian as autonomous, it looks as though we must also deem Amy's voting decision likewise.

Despite all these considerations, however, the conclusion that Amy's decision was autonomous remains dubious. On the contrary, her decision appears to have been non-autonomous, primarily because, as I will argue in the following sections, the influence of the psychological profiling team should be deemed manipulative rather than persuasive. Therefore, our intuitions surrounding the concept of autonomy (and in particular the condition of voluntariness) are insufficient for a complete understanding of autonomy.

2.3. The problem with the intuitive account

Given the ambiguity surrounding the intuition that autonomy requires voluntariness, the intuitive account is unable to accurately distinguish between manipulation and persuasion. In this section, I propose that including the condition of *authenticity* as a necessary component of autonomy enables us to make such a distinction.

Recall that (aside from accidental or subconscious actions) agents usually perform actions on the basis of intentions to act. In this sense, one has an unyielding authority over one's own actions. So

long as one's physical body is not strapped to a robotic frame controlled by a mad scientist, it is impossible for one's actions to be caused by an intention to act that does not belong to oneself. But this necessary authority that one's intentions to act wield over one's actions does not guarantee one's autonomy. This is because the mental states that *motivate* one's intentions to act may themselves have been influenced by external factors. One's *actions themselves* cannot be directly influenced by anything other than one's own intentions; but the beliefs, desires and values that *motivate* one's intentions and thereby govern one's actions *are* capable of being influenced. In this way, one's power to determine *how* to act – one's autonomy - can be compromised (Buss, 2013). For example, Amy was not physically coerced into performing the *action* of voting for the populist, but her *desire to act* in this way was clearly influenced.

Recall that influence occurs when one's beliefs, desires or actions are affected by external forces, rather than oneself. But *all* of the desires that motivate one's actions are influenced, to some extent, by things external to oneself.⁴ Autonomy ascription depends on distinguishing between those influences upon action that are compatible with autonomy, and those that threaten autonomy. More specifically, it involves "distinguishing those ways of influencing people's reflective and critical faculties which *subvert* them from those which *promote and improve* them" (Dworkin, 1988, p.18, emphasis added). We have seen that persuasive influences are compatible with autonomy, while manipulative influences are not. Thus, a satisfactory descriptive account of autonomy must provide us with the tools to distinguish between persuasive and manipulative influences.

This distinction depends in turn upon the ability to draw a clear conceptual line between those desires that we describe as part of an agent's "authentic" self, and those which an agent cannot authentically embrace.⁵ Autonomy requires not only that we are able to consciously act on the basis of reasons, but also that these reasons reflect our true selves. For an agent's decision to be autonomous, it must be "authentic" – it must "genuinely express a person's identity, commitments or goals" (O'Shea, 2012, p.9).⁶ To understand this idea of authenticity, consider the example of a drug addict who sincerely desires to become clean, but continues to use the drug nonetheless. She obviously fails to govern herself, but, under Beauchamp's criteria, it is difficult to explain exactly why this is so. Her action appears, we can imagine, to be performed intentionally, with understanding, and voluntarily. Intentionality is present because the desire to use does not destroy the authority

⁴ For brevity, I will focus on desires, but the reader should bear in mind that my argument applies equally to other motivating mental states such as beliefs and values.

⁵ The connection between autonomy and authenticity was first introduced by Gerald Dworkin (1976, pp.24-25).

⁶ Of course, the requirement of authenticity assumes that people are capable of possessing such "authentic" qualities. Readers sceptical of such an assumption should consider something they care about very deeply; the love for one's family, for example. It would be difficult to describe actions motivated by this value as anything but a genuine expression of one's identity, commitments and goals.

that the addict had over her decision to act. We can also suppose that she fully understood the nature and consequences of the action at the time which she performed it. And the action can be described as voluntary, since it was not the case that somebody forced her to inject herself under duress of physical violence; she did it by herself, alone in her room.

What is missing from the intuitive account is a distinction between authentic and inauthentic desires. As a long term life goal, the addict wants to stop using drugs. This desire can be described as authentically hers; she fully and genuinely identifies with it. But she also has a strong, immediate urge to use again. This desire is inauthentic; she often rejects this desire, and wishes that it did not motivate her actions. Inauthentic desires are often alienating in this way because they are the product of some manipulative or oppressive external force; in this case, the addictive power of the drug. We *would not* ascribe autonomy to the agent if she *continued* to use the drug, since her desire to use is an inauthentic one. We *would* ascribe autonomy to the agent if she *stopped* herself from using the drug, since this is what she authentically wants. Consideration of the authenticity of the addict's desires explains why the fact that the addict's actions are motivated by the desire to use, rather than the desire to stop using, leads us to claim that her autonomy has been undermined.

So, for an action to be autonomous, it must have been motivated by an authentic desire. Introducing authenticity as a necessary condition of autonomy provides us with a tool for distinguishing between legitimate *persuasive* influences upon desires, and unwarranted *manipulative* influences upon desires. Desires influenced by persuasion are authentic, whereas desires influenced by manipulation are inauthentic. For example, if Amy's desire to vote for the populist was *not* influenced by psychological profiling, and was instead formed solely as a consequence of persuasive arguments made by her friends, then she would be likely to authentically identify with that desire. As we have seen, however, Amy felt alienated from her desire to vote for the populist. Since she did not fully and genuinely identify with it, the desire must be deemed inauthentic. And since this inauthentic desire was caused by the psychological profiling team, their influence must be deemed manipulative rather than persuasive.

The problem of authenticity, then, denotes the difficulties theorists face in describing exactly what it is that makes a desire authentic, rather than inauthentic. In the next section, I will argue that a particular theoretical account of autonomy provides a convincing answer to this problem.

3. A hierarchical account of autonomy

The first claim of my thesis is that the use of psychological profiling to influence voting decisions can violate autonomy. To assess this claim, a theoretical account of autonomy is required, which also

captures our pre-theoretical intuitions regarding the concept. In this section, I argue that an answer to the problem of authenticity can be provided by a particular version of a hierarchical account of autonomy. I begin by exposing the flaws in Harry Frankfurt's hierarchical account. This enables us to see the strengths of John Christman's account, which, I argue, provides a convincing solution to the problem of authenticity.

The basic premise of all hierarchical accounts is that our mental lives are structured in a certain way: we have both lower-order desires and higher-order desires. A lower-order desire is a *mere* desire regarding the actions of an agent; *a desire to do X*. A higher-order desire is a more *considered* desire, because it is *about* lower-order desires; *a desire to desire to do X* (Ashley, 2012, p.14). The implication of this is that one's higher-order desires reveal what one authentically wants; that they are the site of one's "true self" (Anderson, 2008, p.10). According to hierarchical accounts, then, actions motivated by lower-order desires must ultimately be guided by authentic higher-order desires, if they are to be deemed autonomous.

3.1. Structural hierarchicalism

Harry Frankfurt (1988) developed the first hierarchical account. His answer to the problem of authenticity is that autonomy involves the "reflective endorsement" of one's lower-order desires in light of one's higher-order desires. If my action is motivated by lower-order desires that I *endorse* in light of my higher-order desires, then I am acting autonomously. If my action is motivated by lower-order desires that *contradict* my higher-order desires, then I am acting non-autonomously. Thus, for Frankfurt, endorsement is a form of authentication. For example, the addict may have a lower-order desire to use drugs, but this desire cannot be authentically endorsed in light of her higher-order desire to not be a drug addict. If at any moment the addict wants to use drugs, then she wants at that moment what she does not authentically want, and is thereby non-autonomous. Thus, autonomy consists of both the ability to authentically identify with one's lower-order desires (by reflecting upon their coherence with one's higher-order desires), and the ability to reject (or refuse to act upon) those lower-order desires that contradict one's higher-order desires.

This approach might appear plausible at first glance. It accounts for the familiar way in which we sometimes, upon reflection, choose to reject our own desires. It also captures the common malaise we experience when deciding between doing what we *want* to do, and doing what we know we *should* do. However, on closer inspection, it becomes clear that Frankfurt's approach cannot explain the difference between *authentic* higher-order desires and *manipulated* higher-order desires.

Although an agent's ability to reflect upon their higher-order desires seems necessary for autonomy

ascription, this does not, as Dworkin (1988, p.18, emphasis added) writes, tell the “whole story of autonomy...*the choice of the kind of person one wants to become* may be influenced by other persons or circumstances in such a fashion that we do not view those evaluations as being the person’s own.” A good account of autonomy must be capable of identifying and excluding higher-order desires that have been shaped by manipulative influences.

Frankfurt’s hierarchical account is labelled “structuralist” (Buss, 2016), since it sees autonomy as resting solely on the way that an agent’s desires are structured. Structuralist accounts are weak because they ignore the fact that agents are “diachronic”; they exist over a long period of time. Frankfurt’s approach assesses the structural relationship between an agent’s desires at a particular time, and does not take into account their historical origins. For example, under Frankfurt’s account, Amy’s voting decision would be described as autonomous. Frankfurt would assess Amy’s desires at the moment of action, and determine that her lower-order desire to mark her ballot paper by putting a cross next to the populist’s name was endorsed by her higher-order desire to vote for the populist. Because of this endorsement, Frankfurt would hold that Amy’s lower-order desire to mark her ballot paper in that way was authentic, and that her action was thereby autonomous.

But this conclusion is counterintuitive. At the moment of action, Amy did genuinely desire to put her cross in the populist’s box. However, we do not accept that the higher-order desire to vote for the populist which caused her to form the lower-order desire to mark the ballot paper was one of her authentic desires. Amy’s power to *decide* how to behave had been manipulated by an antecedent factor, namely, the psychological profiling team, which caused her desire to vote for the populist. Accurate autonomy ascription cannot be achieved solely through an assessment of the conditions surrounding an action at the time which it is performed. Since Frankfurt’s account does not consider how higher-order desires are *generated*, it cannot distinguish between cases in which one’s higher-order desires are authentic, and cases in which one’s higher-order desire has been manipulated by an external force. Thus, Frankfurt’s account cannot solve the problem of authenticity.

3.2. Proceduralist hierarchicalism

Another type of hierarchical account, developed by John Christman, offers a more promising solution to the problem of authenticity. Christman, like Frankfurt, holds that higher-order desires are capable of expressing one’s authentic, autonomous considerations. However, unlike Frankfurt, he does not assume that manipulation is absent if a lower-order desire is endorsed by a higher-order

desire.⁷ In order to account for the possibility of manipulated higher-order desires, Christman shifts the focus away from the endorsement of a desire at the moment of action, and onto the historical process of desire formation. Thus, his account is described as “proceduralist” rather than merely “structuralist” (Buss, 2016). Instead of requiring that an agent must reflectively endorse her desire in order to be autonomous with respect to it, proceduralism requires that an agent must reflectively endorse the *process* by which she ended up coming to have that desire.

For Christman (1991, p.11), an agent is autonomous relative to some desire (at time t) if and only if:

- i) The agent did not resist the development of the desire (prior to t) when attending to this process of development, (or would not have resisted that development, had she attended to the process).
- ii) The lack of resistance to the development of the desire (prior to t) did not take place (or would not have) under the influence of factors that inhibit self-reflection.
- iii) The self-reflection involved is (minimally) rational and involves no self-deception.

The first condition pinpoints the historical process of desire formation as crucial for determining authenticity. The second and third conditions impose limits upon an autonomous agent’s cognitive capacities. An agent who meets all three conditions acts upon the basis of an authentic desire, and must thereby be deemed autonomous.

Considering the first condition, an agent “attends to” the development of a desire when she is able to fully consider a complete description of the causal processes that led her to have this desire. Imagine that our reluctant addict, recently released from rehab, pays conscious attention to the formation of a resurgent lower-order desire to use the drug. She becomes aware of her physical cravings and an inner mental conflict between two opposing desires. She consciously entertains her higher-order desire not to use the drug, and thereby attempts to resist the process by which she forms the lower-order desire to use. Thus, under Christman’s first condition, her lower-order desire to use is inauthentic, and her eventual use of the drug constitutes a non-autonomous action.

It is important to note that the question of whether or not an agent resists the desire formation process can also be considered hypothetically. The addict was able to fully attend to the development of her desire, but in most cases, agents form desires without full comprehension of the causal steps that lead to them. For example, recall that Amy posted xenophobic messages on social

⁷ To give credit to Frankfurt (1987), he identifies that this assumption is flawed. As a solution, he suggests that autonomous endorsement must be “wholehearted,” in the sense that the agent’s endorsement is devoid of all doubt. But this account is unconvincing, because it seems to permit cases of self-deception (Anderson, 2008, pp.11-15; Taylor, 2008, p.7).

media. Suppose that she developed this lower-order desire to share xenophobic content because she had a higher-order desire to be accepted by her peers, who would approve of such posts. Since Amy was not consciously aware of the way in which her desire to post the messages was formed, we must consider whether she *would have* resisted the formation of the desire if she *had* paid conscious attention to it. And given Amy's higher-order desire to gain acceptance among her peers is a genuine and authentic one, we can reason hypothetically that the desire formation process would be acceptable to her, and that her decision to post xenophobic messages was thereby autonomous.

There are other cases, however, where despite the fact that an agent accepts the desire formation process, that desire should nevertheless be deemed inauthentic. I am talking here about cases in which a manipulative influence is responsible for the agent having accepted the process by which the desire in question was formed. Hence the second condition, which holds that a desire is not authentic if the agent's acceptance of it was caused by "factors that inhibit self-reflection".

Ultimately, for Christman (1991, p.11), "autonomy is achieved when an agent is in a position to be aware of the changes and development of her character and of why these changes come about. This self-awareness enables the agent to foster or resist such changes." That self-reflection requires the ability to both become *aware of* and *resist* the desire formation process aligns with our intuitions that autonomy requires both understanding and voluntariness. For a desire to be deemed authentic, we must hypothesise that a minimally rational agent would understand (by being "aware of") and voluntarily accept (by not "resisting") the process by which it was formed. The inference I have made, on the basis of Christman's account, is that desires are only authentic if the agent both understands and voluntarily accepts (or would accept) the influences that led him to accept it. To put the point in another way, we could say that an agent's desire is authentic if they provide (or would provide) their "informed consent" to the influences that caused it.

Here, the advantage of proceduralism over structuralism can be highlighted. Christman, unlike Frankfurt, seems able to distinguish between cases in which one's higher-order desires are *authentic*, and cases in which one's higher-order desire has been *manipulated* by an external force. For example, the prospective vegetarian understood the nature of his friend's attempt to influence his desires, and he voluntarily and consciously chose not object to this influence. Thus, his desire to stop eating meat was authentic. But in other cases, an absence of understanding necessarily precludes the possibility of actual voluntariness. Othello *did reflect* on some of the reasons behind his desire to kill Desdemona (such as his belief that she was unfaithful) and he found those reasons acceptable. However, Othello did not have access to the *full story* behind the development of his desire, because he was *unaware* of Iago's manipulative influence. Since Othello's acceptance of the process by which

he formed the desire to murder was influenced by Iago, his reflection upon that process does not meet Christman's second condition of competence; he was under the influence of factors that inhibit self-reflection. Furthermore, we can hypothesise that if Othello *had* been aware of Iago's deception, he surely *would* have rejected the desire to kill Desdemona. Thus, the case also fails to meet the hypothetical strand of Christman's first condition. We can thereby infer from Christman's account that Othello's desire to kill Desdemona was inauthentic, and his subsequent action was non-autonomous.

Aside from manipulation, there are other factors that could prevent agents from providing their informed consent to the way that their desires are formed. Christman's third condition addresses this issue. The third condition, like the second, ensures that autonomous agents have the necessary cognitive capacities to adequately reflect upon the development of their desires. But while the second condition ensures that the agent's mental states are free from external influences, the purpose of the third condition is to ascertain that the agent employs the correct techniques of reasoning when attending to the desire development process. If an agent attends to that process, but does so *irrationally*, then their acceptance of the process would not suffice for the desire to be deemed authentic.

A common worry with rationality requirements is that they are often too stringent. An overly demanding rationality requirement would imply that agents are not autonomous if they fail to make decisions that maximise their own well-being. But the human capacity to reason analytically is often overestimated (Kahneman, 2011), and stringent rationality conditions would thus contradict our pre-theoretical intuition that most people are autonomous most of the time. Therefore, an acceptable account of autonomy must allow autonomous agents the space to make sub-optimal or unwise choices. Even the seemingly minimal stipulation that a belief must be coherent with one's wider set of beliefs is too demanding; it is not clear whether anyone's set of beliefs is entirely devoid of conflict. Thus, Christman's condition of *minimal* rationality requires only that an agent does not consciously and concurrently entertain desires that obviously conflict; such as desiring to both use and not use a drug. If one is consciously aware of such a "manifestly inconsistent" (Christman, 1991, p.15) contradiction within their reasoning, and yet still acts on the basis of such reasoning, then they

are either reflecting irrationally, or they are partaking in self-deception. Neither of these processes are compatible with autonomy.⁸

I have outlined Christman's three conditions for autonomy. In the next section, I will show how this account helps us answer the question of whether Amy's desire to vote for the populist is authentic.

3.3. A solution to the problem of authenticity

The intuitive account struggled with the question of whether the influence of the profiling team upon Amy's desires was manipulative or persuasive. Since this means that it failed to determine the authenticity of Amy's desire to vote for the populist, the intuitive account was incapable of accurately ascertaining whether or not her voting decision was autonomous. Christman's authenticity conditions enable us to identify the influence of the profiling team upon Amy's voting decision as manipulative, thereby allowing us to make sense of the intuition that her decision was non-autonomous.

Of course, Amy meets Christman's third condition, since it is presumed *ex hypothesi* that she is minimally rational and not self-deceiving. But she does not meet the other two conditions for authenticity. Recall that Amy developed her desire to vote for the populist after the phone call with the psychological profiling team. It is true that she did not resist the development of the desire at the time. But, under Christman's account, a simple absence of resistance is not enough to secure authenticity. It is only the case that Amy's desire to vote for the populist was authentic if she provided (or would have provided) her "informed consent" to the influences that caused it.

It is clear, however, that reflection-inhibiting factors prevented Amy from properly consenting to the influence of the profiling team. Amy did not fully understand the causes of her desire, and was thereby unable to voluntarily resist the influence of the profiling team. As far as Amy was concerned, the caller had access to, at the most, very minimal personal information about her. She was not aware that the team had access to her private Facebook messages, her Twitter comments, her purchase history, and so on. She did not know that this data had been aggregated into a model of her beliefs and desires more accurate than even she would be able to manufacture herself. She was unaware that psychologists provided information and cues to the caller in order to influence her

⁸ The concepts of self-deception and manipulation might appear entangled; both involve a process whereby an agent's higher-order desires are overruled. The difference is that during self-deception, the agent actively and consciously engages in this process, whereas during manipulation, the agent is passive and unaware. Christman's requirement that self-deception must be absent is largely irrelevant for my project. I focus on cases of psychological profiling, in which *passive* agents are *unaware* of an external influence upon their reasoning. Thus, the presumption that agents are not deceiving themselves is implicit in the following examples and discussion.

voting decision. Amy's ignorance of these factors clearly inhibited her potential for reflection, thereby violating the second criterion for autonomy. Furthermore, we can hypothesise that if Amy had been free from these reflection inhibiting factors, and able to fully attend to the process of desire formation, she would have been shocked to discover the way in which her desire had been formed, and felt alienated from it. She would not regard this influence as merely a benign and persuasive one; she would surely deem it to have been malign and manipulative. For these reasons, a fully informed Amy would have repudiated the desire formation process.⁹ Therefore, her desire to vote for the populist must be deemed inauthentic, and her consequent decision to place the vote was thereby not autonomous.

Christman's account of autonomy provides a solution to the problem of authenticity, and thus enables us to explain that the profiling team manipulated Amy's vote and rendered it non-autonomous. So far, I have argued in support for the first claim of my thesis; that psychological profiling violates autonomy. In the remainder of this paper, I will turn my attention to the second claim of my thesis; that psychological profiling should be outlawed.

4. Autonomy and paternalism

I claim that it would be justifiable for the state to intervene in Amy's life in order to prevent the profiling team from violating her autonomy. Such an intervention would be "paternalistic," because it aims to improve Amy's situation by interfering with her actions. In this section, I will show that Christman's account of autonomy permits paternalistic interventions, but only when they do not overrule autonomous decisions. Therefore, under Christman's account of autonomy, although a specific intervention with Amy's non-autonomous decision would be acceptable, a blanket ban on psychological profiling would not. Firstly, in order to clarify Christman's attitude towards paternalism, I will explain the "internalist" and "content neutral" features of his account of autonomy. Secondly, a distinction between "hard" and "soft" paternalism will be made. Then, I argue in favour of a "privacy paternalism" policy, that would prevent psychological profiling teams from manipulating voting decisions. Finally, I concede that such a policy would constitute a "hard" paternalistic intervention, and would therefore be rejected by Christman.

⁹ In section 1, I described how Amy repudiated her desire to vote for the populist *two years after* the event. In that example, one could object that Amy might still have authentically desired to elect the populist *when she placed her vote*, and that she only regretted her decision two years later. However, this was merely a rhetorical, pre-theoretical device, employed to suggest that Amy's attitude towards the desire changed when she was free from factors that inhibit self-reflection. Under Christman's account, we ask instead whether Amy would have rejected the formation of her desire *at the time at which it was formed*, if she were able to properly attend to and reflect upon such a process.

4.1. Internalism, content neutrality and soft paternalism

The way in which we define autonomy has important political consequences, since autonomy is often invoked in the justification or rejection of policies and laws. Christman holds that only a certain “soft” type of paternalistic law is justifiable. To understand why, we must first comprehend two features of his account of autonomy; it is “internalist” and “content neutral.”

According to Christman, if a minimally rational agent accepts (or would accept) the process by which she formed her desire, and her acceptance of this process was itself not constrained by any inhibiting factors, then the desire is authentic, and any decisions she makes on the basis of this desire are autonomous. Autonomy requires only that these internal competency conditions are met. That these conditions are met is determined by considering the agent’s psychology, rather than by reference to the external world. This exclusive focus on the psychological conditions of an agent is thus deemed “internalist.” Under Christman’s “proceduralist internalism” then, the *content* of what we value is irrelevant to autonomy. In assessing whether or not Amy’s voting decision was autonomous, we did not evaluate the merits of voting for the populist; we only considered the way in which the desire to vote for the populist was formed. Christman’s position, then, is also described as “content-neutral,” because autonomy ascription does not assess *what* we value, it only assesses *how* we come to value it.

Content neutrality does not mean that value-talk is avoided entirely. It would be difficult to understand why the prospective vegetarian’s decision to stop eating meat is autonomous, for example, without appeal to his higher-order desire to protect the environment. Instead, Christman’s account is content-neutral in the sense that it refrains from appealing to values other than those that the agent in question endorses. In other words, it does not specify any *particular values* that an agent must endorse in order to be described as autonomous.¹⁰ Christman (2008) believes that an account of autonomy must be content neutral in order to ensure congruity with the aforementioned liberal principle of state neutrality, which holds that the state should refrain from promoting any particular conception of the good. This position implies that a law is only justified if it respects citizens’ autonomy, in the sense that the law does not impose upon citizens values that they reject.

The principle of state neutrality is challenged by a position known as *perfectionism*, which holds that the state must *promote certain values* deemed objectively necessary for a good life. Perfectionism implies that laws based on such values are justified, even if the values are authentically rejected by citizens themselves. Perfectionists thereby also oppose content neutrality, because they claim that

¹⁰ This is why Christman’s third criterion places only very minimal standards of rationality upon autonomous agents – he believes that such standards are acceptable to all.

“some states of affairs are objective components of the person’s good *independently of what the person desires*” (Szerletics, 2011, p.13, emphasis added). For Christman, the perfectionist rejection of content neutrality entails a lack of respect for autonomy, because it legitimises unacceptably “paternalistic” laws and policies.

Paternalism can be loosely defined as the interference, restriction or overriding of an agent’s freely made choices, for their own good. Paternalistic intervention is often thought to be in direct tension with autonomy, because it imposes certain values upon agents, thus denying them the ability to decide for themselves how best to pursue their own conception of the good. Yet there are at least some versions of paternalistic intervention which *are* compatible with a content-neutral conception of autonomy (Dworkin, 1972). Firstly, a paternalistic intervention is certainly reasonable when agents provide their explicit informed consent to it. If an agent has sufficient understanding of the nature and consequences of an intervention, and they voluntarily agree to such an intervention taking place, there does not seem to be any violation of autonomy. Imagine our recovering addict gives the state permission to coercively intervene with her life in the event of a relapse. Since the addict has autonomously and explicitly chosen to have her future influenced in this way, the intervention would not conflict with her capacity for self-governance.

In other cases, paternalistic intervention can be justified when consent is not explicit but rather hypothetical. Hypothetical consent occurs when an agent is unable to provide their informed consent at the time of intervention, but it can be hypothesised that they *would have* consented had they been able to. For example, we can hypothesise that if Amy had been sufficiently informed, she would have wanted the state to prevent the profiling team from manipulating her voting decision. In cases in which either explicit or hypothetical consent justifies paternalistic intervention, it is clear that agents are not autonomous at the time of intervention. Such cases are thus labelled “soft” because they do not conflict with autonomy. Soft paternalism is relatively uncontroversial, and since it refrains from the perfectionist impulse to impose repudiated values upon *autonomous* agents, it is compatible with content-neutrality.

While “soft” paternalism only permits interference with non-autonomous decisions, “hard” paternalism holds that sometimes, autonomous decisions can justifiably be overruled.¹¹ Christman (2004, p.157) rejects hard paternalism. He claims that autonomy marks out the “parameters within which a person is immune from paternalistic intervention.” In other words, laws that interfere with *autonomous* decisions are incompatible with a respect for autonomy. For example, in an effort to reduce the population, China’s one child policy mandated coercive sterilisations and forced

¹¹ The distinction between soft and hard paternalism was first made by Feinberg (1989, p.12).

abortions. In this case, a perfectionist conception of the good was imposed upon people who thought otherwise; that living in a society with a manageable population is more valuable than living in a society in which women enjoy reproductive autonomy. The government claimed that overruling their autonomous decisions was justified because a reduction in population was for the benefit of everyone. Under Christman's content neutral account, such a policy would be unacceptable. As Christman's account appears to defend people's autonomy in this way, he claims that it sets the right sort of barriers against paternalism.¹²

Since we deem Amy's vote to be non-autonomous, any paternalistic intervention into her situation would be soft, because it would not be in tension with her autonomy. Consequently, even Christman would find such an intervention acceptable. In the next section, I propose a paternalistic law that would protect agents like Amy from being manipulated by psychological profiling.

4.2. Anti-aggregation legislation

What sort of intervention would have protected Amy's autonomy? It is clear that her desire to vote was inauthentic because it had been manipulated by the psychological profiling team, and that this was possible because the team had access to Amy's private personal data. Thus, legislation that aims to protect Amy's autonomy would be concerned with the protection of Amy's informational privacy.

Informational privacy refers to one's ability to have control over the access and use of one's personal information. (Fried, 1990; Moore, 2003; Rachels, 1975; Westin, 1967). That others have access to certain types of one's personal information may seem fairly harmless when this information is understood as part of a singular, encapsulated database (such as the information about one's purchases from an online grocery store, one's movements on public transport, or one's Facebook "likes"). But Reiman (1995, p.34) notes the potential for harm once this data is aggregated; "as we look at each kind of information-gathering in isolation from the others, each may seem relatively benign," but "when the whole complex is in place, its overall effect on privacy will be greater than the sum of the effects of the parts." Aggregation makes "our lives as a whole visible from a single point, like the panopticon," and this is what gives psychological profiles their manipulative power. Almost any large and varied combination of databases, containing seemingly unimportant personal data, can be used to create a psychological profile which reveals important private information. This information, as we have seen, can be used to manipulate our desires and thereby influence our actions in ways that render us non-autonomous. Ultimately, the more

¹² The majority of liberal theorists reject hard paternalism, for similar reasons (Arneson, 1980; Dworkin, 1972; Feinberg, 1989; Goldman, 2012; Mill, 2003).

personal information that is available about an individual, the easier it is for their actions to be influenced by psychological profiling. Therefore, informational privacy protects autonomy, because it limits the ability of others to manipulate us.

The profiling team had access to a vast amount of Amy's personal data because, when signing up to online services, she gave her acquiescence to the sale of that data to third parties. It is clear, however, that she did not autonomously approve of the manipulation of her voting decision which occurred as a consequence of that initial agreement. Thus, a soft paternalistic law that prevents agents like Amy from non-autonomously waiving their informational privacy would be justified, since it would protect their autonomy. In fact, it appears that such legislation is already in place, at least within the European Union. Under the Data Protection Directive, "collecting and processing the personal data of individuals is only legitimate...where the individual concerned...has unambiguously given his or her consent, after being adequately informed" (European Commission, 2011).

However, although the notion of explicit informed consent is evidently written into law, it has not been translated into practice. The reader will no doubt be familiar with the way in which digital services obtain the "informed consent" of their users. Before accessing a service, the user is presented with a long, complex body of text setting out the company's data usage policy, and they cannot proceed until they click a button to signify that they have "read and agreed to the terms and conditions." In doing so, they make a legal agreement, akin to signing a contract, regarding the use of their personal data. Legally, then, Amy provided her "informed consent" to the use of her data. But both the service provider and the user alike know that the consent obtained is often ambiguous at best. The complexity of the text means that users are clearly disincentivised from reading and understanding it, and are thereby incapable of providing their actual informed consent. Most privacy policies are above the reading level of the average person (Jensen and Potts, 2004). The average length of a set of terms and conditions on a website is 2,514 words (McDonald and Cranor, 2008, p.554), and the average internet user would have to spend 40 minutes each day for a whole year in order to read all of the policies he has agreed to (Ibid., p.563). These findings suggest that the current model for obtaining informed consent is impossible to achieve in practice, and that the burden of ensuring that privacy is protected cannot rest solely on the user.

I propose an alternative model, inspired by a controversial article written by Anita Allen named "Coercing Privacy" (1999). Allen supports a position that could be labelled *privacy paternalism*; she makes the case that "government will have to intervene in private lives for the sake of privacy and values associated with it" (Ibid., p.755). In solidarity with Allen, my suggestion is that people should be forced into protecting their privacy, via a paternalistic law that bans companies from trading

personal data for the purposes of aggregation. Such a law would make psychological profiling impossible, and thereby protect citizens from manipulative forces which are capable of violating their autonomy.

My proposed anti-aggregation legislation is paternalistic in an *indirect* way, because it punishes companies for trading data, rather than users for providing it. It would be wrong to directly prevent citizens from sharing their private information, because this action is sometimes beneficial rather than harmful. As Roessler (2015, p.143) notes, allowing companies access to one's data "is convenient, helps people in their searches and in their purchases, enormously simplifies communication, and can generally help enrich people's lives." In many cases, when we choose to share our data with digital companies, our privacy and autonomy remains intact. For example, Amy might find it useful for the weather application on her smartphone to collect and store data about her previously visited locations. In this case, manipulation is only possible if the developers of the app are able to sell this data to third parties, who could then combine it with other data to create a psychological profile. If the developers were prohibited from selling Amy's data, then the fact that they have access to it would not threaten her autonomy. Therefore, since we are concerned with protecting autonomy, the appropriate legislation is that which prevents companies from trading the personal data of its users for the purposes of aggregation.

I have argued that a type of privacy paternalism, which I call *anti-aggregation legislation*, would prevent people from having their voting decisions manipulated and their autonomy violated by psychological profiling teams. Next, I will explain why my proposed legislation is a version of hard paternalism, and that it is thereby incompatible with Christman's account of autonomy.

4.3. A hard paternalistic intervention

Application of the anti-aggregation legislation to *Amy's case* would indeed constitute a *soft* paternalistic intervention, because it would not conflict with any of her autonomous decisions. However, in this section I explain that, since the anti-aggregation legislation would affect all members of society, it would conflict with some people's autonomy. This entails that the anti-aggregation legislation is unjustified, at least under Christman's account of autonomy, because it would constitute a *hard* paternalistic intervention into some people's lives.

While the anti-aggregation legislation protects the autonomy of agents like Amy, who do not give their informed consent to psychological profiling, it violates the autonomy of other agents, who make fully autonomous decisions about their informational privacy. To see this, imagine another agent, Jamie, who had his voting decision manipulated in the same way as Amy. But Jamie knows

about the practice of psychological profiling. He understands that by trading his personal information for access to online services, he is very likely to have his decisions manipulated in the future, whether that be by a political campaign team, an advertising agency, or the government. Nevertheless, he decides that he is willing to expose himself to such manipulation in order to access online services for free. For Jamie, the cost is worth the benefit, so he makes the autonomous decision to allow psychological profiling teams access to his personal data.

Many internet users make this decision; we knowingly trade our privacy (and thereby our autonomy) for convenience. A recent global survey on people's attitudes towards privacy identified a "we want it all paradox" when it comes to the tension between privacy and convenience (The EMC Corporation, 2014). Only 27% of respondents said they were "willing to trade some privacy for greater convenience and ease online." Yet 68% of respondents use social media accounts, allowing others access to huge amounts of one's personal data. This discrepancy suggests that these people are either (often contrary to testimony) willing to autonomously trade privacy for convenience, or that they do so non-autonomously. As we have established, Amy would be placed in the non-autonomous category. Jamie, on the other hand, is one of the 27% who proudly and autonomously chooses to trade privacy for convenience. Every time he signs up for a service, he reads and fully understands the terms and conditions, and thus provides his explicit informed consent to having his privacy and autonomy violated by way of psychological profiling.

If the anti-aggregation legislation was enacted, then Jamie would no longer be able to trade his data in exchange for access to online services. This is because a consequence of the legislation would most definitely be that many online services would cease to be "free," since revenue from targeted advertising would dry up considerably. Therefore, users will be required to pay if they want to access online services. This would indirectly prevent people like Jamie from autonomously choosing to trade privacy for convenience. Therefore, for soft paternalists such as Christman, the anti-aggregation legislation would be impermissible, because it legitimates a hard paternalistic interference with the lives of people like Jamie. According to Christman, since Jamie autonomously chose to put himself in a situation whereby the profiling team could manipulate his voting decision, that voting decision itself must be deemed autonomous, because Jamie would accept the process by which the desire to vote was formed. And since we care about protecting Jamie's autonomy, any hard paternalistic interference into this state of affairs would be unjustified.

I will argue, however, that hard paternalism *is* compatible with a respect for autonomy, and that the anti-aggregation legislation is therefore justified. In doing so, I reject Christman's account of autonomy, and develop an alternative.

5. The problem with Christman's account

In this section, I argue that Christman's account is problematic because it defends the ability of agents to govern their lives as they please, even when they choose to put themselves in autonomy-undermining situations. I conclude that Christman's internalist and content neutral account of autonomy is unconvincing, and that an alternative must be developed.

Consider the case of a housewife who has completely relinquished her autonomy to her husband.¹³ She must follow his every command, and is beaten when she disobeys. Due to the imbalance of power, the husband can easily manipulate his wife into doing his bidding. For example, he convinces her that in order to prove her commitment, she must put herself at risk by continuously shoplifting bottles of alcohol for him to drink. The housewife thus has a lower-order desire to steal alcohol for her husband. Now, one might suppose that under Christman's model, the housewife would, if she were minimally rational and free from debilitating influences upon her desires, reject the process by which she formed the desire to steal alcohol. Of course, if this were the case, then the desire would be inauthentic and her action thereby non-autonomous. Consequently, a soft paternalistic intervention to save her from this sorry state of affairs would be permitted, since she did not provide her informed consent to being manipulated by her husband in that way. This example thus mirrors Amy's case, in that Amy did not provide her informed consent to being manipulated by the profiling team.

Imagine further, however, that the housewife *autonomously chose* to enter into this manipulative relationship. This example now mirrors Jamie's case, in that Jamie did provide his informed consent to being manipulated by the profiling team. Imagine that the housewife chose to marry her husband under conditions of minimal rationality, and free from debilitating influences upon her desires. Before the marriage, she understood his extreme patriarchal views and was aware of the implications that these views would have for her future autonomy. In fact, she even shares his higher-order religious belief that a wife must be subservient to her husband; she thus authentically desires to be subservient. Since the housewife's decision to enter the relationship was autonomous, those who support content neutrality and oppose perfectionism are forced to defend this oppressive state of affairs. Only the perfectionist can say that the housewife *should not* value subservience, for this would be to promote a particular conception of the good life, independently of what the housewife herself desires. And only the hard paternalist can make the (intuitively credible) moral case that, for her own good, the housewife should be rescued from the physical and emotional abuse that she suffers on a daily basis.

¹³ This example is inspired by those given by Oshana (1988, pp.89-91) and Kristinsson (2000).

Furthermore, according to Christman, the housewife's lower-order desire to steal alcohol is authentic (even though she came to possess it under conditions of manipulation) because she autonomously chose to place herself in a situation in which she would be manipulated. Even though the husband's manipulative influence was responsible for her forming the desire to steal alcohol, Christman's second condition - that autonomy depends upon an absence of debilitating factors - is not violated in this case. This is because, since she autonomously chose to be in a manipulative relationship, the housewife accepts the causal steps which led to the formation of that desire:

"Insofar as a person has authentically embraced even (what we might call) oppressive social status or subservient roles, that person deserves respect insofar as her judgment about those roles has the same formal features as our own judgment about our own lives" (Christman, 2004, p.153).

But this conclusion is counter-intuitive, because it labels a clearly non-autonomous agent as autonomous, and it prohibits a paternalistic intervention that seems very reasonable. Christman accepts these conclusions as an inescapable consequence of the need to ensure that the state cannot interfere with the autonomous decisions of its citizens. He claims that respect for autonomy is incompatible with perfectionism and hard paternalism. This is why, under Christman's account, the anti-aggregation legislation is impermissible. Even though Jamie's vote was manipulated by the psychological profiling team, intervention would be impermissible, because he autonomously chose to put himself in such a situation.

Why does Christman's account of autonomy entail such counter-intuitive conclusions? Recall that his account is internalist and content-neutral. It focuses solely on the process of desire formation, thereby avoiding evaluative perfectionist judgments about the content of agent's decisions. But the case of the subservient housewife suggests that autonomy cannot simply be a matter of the *internal* mental states of the agent. The housewife met the conditions of authenticity and cognitive competence demanded by Christman, and yet her life is clearly not one of a self-governing, autonomous agent. The reason the internalist conception of autonomy fails to adequately explain the housewife's case is because it does not consider the way in which her subservience limits her autonomy. Despite fitting the internal psychological criteria, the housewife's *external* situation entails that she is non-autonomous. The fact that she *wants* to be subservient is irrelevant; the pertinent fact is that she *is* subservient (Oshana, 1998, p.90). A subservient life, by its very definition, is a life absent of self-governance, so an agent committed to subservience cannot be described as autonomous.

It thus seems that, in addition to internal psychology, autonomy also depends upon the absence of certain *external* factors which restrict an agent's ability to have control over their life as a whole.

Christman's conception of autonomy cannot accommodate external factors. Therefore, in the next section, I develop an alternative account of autonomy, under which hard paternalistic interventions in the lives of agents like the housewife and Jamie are compatible with a respect for their autonomy.

6. An externalist-substantivist account of autonomy

My account of autonomy has two features that make it distinct from Christman's account. Rather than being exclusively internalist, my account is externalist, because it holds that autonomy depends on aspects outside of the psychology of the agent. And rather than being content-neutral, my account is "strongly substantive"; it holds that autonomy ascription must involve an evaluation of the content of an agent's desires. In this section, I will explain both of these features in turn, before presenting my own set of conditions for autonomy.

6.1. Externalism

The example of the subservient housewife exposed the weaknesses of Christman's account; it does not accurately ascribe autonomy, and it does not set the right sort of limits on paternalistic intervention. Therefore, in this section, I propose that we should adopt an alternative conception of autonomy, which avoids counter-intuitive autonomy ascription, and permits certain forms of hard paternalism. But how is it possible that a hard paternalistic intervention, which overrules an autonomous decision, engenders the protection of autonomy? In order to answer this question, a *global perspective* of autonomy must be adopted.

Christman holds that autonomy should be ascribed at the *local* level, in the sense that an agent is autonomous if they accept the formation process of a particular desire at a particular time. The *global* perspective of autonomy, by contrast, considers the complete set of an agent's higher-order desires over a longer period of time. Those who adopt this approach, such as Dworkin (1988, p.16), hold that Christman's focus on the diachronic nature of agents does not go far enough. Agents exist for a lifetime; they form long-term plans and develop deeply held values; these things are constitutive of an agent's identity. Thus, autonomy should be perceived on a *global* level, assessing the capacity of agents to realise their goals over the course of their whole lives.

It must be said that Christman does not reject the concept of global autonomy outright. Rather he believes that it "is parasitic on the property of autonomy for isolated preferences and values" (1991, p.3), and that local autonomy is thereby the correct level for ascription. He points to cases in which locally non-autonomous agents (like Amy) "display all of the level-headedness and freedom of thought characteristic of autonomy" in all other aspects of their lives (Ibid.). Christman's suggestion

appears to be that, given Amy's failure to govern her actions whilst under the influence of the profiling team, those who adopt the global perspective are forced to counter-intuitively label her as a non-autonomous agent. The wider implication of this conclusion would be that a global perspective of autonomy narrows our ability to describe ordinary people as autonomous, and thereby conflicts with our intuitions about what a satisfactory descriptive account of autonomy must do.

In response, one further point regarding the scope of autonomy can be made. Those who seek necessary and sufficient conditions for autonomy are often tempted to label agents in a binary manner; agents are either autonomous, or non-autonomous. But this concept is not so black and white. Autonomy, like rationality or morality, is a skill that individuals develop to varying degrees (Anderson, 2008, p.8). Recall our pre-theoretical, intuitive conditions for autonomy. Granted, intentionality does appear to be binary; acts are either intentional or nonintentional. But the conditions of understanding and voluntariness can be satisfied to *varying degrees* (Beauchamp, 2009, p.70). The consequences of an action can be fully or partially understood; an action can be fully or partially voluntary. It is important to bear these considerations in mind when we partake in autonomy ascriptions. Autonomy is not a label that, once ascribed, sticks with agents for life. At different moments, under different circumstances, agents can be described as more or less autonomous. Similarly, the global and local perspectives on autonomy are not mutually exclusive; an agent who decides to kill himself may be described as autonomous in the local sense, but non-autonomous in the global sense. This more complex understanding of the scope of autonomy enables an improved assessment of Amy's autonomy. When under the influence of the psychological profiling team, Amy is locally non-autonomous. However, since she, *ex hypothesi*, displays all the characteristics of autonomy in other aspects of her life, it is appropriate to describe her as a globally autonomous individual.

Despite these nuances, however, it remains the case that a satisfactory account of autonomy must delineate cut-off points that enable us to describe an agent as non-autonomous (Beauchamp, 2010 p.71). In other words, there must be a threshold under which the classification of an action falls from minimally autonomous to non-autonomous. The cut-off point delineated by Christman's account is that a minimally rational agent is autonomous if she accepts (or would accept) the process by which she formed her desire, and her acceptance of this process was itself not constrained by any debilitating factors. But, as we have seen, this boundary is too lenient; it implies that agents who once made a locally autonomous decision to relinquish their autonomy must continue to be described as fully autonomous individuals, despite their occurrent oppression or manipulation.

Using the global perspective of autonomy, we can set a more rigorous boundary for autonomy ascription. When considering whether an agent's decision ought to count as fully autonomous, we must ask whether that decision puts their global autonomy at risk. Even if an agent like Jamie or the subservient housewife autonomously chose (in the local sense) to put themselves into a manipulative situation, since such a situation violates their global autonomy, their decision is not fully autonomous. This way of ascribing autonomy ascription is perfectionist; it holds that certain values are objectively necessary for autonomy, even if the values are authentically rejected by citizens themselves. For example, an autonomous agent must value a life devoid of subservience. This method of autonomy ascription is also thereby "externalist," because it holds that autonomy involves value judgments to which the agent might not cede; values external to the agent. Externalism thus enables us to *avoid* Christman's counter-intuitive conclusions that the housewife is autonomous when she steals alcohol for her husband, and that Jamie is autonomous when he acts upon his manipulated desire to vote for the populist.

This externalist position, however, does not entail that internalism is false. My argument is only that an *exclusive* focus on internal conditions produces an incomplete account of autonomy. In other words, internal and external conditions are each necessary - but only jointly sufficient - for an agent to be deemed autonomous. External conditions *supplement* rather than replace Christman's internal conditions of authenticity and cognitive competence.

6.2. Substantivism

My position does, however, have the implication that content neutrality must be jettisoned, because I claim that the actual content of what an agent values is relevant to autonomy. Such a position is labelled "substantive", because, in contrast to Christman's content neutral account, it adds normative substance to the conditions for autonomy (Ashley, 2012, pp.15-18). I argue in support of a "strong" substantive account, which evaluates the desires of autonomous agents, rather than a "weak" substantive account, which attempts to preserve the idea that an account of autonomy must be content neutral.

"Weak" substantive accounts impose normative constraints on autonomy *indirectly*, by stipulating that certain internal mental *capacities* are necessary for autonomous judgement (Ashley, 2012, p.17). For example, Myers (cited in Benson, 2008, p.135) proposes that an autonomous agent must be capable of "introspection, communication, memory, imagination, analytical reasoning, 'self-nurturing', resistance to pressures to conform, and political collaboration." The purported benefit of weak substantive accounts is that these normative standards apply to the cognitive *competencies* of

the agent, rather than to the *content* of her values. Thus, autonomous agents are permitted to make choices which may appear to be wrong or bad for them, so long as their desire formation process meets those competency conditions. Supporters of weak substantivism thereby claim that this position achieves two goals (Ibid.). First, it is capable of supporting paternalistic intervention in cases like that of the subservient housewife. Second, it is capable of doing so without making a perfectionist claim about what the housewife should value. For example, under Myers' account, since the subservient housewife failed to "resist pressures to conform", she is in fact non-autonomous, and a soft paternalistic intervention is thereby compatible with a respect for her autonomy. Furthermore, since this intervention is justified without reference to the content of the housewife's desires, it is said to avoid a perfectionist judgment.

Note, however, that the conditions proposed by Myers, instead of being "weak", are in fact rather stringent. Although stringent conditions for autonomy enable us to (correctly) label non-autonomous agents like the subservient housewife as non-autonomous, they also lead us to (mistakenly) label some *autonomous* agents as non-autonomous. Consider Myers' requirement that autonomous agents employ "analytical reasoning" when making decisions. This could entail that agents who partake in risky activities (such as gambling or skydiving) are irrational and thereby non-autonomous. If Myers' account of autonomy was used to inform policy decisions, then the state would be justified in outlawing such risky activities in order to protect autonomy. But this conclusion is counter-intuitive, and it also seems to conflict with the liberal principle of state neutrality.

In an effort to avoid the problems caused by such overly demanding conditions for autonomy, Christman ensured that his rationality condition was extremely minimal (see section 3.2). Yet even Christman's rationality condition imposes an indirect normative constraint upon autonomy, because any ascription of rationality, no matter how minimal, necessarily involves an evaluative judgment. Therefore, despite pretensions to the contrary, Christman's account should also be described as weakly substantive. And, as we have seen, Christman's (very weak) version of weak substantivism runs into problems of its own. Meyers' conditions are too demanding; they label autonomous agents as non-autonomous. Conversely, Christman's conditions are too weak; they label non-autonomous agents as autonomous. We should therefore be sceptical of the ability of accounts that impose normative constraints on autonomy *indirectly* to accurately distinguish between autonomous and non-autonomous agents.

"Strong" substantive accounts, on the other hand, impose *direct* limits on the *content* of an autonomous agent's desires. It may appear that such a perfectionist position is incompatible with respect for autonomy, since it permits hard paternalistic laws which prohibit locally autonomous

agents from acting on the basis of certain values. However, Robert Young's (1981) version of strong substantivism avoids this criticism, because it appeals directly to autonomy itself. He argues that, since autonomy is intrinsically valuable, we are committed to "defending, preserving and enhancing it" (Ashley, 2012, p.16). It is therefore unacceptable to allow agents to relinquish their autonomy, as Jamie does when he allows his personal data to be sold to psychological profiling organisations. The global perspective makes this externalist criterion intelligible; although an agent might be exercising their *local* autonomy by voluntarily entering into manipulative situations, such an act would directly violate their future capacity for autonomous choice, and thus is in tension with their *global* autonomy. Fundamentally, we value autonomy because it enables us to have a holistic control over our own lives, rather than because it enables us to perform singular actions. This means that global autonomy must take priority over local autonomy. Young's account thus motivates what I term *the perfectionist principle* - that to be described as globally autonomous, agents must not make choices which condemn them to a manipulative situation in the future.

The perfectionist principle entails that a hard paternalistic intervention with an agent's locally autonomous choice is justified if it protects their global autonomy. For example, Jamie's desire to enter into an autonomy-undermining relationship was authentic, and his decision to do so was *locally* autonomous. But Jamie's decision enabled a psychological profiling team to manipulate his vote. Therefore, Jamie's decision to trade his privacy for convenience subverted his future capacity for autonomy. Since the state is committed to defending, preserving and enhancing autonomy, it must intervene with Jamie's situation, against his will, for his own good.

6.3. Four conditions for autonomy

At this stage, it would be advantageous to conduct a brief review of our journey towards an externalist-substantivist conception of autonomy. The intuitive account of autonomy struggled to determine whether the psychological profiling team had a manipulative or persuasive influence upon Amy's desires, and thereby failed to accurately determine the authenticity and autonomy of her voting decision. By employing Christman's authenticity conditions, we were able to ascertain that the profiling team manipulated Amy's voting decision and thereby violated her autonomy. However, due to his exclusive focus on the internal procedure of desire-formation, Christman's account fails deal with cases in which a locally autonomous decision condemns an agent to manipulation in the future.

I propose two modifications to Christman's internalist-proceduralist account. The first is to observe that an agent can be autonomous merely in a local sense, but also in a global sense, and that global

autonomy is paramount. The second is to supplement Christman's internalist conditions with an externalist condition; *the perfectionist principle* introduced in the previous section. Both the internalist conditions and the externalist condition are each necessary - but only jointly sufficient - for an agent to be deemed fully autonomous.

So, my alternative account of autonomy states that an agent is *merely locally* autonomous relative to some desire at time t if and only if:

- i) The agent did not resist the development of the desire (prior to t) when attending to this process of development, (or would not have resisted that development, had she attended to the process).
- ii) The lack of resistance to the development of the desire (prior to t) did not take place (or would not have) under the influence of factors that inhibit self-reflection.
- iii) The self-reflection involved is (minimally) rational and involves no self-deception.

An agent is *also globally and thereby fully* autonomous relative to the desire if and only if:

- iv) Decisions motivated by the desire do not condemn the agent to a manipulative situation in the future.

This account is *hierarchical*, in that it sees autonomy as being guided by an agent's higher-order attitudes towards their lower-order desires. It is *historical*, in that it holds the origin of an agent's desires to be crucial for autonomy. It is *proceduralist*, in that autonomy ascription requires an agent to attend (hypothetically, at least) to the process by which their desires are formed. It is *internalist* - in that it holds autonomy to depend upon the psychological state of the agent - but not exclusively so. This is because the account is also *externalist*, as it holds that agents who fail to value their own global autonomy are non-autonomous. It is thus *strongly substantive*, because it makes the content of an agent's desires relevant to autonomy ascription. It is also *perfectionist*, in that it promotes a particular conception of the good. This entails that the account justifies *hard paternalism*, but only in those cases where an agent's global autonomy is under threat of manipulation.

Although my account does reject *content* neutrality, it is compatible with the liberal principle of *state* neutrality, because I assert the primacy of global autonomy over local autonomy. The fundamental purpose of the principle of state neutrality is to enable citizens to direct their lives as they see fit. Global autonomy is essential to such a project. One's capacity to pursue one's own conception of the good life depends more on the ability to form long-term life plans, than on the

ability to make singular decisions, especially when these decisions lead to autonomy-undermining situations. In a sense, the externalist condition does endorse a *particular conception* of what is valuable; namely, global autonomy. However, since the value of global autonomy is supposed to be the foundation of all liberal societies, the externalist condition is one that *every* citizen should be able to accept.

I have argued that this account of autonomy is theoretically plausible. It is also intuitively plausible, because it accords with our pre-theoretical intuitions surrounding the concept. It holds that autonomy is intrinsically valuable; that autonomous actions involve intentionality, understanding and voluntariness; and that most agents, most of the time, are (at least locally) autonomous. Furthermore, since my account retains Christman's conditions for authenticity, it evades the problems encountered by the intuitive account. Specifically, it enables us to distinguish between persuasive and manipulative influences.

I have argued that my alternative account of autonomy should be adopted. In the next section, I show how this account of autonomy is compatible with the proposed anti-aggregation legislation.

7. In support of hard privacy paternalism

In this final section, I offer a defence of the anti-aggregation legislation. First, I reassert why it is compatible with a respect for autonomy. Then, I defend it against a more general objection to hard paternalism. Finally, I argue that those of both libertarian and liberal persuasion have reasons to adopt the legislation.

Under my externalist-substantivist account of autonomy, the proposed anti-aggregation legislation is compatible with a respect for autonomy, even though it constitutes a hard paternalistic intervention. For example, if Jamie decides to give away his personal data, this would violate my fourth condition for autonomy, because he knowingly condemns himself to future manipulation. Since Jamie meets the internal conditions for autonomy, and his desire to waive his privacy is thereby authentic, his decision to give away his privacy could be described as locally autonomous. But this decision also subverts his global autonomy, because it condemns Jamie to a state of affairs in which he is unable to control his own voting decision. Given that we value autonomy because it enables us to have a holistic control over our own lives, rather than because it enables us to perform singular actions, global autonomy takes precedence over local autonomy. It is justifiable to interfere with an agent's locally autonomous decisions in order to protect their global autonomy. Therefore the anti-aggregation legislation respects the autonomy of every citizen, even if it sometimes involves hard paternalistic intervention.

It can be objected, however, that hard paternalism is untenable, because it leads us down a slippery slope, upon which more controversial interferences with people's lives could be defended. Recall the objection I raised to Myers' stringent condition of rationality in section 6.2. The worry was that strong conditions for autonomy permit policy makers to justify paternalism in a whole range of cases, from gambling to skydiving. Similarly, those who object to hard paternalism worry that, by appealing to an agent's global autonomy, almost any policy could be justified as being 'for their own good.' In fact, I suggested that China's one child policy was given such a justification. These concerns reflect one of our pre-theoretical intuitions about autonomy; that conditions for autonomy must not be so stringent that they exclude everyday cases of self-governance. There must, therefore be a threshold of acceptability underneath which paternalistic interferences must not fall. Soft paternalists, such as Christman, propose local autonomy as the appropriate threshold: if an action is locally autonomous then paternalistic interference is impermissible. Christman's reasoning seems to be that although this threshold might lead to the occasional violation of global autonomy, it is best to err on the side of caution, in order to keep the state from interfering with people's lives in unacceptable ways.

In response to this objection, it must be pointed out that my version of hard paternalism enacts a similar threshold. My threshold, imposed by the fourth condition, is that interference with a locally autonomous action is *only* permissible if the threat of *manipulation* is present. This rules out interference with many locally autonomous actions; those which do threaten an agent's global autonomy, but *do not involve manipulation*. My threshold thereby keeps us away from the slippery slope upon which many unacceptable paternalistic laws can be justified. For example, since the locally autonomous decision to gamble or to skydive does not condemn one to future manipulation, hard paternalism is impermissible in these cases. But the locally autonomous decision to give away one's personal data does condemn agents to future manipulation, via psychological profiling. Therefore, hard paternalism is permissible in these cases.

If one finds this defence of hard paternalism unconvincing, however, this does not entail that the anti-aggregation legislation is unjustifiable. In fact, the legislation can even be justified to "libertarians", who are surely the most ardent anti-paternalists imaginable. Since libertarians believe in extremely minimal principles of justice, any legislation that they would approve has a good chance of being approved by others. The basic foundation of libertarian thought is John Stuart Mill's *harm principle* (Hamowy, 2008, p.xxi). In opposition to paternalism, Mill (2003, p.68) argued that "the only purpose for which power can be rightfully exercised over any member of a civilized community, against his will, is to prevent harm to others." In other words, intervention with an agent's action is only permitted if that action will harm *other agents*. Therefore, according to libertarians, even the

soft variety of paternalism is unwarranted, because paternalistic interference, by definition, functions to protect individuals from harming *themselves*. With regards to both Amy and Jamie then, the libertarian would argue that, since their actions violate only their *own* global autonomy, and nobody else's, intervention is unjustified.

In response, an alternative justification for the anti-aggregation legislation can be provided, which shows how Amy and Jamie's decisions about their own privacy cause harm to *others*. Firstly, it is clear that, if we allow powerful agents to manipulate the results of elections, this has the potential to harm everyone. A politician elected by such means would be unlikely to find it immoral to violate the privacy and autonomy of his citizens in other ways. One major worry in such a scenario would be that the surveillance powers of the state could be employed in order to repress his opponents. Cumulative individual decisions to trade privacy for convenience are leading us towards a society in which everyone's privacy, and thereby their autonomy, are drastically reduced.

So, when people like Jamie exhibit indifference to their own informational privacy, the libertarian might presume that the only person Jamie is harming is himself. But Jamie's decision to give up his privacy does not just have consequences for him; it also makes other people, who *do* care about privacy, less able to defend it. As Roessler (2013, p.18) points out, "some people's indifference to their privacy makes it harder for others who care about their privacy to convince firms and officials to institute data protection measures." A case in point is the UK's Investigatory Powers Act 2016, which has been described as "one of the most extreme surveillance laws ever passed in a democracy" (Travis, 2016). The act was passed in an environment of public indifference and "devastatingly poor political opposition" (Carlo, 2016). Criticisms of the act were not publicised because, according to MP David Davis, "the public doesn't care enough about encroachments on their freedom" (Moss, 2015). This example shows that when individuals adopt attitudes of indifference towards their informational privacy, they limit the capacity of others to protect and enjoy their own privacy. And, as we have seen, if an agent lacks informational privacy, their global autonomy is threatened by psychological profiling. Since Jamie's decision can be said to *harm others* in this way, even the libertarian, who rejects all paternalistic intervention, is forced to accept the legitimacy of the anti-aggregation legislation.

Liberals, as well as libertarians, should accept the anti-aggregation legislation, because it is compatible with the liberal principle of state neutrality. Although we want the government to be neutral with regards to competing conceptions of the good, there are certain fundamental values that the liberal state can defend without hypocrisy. Since the state is committed to defending global

autonomy, and informational privacy protects global autonomy, the state must enact privacy paternalism. As Allen (1999, p.752) puts it:

“A conception of the good that permits privacy to be waived...is like a vision of the good that permits freedom to be waived. As liberals, we should not want people to sell all their freedom, and, as liberals, we should not want people to sell all their privacy and capacities for private choices.”

Conclusion

I have argued that the use by political electoral campaigns of psychological profiling in order to influence voting decisions threatens the autonomy of voters, and should therefore be outlawed. To support my claim, I have developed a unique theoretical account of autonomy. As Calabrese (in an interview with Beckett, 2012) warns, if informational privacy is not defended, then “any voter is open to manipulation.” And since the manipulation of one’s voting decision constitutes a severe violation of one’s autonomy, a hard version of privacy paternalism is justified. Although privacy paternalism might limit options in the short term, it protects global autonomy in the long term.

Bibliography

- Allen, A. (1988) *Uneasy access: privacy for women in a free society*. New York, NY, United States: Rowman & Littlefield.
- Allen, A. (1999) 'Coercing privacy', *William & Mary Law Review*, 40(3), pp. 723–757.
- Allen, A. (2013) 'Our privacy rights and responsibilities: replies to critics', *Philosophy and Law*, 13(1), pp. 19–27.
- Anderson, J. (2008) 'Disputing autonomy: second-order desires and the dynamics of ascribing autonomy', *Sats - Nordic Journal of Philosophy*, 9, pp. 7–26.
- Anderson, J. (2010) 'Review of Richard Thaler and Cass Sunstein: nudge: improving decisions about health, wealth, and happiness', *Economics and Philosophy*, 26, pp. 369–376.
- Anderson, S. (2011) 'Coercion', in *The Stanford Encyclopedia of Philosophy*. Available at: <http://plato.stanford.edu/archives/sum2015/entries/coercion/> (Accessed: 11 November 2016).
- Arneson, R.J. (1980) 'Mill versus paternalism', *Ethics*, 90(4), pp. 470–489. doi: 10.2307/2380448.
- Ashley, V. (2012) 'Philosophical models of autonomy', in *Essex Autonomy Project Green Paper Report*. Available at: <http://autonomy.essex.ac.uk/philosophical-models-of-autonomy> (Accessed: 11 November 2016).
- Beauchamp, T. (2010) 'Autonomy and consent', in Miller, F. and Wertheimer, A. (eds.) *The ethics of consent*. USA: Oxford University Press, pp. 55–79.
- Beckett, L. (2012) *How Microsoft and Yahoo are selling politicians access to you*. Available at: <https://www.propublica.org/article/how-microsoft-and-yahoo-are-selling-politicians-access-to-you> (Accessed: 11 November 2016).
- Benn, S.I. (1984) 'Privacy, freedom, and respect for persons', in Schoeman, F.D. (ed.) *Philosophical dimensions of privacy: an anthology*. Cambridge: Cambridge University Press, pp. 223–244.
- Benson, P. (2008) 'Feminist intuitions and the normative substance of autonomy', in Taylor, J. (ed.) *Personal autonomy: new essays on personal autonomy and its role in contemporary moral philosophy*. Cambridge: Cambridge University Press, pp. 124–142.
- Birnhack, M. (2011) 'Review: a quest for a theory of privacy: context and control', *Jurimetrics*, 51(4), pp. 447–479. doi: 10.2307/41307137.
- Bok, S. (1989) *Secrets: on the ethics of concealment and revelation*. New York: Knopf Doubleday Publishing Group.
- Brennan, A. (2012) *Microtargeting: how campaigns know you better than you know yourself*. Available at: <http://edition.cnn.com/2012/11/05/politics/voters-microtargeting> (Accessed: 11 November 2016).
- Buss, S. (2016) 'Personal autonomy', in *The Stanford Encyclopedia of Philosophy*. Available at: <http://plato.stanford.edu/archives/win2016/entries/personal-autonomy/> (Accessed: 11 November 2016).

- Carlo, S. (2016) *The snooper's charter passed into law this week – say goodbye to your privacy*. Available at: <http://www.independent.co.uk/voices/snoopers-charter-theresa-may-online-privacy-investigatory-powers-act-a7426461.html> (Accessed: 30 November 2016).
- Carter, I. (2016) 'Positive and negative liberty', in *The Stanford Encyclopedia of Philosophy*. Available at: <http://plato.stanford.edu/archives/fall2016/entries/liberty-positive-negative/> (Accessed: 11 November 2016).
- Christman, J. (1991) 'Autonomy and personal history', *Canadian Journal of Philosophy*, 21(1), pp. 1–24. doi: 10.2307/40231731.
- Christman, J. (1999) 'Autonomy', in Gray, B. (ed.) *The philosophy of law: an encyclopedia*. New York: Garland Publishing, pp. 72–73.
- Christman, J. (2004) 'Relational autonomy, liberal individualism, and the social constitution of selves', *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 117, pp. 143–164. doi: 10.2307/4321441.
- Christman, J. (2008) 'Procedural autonomy and liberal legitimacy', in Taylor, J. (ed.) *Personal autonomy: new essays on personal autonomy and its role in contemporary moral philosophy*. Cambridge: Cambridge University Press, pp. 277–298.
- Christman, J. (2009) 'Autonomy in moral and political philosophy', in *The Stanford Encyclopedia of Philosophy*. Available at: <http://plato.stanford.edu/archives/spr2011/entries/autonomy-moral/> (Accessed: 11 November 2016).
- Christman, J. (2011) *The politics of persons: individual autonomy and socio-historical selves*. Cambridge: Cambridge University Press.
- Christman, J. (2014) 'Introduction', in *The inner citadel: essays on individual autonomy*. United States: Echo Point Books & Media, pp. 3–26.
- Crisp, R. (1987) 'Persuasive advertising, autonomy, and the creation of desire', *Journal of Business Ethics*, 6(5), pp. 413–418. doi: 10.2307/25071678.
- Davidson, D. (1963) 'Actions, reasons, and causes', *The Journal of Philosophy*, 60, pp. 685–700. doi: 10.2307/2023177.
- Davies, H. (2016) *Ted Cruz using firm that harvested data on millions of unwitting Facebook users*. Available at: <https://www.theguardian.com/us-news/2015/dec/11/senator-ted-cruz-president-campaign-facebook-user-data> (Accessed: 11 November 2016).
- DeCew, J. (2015) 'Privacy', in Zalta, E. (ed.) *The Stanford Encyclopedia of Philosophy*. Available at: <http://plato.stanford.edu/archives/spr2015/entries/privacy/> (Accessed: 11 November 2016).
- Dryden, J. (2016) 'Autonomy', in *The Internet Encyclopedia of Philosophy*. Available at: <http://www.iep.utm.edu/autonomy/> (Accessed: 11 November 2016).
- Dunn, L. and Hoegg, J. (2014) 'The impact of fear on emotional brand attachment', *Journal of Consumer Research*, 41(1), pp. 152–168. doi: 10.1086/675377.
- Dworkin, G. (1972) 'Paternalism', *The Monist*, 56, pp. 64–84. doi: 10.2307/27902250.

Dworkin, G. (1976) 'Autonomy and behavior control', *The Hastings Center Report*, 6(1), pp. 23–28. doi: 10.2307/3560358.

Dworkin, G. (1988) *The theory and practice of autonomy*. Cambridge, MA: Cambridge University Press.

The EMC Corporation (2014) *The EMC privacy index*. Available at: <http://www.emc.com/collateral/brochure/privacy-index-global-in-depth-results.pdf> (Accessed: 29 November 2016).

European Commission (2011) *Collecting & processing personal data: what is legal?* Available at: http://ec.europa.eu/justice/data-protection/data-collection/legal/index_en.htm (Accessed: 29 November 2016).

Eyal, N. (2011) 'Informed consent', in *The Stanford Encyclopedia of Philosophy*. Available at: <http://plato.stanford.edu/archives/fall2012/entries/informed-consent/> (Accessed: 11 November 2016).

Feinberg, J. (1989) *The moral limits of the criminal law volume 3: harm to self*. New York: Oxford University Press.

Frankfurt, H. (1987) 'Identification and wholeheartedness', in Schoeman, F. (ed.) *Responsibility, character, and the emotions: new essays in moral psychology*. New York: Cambridge University Press, pp. 27–45.

Frankfurt, H. (1988) *The importance of what we care about*. Cambridge: Cambridge University Press (CUP).

Fried, C. (1984) 'Privacy [a moral analysis]', in Schoeman, F.D. (ed.) *Philosophical dimensions of privacy: an anthology*. Cambridge: Cambridge University Press, pp. 203–222.

Fried, C. (1990) 'Privacy', in Ermann, M.D., Williams, M.B., and Gutierrez, C. (eds.) *Computers, ethics, and society*. New York: Oxford University Press, pp. 51–67.

Gavison, R. (1980) 'Privacy and the limits of law', *The Yale Law Journal*, 89(3), pp. 421–471. doi: 10.2307/795891.

Gertner, J. (2004) *The very, very personal is the political*. Available at: <http://www.nytimes.com/2004/02/15/magazine/the-very-very-personal-is-the-political.html> (Accessed: 11 November 2016).

Gladwell, M. (2007) *Blink: the power of thinking without thinking*. New York: Little, Brown & Company.

Goldman, A. (2012) 'The refutation of medical paternalism', in Arras, J., London, A., and Steinbock, B. (eds.) *Ethical issues in modern medicine: contemporary readings in bioethics*. New York: McGraw-Hill, pp. 60–67.

Gunter, B. (2016) *The psychology of consumer profiling in a digital age*. Oxford: Routledge.

Hamburger, T. (2015) *Cruz campaign credits psychological data and analytics for its rising success*. Available at: <https://www.washingtonpost.com/politics/cruz-campaign-credits-psychological-data->

and-analytics-for-its-rising-success/2015/12/13/4cb0baf8-9dc5-11e5-bce4-708fe33e3288_story.html (Accessed: 11 November 2016).

Hamowy, R. (2008) *The encyclopedia of libertarianism*. Thousand Oaks, CA: Sage Publications.

Herman, E.S. and Chomsky, N. (1988) *Manufacturing consent: the political economy of the mass media*. New York: Pantheon Books.

Hildebrandt, M. and Gutwirth, S. (eds.) (2010) *Profiling the european citizen: cross-disciplinary perspectives*. Netherlands: Springer-Verlag New York.

Issenberg, S. (2016) *How Obama's team used big data to rally voters*. Available at: <https://www.technologyreview.com/s/509026/how-obamas-team-used-big-data-to-rally-voters/> (Accessed: 11 November 2016).

Jensen, C. and Potts, C. (2004) 'Privacy policies as decision-making tools: an evaluation of online privacy notices', *CHI 2004*, 6(1), pp. 471–478.

Kahneman, D. (2011) *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.

Kranish, M. (2016) *Trump's plan for a comeback includes building a 'psychographic' profile of every voter*. Available at: https://www.washingtonpost.com/politics/trumps-plan-for-a-comeback-includes-building-a-psychographic-profile-of-every-voter/2016/10/27/9064a706-9611-11e6-9b7c-57290af48a49_story.html (Accessed: 11 November 2016).

Kristinsson, S. (2000) 'The limits of neutrality: toward a weakly substantive account of autonomy', *Canadian Journal of Philosophy*, 30(2), pp. 257–286. doi: 10.2307/40232091.

Kupfer, J. (1987) 'Privacy, autonomy, and self-concept', *American Philosophical Quarterly*, 24(1), pp. 81–89. doi: 10.2307/20014176.

Lazaro, C. and Métayer, D.L. (2015) 'Control over personal data: true remedy or fairytale?', *SCRIPTed*, 12(1). doi: 10.2966/scrip.120115.3.

Leetaru, K. (2016) *Data breaches, psychological profiling, voter modeling: inside the big data world of campaign 2016*. Available at: <http://www.forbes.com/sites/kalevleetaru/2016/01/01/data-breaches-psychological-profiling-voter-modeling-inside-the-big-data-world-of-campaign-2016/#7ec40652737d> (Accessed: 11 November 2016).

Lippke, R. (1989) 'Advertising and the social conditions of autonomy', *Business & Professional Ethics Journal*, 8(4), pp. 35–58. doi: 10.2307/27800021.

Lynch, M.P. (2016) *The Internet of us: knowing more and understanding less in the age of big data*. United States: Liveright Publishing.

McDonald, A. and Cranor, L.F. (2008) 'The cost of reading privacy policies', *Privacy Year in Review*, , pp. 543–568.

Mill, J.S. (2003) *On liberty*. Edited by D Bromwich and G Kateb. New Haven: Yale University Press.

Moore, A. (2000) 'Employee monitoring and computer technology: evaluative surveillance v. privacy', *Business Ethics Quarterly*, 10(3), pp. 697–709. doi: 10.2139/ssrn.1973423.

Moore, A. (2003) 'Privacy: Its meaning and value', *American Philosophical Quarterly*, 40(3), pp. 215–227. doi: 10.2307/20010117.

Moore, A. (2013) 'Coercing privacy and moderate paternalism: Allen on unpopular privacy', *Philosophy and Law*, 13(1), pp. 10–14.

Moss, S. (2015) *Tory rebel David Davis: 'we haven't had a Stasi or a Gestapo in Britain, so are intellectually lazy about surveillance'*. Available at: <https://www.theguardian.com/politics/2015/nov/08/david-davis-stasi-gestapo-surveillance-tory-civil-liberties> (Accessed: 30 November 2016).

Oshana, M. (1998) 'Personal autonomy and society', *Journal of Social Philosophy*, 29(1), pp. 81–102. doi: 10.1111/j.1467-9833.1998.tb00098.x.

O'Shea, T. (2011) 'Consent in history, theory and practice', in *Essex Autonomy Project Green Paper Report*. Available at: <http://autonomy.essex.ac.uk/consent-in-history-theory-and-practice> (Accessed: 18 November 2016).

O'Shea, T. (2012) 'Autonomy and value', in *Essex Autonomy Project Green Paper Report*. Available at: <http://autonomy.essex.ac.uk/autonomy-and-value> (Accessed: 18 November 2016).

Parent, W.A. (1983) 'Privacy, morality, and the law', *Philosophy & Public Affairs*, 12(4), pp. 269–288. doi: 10.2307/2265374.

Pasick, A. and Fernholz, T. (2015) *The stealthy, Eric Schmidt startup that's working to put Hillary Clinton in the white house*. Available at: <http://qz.com/520652/groundwork-eric-schmidt-startup-working-for-hillary-clinton-campaign/> (Accessed: 11 November 2016).

Payne, A. (2016) A British firm which helped deliver Brexit is working for Donald Trump's campaign. Available at: <https://www.businessinsider.nl/donald-trump-brexit-us-presidential-election-2016-9/?international=true&r=UK> (Accessed: 10 December 2016).

Rachels, J. (1975) 'Why privacy is important', *Philosophy & Public Affairs*, 4(4), pp. 323–333. doi: 10.2307/2265077.

Regan, P. (2015) 'Privacy and the common good: revisited', in Roessler, B. and Mokrosinska, D. (eds.) *Social dimensions of privacy: Interdisciplinary perspectives*. Cambridge, United Kingdom: Cambridge University Press, pp. 50–70.

Reiman, J.H. (1995) 'Driving to the Panopticon: a philosophical exploration of the risks to privacy posed by the highway technology of the future', *Santa Clara High Technology Law Journal*, 11(1), pp. 27–44.

Roessler, B. (2013) 'Autonomy, paternalism and privacy: some remarks on Anita Allen', *Philosophy and Law*, 13(1), pp. 14–18.

Roessler, B. (2015) 'Should personal data be a tradable good? on the moral limits of markets in privacy', in Roessler, B. and Mokrosinska, D. (eds.) *Social dimensions of privacy: interdisciplinary perspectives*. Cambridge: Cambridge University Press, pp. 141–161.

Roessler, B. and Mokrosinska, D. (eds.) (2015) *Social dimensions of privacy: interdisciplinary perspectives*. Cambridge: Cambridge University Press.

- Rule, J.B. (2015) 'Privacy: the longue durée', in Roessler, B. and Mokrosinska, D. (eds.) *Social dimensions of privacy: interdisciplinary perspectives*. Cambridge: Cambridge University Press, pp. 11–31.
- Scherer, M. (2012) *Inside the secret world of the data crunchers who helped Obama win*. Available at: <http://swampland.time.com/2012/11/07/inside-the-secret-world-of-quants-and-data-crunchers-who-helped-obama-win/> (Accessed: 11 November 2016).
- Schwartz, P. (2004) 'Property, privacy, and personal data', *Harvard Law Review*, 117(7), pp. 2055–2128.
- Scoccia, D. (2007) 'In defense of hard paternalism', *Law and Philosophy*, 27(4), pp. 351–381. doi: 10.1007/s10982-007-9020-8.
- Shakespeare, W. (1999) *Othello*. New York, NY, United States: Classic Publishers.
- Simon, G.K. (2010) *In sheep's clothing: understanding and dealing with manipulative people*. 2nd edn. United Kingdom: Parkhurst Brothers Publishers.
- Simonite, T. (2014) *Sell your personal data for \$8 a month*. Available at: <https://www.technologyreview.com/s/524621/sell-your-personal-data-for-8-a-month/> (Accessed: 28 November 2016).
- Sneddon, A. (2001) 'Advertising and deep autonomy', *Journal of Business Ethics*, 33(1), pp. 15–28. doi: 10.1023/a:1011929725518.
- Solove, D. (2006) *The digital person: technology and privacy in the information age*. New York: New York University Press.
- Solove, D. (2009) *Understanding privacy*. Cambridge, MA: Harvard University Press.
- Stoljar, N. (2013) 'Feminist perspectives on autonomy', in *The Stanford Encyclopedia of Philosophy*. Available at: <http://plato.stanford.edu/archives/fall2015/entries/feminism-autonomy/> (Accessed: 11 November 2016).
- Stroud, S. (2014) 'Weakness of will', in *The Stanford Encyclopedia of Philosophy*. Available at: <http://plato.stanford.edu/archives/spr2014/entries/weakness-will/> (Accessed: 11 November 2016).
- Szerletics, A. (2011) 'Paternalism', in *Essex Autonomy Project Green Paper Report*. Available at: <http://autonomy.essex.ac.uk/paternalism-2> (Accessed: 18 November 2016).
- Tavani, H.T. (2008) 'Informational privacy: concepts, theories, and controversies', in Himma, K.E.E. and Tavani, H.T. (eds.) *The handbook of information and computer ethics (Kingfisher leisure guide)*. United Kingdom: Wiley-Blackwell, pp. 131–164.
- Taylor, J. (2008) 'Introduction', in Taylor, J. (ed.) *Personal autonomy: new essays on personal autonomy and its role in contemporary moral philosophy*. Cambridge: Cambridge University Press, pp. 1–32.
- Travis, A. (2016) 'Snooper's charter' bill becomes law, extending UK state surveillance. Available at: <https://www.theguardian.com/world/2016/nov/29/snoopers-charter-bill-becomes-law-extending-uk-state-surveillance> (Accessed: 30 November 2016).

Van den Hoven, J. (2008) 'Information technology, privacy, and the protection of personal data', in Van den Hoven, J. and Weckert, J. (eds.) *Information technology and moral philosophy*. Cambridge: Cambridge University Press, pp. 301–321.

Van den Hoven, J., Blaauw, M., Pieters, W. and Warnier, M. (2014) 'Privacy and information technology', in *The Stanford Encyclopedia of Philosophy*. Available at: <http://plato.stanford.edu/archives/spr2016/entries/it-privacy/> (Accessed: 11 November 2016).

Vavreck, L. and Sides, J. (2014) *Obama's not-so-big data*. Available at: <https://psmag.com/obama-s-not-so-big-data-a2546a5c8b28#.onosptq6b> (Accessed: 11 November 2016).

Velleman, D. (1992) 'What happens when someone acts?', *Mind*, 101(403), pp. 461–481. doi: 10.2307/2253898.

Wall, S. (2003) 'Freedom as a political ideal', *Social Philosophy and Policy*, 20(2), pp. 307–334. doi: 10.1017/S0265052503202132.

Westin, A. (1967) *Privacy and freedom*. New York: Atheneum.

Wikipedia (2016a) 'Information privacy', in *Wikipedia*. Available at: https://en.wikipedia.org/wiki/Information_privacy (Accessed: 11 November 2016).

Wikipedia (2016b) 'Internet privacy', in *Wikipedia*. Available at: https://en.wikipedia.org/wiki/Internet_privacy (Accessed: 11 November 2016).

Wikipedia (2016c) 'Marketing and artificial intelligence', in *Wikipedia*. Available at: https://en.wikipedia.org/wiki/Marketing_and_artificial_intelligence (Accessed: 11 November 2016).

Wikipedia (2016d) 'Marketing ethics', in *Wikipedia*. Available at: https://en.wikipedia.org/wiki/Marketing_ethics (Accessed: 11 November 2016).

Wikipedia (2016e) 'Privacy', in *Wikipedia*. Available at: <https://en.wikipedia.org/wiki/Privacy> (Accessed: 11 November 2016).

Wikipedia (2016f) 'Propaganda', in *Wikipedia*. Available at: <https://en.wikipedia.org/wiki/Propaganda> (Accessed: 11 November 2016).

Wikipedia (2016g) 'Psychological manipulation', in *Wikipedia*. Available at: https://en.wikipedia.org/wiki/Psychological_manipulation (Accessed: 11 November 2016).

Young, R. (1981) 'Autonomy and paternalism', *Bulletin of the Australian Society of Legal Philosophy*, 21, pp. 32–54.

Zimmer, M. (2016) *OkCupid study reveals the perils of big-data science*. Available at: <https://www.wired.com/2016/05/okcupid-study-reveals-perils-big-data-science/> (Accessed: 11 November 2016).