A RELATIONAL APPROACH TO THE MORAL VALUE OF PRIVACY

by

Charles Gordon Hudgins

A dissertation submitted to the faculty of
The University of Utah
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Department of Philosophy

The University of Utah

August 2013

Copyright © Charles Gordon Hudgins 2013

All Rights Reserved

The University of Utah Graduate School

STATEMENT OF DISSERTATION APPROVAL

The dissertation of **Charles Gordon Hudgins**

has been approved by the following supervisory committee members:

Cynthia Stark	Chair	**5/8/13** Date Approved
Mariam Thalos	, Member	**5/13/13** Date Approved
Ella Myers	, Member	**5/13/13** Date Approved
Leslie Francis	, Member	**5/14/13** Date Approved
Margaret Battin	, Member	**5/13/13** Date Approved

and by **Stephen Downes** , Chair of

the Department of **Philosophy**

and by Donna M. White, Interim Dean of The Graduate School.

ABSTRACT

This dissertation constructs a theoretical approach to understanding the moral value of personal privacy. In its current state, the philosophical literature on the moral value of privacy is fractured in that there are several differing approaches, each emphasizing different aspects of the problem. Some approaches, referred to as “functionalist,” consider the value of privacy as being based on the goods that it brings about. Others emphasize respect for the autonomy of the individual in question, referred to as “autonomy-based.” The view developed in this dissertation combines the central intuitions behind both of these approaches by basing the moral value of privacy on the notion of relational autonomy. Relational autonomy is a conception of autonomy that emphasizes the interpersonal and social embeddedness of individuals, instead of treating autonomy as a singular ideal of independence from such influences. By understanding the value of privacy as a kind of respect for relational autonomy, the relational approach is technically a kind of autonomy-based approach, but one that manages to incorporate consideration of the socially oriented goods emphasized by functionalist accounts. The examples of social media (such as Facebook) and lifelogging are used to explain this approach to the value of privacy.

To persistence.

CONTENTS

ABSTRACT.. iii

Chapters

1 INTRODUCTION 1

2 AUTONOMY-BASED AND FUNCTIONAL ACCOUNTS OF THE VALUE OF PRIVACY 16

- Introduction 16
- Forms of Privacy 18
- Autonomy-Based Accounts...21
- Functionalist Accounts 24
- Reconciling Functionalist and Autonomy-Based Accounts.....................................27

3 RELATIONAL AUTONOMY 34

- Autonomy Simpliciter 34
- Three Questions Regarding Autonomy...38
- What “Relational” Means: A First Pass 39
  - Relational Agency 40
  - Relational Standards of Autonomy 43
  - Relational Accounts of Respect for Autonomy..48
- What “Relational” Means: A Second Pass...52

4 AUTONOMY: RETHINKING THE QUESTION..58

- Introduction 58
- Meyer’s Procedural Standard of Autonomy...62
- Stoljar’s Argument for a Strong Normative Competence Theory of Autonomy 65
- Rethinking the Question: Respect for Autonomy 76
  - Going Beyond Autonomy to Understand Autonomy 76
  - Respect for Autonomy 79
- Conclusion...85

5 PRIVACY AND RELATIONAL AUTONOMY...88

- Introduction 88
- Privacy and Relational Processes 90

Putting the Relational Approach to the Value of Privacy in Context 95
    Nissenbaum's Contextual Integrity...95
    Contextual Integrity and Walzer's Spheres of Justice 103
Conclusion..110

6 APPLYING THE RELATIONAL APPROACH TO PRIVACY 112

Introduction 112
Facebook 118
Contextual Integrity and Facebook 120
Facebook and the Relational Approach to Respect for Privacy............................124
Lifelogging and the Relational Approach to Respect for Privacy 128
Confidentiality, Facebook, and Lifelogging 132
Conclusion...138

7 CONCLUSION...139

SELECTED BIBLIOGRAPHY...142

CHAPTER 1

INTRODUCTION

On November 9th 2012, Gen. David Petraeus resigned as director of the CIA. During the investigation that led to his resignation, the FBI accessed the emails sent by Petraeus to his mistress, Paula Broadwell, as well as emails sent by Broadwell to other parties. Public response to the event predictably raised the issue of privacy. Was Gen. Petraeus' privacy violated? What about Broadwell's or some of the other individuals involved? Did the FBI have the right to access personal email accounts? If so, was it because of the nature of Petraeus' position? Given his position, does he have a right to any privacy at all when it comes to his personal communications?

Clearly, context matters a great deal here. The Director of the CIA is a position holding a great deal of power and responsibility, mostly pertaining to national security. Does that fact not mitigate his or her right to privacy? Or is it still more important to respect his right to pursue goods in his life, such as relationships he may choose to cultivate, by providing some measure of privacy? Yet, in order to answer that question, would we not need to understand how privacy's value relates to other goods, other objects of moral concern? That is, the answers are dependent upon how we understand the relationships between the value of privacy, respect for personal autonomy, and the values of other goods pursued by individuals.

It is cases like these that make understanding the value of privacy so difficult. Clearly we need to understand the contexts involved. We need to understand the value that privacy has relative to that context, as in what we would like privacy to do for individuals in such situations. But we also need to understand the limitations of that value and our corresponding rights and obligations towards the people involved. Furthermore, we need to understand how privacy relates to other issues such as autonomy, liberty, and our larger sense of justice.

At first glance, this may seem like an intractably complicated mess, and hence a lost cause. How are we supposed to untangle this knot of abstract ideas in such a way that our resulting understanding leaves us with a sense of practical direction? After all, if we are going to take the time to think through the value of privacy, we want the resulting account to be able to provide practical and prescriptive recommendations for dealing with complex privacy issues. The goal of this dissertation is to develop an approach to the value of privacy that captures why we take privacy to be important in the first place, and does so in a way that enlightens our understanding of how privacy's value connects with other concerns. The approach that I will describe and advocate is what I call the relational approach to the value of privacy. Before providing an initial description of what the relational approach looks like, I would like to make two points that are important for setting up this project.

First, I will not be making any attempt to define privacy. Privacy has been defined in a number of different ways over the years, from being let alone to more recently a kind of Wittgensteinian family resemblance concept; but my concern here is to understand why we should really care about it in the first place. The task of defining privacy has proven to be a very difficult one. [1] For every definition offered, there seems to be at least one counterexample that has come along with it. That said, I can offer, for explanatory purposes only, a definition provided by Anita Allen.

> [P]ersonal privacy is a condition of inaccessibility of the person, his or her mental states, or information about the person to the senses or surveillance devices of others. To say that a person possesses or enjoys privacy is to say that, in some respect and to some extent, the person (or the person's mental state, or information about the person) is beyond the range of others' five senses and any devices that can enhance, reveal, trace, or record human conduct thought, belief, or emotion.[2]

This is a fairly broad definition that is meant to capture many of the ways in which the term tends to be used. It claims that privacy is a state of restricted access relative to some aspect of the person in question. The access in question can be relative to information about the person, physical access, her conduct, etc. It is important to keep in mind that Allen does not intend this definition to constitute a set of sufficient conditions for what privacy is. Something can be inaccessible but not be private in the usual sense, such as an obscure work of art, long lost texts, Tantalus' grapes, etc.

The second comment I would like to make is regarding motivation. Why undertake a project like this? The main reason is probably the most straightforward. Privacy issues are ubiquitous in much of Western society; and it is clear from public discussion that there are wildly disparate understandings of what goods are really at stake in issues involving privacy. It is also clear that there are very different ways of making sense of the right to privacy. As it turns out, this same distinction between the goods made possible by privacy and the right to privacy runs through much of the academic and legal scholarship on the subject of privacy's value.

---

[1] The literature here is massive. The earliest and canonical reference is Warren and Brandeis (1890), with one of the most recent and extensive attempts being by Daniel Solove (2008) where he appeals to Wittgenstein's family resemblance concept. There the idea is that there is not one single feature that defines privacy, but a cluster of concepts members of which may or may not be present in any one particular instance of privacy.

[2] Allen, Anita L. *Uneasy Access: Privacy for Women in a Free Society*. Totowa, NJ: Rowman & Littlefield, 1988. p. 15

Along these same lines, there are two main schools of thought regarding the value of privacy. The first is made up of instrumental or functional accounts that identify the goods that privacy can enable. These goods include such things as personhood and relationship creation (as well as maintenance), creative endeavors, mental health, and more.[3] When we are "let alone" or enjoy varying "degrees of inaccessibility," we are often times more able to pursue such goods. Hence, we can consider privacy as being instrumental in the achievement of them.

To offer a simple example, consider someone learning a new form of artistic expression, such as playing the cello. The cello is an easy instrument to play, but a difficult one on which to cultivate a pleasant sounding solo voice. It requires a great deal of practice to develop the kind of technique required to make something sound musical and to express emotion through the instrument. Because of this many people might prefer to practice in private, thereby worrying less about the inevitable mistakes and occasional banshee-like sounds.

Privacy in this context can provide the opportunity for the individual to learn to develop a skill and develop a form of self-expression. But notice that in this example we can also see the importance of encouraging the student to play in an orchestral setting, and eventually to try her hand at a solo in front of an audience. Here the idea would be to expand her sense of the instrument and its potential; as well as to enjoy the potential improvements gained through playing with other musicians. This point brings out the sense in which privacy's value on the functionalist view is intended to be instrumental in that it *can* bring about good things, not that it always does. If our cellist insists on always playing in private, she misses out on so much of what there is to learn about the instrument and what can make learning a musical instrument so rewarding.

---

[3] Several authors have constructed such lists. For two examples, see Allen 1988 and Gavison 1980.

The second school of thought on privacy's value is comprised of autonomy-based accounts that link the right to privacy to respect for autonomy. Here the idea is that respect for privacy can oftentimes be tantamount to respect for autonomy. But when we say respect for autonomy we do not just mean respecting wishes or desires. Respect for autonomy is respect for the capacity to form desires and preferences. Hence the relationship between respect for privacy and respect for autonomy speaks to the fact that privacy can impact individuals at the level of their ability to form preferences and make decisions. The idea is that the *kinds of things* we give up or have violated can be very fundamental to our quests for self-creation, our attempts at maintaining our public persona, and more.[4] The loss of opportunities to spend time alone, or the ability to control the extent of our exposure to others, can in turn affect our ability to flourish as people. Therefore respecting the privacy of an individual can be a matter of respecting her autonomous capacity.

One of the core goals of this dissertation is to provide an account of the value of privacy that combines these two schools of thought on privacy's value. On the one hand, we do not want to lose sight of the fact that privacy as a state of affairs can do a lot of very beneficial things for us. This seems fairly obvious and it also seems like an intuition that is worth preserving in our final account of privacy's value. On the other hand, it seems obvious that individuals collectively have a wide variety of views and preferences regarding privacy; and those preferences in turn can be a significant part of what we take to be fundamental to our autonomy. The question is of course, how does one go about combining these two views?

---

[4] See Benn 1971, Beardsley 1971, and Beauchamp and Childress 2001 for examples. See Chapter 2 of this work for an expanded discussion.

What I attempt to do is to adopt an autonomy-based account of the value of privacy that is able to capture the intuitions that drive the functionalist account. After hearing how autonomy-based accounts generally work, it would seem a fairly sensible question to ask: but what exactly do you mean by autonomy, and respect for it? Yes, it is the capacity to form preferences and make decisions (develop conceptions of the good, adopt or reject values, and so on); but is there not a great deal more detail we would want in addition to such a basic definition? Furthermore, would not providing that detail offer some insight into autonomy's possible connections with privacy?

This is the line of inquiry that I will pursue in this dissertation; and the basic thought process can be described in the following way. Privacy is an inherently social concept. Even though it is usually understood as the absence of something, that something is always in some way, shape, or form related to other people, groups, institutions or even governments. Hence the goods enumerated by functionalist accounts, insofar as they are enabled by privacy, are brought about by social factors. Somewhere in the causal chain, some kind of social element will come into play. It would seem to follow from this conclusion that if our account of respect for autonomy is based on social factors, then it has an excellent chance of tracking the same intuitions that connect privacy's value to the goods described by functionalist accounts.

This line of reasoning should raise questions. First, what exactly does it mean for an account of respect for autonomy to be "based on social factors?" And second, just because an account incorporates social factors, does that enable it to capture the goods we tend to associate with privacy? A good portion of this dissertation is devoted to answering these questions. But for now, I will simply appeal to what is known as relational autonomy. Relational autonomy is a conception of autonomy which treats as conceptually necessary the idea that individuals are to varying degrees influenced by cultural and interpersonal factors. That is, the decisions and preferences we make are shaped by such influences.

The idea that we are influenced by culture, our relationships with others, norms, etc. may seem patently obvious; but the devil is in the details. More specifically, I address relational conceptions of three different questions within the autonomy space. First is the question of agency. What is the constitution of the self? How is it created and maintained? Is it a static thing, a process, or should it be reified at all? Second is the question of standards of autonomy. What are the proper criteria for determining a particular decision or preference to be autonomous? Third is respect for autonomy. How exactly do we go about respecting the autonomy of others? How do we understand rights and obligations based on respect for autonomy in the sense of preserving and even promoting the autonomous capacity of others?

It is the third question, respect for autonomy, which will ultimately do the heavy lifting when it comes to real-life situations involving the value of privacy. According to the relational approach to the value of privacy, respect for privacy is one form of respect for autonomy. Remember that relational conceptions of autonomy focus on the way in which individuals are influenced by interpersonal and social factors. Privacy or the lack thereof can oftentimes play a role in *how* that influence plays out.

Consider a brief example. One might argue that the right to privacy in cases involving mental health should be upheld strictly, at least in the sense that patient anonymity should be reinforced throughout the entire lifecycle of care. Elena Premack Sandler discusses a clinic that operates along these lines in an article in Psychology Today:

> The case I examined is that of a health care system in Maine that created a clinic where clients have the utmost privacy. There's no public waiting room and sessions are scheduled to avoid encountering other patients. The clinic is for a specific population: high-level executives, community leaders, attorneys, doctors, and clergy - all who may be less likely to seek treatment if they have to do so in a public clinic.[5]

The idea is of course that patients are more likely to seek and pursue a course of treatment on their own terms if they have some reassurance that the entire process will remain private. Mental health issues can have strong stigmas attached to them, and reducing the likelihood of experiencing those stigmas would naturally make one more likely to seek treatment.

Where the functionalist view might use this example to show that “the primary benefit of privacy is really mental health,” the autonomy-based account might use it to say, “this is really about the rights of patients to choose the nature and terms of their treatment.” But notice that insofar as privacy is seen as promoting the good in question, namely mental health, it does so in virtue of *how* individuals function in an embedded context. This would include their views on stigmas associated with mental illness, particular relationships that they have with others, how such information (if released) could affect their careers, and so on. Hence, when we think about the right to privacy in a relational context (i.e., as an instance of respecting relational autonomy) we consider the rights of individuals involved in a way that speaks directly to *how* privacy can promote the goods involved.

---

[5] http://www.psychologytoday.com/blog/promoting-hope-preventing-suicide/201107/privacy-and-mental-illness-part-ii

In this case involving patient anonymity, we consider how deciding on a policy either way would impact the ability of the individuals affected to function as autonomous agents in a relational context, and in turn how those goods may or may not be promoted. The good in question here is mental health, and privacy enables that good by allowing patients to seek it more on their own terms, free from the fear of social stigma. Thinking about patient autonomy in a relational sense requires that the individuals in question be thought of as embedded in some social context and thereby seeking those goods in that context. It is in this way that we are able to capture the intuitions driving the functionalist view. As stated before, insofar as the goods described in that view are actually caused by privacy, they are done so in a way that incorporates social factors.

What the relational approach suggests is that we should understand how privacy functions in terms of preserving and/or improving the autonomous capacity of individuals. Hence, while privacy itself is a very distinct concept in a non-normative sense, respect for privacy and the rights and obligations it entails are ultimately best understood as a kind of respect for autonomy. Specifically, they are best understood as a kind of respect for relational autonomy.

Hence, the general strategy here is to adopt a particular conception of autonomy that is based on social factors, flesh it out in more detail by addressing more specific questions within the autonomy space, and then bring that understanding to bear on privacy issues. There is, however, one more step necessary before being able to apply a relational conception of autonomy to the value of privacy in particular cases.

The relational approach to the value of privacy is a “global” one in that it talks about privacy in its most general sense. The difficulty behind taking this approach is that many of the situations involving privacy and questions about its value involve details associated with the particular context in which that situation occurs. Context matters, as they say. Why we think a piece of information is important, or why it is important that an individual enjoy some form of inaccessibility depends greatly on the details behind what that information actually is and why people normally seek to keep it private. If the goal of coming up with an account of the value of privacy is to be able to make prescriptive recommendations, then it is important to bring these details into consideration. However, drawing a connection between respect for relational autonomy and the value of privacy is not in and of itself sufficient to capture these details. We need some addition to the relational approach that enables it to dig into the details of situations and understand what normative impact those details really have. This is why I will incorporate the work of Helen Nissenbaum, specifically her notion of contextual integrity. I will discuss contextual integrity in much greater detail in Chapter 5, but the core idea is that for each context we have a particular understanding of how privacy actually works and a set of expectations as to how privacy is supposed to work. Furthermore, according to Nissenbaum’s contextual integrity approach, there is a presumption in favor of maintaining how things work, and any change is considered a violation of contextual integrity unless it can be shown to be beneficial relative to the context in question. In other words, there is a presumption in favor of maintaining the status quo, unless a change in how information is managed can be shown to improve how things work in that context.

The case of Gen. Petraeus is a good example of the importance of context and how our understanding of the context in question can drive our thoughts regarding privacy’s value. The public debate that has arisen over his case often factors in the nature of his job as Director of the CIA, and how the responsibilities inherent in such a role might preclude the right to personal privacy, at least in the way that we think of it. But in order to get to the bottom of that intuition, we need to understand in a more robust and theoretical sense how privacy relates to certain core concepts in political philosophy. My claim is that, in order to achieve this understanding, we first need a more thorough view of how privacy connects with privacy in a way that is sensitive to contextual considerations.

Towards the end of this dissertation I offer two applications of the relational approach to the value of privacy: Facebook and lifelogging. I chose Facebook as one of the examples because it has two key features. First, the functionality of the site is fairly straightforward, at least from an end user perspective, so the phenomenon in question is simple, *on that level*. Second, because of the vast variety of interactions that occur on the site, and the vast array of personal information that is both shared and generated on the site, it offers a wealth of examples and richness despite the simplicity of the functionality.

Lifelogging was chosen as an example because it is a reasonable extrapolation of Facebook and current social media functionality. Lifelogging is the practice of recording and keeping a journal of multiple aspects of one’s every day experiences through video, audio, and assorted other means of surveillance. The information collected can be essentially anything, from aspects of one's physiological states (such as blood pressure, pulse, etc.) to experiences like face-to-face conversations with friends or strangers. Also, the information recorded can be stored and/or disseminated on the Internet. To be accurate, lifelogging as I am defining it here has only partly come into existence; hence, it has also been chosen in an effort to be forward-looking.

In both cases, Facebook and lifelogging, the goal is to understand the impact of privacy or the lack thereof on autonomous capacity understood relationally. While I will have much more to say about this in Chapter 6, the gist of the analysis is that both cases have a great deal of potential to impact the autonomy of individuals utilizing them. Hence it is important that users be made aware of those potential impacts and be given opportunities to take control of their experiences to the extent that is practical given the nature of the technology involved. Furthermore, users should be made aware of what the limitations are in terms of controlling their experiences while using these services.

To summarize, this dissertation is an exploration of the conceptual consequences of inserting a particular conception of autonomy (i.e., relational autonomy) into the claim that respect for privacy is deeply intertwined with respect for autonomy. I propose understanding the value of and respect for privacy in terms of respecting the autonomy of individuals, where that autonomy is understood within a conceptual framework that incorporates cultural and interpersonal influences and hence the intuitions driving the functionalist view.

To be clear, I am not attempting to provide a means of resolving conflicts concerning privacy in applied ethics directly and finally. Given the vast variety of instances involving privacy, I do not think that it is truly feasible to come up with any sort of method that guarantees or even comes close to guaranteeing a direct resolution in all cases. Instead I am seeking to establish a framework for how such conflicts can be understood and a context in which they can be discussed that focuses on what is genuinely at stake while preserving the central intuitions behind both functionalist and autonomy-based approaches to the value of privacy.

In Chapter 2, I will provide a more detailed exegesis of the functionalist and autonomy-based accounts of privacy's value, and discuss in more detail how to combine the two such that the primary intuitions that drive each are preserved. What I propose is that autonomy-based accounts are ultimately better suited to capturing privacy's value, but only if we understand autonomy in the relational sense and thereby capture many of the goods that motivate the functionalist view.

In Chapter 3 I discuss what is meant by relational autonomy, and how it can be contrasted with the more traditional or received view of autonomy. More specifically, the received view does not consider relational factors as a necessary part of the conceptual picture when it comes to how we form our beliefs or preferences. Whereas the relational view, on the other hand, considers those things as being conceptually necessary. Again, there are several additional distinctions to be made in order to understand just what “relational” means and how it functions as an umbrella term; and that will be the focus of Chapter 3.

In Chapter 4 I discuss a debate regarding autonomy focused on what should be the proper standard for an autonomous decision or preference. That debate, as it stands now, is between procedural accounts of autonomy and substantive accounts. Procedural accounts focus on how preferences are formed, whereas substantive accounts focus on the content of those preferences. In lieu of discussing the details at this point, my contention is that once we make the common sense assumption that interpersonal and cultural factors play a role in both how we form our beliefs and the ultimate content of those beliefs, the debate over which standard is correct becomes intractably complex.

What should be our focus, instead of a standard of autonomy, is the proper way of respecting our autonomous capacity. That is, the ability to recognize an interpersonal or cultural influence on our decision making process, as such, is ultimately a skill, and one that can be nurtured or diminished depending on our experiences and how we are treated by others. Once we make this shift from worrying about a standard of autonomy to worrying about an understanding of how best to respect autonomous capability, it is much easier to understand respect for privacy as it can be a part or component of respect for autonomy, specifically relational autonomy.

Chapter 5 continues this point by discussing what I take to be the relationship between relational autonomy and the value of privacy. This is done in two stages. First I discuss the non-normative dynamic between privacy and what I refer to as relational processes. Here, I will discuss what it is that privacy actually does in the context of relational autonomy. Second, I will introduce Helen Nissenbaum’s contextual integrity approach mentioned above. Here, I will go into more detail as to how and why context matters when seeking to make practical recommendations in situations regarding privacy.

In Chapter 6 I will discuss both Facebook's privacy policy and lifelogging. This discussion will offer more details of both and discuss the potential impact that each has on the autonomy of individuals utilizing the technologies. These impacts are what form the basis for my analysis of how and why privacy should be respected in these contexts.

In Chapter 7 I will conclude with some remarks on what I take this dissertation to have accomplished and what it has not. As mentioned above, the primary goal is to understand the value of privacy in such a way that both captures the two primary intuitions currently identified in the literature, and do so in such a way that offers at least some recourse for resolving issues involving privacy in virtue of understanding how it connects with other concerns.

CHAPTER 2

AUTONOMY-BASED AND FUNCTIONALIST ACCOUNTS OF THE VALUE OF PRIVACY

Introduction

Early treatments of privacy tended to take on the task of defining privacy as well as understanding its value. Attempts at defining privacy deal with the complicating fact that "privacy" is used in wide variety of ways. Judge Thomas Cooley's famous definition described privacy very broadly as "being let alone."[6] Allen and Gavison consider privacy to be a matter of inaccessibility or limited access, respectively.[7] Westin defined privacy as a claim of individuals or groups against others regarding information about themselves.[8] The challenge involved in defining privacy, according to some thinkers, is not necessarily to create a definition that encompasses every conceivable use of the term, but to find a sort of balance between that extreme and defining privacy so narrowly as to eliminate important aspects. For example, Allen's approach was to offer a definition that both made clear the applicability of the term and provided some prescriptive guidance for its use.

---

[6] The most famous discussion of this can be found in: Warren and Brandeis "The Right to Privacy" *Harvard Law Review* Vol. IV December 15, 1890 No. 5

[7] See Allen, Anita *Uneasy Access: Privacy for Women in a Free Society*. Rowman & Littlefield 1988 New Jersey; and Gavison, Ruth "Privacy and the Limits of Law" *The Yale Law Journal* Vol. 89, No. 3 January 1980 p421-471

[8] Westin, A. Privacy and Freedom Atheneum 1967.

It seems that there is a similar problem associated with accounts of the value of privacy. Because the use of “privacy” varies, the sense in which privacy can be thought of as valuable varies as well. Hence, one could offer an account of the value of privacy by attempting to encompass every conceivable way in which privacy can have value. Alternatively, one could focus on what one takes to be fundamental or central to the value of privacy, disregarding certain complexities. The method that will be adopted here is a kind of compromise. I will explain and assess the two approaches to understanding the value of privacy introduced above, namely the autonomy-based and functionalist accounts. The first approach is that of understanding how the value of privacy is associated with respect for autonomy. The second approach is to understand the goods that are enabled by privacy, such as personhood enhancement and relationship enhancement. The chapter will conclude by discussing the need to accommodate the intuitions behind both of these approaches and the difficulties faced by them. In turn, the remaining chapters will be devoted to describing what I refer to as the relational approach to the value of privacy and how it addresses the goal of combining the intuitions behind both the autonomy-based and functionalist accounts.

It may also be helpful to point out that while the primary purpose of this chapter is exegetical, it has the additional goal of providing some sense of what it means to offer a more complete picture of the value of privacy. To be clear, my goal here is not to offer an account of the value of privacy that captures every conceivable way that it can have value. Instead, my intent is simply to capture the main intuitions found in the literature and then, in the following chapters, to explain how this view of privacy’s value can be supplemented by a more detailed understanding of autonomy. But before getting into the accounts of privacy’s value, it will be helpful to say a bit more about the concept of privacy itself by describing some of the different types of privacy.

Forms of Privacy

The first thing to mention here is that there is not much agreement among scholars as to the definition of privacy, much less any real agreement on what the various kinds of privacy actually are. As a general rule, it is helpful to be able to distinguish between different forms of privacy so that we can be more precise about the source of our concern in cases involving privacy. That said, this section will provide a brief overview.

Drawing somewhat from Allen Westin’s analysis as many scholars have done, Anita Allen lays out a few different categories of privacy and synonyms for each: seclusion and solitude, anonymity and limited attention, nondisclosure and confidentiality and secrecy.[9] Seclusion and solitude are the various ways in which different kinds of observation (or other forms of sensing) can be impaired.[10] The impairment of observation can of course occur in a number of different ways, such as a fence, a wall, or sheer distance.

Anonymity and limited attention are just as they sound. Sometimes individuals may enjoy privacy in virtue of not being the focus of attention at all, or only in a limited way. This form of privacy tends to be of more significance when the individual or group in question is part of a larger group in a large crowd. Allen mentions the examples of public places and the homes of large families.[11] To that list we could conceivably add large online social networks such as Facebook or LinkedIn.[12]

---

[9] See Allen,1988, as well as Westin1967. While the literature on the forms of privacy is fairly extensive, it will only be necessary here to get the basic ideas across. Furthermore, the bulk of the variation in the literature centers around how these ideas relate to one another in the sense of collectively forming a definition of privacy. Since that is not my concern in this dissertation, I will keep the exegesis short.

[10] Allen, 1988, p. 23

The cluster of nondisclosure, confidentiality and secrecy is a bit vaguer. Allen describes them in the following way:

> As species of information non-disclosure, confidentiality and secrecy are forms of privacy. Confidentiality is achieved where designated information is not disseminated beyond a community of authorized knowers. Secrecy is information non-disclosure that results from the intentional concealment or withholding of information.[13]

Secrecy is a matter of intentionally preventing information from being discovered by either particular parties, or anyone in general depending on the context. There are numerous examples one might think of, such as concealing parts of one's past, a medical condition, or even ideological or religious beliefs.

Confidentiality is essentially the idea of confining information in a particular way such that only a particular person or group of people, usually predetermined in some way by agreement or convention, has possession of it. The most common example would of course be doctor patient confidentiality. When a patient shares information with a doctor for the purposes of treatment, professional obligation dictates that the medical professional in question not share that information, at the very least not without the patient's consent.

While most scholars would agree that the concepts of privacy and confidentiality are distinct, there is some disagreement as to whether confidentiality is really a form of privacy, or something completely independent. As noted above, Allen considers it a form of privacy; however, Leslie Francis has argued that there is an important sense in which the two can be seen as independent.[14]

---

[11] Allen, 1988, p. 24
[12] I will have more to say about these examples at a later point. For now it is sufficient to point out that an individual may seek out interaction with a subgroup of individuals within an online social network, yet wish to limit interaction with others in that same network.
[13] Allen, 1988, p. 24

The distinction goes roughly as follows. One may suffer an invasion of privacy in virtue of being intruded upon or seen, monitored, etc. Alternatively, one may suffer a breach of confidentiality by having information that was already possessed by another party, whether voluntarily or involuntarily, disseminated to a third party without the individual's consent. So the first leg of the distinction is that when we consider concerns about information, it is important to keep separate the initial gathering of the information (and whether that constitutes an invasion of privacy) and the stewardship of the information after it has been gathered (and whether there has been a breach of confidentiality).

The second leg of the distinction is whether our concern about an invasion of privacy may trump our concern about a breach of confidentiality or vice versa. Francis argues that there are cases where confidentiality is the more important of the two.[15] She cites examples, primarily in medicine and public health, where the taking of a biological sample is often the lesser of the two concerns. In other words, the impact of having one's mouth swabbed is far less than the possible impact of the test results, especially in cases like infectious diseases such as HIV. Hence when considering whether a particular case of testing is justified (forced or otherwise), it is important to keep these two notions distinct and to understand what is really at stake.

For present purposes, I am not concerned with the debate over whether confidentiality should rightly be considered distinct from privacy. I will, for the sake of simplicity, take Allen's approach and consider it a form of privacy. However, I take Francis' point to be an important one, and will return to it later during the discussion of Facebook's data use policy.

---

[14] Francis, Leslie Pickering "Privacy and Confidentiality: The Importance of Context" *The Monist* vol. 91 no 1 p 52 -67 2008

[15] Francis, 2008

Autonomy-Based Accounts

"Autonomy-based" accounts of the value of privacy attach privacy in some way to respect for autonomy. Autonomy-based accounts vary in exactly how they understand the link between autonomy and privacy. Beauchamp and Childress offer a straightforward example of an autonomy-based account.

> The primary justification [of rights to privacy is] based on the principle of respect for autonomy. We often respect persons by respecting their autonomous wishes not to be observed, touched, or intruded upon. On this account, rights of privacy are valid claims against unauthorized access that have their basis in *the right to authorize or decline access*. These rights are justified by rights of autonomous choice that are correlative to the obligations expressed in the principle of respect for autonomy.[16]

The key word here is "correlative." On Beauchamp and Childress' view, the principle of respect for autonomy lays out obligations "to build up or maintain others' capacity for autonomous choice while helping to allay fears and other conditions that destroy or disrupt their autonomous actions."[17] So what we are supposed to be concerned with is the correlative rights to these obligations. The right to authorize or decline access to oneself is a right of autonomous choice that corresponds to the obligation to "build up or maintain" autonomous capacity. Therein lies the basis of rights to privacy. People need to have some semblance of control when it comes to declining or granting access to their bodies and information about themselves. Without such control, their ability to shape their lives and to ultimately pursue their conception of the good is hindered.

---

[16] Beauchamp, T. and J. Childress (2001). *Principles of Biomedical Ethics*. Oxford, Oxford University Press p. 296

[17] Beauchamp and Childress, 2001, p. 63

Stanley Benn offers a similar account of the right to privacy, based on respect for persons as autonomous choosers. The "principle of privacy" that he proposes is:

> [T]hat any man who desires that he himself should not be an object of scrutiny has a reasonable prima facie claim to immunity. But the ground is not in the mere fact of his desiring, but in the relation between himself as an object of scrutiny and as a conscious and experiencing subject.[18]

Individuals, Benn claims, are engaged in "self-creative enterprises" which can "be disrupted, distorted, or frustrated even by so limited an intrusion as watching."[19] On the one hand, then, Benn is noting that privacy enables or makes possible the sorts of things or "enterprises" that we pursue as autonomous beings. On the other hand, what follows from this is that our right to privacy is grounded in the respect that we are due as *autonomous* choosers.

The core idea here is that an invasion of privacy can disrupt a "self-creative enterprise" by changing the goal of the enterprise itself. An invasion of privacy can cause a change in what is seen as worth pursuing and what is not in a wide variety of ways. It can even affect how an individual sees themselves, touching on such things as self-worth and self-trust. However, saying more about the relationship between respect for privacy and respect for autonomy will inevitably require saying more about autonomy itself and how our understanding of autonomy ultimately informs respect for privacy. I will take up that issue in later chapters.

Discussing privacy in the guise of "selective disclosure," Elizabeth Beardsley provides yet another autonomy-based account that construes the right to privacy as a matter of respecting autonomy. What she calls the "norm of selective disclosure" is defined as "Do not seek or disseminate information about X which X does not wish to have known or disseminated."[20] In turn, she understands the basis for this norm, this right, as being an expression of the principle of autonomy.

---

[18] Benn, S. "Privacy, Freedom, and Respect for Persons," *Nomos XIII: Privacy*, Pennock J. & Chapman J. ed.'s Atherton Press NY 1971 p. 13

[19] Benn, 1971, p. 26

> The first question to consider is whether this principle is anything more than a special case of the principle of autonomy. Are violations of selective disclosure with respect to X morally offensive strictly qua violations of X's autonomy with respect to his determining for himself whether or not he will have known or disseminated a certain fact about himself in situation S? They are indeed morally objectionable for this reason…[21]

Much like the previous two accounts, Beardsley is straightforwardly associating the value of privacy with respect for autonomy. It is important to note, however, that privacy is still, on this and the previous two accounts, a distinct concept. Beardsley puts the point this way:

> The norm of autonomy is, I have argued, what gives our obligation to respect another's right of selective disclosure its moral rationale. But selective disclosure constitutes the conceptual core of the norm of privacy. A manageable distinction can be drawn between those violations of autonomy which violate X's right to selective disclosure and those which do not.[22]

The idea, then, is that violations of the right to privacy or selective disclosure form a kind of subset of violations of autonomy. In this way, privacy as a non-normative concept is still quite distinct from autonomy, but the right to privacy is in effect "reduced" to respect for autonomy.

To summarize, each of these accounts considers the value of privacy to be related to, or in some cases, partly constitutive of respect for autonomy. The underlying theme here is that autonomy-based accounts constitute a common intuition regarding the value of privacy, and one that should be addressed in any account of privacy's value. The intuition can be summarized in the following example. Suppose that I have a close friend who wishes that a certain fact about himself be kept secret, or at least be kept between the two of us. However, I do not understand *why* he would want this fact kept secret. Say, for example, he wants to conceal the fact that his middle name is Robert. Now, it seems quite absurd that my friend would keep something like this a secret, but I can at least respect his wish to do so. In this case, I have no inkling of the instrumental value of privacy (i.e., what the benefit is to my friend in keeping this secret), but I can understand my friend's right to privacy as a matter of respecting his autonomy.

---

[20] Beardsley, Elizabeth "Privacy: Autonomy and Selective Disclosure" *Nomos XIII: Privacy* Pennock J. & Chapman J. ed.'s Atherton Press NY 1971 p. 57

[21] Beardsley, 1971 p. 65

[22] Beardsley, 1971 p. 70

Furthermore, it is not just that I respect his autonomy insofar as I respect an autonomous wish. The idea is that I take my friend's preference to be a part of his broader concerns about his privacy and how he wishes to manage it for himself. Those concerns are in turn partly constitutive of a broader self-creative and self-governing enterprise that hinges significantly on how and to what extent he is able to function on an interpersonal and/or cultural level. As will be shown in the following two sections, this intuition about autonomy and the value of privacy is not accounted for by functionalist or instrumental accounts.

Functionalist Accounts

The example above highlights the intuition underlying the association between autonomy and the value of privacy. Suppose, however, that I were to "dig deeper" as it were and inquire as to *why* my friend Robert wants to keep the fact of his middle name a secret. Does the secrecy provide some benefit or enable some good for my friend?

Functionalist accounts focus on this aspect of privacy in order to understand its value. In short, functionalist accounts attribute instrumental value to privacy by identifying the goods privacy makes possible.

Ruth Gavison offers perhaps the broadest version of a functionalist account. Privacy, Gavison argues, can make possible such things as "a healthy, liberal democratic and pluralistic society; individual autonomy; mental health; creativity; and the capacity to form and maintain meaningful relations with others."[23] As an example, consider Gavison's comments regarding "freedom from physical access":

> By restricting physical access to an individual, privacy insulates that individual from distraction and from the inhibitive effects that arise form close physical proximity with another individual. Freedom from distraction is essential for all human activities that require concentration, such as learning, writing, and all forms of creativity.[24]

Privacy, in the form of physical privacy, can allow an individual to pursue certain activities that are otherwise difficult if not impossible to pursue in the company of others. Hence, privacy is instrumental for such goods that require some "alone time" as it were.

Anita Allen's account of the value of privacy is functionalist as well, however, she argues for a more focused version than Gavison. Allen claims that approaches such as Gavison's "underemphasize the close and special connections moralists have stressed between and among privacy, personhood, and fitness for social participation and contribution."[25] According to Allen, the most complete account would hold as central to the goods enabled by privacy what are referred to as personhood creation and enhancement as well as relationship creation and enhancement.

---

[23] Gavison, 1980, p. 442

[24] Gavison, 1980, p. 446-447

[25] Allen, 1988, p. 43

By personhood creation and enhancement, Allen intends that privacy as a state of inaccessibility (whether as informational, physical, or decisional) can, in part at least, allow individuals to pursue conceptions of who they think they should be or want to be.[26] For example, if one were to consider her financial activities and records (such as what and when she buys, her credit record, etc.) to be a significant statement about whom she is, then she would be sensitive to whom was privy to such information. Furthermore, knowing that someone, anyone for that matter, were able to access that information would affect her actions and therefore affect her ability to, in effect, be herself. She might not purchase things that she wanted, take trips, or engage in certain hobbies if she felt that those activities were being monitored by individuals, companies or institutions. The truth of personhood theories does not require that everyone think or feel this way about financial records or any one thing in particular, only that privacy in a general sense, can have this type of "big brother" effect on individuals.

Relationship creation and enhancement works in a similar fashion. Privacy enables individuals to pursue relationships in ways that are more suitable to them. Knowing that individuals aside from one's "significant other" would become aware of certain facts about either ourselves or the relationship itself changes how we act and feel. For instance, offering that first "I love you" in a relationship can be daunting enough "in private," and well-nigh impossible in front of a crowd. It is this function of privacy, in combination with personhood theories that, according to Allen, best gets at the value of privacy.

> On personhood theories, privacy is described as a condition or set of social practices constituting, creating, or sustaining boundaries that should be drawn between ourselves and others in virtue of our status or potential as persons. In

---

[26] Allen, 1988, p. 43

> short, personhood theories maintain that the value of privacy is that it creates, sustains, or enforces personhood. Not all personhood accounts go the step further to explain the moral value privacy has relative to fitness for social participation. To do so would be to provide the most complete, fundamental, and explanatory account of the moral value of opportunities for individual privacy.[27]

I think that Allen is right to emphasize these aspects of the instrumental value of privacy. It seems that when we think of *why* an individual has a preference or a wish that something remain private, we tend to think in terms of such things as social participation. The idea is that, first, privacy is doing something for that individual, and that, second, what that something is involves either being who one wants to be or pursuing a relationship in a particular way. This is not to say that this is strictly all that privacy can do for us, or that these are necessarily good things in all cases. The claim, then, is that privacy *can* enable individuals to pursue and develop individual conceptions of self as well as enact preferences regarding relationships with others. Furthermore, these things form the core of the goods made possible by privacy.

Reconciling Functionalist and Autonomy-Based Accounts

Up to this point I have explained two views of the value of privacy. The first, autonomy-based accounts, associate the value of privacy with respect for autonomy. These approaches to the value of privacy represent the intuition described above, that individuals are prevented from governing themselves when they are prevented from controlling other people's access to them. The second approach, functionalist accounts, state that the value of privacy can be understood in terms of the goods that it can enable. Functionalist approaches, such as Allen's, represent a second common intuition that we have regarding privacy. Namely, that we find privacy important simply because of the goods it enables.

---

[27] Allen, 1988, p. 43

The challenge, then, is to offer an account of the moral value of privacy that is able to make sense of both of these intuitions. It is important to note here that I am not arguing that these intuitions are in some sense fundamental, incommensurable with one another, or irreducible. The point is simply that some way should be found to accommodate both intuitions. In addition, the account offered should also be able to deal with the difficulties faced by each.

Gavison provides an analysis of one such problem. She describes autonomy-based accounts as want-satisfaction accounts:

> The want-satisfaction argument posits the desirability of satisfying wishes and thus provides a reason to protect all wishes to have privacy. It does not require empirical links between privacy and other goals. Moreover, the notion that choice should be respected is almost universally accepted as a starting point in practical reasoning. The want-satisfaction argument cannot carry us very far, however. It does not explain why we should prefer X’s wish to maintain his privacy against Y’s wish to pry or acquire information.[28]

Notice that this is a different sort of respect for choice than what I have outlined above via the accounts of Beauchamp and Childress, Benn, and Beardsley. The relationship, as it is outlined here, is strictly a matter of respecting autonomous wish. Gavison does however mention both Beardsley and Benn in a footnote from this passage. She states that “to some extent, Benn’s discussion goes beyond the want-satisfaction argument when he suggests that there is something especially disrespectful in certain invasions of privacy.”[29] However, she offers no response to this idea of “especially disrespectful.”

---

[28] Gavison, 1980, p. 441

[29] Gavison, 1980, p. 441 n. 64

So the problem is that just looking at the desires of the parties involved offers us no guidance as to whose desire should be honored, and whose should be set aside. And in fact, Gavison sees functionalist accounts as essentially filling in this gap. Now, despite the disparity between Gavison's account and the somewhat deeper notion of respect for autonomy, it is possible to see how the difficulty still applies.

Consider an example. A computer hacker would presumably have a desire for privacy in the form of anonymity, while his intended victim would have a desire to keep certain information private. Clearly these desires are in conflict with one another. Furthermore, most have strong intuitions to the effect that the hacker who is attempting to steal valuable information from the victim is not deserving of privacy in this case. It is this latter intuition about desert that Gavison sees as being supplied by functionalist accounts.

It would seem that autonomy-based accounts do not offer enough guidance in order to make sense of such cases. What is it about the victim's legitimate interest in self-governance that trumps or overrides the hacker's legitimate interest in self-governance? Simply stating that the right to privacy is based on respect for autonomy tells us very little about many of the real world cases involving privacy. In fact, many of the more difficult issues associated with privacy involve conflicting interests in some way or another. As mentioned in the introduction, when someone gains accessibility, someone else usually loses inaccessibility.

This objection is not necessarily fatal to autonomy-based accounts. It does, however, point out the need for further explanation or refinement of these views. More specifically, it highlights the need to understand the underlying account of respect for autonomy in such a way that will drive our understanding of how to adjudicate competing claims to privacy. As has been stated a few times above, this is more a matter of understanding how a particular case involving privacy connects to respect for autonomy. And again, we will need to have an independent account of respect for autonomy in order to make progress here. This is an issue to which I will return in later chapters.

As Gavison points out, functionalist accounts offer one solution to the problem of conflicting autonomies. By identifying the goods associated with privacy, these accounts offer some recourse for determining when privacy has positive value and when it does not. It could be argued, for instance, that the reasons the hacker has for desiring privacy do not, in effect, qualify as the kinds of goods identified by functionalist accounts. He desires anonymity in order to achieve his goal of stealing information from someone, and “stealing” simply is not the kind of thing that can be thought of as a good associated with personhood creation or enhancement. Assuming Allen’s account, however, it is not clear if privacy, in this case, is not actually supporting something like personhood creation. What if this individual identifies the activity of hacking as something central to whom he is?[30] If so, what is to be said in defense of the victim, who might consider the information (and control over it) as being important to maintaining her sense of self? Again, some might have strong moral intuitions to the effect that the hacker is not deserving of privacy in this case, but when we attempt to apply such notions as personhood enhancement and relationship enhancement to this case, it becomes difficult to discern a practical recommendation. In short, it is not clear which instance of personhood enhancement gets respected and for what reason.

---

[30] This is not much of a stretch. Hackers often think of themselves as hackers first, and whatever else second. What is more interesting, is that their accomplishments (i.e., who and what they have been able to hack) go a long way in establishing their online identities and reputations.

It might be argued that Allen's functionalist account simply does not speak to such cases, and hence, its shortcoming is more a matter of a lack of comprehensiveness than a matter of identifying something important about the value of privacy. This response, however, is not a very appealing one. Since this is clearly a case involving privacy, any account of the value of privacy should have something productive and insightful to say about it. The problem, then, is not so much a matter of applicability, as it is a matter of underspecification. Just as autonomy-based accounts do not say enough about respecting preferences in the case of privacy (per Gavison's objection above), functionalist accounts do not say enough about the goods associated with privacy.

Consider a second example. An abusive husband has an interest in keeping the fact of his abuse private. He beats his wife when he feels the need to assert his control over the relationship and to in effect be the "man" of the house. He does this because this was how he was raised. The men are in control and the women do as they are told. Suppose further that if a friend of the family were to find out about this abuse and were to inform the authorities, the husband would no longer be able to pursue his relationship with his wife in the way he sees fit. In other words, a loss of privacy in this case would have just the effect that Allen's account says it would. It would hamper the processes of personhood and relationship enhancement, at least from the point of view of the abusive husband.

Just as in the previous example, we have strong moral intuitions to the effect that the husband's privacy deserves little if any respect. At least, not if what he is going to do with that privacy is physically and mentally abuse his wife. The problem for Allen's account, then, is that there is not enough guidance in terms of what is meant by personhood and relationship enhancement. If her account is to offer any practical guidance, more needs to be said as to what constitutes personhood and relationship enhancement. We could, for example, offer an objective account of personhood and relationship enhancement wherein what the husband is doing here simply does not meet the requirements or falls short in a particular way.

We are now in a better position to describe the challenge posed by these conflicting approaches. This challenge can be broken down into two parts. First, privacy enables the goods enumerated by functionalist approaches as a matter of interpersonal and social dynamics (i.e., that the fact of being inaccessible at times allows one to pursue such things), and those goods were in fact chosen as goods to be pursued. Not only do I need the state of inaccessibility, I need it to pursue a conception of self or a relationship in the way that *I want to*. Hence, the value of privacy comes not only from the goods that it enables but also from the fact that some version of those goods was in fact chosen by an autonomous individual and bare directly on the enterprise of self-creation in an embedded context.

Second, neither autonomy-based nor functionalist accounts offer, on their own, enough specification to be of practical use. While they both identify what are important aspects of the value of privacy, neither is complete in the sense of being able to make an actual recommendation for action. It is not clear, however, what further specification is called for in such cases.

In short, the two primary intuitions regarding the value of privacy create friction with one another. Should we simply respect privacy as an autonomous preference, regardless of what is chosen? Or should we let our intuitions about what should be pursued trump our concern about respect for autonomy?

What falls out of identifying this problem is a kind of requirement for any account of the value of privacy. It should be able to accommodate the intuitions associated with both kinds of approaches in a way that still offers some practical guidance. In order to accomplish this, it will be necessary to specify the goods associated with privacy in such a way that also informs our understanding of respect for autonomy itself. In other words, if I want to claim that the "goods" desired by the abusive husband, for example, do not merit respecting his privacy, I need to show how those goods not only do not fall under the functionalist account but also that they do not qualify under the more general account of human flourishing that underlies respect for autonomy. I need to show why both his privacy (in the functionalist sense) and his autonomy do not deserve respect in this case.

In the following chapters, I will argue that the best way to approach this problem is to understand the value of privacy in terms of respect for *relational* autonomy. Relational autonomy understands autonomous capacity in much the same terms as the functionalist approach understands the goods associated with privacy. That is, the functionalist approach picks out goods that are socially and culturally oriented; and the relational approach to autonomy considers social and cultural factors to be a necessary component to such things as agency and conditions for autonomous decision making. In this way, there is at least some recourse for moderating claims both to respect for autonomy and the functional value of privacy.

CHAPTER 3

RELATIONAL AUTONOMY

Autonomy Simpliciter

Before discussing relational autonomy, it will be necessary to make some remarks about autonomy in the most general sense. “Autonomy” has been applied to states or political entities, as well as “to actions, to persons, to the will, to desires, to principles, to thoughts.”[31] Given that the term has such a long history and a diverse usage, it is useful to, at the very least, state the approach to characterizing autonomy that I will use. That is, is the task here to argue for the proper use of the term, to attempt to describe every conceivable application, or merely find a definition that works given current purposes? Dworkin offers the following approach for characterizing “autonomy.”

> My own view is that as a term of art, one cannot look to the ordinary uses of the concept. What a theorist must do is construct a concept – given various theoretical purposes and some constraints of normal usage. But the construction of the concept must be relative to a set of problems and questions.[32]

In short, the theorist must identify certain aspects of basic usage, and then attempt to connect those aspects to the theoretical issue at hand. This is the approach that will be adopted here, and the two aspects of the term that I will identify are:

---

[31] Dworkin, Gerald. “Autonomy” *A Companion to Contemporary Political Philosophy*. Blackwell Publishing 1995 p. 359.

[32] Dworkin, 1995, p. 359

a) autonomy in the sense of the mental capacity to make decisions and form preferences; and
b) autonomy in the sense of being free, or at liberty to act on those decisions and preferences.

Before describing each of these in more detail, it is important to note that the point here is to understand autonomy in a morally satisfying way. If the larger project is to understand the moral value of privacy via its relationship with the moral value of autonomy, then we need to understand autonomy in a way that is sensitive to at least the more significant moral intuitions we have regarding the term.

Consider an example. If I were to consider whether Adolph Hitler was autonomous while composing *Mein Kampf* in prison, several questions arise. First, regarding autonomy in the sense of a mental capacity, should I consider someone like Hitler to be truly capable in this way? Was he mentally ill, or strongly influenced by society? To what extent do these things matter when making such an evaluation? Second, although he was able to compose notes for the work, was he really free to write in any significant sense? Was the fact that he was in prison something that affected his liberty regarding the work? Would the work have turned out differently? Should I not consider anyone in prison to be autonomous because they are physically confined?

Clearly there are several questions of these sorts that bleed into one another and are interdependent in some sense. Furthermore, in a case like this, there is an additional element of desert, or more specifically whether his autonomy is actually deserving of respect. Prior to this consideration, however, the idea is that in order to feel that the label of “autonomous” applies to some individual in a morally satisfying way (i.e., that his autonomy has been respected, not that it is deserving of respect), questions arise that involve both autonomy as a kind of mental capacity as well as autonomy in the sense of liberty. In other words, if my goal were to respect someone's autonomy, understanding what that would mean would typically lead me to consider both kinds of issues.

For the time being, then, I will be splitting the concept into what are (for the present purposes) the two most relevant components: autonomous capacity and liberty. To reiterate, what I mean by "autonomous capacity" is the mental ability to make decisions, to form preferences and values. It is the capacity, as it were, to make up one's mind. I mean this in only the most minimal sense so as to separate this concept from the question of whether a particular decision or preference should be labeled as autonomous. There is of course an independent debate over the proper standards of autonomous decisions and agency as well. I will address that debate at a later point.

The second component I will refer to as liberty. Here, I will be adopting a simplified version of the definition of liberty offered by several theorists, such as Rawls, Feinberg, and Benn. Theirs is a triadic notion of liberty, taking the form of "x is (is not) free from y to do (not do) z."[33] Hence, liberty is a state of affairs wherein one is free to perform (or not to perform) some action.[34]

Before continuing, it is important to note that I am not arguing that this is how "autonomy" should be understood in all contexts. Instead, and as will be made clearer below, the idea is to identify the aspects of the term that are both relevant to moral considerations (i.e., respect for autonomy) and to how the moral value of privacy can conceivably relate to those aspects. Recall from the previous chapter, that autonomy-based accounts of privacy are not properly understood as want-satisfaction accounts. Instead, they implicate autonomy in the sense of self-creative enterprises and the formation and pursuit of a conception of the good. For that reason, it is this more developed notion of autonomy that will be my focus. For the purposes of starting the discussion and distinguishing separate questions about autonomy, I will first treat it in a sort of simplified sense: the formation of preferences and values. I take these notions in turn to make up the grander notions of self-creation and an overall conception of the good.

---

[33] Kukathas, Chandran. "Liberty," *A Companion to Contemporary Political Philosophy*, Blackwell Publishing 1995 p. 535-536

[34] In this sense, the conception of autonomy being offered here is similar to Berlin's distinction between positive and negative liberty. The metaphorical captain can decide to guide his ship port or starboard, but in order to enact such a decision certain conditions must be in place.

This chapter will be divided up into several parts. First, I will discuss briefly three different questions one could ask about autonomy. The first is what constitutes agency, in the sense of to whom the concept of autonomy is applying. When we say that an agent is autonomous, how should we understand what an agent is? The second has to do with standards of autonomy. Are there necessary and sufficient conditions that determine if a particular agent is autonomous? And third, what does it mean to respect the autonomy of another person?

After laying out those questions, I will take a first pass at understanding how a relational notion of autonomy would answer them. And finally, I will take a second pass at characterizing relational autonomy, given a more nuanced understanding of these questions. I will also identify the notion of embeddedness as the main criterion of relational autonomy.

Three Questions Regarding Autonomy

Thus far I have provided a fairly simple notion of autonomous capacity, essentially the ability to form preferences and values. It would be quite natural, however, to ask a few follow up questions. First is the question of the self or agency. How should we understand the concept of the agent that underlies this concept of autonomy? In other words, we are saying that *someone* is autonomous. Is this someone properly understood as completely distinct from her surroundings, as in something “atomistic” and strictly bounded? Or do we think of agency as more like a process that is in fact defined and/or constituted by its interactions with certain things, such as other agents or cultural factors?

Second is the question of what constitutes an autonomous preference or agent. Is there a set of necessary and sufficient conditions that we can apply to a particular agent, or even a particular preference held by an agent, that qualifies her or it as autonomous? Such conditions are often referred to as standards of autonomy. To what extent do we factor in concerns like mental illness, drug addiction, or such things as cultural and interpersonal issues?

Third is the question of what constitutes respect for the autonomy of others. Given our understandings of autonomous agency and our standards of autonomy, what is required to respect someone else’s autonomy? Are there certain obligations that we have, certain ways that we are supposed to act towards them that constitute being respectful?

There are of course many other questions related to autonomy, such as its relationship to moral responsibility and accountability, whether or not it is really a marker for consideration in certain political issues, and so on. I have identified these three issues here because they are most pertinent to respect for privacy and because it is necessary to have a basic grasp of them in order to understand what distinguishes relational autonomy from other conceptions.

What "Relational" Means: A First Pass

"Relational" is first and foremost an umbrella term, denoting several different attempts in the philosophical literature to develop an account of autonomy that is sensitive to the various objections leveled against what is sometimes referred to as "the received view of autonomy."[35] As described by its detractors, the received view is based on the Western notion of rugged individualism, wherein agents are seen as completely independent and not necessarily embedded in or dependent upon a cultural or interpersonal context. To make better sense of this claim, it would be helpful to put this in terms of the three questions about autonomy outlined in the previous section and to understand how each of these questions may be approached in a relational way. In this section, I will describe various accounts of agency, standards of autonomy, and respect for autonomy; and then make some initial remarks about how they can be understood as relational. In the following section, I will dive a bit deeper into that question, and settle in a working definition of what makes a particular conception of autonomy a relational one regardless of which of the three questions it is trying to answer.

---

[35] For a more detailed discussion of relational autonomy, see the "Introduction" in *Relational Autonomy*, Stoljar & Mackenzie ed.'s Oxford Univ. Press 2000. The term "relational" most likely stems from the work of psychologists such as Carol Gilligan and Jean Baker Miller (among others) who were pioneers of what is referred to as "relational psychology." For an historical account of how this branch of psychology developed, see Robb, Christina *This Changes Everything: the Relational Revolution in Psychology*, Farrar, Straus & Giroux 2006.

Relational Agency

Regarding the question of agency, relational approaches see the self as either embedded in, constituted by, or caused by cultural and interpersonal factors. The received view rejects such interdependence in favor of an atomistic account of the self, wherein such connections are not ultimately considered to be constitutive of agency. The problem is that such atomistic views promote a kind of character ideal wherein individuals "are capable of leading self-sufficient, isolated, independent lives."[36] This ideal is based on a conception of agency that presupposes a great deal of independence from cultural and interpersonal influences. In turn, this ideal constitutes a prescription for promoting individual autonomy in a particular way. The fundamental objection, however, is a matter of which is the proper conception of agency. The claim being made is that the character ideal is based on an incorrect conception of agency, and hence the prescription that falls from it is incorrect as well.

Judith Jordan offers one account of a relational understanding of agency. Jordan describes a conception of agency that is essentially the very opposite of the received view, what she refers to as "relational being."

> To summarize, from a relational perspective, human beings are seen as experiencing a primary need for connection and essential emotional joining. This need is served by empathy, which in authentic relatedness, is characterized by mutuality. Further, in relationships one comes to experience: clarity about one's own experience and others; the capacity for creating meaningful action; an increased sense of vitality; and capacity for further connection.[37]

In Jordan's view, agents have a fundamental need for empathy with others. It is a need to understand how others function and feel, and to be understood in a similar way.

---

[36] Mackenzie and Stoljar, 2000, p. 6

[37] Jordan, Judith V. "A Relational Perspective," *Women's Growth in Diversity.* Jordan, J. ed. 1997 The Guilford Press NY, NY p. 20

Furthermore, the satisfaction of this need is what allows individuals to develop as persons. Jordan describes empathy in more detail:

> Empathy, the dynamic cognitive-affective process of joining with and understanding another's subjective experience, is central to this perspective. Mutual empathy, characterized by the flow of empathic attunement between people, alters the traditional boundaries between subject and object and experientially alters the sense of separate self in a profound way. In true empathic exchange, each is both object and subject, mutually engaged in affecting and being affected, knowing and being known. In interpersonal language, in a mutually empathic relationship, each individual allows and assists the other in coming more fully into clarity, reality, and relatedness; each shapes the other.[38]

Here, Jordan is rejecting the notion of the self as a singular entity, and even the reification of the self. That is, on the relational being conception, there is an intentional ambiguity between the self as a thing and the self as a process. Agency here is the process of interacting with others in order to better understand oneself. Ideally, this process is marked by mutual empathy, wherein each understands the other in such a way that the traditional lines between self and other become less distinct.

Since this account is a psychological one, it would appear that the conception of agency being offered is a contingent one, and not metaphysical. That is, the claims being made here regard the nature of psychological interaction and what the proper conception of self is given that interaction. I therefore take this not as a metaphysical claim in the sense of the *necessary* constitution of agency. However, that said, Jordan is concerned with the distinction between self as process and self as thing. One of the primary motivations behind this account is the idea that the reification of the self is ultimately a misleading tendency *because* it downplays the importance of the effects of social interaction. Jordan even goes so far as to cite the influence of Newtonian physics.

---

[38] Jordan, 1997, p. 15

> The prevailing concept of self is modeled on the now outdated Newtonian physics, a paradigm which posited separate objects possessing clear identities whose interactions were secondary to their atomistic and bounded structures.[39]

Instead the interactions are intended to be primary, while there should be no talk of bounded structure at all. Hence, it is not completely clear if the kind of atomism to which she refers constitutes an ineffective and misleading heuristic when thinking about the self, or if something deeper and necessary is intended. But again, because the analysis is primarily psychological and therefore focuses on contingent factors, I take the general claim regarding agency to be contingent as well.

While this account of agency does not directly address the second question above (i.e., what constitutes an autonomous decision or preference), it does speak to the issue of respecting autonomy by providing an account of human flourishing. To flourish, according to this conception, is to experience mutual empathy (mutuality) with others. The idea is that individuals work towards varying states of mutuality and that this process is not a matter of achieving two different goals (i.e., the flourishing of the self as well as maintaining distinct relationships with others), but instead it is the natural process of relational being. It is the process itself that constitutes who we are and what we do.

Hence, flourishing means achieving or being allowed to achieve varying states of mutuality. It means understanding that to be autonomous is first to acknowledge this kind of interdependence and second, as strange as it may sound, to support autonomous capacity through a particular kind of dependence. It means supporting certain interpersonal processes that in turn facilitate self-knowledge and understanding.

---

[39] Jordan, Judith V. “Some Misconceptions and Reconceptions of a Relational Approach ” *Women’s Growth in Diversity.* Jordan, J. ed. 1997 The Guilford Press NY, NY p. 30

At this point, it is important to note that what is meant by autonomy by those who subscribe to the mutual empathy account of agency is fundamentally different than autonomy in the received sense. And to be clear, I am offering Jordan's view as an example of a relational view. Recall from the beginning of this chapter "autonomy" simpliciter was divided between the notion of the mental capacity to form preferences and make decisions and liberty, with my focus being on the former. The point behind defining autonomous capacity so broadly is first to capture the root idea behind its common usage, but at the same time not to make any assumptions as to the nature of how it works. On the relational being conception of agency, autonomy as a mental capacity is a process that is developed and nurtured through interactions with others. It is the very antithesis of the inner citadel.[40]

Relational Standards of Autonomy

The second question mentioned above is the question of what constitutes an autonomous preference or decision. Are agents supposed to be aware of cultural and interpersonal influences in order for their choices to count as autonomous? If so, to what extent? This is where the debate between procedural and substantive accounts of autonomy comes into play.

Procedural accounts of autonomy attempt to establish a set of standards for the degree of internal reflection to which agents should subject themselves in order for their choices to count as autonomous. Such accounts are often referred to as "internalist" because they focus on the internal psychological states of the agent, instead of any external conditions she may be experiencing.[41] These accounts are also content-neutral in the sense that an autonomous preference or decision need not have a particular subject matter in order to count as autonomous, but must have been subjected to a certain degree of "critical reflection."[42] Finally, procedural accounts of autonomy are value neutral in the sense that understanding how a preference was generated is not supposed to affect the content of any value laden judgments. As Paul Benson puts it, "Understanding one's freedom as an agent is not supposed to reveal anything about how it is good to live or right to act."[43]

---

[40] See The Inner Citadel: Essays on Individual Autonomy, Christman, John,(ed.) for an oft cited contemporary source of the notion of an inner citadel as the notion of a bounded conception of agency. Also the phrase is linked to the thought of Marcus Aurelius by Pierre Hadot.

One of the better known examples of a procedural approach is Christman's counterfactual standard. An action or decision is considered to be autonomous if it is the case that, had the agent critically reflected upon the action, she would have chosen it.[44] The idea is fairly straightforward. Would she have still chosen x, if she had critically reflected on x? If so, then by this standard x is chosen autonomously. Another part of Christman's account is that the agent should be able to endorse the process by which the preference or value was formed.[45] So it is not just a matter of reflecting upon the preference itself, but accepting how the preference came about.

Consider an oft cited example from Benson to help clarify this second part of the standard.[46] In the 1944 movie Gaslight, Ingrid Bergman plays a woman (Paula Alquist) being victimized by a con man (Gregory Anton, played by Charles Boyer). Anton is attempting to undermine Alquist's sanity and confidence in her judgment in an effort to

---

[41] Oshana, Marina "Personal Autonomy and Society," *Journal of Social Philosophy*, vol 29 No. 1, Spring 1998 81 - 102

[42] Oshana, 1998, p. 14

[43] Benson, Paul "Feminist Second Thoughts about Free Agency," *Hypatia*, Vol. 5, No. 3 (Autumn, 1990), p. 49

[44] Christman develops his theory at multiple points. See 1991and 2004 especially.

[45] Christman 1991

[46] Benson, 1994

"procure" estate jewelry left to Alquist by her recently deceased aunt. Benson describes the villain's tactics as follows:

> Boyer's scheme is to reduce Bergman to a state of such apparent confusion and disorientation that she will be unlikely to realize what he is up to. He isolates her from other people, not by force but through suggestions that seem innocent enough to Bergman. He makes her think that she is losing things, that she cannot remember things she recently has done, that she is subject to hallucinations.[47]

This example has received quite a bit of attention in the literature on autonomy, and I will have more to say about it later, but for present purposes the question is whether Alquist should be considered autonomous, and why or why not. Keep in mind that Alquist is still very much capable of considering her preferences and she is most definitely not ambivalent about what is going on around her. Furthermore, she is still capable of ranking her concerns and goals, essentially prioritizing them.

On Christman's view, she would not qualify as autonomous because although she may even endorse some of her beliefs in the sense of reflecting on them and consciously accepting them, she would not endorse the way in which they were formed (due to the devious intentions and actions of Mr. Anton). In the parlance, they are authentic but only on a superficial level.

The alternative to procedural standards of autonomy is substantive standards. Natalie Stoljar advocates an example of this kind of standard with what she refers to as a strong substantive view.[48] An individual who is autonomous in the strong substantive sense is able to criticize a given course of action by what Stoljar refers to as "relevant normative standards."[49] As will be discussed in more detail in the following chapter,

---

[47] Benson, 1994, p. 655

[48] Stoljar, Natalie. "Autonomy and the Feminist Intuition," *Relational Autonomy*, Stoljar & Mackenzie ed.'s Oxford Univ. Press 2000.

[49] Stoljar, 2000, p. 107

Stoljar focuses her analysis on the oppressive norms of femininity. Essentially, these are the damaging norms or cultural gender-based biases that hinder and limit women's opportunities for self-fulfillment and flourishing. Women who have accepted such norms, Stoljar argues, are not autonomous because they have accepted a false norm and are essentially incapable of seeing it as such.

This kind of strong substantive view is sometimes distinguished from a weaker version, which has been advocated by both Benson and Mackenzie.[50] In the weaker substantive view, there is intended to be far less constraint on the actual content of the values held by agents, save for the idea that agents must demonstrate certain core values or capacities, such as self-esteem, self-trust, and self-respect. Here the idea is that agents must possess the core capacity to hold themselves accountable to relevant normative standards, and worthy of them as well.[51]

These substantive views are also sometimes labeled as normative competence views; however, sometimes this phrase is reserved just for the weaker version. The idea behind normative competence itself is just as has been described, that agents must have the ability to identify relevant or genuinely applicable normative concepts when evaluating and possibly endorsing some value. So in the case of both weak and strong substantive views, the content of the preferences matters, hence the label substantive; and there must be some additional capacity on the part of the agent to identify relevant (or "correct" on Stoljar's view) standards, hence the notion of normative competence.

As mentioned previously, it is not my goal here to provide a complete taxonomy of standards of autonomy, but instead my intent is to lay out the core constitutive concepts. Also, it is important to note that the ways in which various philosophers have presented these views has varied over the last twenty years of scholarship, and is still very much evolving. Hence, pinning down precise versions of each is not generally feasible given the level of variation in which they have been understood by scholars.

---

[50] See Benson's "Feminist Intuitions and the Normative Substance of Autonomy" in Taylor 2005, as well as Mackenzie 2008

[51] Mackenzie p.527

The next question is in what sense are these different standards to be considered relational, if at all? In this section I will make some initial observations about this question, and will return to it in the following section in more detail. Recall that procedural standards are sometimes referred to as “internalist.” That is, they focus on the internal psychological states of the agent in question. So on a simple reading of procedural accounts we might lean towards saying that they are not relational. However, a richer reading of Christman’s procedural account reveals some ways in which it can capture cultural and interpersonal factors that often underlie relational views.

As Mackenzie points out, Christman’s view includes what can be considered relational competency conditions, such as care, empathy, intimacy and social cooperation.[52] These are certain competencies that agents must possess and exercise in order to qualify as autonomous; and they are of course outward looking in the sense that they do not simply disregard the agent’s relations with other individuals or certain norms.

In addition to these competencies, Mackenzie points out that some of the authenticity requirements in Christman’s account can be understood as relational:

> He argues that any adequate account of authenticity must recognize that not only is a person’s practical identity shaped by complex, intersecting social determinants and constituted in the context of interpersonal relationships; the process of reflection is similarly shaped by these factors.[53]

---

[52] Mackenzie, 2008, p. 520

[53] Mackenzie, 2008, p. 520

Authenticity requirements are basically the notion that the agent must endorse the values and preferences in question. So in this sense we can see how a procedural theory can make space for relational concerns about how authenticity is not a strictly internal notion.

While these considerations offer some room for viewing procedural accounts as relational, recall that they do not possess elements of normative competence nor substantive constraints. These two components are what relational and feminist theorists most often focus on when attempting to construct standards of autonomy that are sensitive to the experience of oppressed individuals. They are also typically identified as the missing elements when those same theorists worry about the ability of procedural accounts to capture the ways in which oppressive norms can impact autonomy. I will return to this issue in the next section, but for the time being, it is sufficient to point out that there is a great deal of debate over whether normative competence and/or substantive constraints are actually necessary in order to capture the ways in which oppressive norms can affect autonomy and whether they are a necessary part of any relational conception of autonomy.

Relational Accounts of Respect for Autonomy

The third issue mentioned at the beginning of this section is the question of how to make sense of respect for autonomy. Given the nature of agency and what constitutes an autonomous preference or decision, how does one go about respecting autonomy? What does it mean to say that we value and respect the personal autonomy of another individual? What are the actions that we typically take? What are the sorts of things that we should do to maintain the autonomous capacity of individuals?

Consider, as an example of how these questions come into play, the issue of informed consent in bioethics. On the received view of autonomy, the idea is to enable the individual to make an autonomous decision by providing as much information as is practical and relevant to the decision at hand. For instance, if the decision is whether or not to undergo some form of medical treatment, the patient presumably needs to know such things as the possible side effects, the anticipated chances of success and so forth. However, as one critic points out, this approach tends to view both patient and medical practitioner as independent contracting agents.[54] That is, it treats both parties as entering into a contract that is primarily structured and motivated by the need for information unconnected to relational concerns. There is no talk of the importance of familial relations, cultural influences upon either party, and how those factors influence the decision whether to undergo treatment. In short, in order to respect a patient's autonomy in such cases, one must acknowledge that the decision making process will incorporate such factors in various ways, and hence not addressing such issues is tantamount, on the relational view, to disrespecting autonomy.

Mackenzie also points out a very important concern about respect for autonomy, one that falls out of the debate over standards of autonomy. If our standard of autonomy attempts to incorporate relational factors in some way, then it would seem to follow that those factors would also inform our understanding of what it means to respect autonomy. For example, in the case of Mackenzie's normative competence view (i.e., her weak substantive view), the ideas of self-esteem, self-trust and self-respect are of central importance. They make up the core notions of what it means for an agent to be minimally functional and competent to handle the myriad ways in which she might interact with relational factors, as well as form and pursue her conception of the good.

---

[54] See Donchin, Anne "Autonomy and Interdependence: Quandaries in Genetic Decision Making" in Mackenzie and Stoljar 2000

So how do we incorporate these notions in our account of respect for autonomy such that it too is genuinely relational? Mackenzie offers three points in this regard. First, we must recognize the agent's humanity. That is, we must treat agents as someone with a specific set of beliefs and preferences, someone to whom certain things matter more than others.[55] Second, we must try to understand the agent's subjective perspective. While discussing the case of a hypothetical hospital patient suffering from a potentially terminal illness, Mackenzie states that if she "feels that her carers do not understand her and have made no effort to do so then she is likely to feel that the medical team is just being coercive if it continues to discuss treatment options with her."[56]

And third, there is the need to promote the actual competencies of the agent. Again this would be understood in terms of what the particular conception or standard of autonomy has to say about normative competence and/or procedural competencies. But in the case of Mackenzie's conception of autonomy, the obligation is a matter of actually improving an agent's sense of self-value:

> [B]ecause [the agent's] perspective is shaped by attitudes toward herself that undermine her flourishing, respect for her autonomy involves an obligation not just to understand but to try to shift her perspective and to promote her capacities for autonomy. This involves trying to find ways to change [the agent's] attitudes toward herself—to try to counter her sense of personal worthlessness, to promote her sense of self-respect, to assist her to find reasons for living and to envisage a possible future in which she would find her life meaningful.[57]

I will not go into too much more detail at this stage, but this idea that Mackenzie is describing will be a crucial one when I connect respect for relational autonomy to respect for privacy in later chapters. What she is describing here is an articulated obligation to enhance the sense of self-respect, self-esteem and self-trust of agents in cases where we have reason to suspect that such capacities are diminished. I will return to this issue at a later point.

---

[55] Mackenzie, 2008, p. 528

[56] Mackenzie, 2008, p. 528

[57] Mackenzie, 2008, p. 528

Now I would like to summarize this first pass at understanding what it means for an understanding of autonomy to be relational. As stated, there are three main questions concerning relational autonomy. First, there is the debate concerning the proper conception of agency. To what extent are agents embedded or dependent upon cultural and interpersonal factors? Second, what constitutes an autonomous decision or preference? Are there procedural or substantive requirements for autonomy? Third, how does one go about respecting autonomy? What are the sorts of things that need to be done to enable or provide for autonomous decision making? Though conceptually distinct, these issues are clearly cognate and do not pull apart cleanly.

Given that there are a number of different questions about autonomy as I have described it here, is there a set of necessary and sufficient conditions that distinguish relational accounts from more traditional versions of autonomy? As Christman points out, the notion that cultural and interpersonal factors (what I refer to as “embeddedness”) play a role in autonomy seems to function as a “conceptually necessary” part of relational accounts, regardless of whether they are addressing the issue of the nature of the self, what constitutes an autonomous decision, or the proper conception of respect for autonomy.[58]

What should be clear from the previous discussion of the three core questions is that embeddedness is not necessarily the notion that individuals are completely dependent upon cultural and interpersonal factors in terms of what decisions they make and what preferences they develop. It is instead, at a minimum anyway, the notion that individuals respond to and in some cases even desire such influences. Embeddedness is the concession that these kinds of factors play a legitimate role in how we make decisions and develop preferences, not necessarily that these factors completely control those decisions and preferences.

---

[58] Christman, 2004

So for now, what we have is a more minimalist version wherein embeddedness is a conceptual consideration throughout, and wherein it directly informs attempts to formulate answers to particular issues or questions within autonomy. There is however, a separate question of whether some element of substantive or normative competence elements are genuinely a part of what we take to be relational conceptions of autonomy. This is the issue I turn to in the second pass at what relational means.

What "Relational" Means: A Second Pass

I ended the previous section asking the question, to what extent are substantive and normative competence elements a necessary part of relational conceptions of autonomy. Another way to think about that question, as I have framed things here, is to ask what exactly we mean by embeddedness when we say it is conceptually necessary for relational autonomy. I take embeddedness to be an obvious fact of the human condition, at least on a simplified reading. It is hard to argue against the claim that individuals are to varying degrees influenced by interpersonal and cultural factors. So it would seem natural to have embeddedness as a consideration when formulating any particular conception of autonomy. The question, however, is how far do we take this consideration; and how do we apply it to gain a better understanding of what constitutes *relational* autonomy? I will, for the most part, adopt Christman's stance on this issue:

> It is one thing to claim that social conditions that enable us to develop and maintain the powers of authentic choice and which protect the ongoing interpersonal and social relationships that define ourselves are all part of the background requirements for the development of autonomy…It is another thing, however – and a more dangerous and ultimately problematic move, I have argued – to claim that being autonomous *means* standing in proper relations to surrounding others and within social practices and institutions.[59]

In short, Christman is arguing that embeddedness is part of the explanation for how we develop the capacity for autonomy, but it is not part of what *it is* to be autonomous. Here, Christman is essentially referring to strong substantive views of autonomy. To be clear, it is not Christman's intent here to pick out a standard for when a particular conception of autonomy qualifies as relational; he is instead pointing out a difficulty associated with conceptions that are both substantive and relational. Setting aside the question of whether his objection to substantive views is truly effective or even fatal, I take his assertion to constitute an effective way of rounding out what it means to take embeddedness as conceptually necessary for relational approaches to autonomy.

To explain what I mean by this, it is first necessary to understand in more detail the target of Christman's criticism. He is aiming this objection towards strong substantive accounts, and specifically the one presented by Marina Oshana.[60] Oshana's view incorporates procedural or internalist competency and authenticity conditions, but also takes the further step of insisting that agents stand in particular relations with others in order to qualify as autonomous. She outlines four criteria for this part of her standard.

---

[59] Christman, 2004

[60] Oshana, 1998

First, individuals should be able to defend themselves against this kind of psychological or physical assault when necessary. One should be able to counter, in some way, actions of others that directly threaten psychological or physical capabilities. Second, individuals should be able to defend themselves against attacks on their “civil and economic rights.”[61] Third, individuals should not be required to take responsibility for another person’s “needs, expectations or failings” unless she has agreed to do so or it can be reasonably expected of her to do so as per a matter of the relationship or performing some function.[62] And fourth, agents must be able to pursue goods and values that are independent from those of the individuals who otherwise have authority over her.

It is not necessary to go into too much detail on these criteria, but just to note that they each pick out some form of external relationship an agent may or may not have with another person, group of people, or even an institution or government. It is this kind of constraint to which Christman is objecting to as perfectionist. The worry is that incorporating these kinds of requirements into a standard of autonomy actually goes against what is the spirit of relational conceptions in the first place. In order for an agent to develop or maintain such relations, she would in some cases have to demonstrate a level of independence from social factors that flies in the face of the original motivation behind the relational approach.

I think that Christman is right to point this out, and that it constitutes a genuine difficulty for substantive or perfectionist views of relational autonomy. But again, he is not concerned so much with what specifically distinguishes a relational conception of autonomy from other conceptions. But what falls out of this is that it is not enough to classify relational views as any conception of autonomy that takes embeddedness as conceptually necessary. Instead, we must add the criterion that relational conceptions do not posit as ideals any sort of individualist views of agency and respect for autonomy. What is more, we must be careful to determine whether the inclusion of certain substantive constraints on standards of autonomy do not ultimately lead to a promotion of individualist ideals, even if just inadvertently. This is the possibility alleged by Christman regarding Oshana's view.

---

[61] Oshana, 1998, p. 94

[62] Oshana, 1998, p. 94

To explain, suppose I were to advance a conception of respect for autonomy wherein I acknowledge embeddedness as a state of affairs, but then claim that the goal of respect for autonomy should be to promote greater independence from interpersonal and cultural factors. Suppose that I also claim that it is only through an exhaustive study of embeddedness that we can really understand respect for autonomy. In at least that sense, embeddedness is a part of my view throughout the conception. Everything that I might say about respect for autonomy at that point would center on my understanding of embeddedness (and how to minimize it).

On the minimal view of relational autonomy, this conception of respect for autonomy might seem to qualify as relational. However, on a normative level this would seem to contradict the explicitly stated purpose of the approach altogether. The various symbolic and metaphysical critiques of the received view of autonomy, which are in large part the primary motivation for the approach in the first place, are typically based on the claim that individualism is either based on an incorrect view of agency, or a view that is at the very least not inclusive of other possibilities. Hence to put individualism forth as a goal exclusive of other views simply is not consistent with the relational approach.

This hypothetical conception of respect for autonomy has a normative component which states that individualism, or a departure from embeddedness, is essentially a goal that should drive our understanding of respect for autonomy. When we respect autonomy, we respect the need to gain independence from interpersonal and cultural influences on the individual's capacity to make decisions and form preferences. We would not in this conception understand any positive obligations (based on the principle of respect for autonomy) to respect the need to preserve interpersonal or cultural relations that are premised on dependency as a positive thing. In this sense at least we would not be in any significant way acknowledging alternative conceptions of agency.

Based on the minimal account of what constitutes a relational view of autonomy that I am here proposing, this conception of respect for autonomy would not qualify as relational because it promotes a particular conception of agency to the exclusion of others. Regarding Oshana's substantive view, it is often cited as a paradigmatic example of a relational conception of autonomy. So to be clear, I am not here claiming that her view is nonrelational. Her account does not posit, as an ideal, the sort of independence from relational factors as the hypothetical view described previously. However, Christman's objection to Oshana's view raises the possibility that she may be committed to nonrelational elements. That said, it is a separate question as to whether Christman's worry about perfectionism actually sticks, and if so, whether that would create an inherent contradiction in Oshana's view. Those questions are not ones I will take up in this dissertation.

To summarize, the criteria I am proposing for a conception of autonomy to count as relational are first that embeddedness is treated as conceptually necessary, and second that it not posit an individualistic view of agency, standard of autonomy, or conception of respect for autonomy *as an ideal*. It is important to provide a more inclusive account of the relational approach to autonomy so that a wider variety of views about agency, standards of or respect for autonomy can find space under the umbrella as it were. More specifically, it is important to be able to make sense of an account of relational autonomy in general that can accommodate procedural conceptions of standards of autonomy and/or respect for autonomy.

On the other hand, however, it is still important to derive an understanding of relational autonomy that is still distinct enough so that there is something concrete to which we can connect our understanding of respect for privacy. The ultimate goal, remember, is to understand how relational autonomy can help us to track the intuitions that drive the functionalist views of the value of privacy; and on a deeper level, understand how debates about relational autonomy can ultimately inform discussions about the value of privacy. To achieve that, it will be necessary to have on hand an account of relational autonomy that is detailed enough to motivate concerns about what it actually means for an account of the value of privacy to be committed to a relational view of autonomy.

In the following chapter, I will discuss the debate between procedural and substantive standards of autonomy in more detail. And now that we have a stronger sense of what really constitutes a relational approach to autonomy, I will also discuss how the proceduralist/substantive debate itself becomes inherently problematic assuming the accounts in question meet the minimal requirements of the relational approach discussed above.

CHAPTER 4

AUTONOMY: RETHINKING THE QUESTION

Introduction

In the previous chapter, I devoted a fair amount of discussion to standards of autonomy. In this chapter, I will conclude my discussion of autonomy and provide two reasons for why it would be beneficial to shift our focus from standards of autonomy to respect for autonomy.

The first reason is due to the assumption of embeddedness and is broken down into two parts. The first part is that there is a difficulty in reconciling procedural standards with considered intuitions about the internalization of oppressive norms. The assumption of embeddedness adds a layer of complexity to our understanding of how norms are internalized. This added complexity makes it more difficult for procedural accounts to capture all of the ways in which agents might internalize a norm. This in turn makes it more difficult to reconcile procedural accounts with intuitions about the autonomy of oppressed individuals. In other words, procedural accounts offer no reliable way to connect the acceptance of an oppressive norm and nonautonomy. The second part has to do with the plausibility of agents qualifying as autonomous under substantive accounts. Substantive constraints on the content of autonomous beliefs and preferences sometimes require a level of independence from interpersonal and cultural influences that contradicts the assumption of embeddedness itself.

The second reason is that there is a need for our understanding of autonomy to cohere with how it is used in other theoretical contexts. As Christman has argued, adopting a particular stance on the issue of standards of autonomy implies certain commitments in how autonomy can be employed in related debates. At first blush, this may seem like an unfair and unnecessary complication of the standards debate; however, an analysis of Stoljar's argument against procedural accounts will reveal it to be an important consideration. In response to this call for greater coherence with cognate issues involving autonomy and in response to the difficulties associated with the assumption of embeddedness mentioned above, I will suggest a respect-based approach that, as I will argue, is better able to handle these complications.

I will start this chapter first with a presentation of another standard of autonomy, which is the skills-based account provided by Diana Tietjens Meyers. This standard, as I will show, is primarily procedural, but as Paul Benson points out, has certain very minimal substantive elements. Second, I will work through Natalie Stoljar's strong substantive view and her argument against procedural accounts based on the importance of what she calls the feminist intuition. Her argument is important because it very effectively shows procedural accounts to be content neutral not just as a matter of stipulation, but as a *necessary* consequence of the assumption of embeddedness.

Third, I will discuss in more detail Christman's idea that resolving the procedural/substantive debate may very well rely on how we think of autonomy operating in other theoretical contexts. The "other context" he has in mind is the debate between perfectionism and political liberalism. And finally, I will try to bring these points together to show that respect for privacy may be a suitable candidate for dealing with these difficulties and complications. Ultimately, what I hope to show is that, given the assumption of embeddedness, autonomy is not something that is best understood as an isolated concept. Its relationship with political considerations and other issues such as the value of privacy is one of reciprocal illumination. How we understand the relationships between them can help us to understand each individually.

Before beginning, it would be helpful to lay out more specifically what I have in mind in terms of the shift from standards of autonomy to accounts of respect for autonomy that are informed by those standards. As discussed in the previous chapter, a standard of autonomy is a set of necessary and sufficient conditions that must obtain for an agent or a particular value held by an agent to be considered autonomous. In the case of procedural standards, there are sets of authenticity and competency conditions that need to be obtained. Many of these conditions, however, are essentially a matter of degree. For example, Christman makes use of certain relational competency conditions such as care, empathy, intimacy and social cooperation. Each of these is essentially a comparative notion in the sense that an agent can demonstrate varying degrees of caring for another, empathizing with another, and so on.

In this sense, procedural standards of autonomy offer both an account of what constitutes autonomy, and what counts as being autonomous in the sense of qualifying for a standard. However, it is possible to separate these two claims once we acknowledge that they are a matter of degree. In other words, there is the claim that autonomy is constituted by an individual exercising certain competencies and demonstrating authenticity by a degree of self-evaluation. There is also the claim that once having done so, that agent qualifies as autonomous. Furthermore, when we bestow the label of autonomous, we attach certain moral significance to that agent. As will be discussed in more detail below, being autonomous is a "marker" for participation and consideration in certain political issues, as well as qualification for earning a certain kind of respect from other agents.[63] In short, we can consider what generally makes up autonomy, and we can separately consider when enough of those things are present (or are present to a high enough degree) to trigger certain moral intuitions.

The same general point can be made of substantive standards of autonomy. First, substantive views often incorporate elements of procedural accounts.[64] Second, they often require that agents have the ability to identify or incorporate relevant normative elements to varying degrees. Self-respect, self-trust and the like are of course matters of degree, as well as the extent to which an agent may or may not enjoy certain sociorelational statuses such as being free from certain types of coercion.

What I will argue in this chapter is that the question of what constitutes a proper standard of autonomy may not be resolvable as an isolated issue; at least not once we incorporate other factors such as embeddedness. Instead, looking at these issues from the point of view of respect for autonomy may in fact be a better way to determine to what extent our intuitions about whether one qualifies as autonomous cohere with what we take to be constitutive of autonomy.

---

[63] Christman, "Procedural Autonomy and Liberal Legitimacy" from Taylor 2005.

[64] Marina Oshana's view would be an example of this. See Oshana (1998) for her view that incorporates both procedural competency and authenticity criteria as well as strong substantive criteria involving socio-relational conditions.

Meyer’s Procedural Standard of Autonomy

Meyers’ account identifies a list of skills that individuals must possess in order for their choices to qualify as autonomous.[65] Examples of such capabilities include: introspective skills (having a sense of how well one understands oneself), communication skills, memory skills, imaginative skills (the ability to see oneself otherwise and thereby develop goals and aspirations), and others. These skills, collectively at least, are sensitive to the sorts of concerns responded to by relational approaches. That is, if agents are to some degree embedded in and dependent upon cultural and interpersonal factors, then an autonomous individual must possess the skills necessary to identify those influences and be able to form preferences while cognizant of them. Otherwise she is merely pushed and pulled by these influences and so her choices are not autonomous.

As an example, consider whether Carol Gilligan, one of the early pioneers of relational psychology, was autonomous in her professional work. Gilligan was trained (in graduate school and in her earlier professional life) in an environment where women were not considered as equals to men, both intellectually and as a matter of what was expected in terms of professional accomplishments.[66] The bias extended even to how men were treated differently than women as subjects of psychological study. Robb describes how Gilligan, and those who were interested in her work (either professionally or casually), would work in ways outside of the norms of academia.

> They were friends who’d been neighbors and mothers together in a suburb where their kids came home from school for lunch every day. Now they sat around Carol’s kitchen table and brainstormed a way of really listening that could also pass as psychological research.[67]

---

[65] Meyers, 2001, p. 741

[66] Again, see Robb 2006.

[67] Robb, 2006, p. 21

In short, Gilligan and her coworkers did the work in any way that they could, and were willing to collaborate with one another in order to develop questions that fell outside the sort of paradigm of male-dominated psychological theory.

Looking at Gilligan's situation, it is possible to see how some of the skills enumerated by Meyers were applied. For instance, she describes communication skills that "enable individuals to get the benefit of others' perceptions, background knowledge, insights, advice and support."[68] The kitchen table research sessions described by Robb offer a good example of this sort of process. The skills that are perhaps most relevant to this example, however, are interpersonal skills. Meyers describes these in the following way:

> [I]nterpersonal skills that enable individuals to join forces to challenge and change cultural regimes that pathologize or marginalize their priorities and projects and that deprive them of the discursive means to represent themselves to themselves and to others as flourishing, self-respecting, valuable individuals.[69]

Clearly, the ability of Gilligan and her peers to have these discussions and to build upon them allowed them to advance their particular vision of psychological development and moral deliberation. Hence, in the sense of their professional lives at least, having and applying such interpersonal skills contributed to their autonomy in Meyers' sense.

So one might ask, in what sense was Gilligan autonomous? Most notable is the presence of certain mental capacities. Faced with various discriminatory practices, Gilligan managed to advance her work through a number of different means. In short, she was able to apply the sorts of skills that Meyers describes. She was able to imagine or see herself as pursuing research projects independently of male mentors. She was able to communicate with others who shared similar interests and concerns, and (perhaps most importantly) she was able to apply skills of introspection to the effect that the vision of the future presented to her by male professors and colleagues was something she knew she did not want. In these ways, despite the oppression, she was autonomous in the sense of being able to exercise certain mental capacities.

---

[68] Meyers, p. 741
[69] Meyers, p. 742

While I am not here arguing in favor of Meyers' approach as a standard of autonomy, the above example points out one of its attractive qualities as a procedural account. It is in such cases as the one just described that we feel a strong intuition to the effect that women such as Gilligan epitomize autonomy. Despite their condition, as in their lack of certain liberties, they are able to choose goals independently of the influences that they face. The more important thing to note, however, is that Meyers' approach is capable of acknowledging the embeddedness of individuals. Individuals such as Gilligan can be seen as being dependent upon their interpersonal relationships and cultural conditions (i.e., that they are women in a male dominated profession and society, but with a network of friends upon whom they can rely for support and discussion), but they can also be seen as applying certain skill sets to function autonomously within those conditions.

While Meyers' account is primarily procedural, Paul Benson has made the claim that the skills she enumerates could be seen as importing substantive elements.

> The agentic skills that she describes seem to import specific values into the account. For instance, Meyers defines the skills of "self-nurturing" in relation to the value of self-worth. She says that these skills enable agents "to appreciate the overall worthiness of their self-portraits and self-narratives" and to "sustain their self-respect." … Therefore, although Meyers' approach eschews direct restrictions on what autonomous agents can prefer or value, it carries normative content indirectly, through the values subsumed in its descriptions of autonomy competencies.[70]

---

[70] Benson, 2005, p. 135.

In short, Meyers' view fits cleanly in neither the procedural category nor the weak normative competence category. The skills she describes are difficult to understand without importing in some form the ideas of self-worth and self-respect. As mentioned previously, these ideas are central to the weak normative competence views presented by Mackenzie. The other thing to note is that this puts Meyers' view in the middle of the spectrum of standards in at least one sense. If we view all of the accounts discussed thus far in terms of the extent to which they entail constraints on the content of agents' values, with procedural accounts being on one end and strong substantive views being on the other, then Meyers' and Mackenzie's views would seem to occupy a middle ground.[71] For reasons I will provide later in this chapter, this middle ground constitutes the conceptions of autonomy (albeit in the form of respect for autonomy) that I will focus on in subsequent chapters when considering relational autonomy's connection with the value of privacy.

Stoljar's Argument for a Strong Normative Competence Theory of Autonomy

In this section, I will discuss Natalie Stoljar's argument against procedural autonomy as well as two difficulties faced by the theory of strong normative competence that she proposes. I use Stoljar's discussion to show that there are certain difficulties resolving the debate between procedural and substantive views by appealing to autonomy in and of itself, if we also adopt a relational approach to autonomy and thereby take embeddedness as conceptually necessary.

---

[71] Benson, 2005, p. 136

The feminist intuition, as Stoljar describes it, states that preferences or choices based on the norms of femininity are not autonomous. Here “norms of femininity” refer to the sorts of cultural practices and beliefs that are specifically oppressive and damaging to the well-being and life prospects of women. Basic examples include the notion that a woman’s social standing or moral worth is contingent upon her making the “right” reproductive decisions and the insistence that a woman’s place is in the home. In “Autonomy and the Feminist Intuition” Stoljar presents an argument that, if one wants to maintain the feminist intuition, then one must adopt a strongly substantive view of autonomy, what she refers to as a strong normative competence view.

Focusing on norms involving reproductive practices, Stoljar considers the rationales of a group of women who took contraceptive risks. Utilizing a study performed by Kristin Luker in the early 1970’s, Stoljar applies several procedural standards to these rationales and argues that, despite their having internalized (i.e., accepted as true) certain norms regarding contraception, the subjects qualify as autonomous according to those standards. Yet, she thinks, the “feminist intuition” tells us that the choices of many of these women to risk unwanted pregnancy were not autonomous. From this it follows that some other standard of autonomy is needed to preserve the feminist intuition. To that end, Stoljar claims it is necessary to adopt what she calls a strong normative competence theory of autonomy. She borrows a view based on the earlier work of Paul Benson.[72] He claims that the rejection (or at least nonacceptance) of false or irrelevant norms is a necessary condition for autonomy. The idea is that, in cases where a course of action, a preference, or a decision is based on a norm that either does not properly apply to the situation or is simply false, the autonomous capacity of the agent is compromised.

---

[72] To be clear on this point, Stoljar states that her version of a strong substantive view of autonomy relies heavily on the earlier works of Benson, primarily “Freedom and Value” (1987) and “Feminist Second Thoughts about Free Agency” (1990). However, in “Feminist Intuitions and the Normative Substance of Autonomy” (2005) Benson advocates a weak substantive view similar to the view of Mackenzie described in the previous chapter. Interestingly enough, he argues for this view while responding to Stoljar’s “Autonomy and the Feminist Intuition” which is the work I am discussing here.

A discussion of three of the procedural views canvassed by Stoljar will show the general argument *form* she uses to critique proceduralism. The first standard she rejects is offered by Christman. He has developed what he calls "the counterfactual condition." In order for a decision to be autonomous on this standard it must be the case that had one critically reflected upon the decision, she would have consented to it and the way in which the decision was formed. This is essentially the notion of authenticity discussed in the previous chapter. Certain habits individuals have, while not really given much thought, are still considered autonomous actions or decisions so long as they would have been consented to had such deliberation occurred. The problem with this approach, Stoljar argues, is that the individuals in question have already internalized certain norms. They have accepted them as true; hence their presence in the belief set of the subject in effect "blocks the capacity of the agent to resist the development of preferences based on the norm."[73] For instance, if a woman accepts the notion that her worth is tied to her ability to have children, then the pressure to take contraceptive risk is there even if she evaluates the desire to take the risk in the first place because she has already accepted the norm as true. It is the acceptance of the norm that is dictating the terms of the internal debate for her, and thereby "blocking" her ability to resist the preference to take the contraceptive risk.

---

[73] Stoljar, 2000, p. 101

The second criterion for autonomous decision making that Stoljar considers is that of self-knowledge. This criterion is an addendum to the counterfactual condition. For an action to be autonomous, one's hypothetical critical deliberation must not involve self-deception. For example, one subject in the study discussed by Stoljar seemed to be in a state of denial regarding the very existence of her sexual activity. Sex was not something that she and her partner ever planned or talked about, it was as she saw it something that just happened.[74] This way of thinking about her sexual activity allowed the subject to pretend that it never did happen.

While the self-knowledge criterion captures this case, it is not able to capture cases that involve other forms of deception. A second subject discusses how she decided not to use contraception because her doctor was likely to tell her father (who was also a doctor), and that this knowledge would likely hurt or deeply offend him. In this case, there is the norm, having been internalized by the subject, that premarital sex is immoral. In turn, the subject decides not to make use of contraception because doing so would inform her father of her sexual activity and cause some sort of emotional distress.[75] Since this is not a kind of self-deception, and is instead the deception of another (i.e., the subject's father), Stoljar claims that the self-knowledge criterion fails to reconcile with the feminist intuition as well.

A third procedural approach to standards of autonomy considered by Stoljar is that of "weak normative competence." Here the idea is that, in order for a preference to be autonomous, it must reflect the notion of self-competence in the following way. One must be able to justify one's actions in terms of normative standards that are seen as being relevant by other individuals. Agents must have a certain degree of "self-confidence, self-trust, or self-esteem" and thereby see themselves as needing to live up to certain normative standards.[76] This is of course the same kind of view described in Chapter 3, one advocated by Mackenzie. Stoljar uses the example of a woman who, because she is prone to emotional outbursts, is considered by her husband and others to be emotionally unstable or "just plain crazy." Because she has accepted this norm or belief (i.e., the idea that such outbursts are indicative of women's' inherent emotional instability and incompetence), she does not, according to the weak normative competence theory, demonstrate the ability to answer to certain normative standards.

---

[74] Stoljar, 2000, p. 102
[75] Stoljar, 2000, p. 103

Now consider the case of a young woman who believes that her body type, which is typical, is unappealing, and has internalized the norm that women should strive to look like the atypical runway model. This individual, according to Stoljar, would qualify as autonomous on the weak normative competence approach, but not on the strong normative competence approach. The reason for this is that she has internalized a norm that effectively blocks her ability to understand her own value in the proper way. In other words she has accepted a normative claim that is false as true, and in turn bases certain decisions on that claim. She is considered autonomous on the weak conception because she understands herself as competent to meet certain standards set forth by others, which is just what she is doing in the case of her preferences regarding body image. However, she is *not* autonomous on the strong normative competence view because she has internalized a norm that is essentially false.

I turn now to what I take to be the general form of the examples used by Stoljar. In each case, the problem is that a certain norm has been internalized or accepted as true by the agent. Hence it is not the process of internalization that is the object of Stoljar's analysis, but instead it is the falsity of the norms internalized that is the problem. The first thing to note is that this squares neatly with the form she provides for the feminist intuition. That is, *if* an agent has accepted as true one or several of the norms of femininity, *then* the preferences or decisions based on those norms are nonautonomous.

---

[76] Stoljar, 2000, p. 107

The second point to note about the form of Stoljar's criticism is that the deliberative process, given the internalization of a false norm, possesses a great deal of potential complexity and variation. This complexity leads to what I take to be the upshot of Stoljar's analysis. Due to this complexity, procedural accounts of autonomy are not capable of capturing every conceivable way that an agent might deliberate given the internalization of a false norm. For instance, in the example above regarding the denial of sexual activity, such deliberation might involve self-deception and thereby be captured by the procedural standard (of self-knowledge), or it might involve the deception of another and thereby not be captured by the procedural standard in question. What falls out of Stoljar's analysis is that the only way to claim that such deliberation is non-autonomous is to stipulate that the adoption of the norm itself is what constitutes a lack of autonomy (i.e., because it is a false norm). Note that I am not here claiming that Stoljar has produced an "in theory" argument to this effect. Hence it is in some sense still possible to construct a set of criteria that can deal with the degree of variation and complexity of this kind of deliberation. However, I am skeptical that such a thing can be done.

To put this in terms of the concepts discussed in the previous chapter, what Stoljar has done is provide an argument for the notion that procedural views are content-neutral in a more significant way. In other words, this is no longer a matter of stipulating a feature of procedural theories, but a consequence both of their structure and of viewing autonomy as relational (i.e., embedded). If we accept her conclusion, there is no guarantee of agents accepting any particular values given procedural constraints. And of course the worry from Stoljar's point of view is that this leaves no consistent connection between agents who qualify as autonomous on procedural accounts and whether they have internalized the norms of femininity.

Furthermore, she has presented strong reasons to believe that there is a similar disconnect between those norms and weak normative competence views. So even though I have presented those views as being weakly substantive, the substantive constraints are not of the sort that will guarantee the kind of connection to the feminist intuition Stoljar claims that feminist should insist upon.

To this extent, I would agree with Stoljar that the strong normative competence theory is the approach consistent with the feminist intuition. However, I believe the approach suffers from a separate problem insofar as it is intended as a general standard of autonomy. It is not clear why the falsity of the norm in question should mark the absence of autonomy.[77] Consider an example. An individual living in ancient Greece decides to make an animal sacrifice to Poseidon before taking a voyage at sea. Now, assuming that there is no actual causal connection between tossing a domesticated animal over the side of a cliff and one's chances of surviving a trek on the ocean, it seems fair to say that the norm in question is false. Of course the nature of the norm in question here is quite different from the norms of femininity, but that is the point. The only thing that this example has in common with Stoljar's is it involves acting on a false norm. Yet it does not seem correct to claim that the individual performing the sacrifice is nonautonomous simply because his action is based on a false norm.

---

[77] Benson makes a similar point (2005). He sets it out as a distinction between autonomy and orthonomy, or self-rule and right-rule, respectively. Just because an agent has gotten something wrong, at least in a moral sense, does not mean that she is not autonomous. I am also trying to argue here that the addition of embeddedness makes strong substantive notions implausible because of the extent to which agents may be influenced by relational factors.

Notice that the reason why one might have this intuition is because of the larger historical and social context in which the example occurs. To say that someone living in Greece a few hundred years BC is not autonomous because he accepts the dominant norms of his time and culture is to, a certain extent, reject the central intuition behind relational autonomy itself. Individuals are, on the relational account, embedded to at least some degree, hence to insist on such a degree of independence from embeddedness is to in fact embrace the received view of autonomy to an extent that even some of its proponents might find unreasonable.

One might, at this point, claim that it is not the empirical falsity of the norm in question that strong normative competence is meant to pick out, but its lack of moral justification or moral falsity. Even with this distinction, it is still not difficult to construct a similar counterexample. Take instead the Mayan game and ritual of Pitz. This game was played between teams of two to five players, and the object was to hit a marker on the court with a ball (almost like an early form of basketball). Frequently Pitz was played just for fun or recreation; however, there are historical accounts of cases where the game was used as part of a much more serious ritual, and the captain of the losing team was put to death.[78] Some anthropologists believe that the game itself would symbolize the struggles between the gods of the Mayan pantheon, or interpretations of the movement of celestial bodies and transitions between seasons. In addition, there is some evidence to the effect that the heads of losing captains were wrapped up and used as balls for future games.

---

[78] See http://library.umaine.edu/hudson/palmer/Maya/ballgame.asp, which describes part of the William P. Palmer III Collection at the University of Maine Orono Library.

With an example such as this one, we might think it immoral to allow the sacrifice of members of a losing team. Of course, the idea behind this example is to show that there are ideas that can be deeply entrenched in a culture yet still seem to lack moral justification; hence, adding moral falsity to the conception of the falsity of a norm is not sufficient because we are still left with reason to be skeptical of the ability of anyone in such cultures to meet the substantive standard.

There is a second problem with the strong normative competence approach as a general standard of autonomy. Recall the distinction made above regarding the difference between internalizing a particular norm and deliberation based on that norm. Assuming a certain degree of embeddedness of agency, it seems that any relational account of autonomy would need to say something about *both* of these processes. The strong normative competence approach claims that autonomy is dependent upon the agent's success in resisting the internalization process. In other words, it is the end result that matters on this account. My contention here is that this approach misses an important opportunity to respect autonomy by not speaking to the *way* that an agent could resist a false norm.

The story of Carol Gilligan described in the previous section offers an example of what I have in mind here. It is clear that Gilligan was confronted with a false and oppressive norm through the actions and words of her professors. It is also clear that she was successful in resisting the internalization of that norm. By applying the sorts of skills that Meyers describes, Gilligan was able to pursue the research program that she preferred despite the resistance she encountered. Even if one maintains the intuition that it is *because* Gilligan was successful in resisting the false norm that she should be considered autonomous, saying nothing as to *how* that process of resistance played out leaves the corresponding conception of respect for autonomy crucially underdeveloped. The problem for the strong normative competence approach, then, is that, without any sort of account of how such resistance occurs or more importantly should occur there is little or no guidance as to how one should respect autonomy.

In other words, it seems problematic that one’s *standard* of autonomy in no way informs one’s account of *respect* for autonomy. This problem would seem more applicable to the strong substantive views than weak normative competence views. As mentioned in the previous chapter, Mackenzie describes a few ways that values such as self-trust and self-respect could be promoted within other agents, and how we might have an obligation to do so as a matter of respect for autonomy. That said, strong substantive views simply insist on the rejection of damaging or false norms such as the norms of femininity. This is focusing on the end result as it were, instead of identifying some of the underlying values that could conceivably prevent such internalization or more directly inform opportunities for respecting autonomy as a capacity.

One might object that the same kind of difficulty Stoljar points out regarding deliberation based on norms applies also to the resistance to those norms. That is, given that the process of resistance to false norms is potentially as complex as the deliberation based on false norms, procedurally based accounts of respect for autonomy will fall victim to the same sorts of counterexamples. In short, I think that this objection is correct. However, what we are left with when we add this consideration is simply a high degree of skepticism regarding the ability of both procedural and substantive accounts to do the work that we want them to do. On the one hand, there is always room to generate a counterexample to any given procedural standard, Stoljar's analysis showed us this. On the other hand, the strong normative competence standard, insofar as it requires all norms an agent uses in his deliberation to be true in order for his choice to be autonomous, is susceptible to counterexamples as a standard of autonomy as well. Furthermore, the strong substantive standard generates a disconnect between itself as a standard of autonomy and what is supposed to constitute respect for autonomy. What is said about what constitutes an autonomous preference does not properly inform how we should go about respecting autonomy.

To summarize, I am skeptical that the debate over what constitutes an autonomous preference can be resolved by appeal to standards of autonomy alone, at least as it is conceptualized here. Given the complexity of the human deliberative process and the assumption of embeddedness, both kinds of standards seem fundamentally subject to counterexamples. In lieu of rejecting the debate altogether as intractably complex, I think the more sensible move to make is to point out that our understanding of the proper standard of autonomy will ultimately be informed by our understanding of respect for autonomy and how autonomy is intended to function in other theoretical contexts. This is the idea of reciprocal illumination mentioned in the introduction of this chapter, and what I now turn to in the following section.

Rethinking the Question: Respect for Autonomy

Going Beyond Autonomy to Understand Autonomy

To be clear, I am not arguing that the task of determining a proper standard of autonomy is a lost cause as it were, but that our best chances for deriving such a standard is to appeal to more than just ideas of what constitutes autonomy. While I have provided reasons to be skeptical of the project of deriving a standard of autonomy given embeddedness, there are the separate notions of respect for autonomy and how we go about using or treating autonomy as a marker for other moral notions and inclusion in certain political ideals. My claim is that understanding these other notions can help, at least to some degree, understand what a proper standard of autonomy might look like. This is the notion of reciprocal illumination, and what I will try to unpack in this section.

Although I touched on the ideas of perfectionism in the previous chapter when discussing what properly counts as a relational conception of autonomy, it will be necessary to go into a little bit more detail here. Bearing in mind that this debate is complex and well discussed in its own right, my goal here is just to give the reader a sense of the notions of political liberalism and perfectionism in order to show how they connect with the procedural/substantive distinction in autonomy.

John Christman has argued on more than one occasion that identifying with political liberalism instead of perfectionism may in fact require accepting a procedural standard of autonomy.[79] To provide a very general definition, perfectionism "is the view that values and moral principles can be valid for a person independent of her judgment of those values and principles."[80] From this it follows that there are certain values that should drive the structure and aims of political institutions, independent of whether a reasonable person might be expected to accept those core values.

---

[79] See both "Relational Autonomy, Liberal Individualism, and the Social Constitution of Selves" (2004) and "Personal Autonomy and Liberal Legitimacy" (2005).

Political liberalism, in contrast, is the notion that respect for liberty should drive the structure of political institutions; and that, in pluralistic societies, the rationale for "specific values or comprehensive moral views" should be something that individuals may reasonably accept; not something that must be accepted because they are based on some sense of moral objectivity (again in contrast with perfectionism). In other words, when a government institutes some law or policy, the reasoning for it should not be dependent on any one comprehensive moral view. The idea is that principles should be grounded "neutrally" (i.e., not as a result of appealing to a particular moral value), and this neutrality will provide the best chance at acceptance from individuals in a deeply pluralistic society.

Again, my aim here is just to provide a rough sketch of the debate, and not to describe any of the strengths or weaknesses of either point of view. Instead, the idea here is to point out, as Christman does, that adhering to a version of political liberalism commits us to a procedural version of autonomy. The reason for this plays out in two stages. First, autonomy constitutes a marker for identifying individuals who will be "bound by" the principles of the institutions and/or laws of a society. Second, *how* those principles are justified to those individuals depends on how we see agents deliberating on and possibly accepting those principles. In other words it depends on how we understand the nature of their autonomous deliberation. This is the crucial part, and worth quoting Christman at length:

---

[80] Christman (2004) p. 152

> For if by “autonomous” we mean only those persons who have particular value commitments – say, to relations of independence of a certain sort – then the principles of justice that shape social institutions will automatically rule out the perspective of those who reject those values. This perspective will not be considered and rejected in the mechanisms of public reason; rather, the viewpoint will not even be *considered*. Such “persons” will be viewed like children or the insane, as “unreasonable” and politically irrelevant.[81]

So the point is that autonomy is meant to pick out not only who is included in participation in public institutions or policies, but to whom those policies should be justified (and this part is true for both sides of the debate). But it also picks out how those institutions should be justified to those individuals. *What counts as legitimate forms of* deliberation regarding a given aspect of political institutions is determined in part by whether it includes some particular value, in the case of perfectionism. Hence, a procedural account of autonomy would not be consistent with this.

Again, the point here is not to take one side or the other, although Christman is clearly a proponent of political liberalism. Instead, the goal is just to point out there is a need for coherence between our understanding of conceptions of autonomy and conceptions of political legitimacy and modes of justification. From this, it would seem to follow that the two questions can inform one another as well.

In a similar way, Stoljar’s appeal to the feminist intuition can be seen as informative about standards of autonomy. To understand why this is, it is necessary to note that the feminist intuition is not just another instance of what is sometimes referred to as a “happy slave” counterexample. With these, the idea is to attack a particular standard of autonomy with an example of an agent who qualifies as autonomous under the standard, yet who has made a choice (or who has found herself in a situation) that we intuitively take to be an instance of nonautonomy. While Stoljar’s argument could be viewed as having this structure, and on a basic level it does, the difference is that the motivation behind the feminist intuition itself is embedded in a much more complex theoretical context informed and motivated by feminism itself. Appealing to the intuition is not just a matter of appealing to something "given" as it were, but something that can be reflected upon in its own right.

---

[81] Christman (2005) p. 293

As a kind of oppression, there are different ways of understanding the moral and political implications of an agent having internalized the norms of femininity. This is of course an independent and highly complex question in its own right, and one that is part of the larger literature generated by feminism itself.[82] Hence, I take this to be an example of the kind of issue Christman has in mind when he makes the claim that our understanding of autonomy relies in part on how it functions in other contexts. Stoljar's argument is insufficient in the sense that it makes a simple appeal to the feminist intuition in the same sense of "happy slave" examples, when in fact the grounding for the intuition is much more complex. This is not to invalidate the connection she is making overall, but simply to point out that more would need to be said about the grounding of the feminist intuition independent of anything she can argue about autonomy itself.

Respect for Autonomy

The alternative I propose is, in short, to understand the problem constituted by the norms of femininity and other such appeals to intuitions in terms of *respect* for autonomy. This requires adopting procedural relational accounts of autonomy (such as Meyers'), not as standards of autonomy, but instead as means of adopting practices and policies that nurture capacities that are generally conducive to autonomy. Part of the goal then is not to rely on the feminist intuition strictly as an intuition that must be accepted wholesale in order to inform our discussion of standards of autonomy. Instead, it is a sort of trigger for considerations about respect for autonomy.

---

[82] Benson makes a similar point (2005). Not all feminists would view the subjects in Lukers' study as non-autonomous.

The solution I am presenting requires understanding how the promotion of the norms of femininity may constitute disrespecting autonomy. Not only should we provide, whenever possible, for the skills and capacities outlined by procedural relational accounts, but respect for autonomy also demands that we seek out and eliminate instances of the promotion of the norms of femininity. Furthermore, we do this with an eye towards developing a practical solution to what is essentially the real problem: the intuited lack of autonomy on the part of women who internalize or act from the norms of femininity. In this sense, the feminist intuition becomes, not something to be accommodated by a *standard* of autonomy, but instead one part of a understanding of progress for feminism as a social movement. It offers one way of understanding whether or not autonomy has been respected, and whether there is more that needs to be done.

Reconsider the example discussed above regarding the difficulties faced by Gilligan. The initial point of this example is that procedural standards address both stages of the deliberative process. That is, they address the process of the internalization of norms (or the resistance of those norms as the case may be) as well as deliberation based on those norms. In this sense, they address a fundamentally important aspect of the relational view of agency. Hence, if embeddedness is a conceptually necessary part of relational agency and the potential for internalizing norms is a necessary part of embeddedness, then procedural accounts must be incorporated into a relational conception of respect for autonomy in some way. Otherwise the conception has no way of addressing a fundamental part of embedded autonomous agency.

The most significant problem with this account of respect for autonomy is that, given Stoljar's objections to procedural accounts of standards of autonomy, why should we expect it to be up to the task of satisfying the feminist intuition? Why should we expect the capacities and criteria of procedural accounts, even when satisfied, to meet the practical goal set by the feminist intuition? It seems as if there will still be imaginable cases wherein an individual has had every applicable capacity respected, yet still has internalized some false norm of femininity. In such a case, it would seem that our account of respect for autonomy has failed us for not accommodating the feminist intuition, in the same way that procedural accounts of standards of autonomy have failed us.

There is an important distinction to be made in order to respond to this problem. Recall that the form of the objection generated by Stoljar to procedural accounts of standards of autonomy was that the deliberation based on false norms is essentially too complex for any one procedural standard to capture. There always seems to be some way in which the deliberation can go that is missed by any one procedural standard. In the case of respect for autonomy, however, we are looking at embeddedness from the opposite direction, as it were. We are considering the ways in which a norm of femininity can be internalized, not how it can be used in deliberation.

With this distinction in hand, one can see that respect for autonomy forces us to ask the following question. Were I to promote a particular norm, is there any conceivable competency or skill associated with autonomy that would be diminished on the part of the agent, were she to internalize the norm? If so, then the promotion of that norm constitutes disrespect for autonomy. Notice that approaching the issue from this direction places us in a much different situation. In the case of accommodating the feminist intuition via a procedural *standard of autonomy*, the task we were forced into (as a result of Stoljar's objection) was to construct an exhaustive list of skills or formal requirements that could capture every conceivable form of deliberation based on a false norm of femininity. Instead, the task we now have is to ask whether, *given a particular norm*, is there *any* sense in which the promotion of that norm diminishes *any* skill associated with autonomy on the part of the agent.

I take this to be a much more manageable task than the one previously described, primarily because we have particulars "in hand" as it were. Given the facts of some particular situation, it is much easier to understand how some agent (whether hypothetical or real) can conceivably suffer the diminishment of some skill that we intuitively take to be constitutive of autonomy.

To highlight the difference between the two approaches, I offer the following example. A young woman, say in her mid-twenties, lives in Las Vegas with her husband. We will call her Sarah. Sarah's husband and his friends frequent strip clubs, and often speak of the dancers that they see perform. They comment, albeit in an approving way, of how attractive these women are and how much money that they must make from tips. These dancers are lauded by Sarah's husband and his friends for being everything a man would really want in a wife: attractive, financially independent and they know "how to have a good time."

In some cases, the element of financial independence can be seen as part of what it means to be an ideal wife simply because there are men who feel that a wife should not be a "financial drain." In this way, some norms can seem to be a mixture of things that, in other contexts, would otherwise be taken as a positive norm. In virtue of making such comments and putting these women up on a pedestal as it were, Sarah's husband and his friends essentially promote a norm to her that this is what a woman should be.

Per Stoljar's analysis and the feminist intuition, assuming that Sarah accepted such a norm and thereby sought to be this type of woman, we would take her to be non-autonomous. Furthermore, we would have difficulty explaining exactly why assuming a procedural standard of autonomy. Suppose for example, that Sarah decides to pursue a career in exotic dance in an effort to prove to her husband that she can be an ideal woman to him. Her deliberation in this case would first be based on the norm presented to her, but exactly how it goes could vary in a myriad of ways. We could say that her decision lacks a kind of self-respect, but that it still qualifies as autonomous according to a counterfactual standard (such as Christman's). Or, it is possible that Sarah demonstrates a sufficient level of self-knowledge by acknowledging that she does in fact have the ability to be one of these women, but must deceive others by claiming that she actually has the desire to do so for a purpose other than satisfying her husband.

If what it means for a standard of autonomy to accommodate the feminist intuition is to avoid all such counterexamples, then as I have stated previously, Stoljar is correct to claim that procedural accounts are not up to the task. However, the kind of counterexample that concerns us in the case of *respect* for autonomy has the following form: agent X has had all rights/obligations respected regarding her autonomy, yet agent X has still internalized a false norm. For an account of respect for autonomy to accommodate the feminist intuition, it must not be subject to *these* kinds of counterexamples.

Returning to Sarah's case, it would not be difficult for one to conceive of a number of different ways in which promoting this false norm of the "ideal wife" could have negative effects on any number of the skills that constitute Sarah's autonomy. Primarily, such promotion would diminish Sarah's ability to conceive of and pursue conceptions of herself that are more conducive to her flourishing and might very well diminish her capacity to envision other alternative long-term goals in her life. In short, by attempting to convince Sarah that this is how women should be, we limit her ability to understand her real options in a damaging way. This being the case, the promotion of that norm should be deemed disrespectful of Sarah's autonomy.

It is important to note that I do not take this to be an "in theory" argument that this approach will always accommodate the feminist intuition. I do contend, however, that cases wherein individuals seem to possess the relevant capacities (that is, the capacities that would make them autonomous on procedural accounts) but have still internalized some false norm are theoretically possible but empirically very rare. The advantage of this approach is therefore an epistemic one. It is simply much easier to conceptualize the damage to an agent's autonomy *given* a particular norm and situation than it is to construct an exhaustive list of every conceivable possibility. With this approach, one is presented with a particular scenario and can then deliberate as to the conceivable effects on the agent in question.

Keep in mind, that this is only one aspect of the shift I am proposing. This analysis helps us to see how shifting the question from standards of respect for autonomy changes how we view particular instances of oppression or a lack of autonomy. What is also being recommended, as stated above, is that there is a great deal of analysis regarding the feminist intuition itself that can be brought to bear on these examples.

Up to this point in the conversation, respect for autonomy has been discussed in terms of oppressive norms and agents' decisions based upon them. I believe that this notion of respect for autonomy can be made more general in the sense of viewing any action, not just the promotion of a particular norm, in terms of whether it could diminish the agentic skills described by Meyers or the competencies described by Mackenzie. In this way we would have a general sense of what constitutes a generally relational conception of respect for autonomy. Were I to commit a particular action towards an agent, is there any conceivable competency or agentic skill associated with autonomy that would be diminished on the part of the agent as a result? This is the test that I will be applying in Chapter 6 in the cases of social networks and lifelogging.

Conclusion

I have covered a great deal of ground in this chapter, and so some recapitulation is in order. First, I presented Diana Meyers' skills-based account of autonomy. Although I take it to be primarily procedural in nature, it has, as Benson has suggested, some substantive content in terms of making sense of the skills themselves. Hence, her view has been presented as occupying a sort of middle ground in the spectrum of standards of autonomy and the level of substantive constraints that they entail. For this reason, Meyers' and Mackenzie's views will be the basis of my analysis in subsequent chapters outlining examples of how relational autonomy connects with the value of privacy. To be clear, I do not consider the analysis I have provided here to be an argument to the effect that their views are conclusively better than the others mentioned, such as Christman's or Stoljar's. I am simply taking their occupation of a middle ground as sufficient reason to use them as examples of relational conceptions of autonomy. However, I will offer some evidence to the effect that theories such as these that strike a sort of balance between content neutrality and substantive constraints connect well with pre-existing claims in the literature on the value of privacy.

In the second section, I presented Stoljar's objection to procedural accounts of autonomy based on the feminist intuition. I take her objections to demonstrate that there is more than just a stipulated element of content-neutrality with procedural accounts of autonomy. Given embeddedness, there is essentially no way to guarantee the substance of values held by agents who qualify as autonomous on those accounts. Hence there is no sense in which they are consistent with the feminist intuition. In addition, I also pointed out some of the difficulties associated with strong substantive views, again, given embeddedness. These included the notion that they incorrectly link the falsity of norms to nonautonomy, and that they generate a disconnect between standards of autonomy and respect for autonomy.

What falls out of the analysis of Stoljar's argument (and my objections to it) is that the general project of deriving standards of autonomy is problematic. Specifically, there are difficulties reconciling standards of autonomy with intuitions about certain values, that is, *if* we run the analysis strictly in terms of the core concepts constitutive of autonomy.

In the third section I turned my focus to Christman's assertion that part of how we understand autonomy is dependent upon how it is used in other theoretical contexts. He mentioned the debate between perfectionism and political liberalism, and contends that analyzing standards of autonomy from such external perspectives offers insights that cannot be gained from analysis of autonomy alone. For these two reasons, I advocate a shift from standards of autonomy to focusing on respect for autonomy.

In the next chapter, I will discuss how accounts of respect for autonomy, at least those informed by either Meyers’ or Mackenzie’s relational standards of autonomy, connect with the value of privacy.

CHAPTER 5

PRIVACY AND RELATIONAL AUTONOMY

Introduction

In this chapter I will attempt to explain what I take to be the relationship between the relational conception of respect for autonomy discussed in the previous chapter and the moral value of privacy. I will do so in two stages. First, I will discuss the non-normative dynamic between privacy and what I refer to as relational processes. Whether or not we enjoy a particular moment of inaccessibility can have an impact (good or bad) on agentic skills and competencies. The first section of this chapter will be devoted to discussing how this is the case. This is the core dynamic constituting what I call the relational approach to the value of privacy.

Second, I will situate the relational approach to privacy in a larger context of other kinds of theories about privacy. The ultimate goal here is to come up with an approach to the value of privacy that can actually make recommendations in particular situations. To this end, it will be necessary to connect the relational approach to privacy to approaches that are able to make sense of the details involved in those situations. This discussion will focus on some of the recent work by Helen Nissenbaum. Nissenbaum's approach, what she calls contextual integrity, focuses on the context in which privacy issues exist. Context matters because it illuminates the social and cultural nuances that are part of how we naturally evaluate what privacy can or should do for us.

More specifically, contextual integrity is the idea that privacy issues occur in specific contexts which are partly constituted by expectations and conventions that in turn help to determine what should be reasonably expected of the parties involved. In addition, Nissenbaum has developed the “contextual integrity decision heuristic” which is essentially a decision making procedure for determining how best to understand what the context is, and then how to determine whether a new practice or technology actually constitutes a significant change, or one that is worth analyzing as a possible violation of privacy. The primary point behind including Nissenbaum at this stage is to help to understand both the contributions and limitations of my theory of the value of privacy in terms of how capable it is of making practical recommendations in actual privacy issues. While I will discuss this in more detail, the gist of it is that my theory occupies a more abstract space in the overall continuum of theories about privacy; however, it interfaces quite well with more practical approaches like Nissenbaum’s, specifically with the heuristic she has developed.

Collectively, this discussion of Nissenbaum’s theory and the discussion of relational processes will set up Chapter 6. There I will discuss certain examples in more detail, with the focus being on privacy issues in the information technology space. What I hope to be able to show in this chapter and the next is that a more nuanced approach to understanding relational autonomy can attach relatively seamlessly to more practically oriented theories about privacy, such as Nissenbaum’s. Furthermore, I hope to show that this can be done in a way that satisfies the original criterion laid out in Chapter 2, which is that we can capture the intuitions that drive functionalist accounts of the value of privacy via autonomy-based accounts and thereby offer a more robust account of privacy's value.

Privacy and Relational Processes

As mentioned in the introduction of this chapter, the first part of the discussion here has to do with the relationship between privacy and relational processes. Privacy can either promote or hinder relational processes. By "relational processes" I mean the interpersonal and cultural interactions that make up the condition of embeddedness. Notice that one of the common elements of relational approaches to autonomy is that cultural and interpersonal factors and interactions (i.e., our social activities, relationships with others, norms with which we interact, etc.) are what drive or hinder the development of autonomous capacity. Respect for autonomy, as described in the previous chapter, is the process of understanding the rights and obligations we have in terms of promoting and not hindering this development. Privacy, in turn, can play a role in these processes in a variety of ways. To understand the dynamic that I am here suggesting, I will revisit both Mackenzie's normative competence account and Meyer's procedural standard of autonomy in the context of how they relate to privacy. The reason for discussing both is to show that privacy (as a state of affairs) can function within a range of conceptions of respect for autonomy. As stated in the introduction, describing how privacy can promote or hinder the processes described in procedural conceptions of autonomy is the first step in understanding its value.

The dynamic between privacy and autonomy according to Meyers' conception of procedural autonomy is based primarily on the effects that privacy can have on the development and/or preservation of the skill set that Meyers sets out. To offer a sense of these effects, I will provide a detailed example.

The Adam Walsh Child Protection and Safety Act (sponsored by Orrin Hatch R-Utah) was signed into law by President George W. Bush in 2006. It requires that juvenile sex offenders (fourteen years of age and older) be registered "with a national sex-offender registry every three months and every time they change their name, address, student status or employment – for the rest of their lives."[83] There has been some debate regarding whether such a requirement violates the privacy of juvenile offenders in the sense that the goal of the juvenile justice system has traditionally been rehabilitation. To that end, juvenile records are not typically available to the public. The Walsh Act of course changes that with the intent of reducing the rate of repeat offenses by equipping others with the knowledge of who has committed sexual offenses in the past.

Consider the case of Shaun, who has gone through the rehabilitation process for juvenile sex offenders and is currently juggling school, a full-time job, and an active social life.[84] Shaun raped an older woman at the age of fifteen, and spent three years in juvenile lockup undergoing an intensive treatment program. The question to consider here is the following: What would be the impact of forcing Shaun to register as a sex offender, and thereby experiencing a loss of privacy, given Meyers' conception of relational autonomy?

Answering this question first requires reconsidering the skills that Meyers describes. Take, for example, the skill of introspection:

---

[83] Peterson, Eric "Branded," Salt Lake City Weekly April 26, 2007 Vol. 23 No. 49

[84] Peterson, 2007, p. 24. Shaun is the name given to the primary interviewee of the article by Peterson.

> [I]ntrospective skills that sensitize individuals to their own feelings and desires, that enable them to interpret their subjective experience, and that help them judge how good a likeness a self-portrait is[85]

Here the idea is that autonomous individuals should be able to place their subjective experiences in the proper context of interpersonal and cultural influences. Individuals naturally develop self-portraits or self-images that may or may not be accurate. Such self-portraits can be influenced by interpersonal experiences as well as by norms that generate certain expectations. Meyers' point here, then, is that autonomous individuals should be able to discern how accurate their self-image is in light of certain influences, at least to a certain extent. While Meyers acknowledges that complete independence from such influences is something of an illusion, the point is to consider the sorts of skills that can be brought to bear on them.

In Shaun's case, one might ask how his introspective skills would be affected were he forced to register as a sex offender. What would the experience of being "branded" as a sex offender do to his ability to understand who he is, in the sense of having an accurate self-image? In conjunction with the Walsh Act, the Department of Justice has provided a web site wherein one can search the databases of all states in compliance with the law.[86] It allows for any member of the public to search by name or location (i.e., by city, county, state, zip code, etc.) for registered sex offenders. The site provides the name of the offender, a picture, a description of the crime, all of the aliases associated with the offender, descriptions of vehicles that the offender has, and the offender's current address. How would the awareness of such a degree of exposure affect Shaun's ability to formulate an accurate self-image?

---

[85] Meyers, 2001, p. 741

[86] See http://www.nsopr.gov

While I take it that there is no simple answer to this question, this would seem to be exactly the kind of potential influence that Meyers has in mind. Knowing that there are strangers out there who possess this kind of knowledge of one's past would seem to make it more difficult to construct the sort of positive self-image associated with rehabilitation. It might very well cause individuals like Shaun to place an undue emphasis (in terms of his self-image) on the crime that he has committed, instead of the individual that he can become. In turn, if the idea behind rehabilitating juvenile sex-offenders is, in part at least, to help them to construct a nonviolent pathology and sense of self, then forcing them to register and thereby experience such a loss of privacy would seem to be counterproductive in this sense. It simply makes the task of rehabilitating Shaun more difficult.

I turn now to the same kind of analysis applied to Mackenzie's weak normative competence account of autonomy. Sticking with the example of Shaun, adopting a view of autonomy that focuses on goods such as self-esteem and self-worth leaves us in a strong position to appeal to the same sorts of intuitions about the rehabilitation of juvenile offenders. Any curtailment of privacy has the potential to limit the opportunities to develop a sense of self-worth thereby hindering the process of building a nonviolent pathology. If Shaun believes that society has deemed him untrustworthy by requiring that everyone else has an enhanced opportunity to find him, know what he did as a minor, and act accordingly (without him every being aware of who knows and who does not), then it is difficult to see how that sort of arrangement will help him to see himself as worthy of full respect typically accorded full citizens.

Notice that this lack of knowledge of who is aware of his past and who is not (the anonymity afforded visitors to the sex offender registry site) makes evaluating Shaun’s relationships with other more difficult for him. When he interacts with someone, say in some way more involved than a passing interaction, he is left wondering what they know about him and what they do not. So it becomes more difficult for him to understand exactly why they might act the way that they do. With this extra variable in place, it becomes, literally, exponentially more difficult to build a sense of social worth and belonging.

Forcing Shaun to be registered on the site simply says that he is not to be trusted. It is a scarlet letter in electronic form. Notice that in the preceding discussion about Meyers’ approach, there is a question about how the awareness of this kind of exposure in turn affects his self-image. So the question is very much the same, but the difference here is that goods pointed out by Mackenzie’s approach speak more directly to that idea by considering the sorts of things we might take to be constitutive of self-image, and how it might impact autonomy. With these sorts of considerations in mind, it is a bit easier to see how privacy can enable the goods we are concerned about (in this case the rehabilitation of a minor) in a way that is tracked by the relational elements of a weak normative conception of autonomy.

What I am hoping to have shown in this section is the dynamic between relational processes and privacy. Privacy plays a role in these processes in some cases by facilitating the development of skills or certain goods such as self-esteem, and in other cases by diminishing both.

Putting the Relational Approach to the Value of Privacy in Context

Nissenbaum's Contextual Integrity

I will start this section by discussing how the relational approach to the value of privacy can be placed into a larger context in terms of what *kind* of theory of privacy it is. The relational approach to the value of privacy, as described thus far, is the claim that the value of privacy is best understood through the lens of respect for autonomy understood relationally. That claim needs more specification if it is to be of any use in real life situations. Hence the relational approach is a more general or global theory about the value of privacy that can be contrasted with Nissenbaum's contextual integrity.

In the second part I will discuss why it is important to understand how more general theories such as mine can be connected with theories that are intended to offer very specific recommendations regarding privacy issues. Specifically, I will discuss Nissenbaum's contextual integrity approach and lay out some ways that my theory can be connected with it, as well as discuss a concern about this connection. In particular, I will discuss an issue that arises from Nissenbaum's reliance on and incorporation of the ideas of political conservatism and Walzer's spheres of justice. Ultimately, I claim that neither is of immediate significance. However, it is important to examine these connections to see exactly how the theories relate to one another when it comes to their normative recommendations

In the introduction of her work *Privacy in Context*, Nissenbaum makes use of a spatial metaphor to categorize different approaches to understanding privacy's value. At the top, in the "stratosphere" as she puts it, are theories such as mine that seek to understand the fundamental moral and/or political underpinnings of the value of privacy.

These views rely on “appeals to universal human values and moral and political principles,” while what happens on the “ground floor” is a matter of interest politics.[87] Here, we have what people commonly think of as a privacy dispute between individuals and/or groups, companies, the government, etc. As Nissenbaum characterizes it, in most cases where there is an imbalance of the interests involved in a particular practice, there is likely to be some dispute.

For example, there may be little dispute in the case of certain types of medical surveillance going on in a hospital such as monitoring blood pressure electronically, and this is because of the obvious benefits to the parties involved. With other practices such as the tracking of consumer behavior, the interests are lopsided (depending on certain factors such as what is done with the information). In these cases, interests clash, and lines are drawn between various parties to settle the dispute. Nissenbaum refers to these as “interest-brawls.”[88]

In the middle, between moral theories and interest-brawls, we have Nissenbaum’s notion of contextual integrity. Contextual integrity is, in the most general terms, an account of what constitutes the “appropriate flow” of information. As Nissenbaum claims, what people are most concerned with is not necessarily restricting the flow of information, but that information flows in ways that they understand, expect, and approve of to at least some extent. She describes it as follows.

> The framework of contextual integrity provides a rigorous, substantive account of factors determining when people will perceive new information technologies and systems as threats to privacy; it not only predicts how people will react to such systems but also formulates an approach to evaluating these systems and prescribing legitimate responses to them.[89]

---

[87] Nissenbaum, 2010, p. 10
[88] Nissenbaum, 2010, p. 8
[89] Nissenbaum, 2010, p. 2

This may seem like a series of unrealistic claims, but the idea is just that there is a way to analyze particular situations involving privacy that brings out what these concerns are and how to address those concerns that is highly sensitive to the particular context (and is effective for partly that reason). Also bear in mind that Nissenbaum's focus here is on information privacy primarily as it relates to information technology. This will also be the focus of my extended examples in Chapter 6.[90]

This is where the "contextual integrity decision heuristic" (hereafter just "decision heuristic") comes into play.[91] Nissenbaum has developed a multistep process for evaluating particular cases of violations of information privacy and then offering possible means of resolving those cases. What she has in mind here is essentially resolving the aforementioned interest-brawls via analyzing how and why the information is flowing the way that it is, understanding the existing relevant norms or conventions in virtue of analyzing the context of the particular case, and then identifying whether a change has occurred. If a change has occurred, it is flagged as a prima facie violation of contextual integrity. Let us outline the steps and core concepts in more detail.

Step one is to "establish the prevailing context." By context, Nissenbaum means "structured social settings characterized by canonical activities, roles, relationships, power structures, norms (or rules), and internal values (goals, ends, purposes)."[92] Here the idea of context as she presents it is generalized from sources in both political and social philosophy including such thinkers as John Searle and Suemas Miller.[93] The core notion is to understand what normative ideas and typical practices make up a generalized version of situation in which the particular privacy issue occurs. Examples of different contexts would include education, healthcare, psychoanalysis, etc.

---

[90] Note that in the block quote Nissenbaum makes the claim that the decision heuristic is able to anticipate the reactions of individuals to changes in privacy policies and practices. This may seem overstated, however, she does offer several examples based on surveys that support her claim. I will not evaluate this part of her theory, other than just to state that I believe her claim is well supported by the evidence she provides.

[91] Nissenbaum, 2010, p. 181

[92] Nissenbaum, 2010, p. 132

To flesh the idea of context out a bit further, we can look at the components of context. Context is made up of roles, activities, norms, and values. Roles are “typical or paradigmatic capacities in which people act in contexts.” Activities are “canonical activities and practices in which people, in roles, engage.” Norms are prescribed actions and practices that in turn “define the relationships among roles and, in this way, the power structures that characterize many familiar social contexts.” And finally, values are made up of any “goals, purposes, or ends; the objectives around which a context is oriented.”[94]

Step two is to determine which “attributes” are affected. Attributes are defined by Nissenbaum as simply the kind and degree of information involved. After identifying the context, we need to look at the actual information involved.

Step three is to establish any changes in the “principles of transmission.” These principles are simply the rules by which the information is changing hands or how it is being managed. In short, they are any constraint on the flow of information from one party to the next.

Finally there is step four, which is to (red) flag any change that has been detected by going through this analysis. Nissenbaum describes it as follows.

---

[93] Nissenbaum, 2010, p. 131
[94] Nissenbaum, 2010, p. 133

> If the new practice generates changes in actors, attributes, or transmission principles, the practice is flagged as violating entrenched informational norms and constitutes a prima facie violation of contextual integrity.[95]

This last step may seem counterintuitive because, on the face of it, it would seem to arbitrarily set the burden of proof on anyone arguing in favor of any change to established practice. However, this step gives the decision heuristic what Nissenbaum sees as an element of conservatism; and I will return to this point shortly.

Let me summarize what I have presented of Nissenbaum's account so far. Focusing on informational privacy, contextual integrity seeks to evaluate particular real world issues regarding privacy in terms of the contexts (as defined above) that surround them. The goal is to evaluate the norms, actors, information (and that information's associated transmission principles), and values that make up that context; and based on that evaluation, determine if a change has occurred. If so, the change is red flagged as a prima facie violation of contextual integrity.

Consider one of the examples provided by Nissenbaum. Suppose a public school administrator is considering whether to implement a new student records management system. This system would provide an increase in the ability to track, organize and present actionable data about the student body and particular students. In applying contextual integrity approach, the administrator would need to identify the context in which this example occurs, education, and try to understand its components. This would involve asking questions about attributes, e.g., what kinds of information would be managed by the system such as academic performance, disciplinary information, information about the student's family of activities outside of school, etc. Other questions would include those involving actors and transmission principles such as who will have access to the system, how and why it is shared with parties outside of that particular school, etc. The administrator would initially flag the new system because it represents new capabilities that could potentially change attributes, transmission principles, and actors. The step essentially completes the basic version of the decision heuristic.

---

[95] Nissenbaum, 2010, p. 150

In the “augmented” version of the decision heuristic that Nissenbaum proposes, there is the additional step of determining whether the “violation” in question is actually warranted or justified. To do this, there of course needs to be some sort of normative component added to the decision heuristic. Because any change is flagged as a prima facie violation of contextual integrity, the heuristic is, on the face of it anyway, biased in favor of maintaining existing contexts, practices and norms contained within those contexts. In this respect, she considers the normative component of contextual integrity to rely in part on the notion of conservatism. By “conservatism,” Nissenbaum generally means the notion that there is at least some moral authority in entrenched social practice. She discusses this idea at length, citing both Jeremy Bentham and Edmund Burke. However, I will not endeavor to track her complete discussion here. For now, it is sufficient to note that she makes an initial appeal to the core intuition behind conservatism, which is that there is value in retaining existing social institutions, instantiated in this case by the contexts surrounding particular privacy cases.

As Nissenbaum points out, this initial reliance on conservatism creates a kind of dilemma. As she describes it, adopting the approach of contextual integrity leaves one with three options in regards to handling its normative component. The first option is to reject the component altogether and simply make use of contextual integrity as an effective means for identifying the significant components of situations involving privacy. In this way it would be effective as an initial tool for analysis, but it would not contribute directly to any prescriptive recommendations.

The second option would be to “dig in our heals, insisting that contextual integrity carries moral weight, committing us to a justificatory framework with a rigidly conservative basis.”[96] In other words, one could stick to her guns so to speak and just rely on the notion of conservatism to do the heavy lifting after the contextual analysis, but thereby being forced to contend with objections to conservatism presented by other competing political views such as liberalism. In other words, we would shift the debate regarding the normative component away from privacy insofar as we rely on conservatism (and all that it entails in the most general sense) to do the normative heavy lifting.

The third option for managing the normative component of contextual integrity, the one she adopts, plays out in two parts. First, she would “bring to bear general moral and political considerations.”[97] These would incorporate the kinds of theories mentioned before, such as mine, which occupy the stratosphere according to her spatial metaphor described earlier. These are the theories that “link privacy with critical individual and social values, including security against harms… protection of individual freedom and autonomy, promotion of fairness, justice, and equality, [etc.].”[98]

So the first part of the third option is essentially the same as the first option described above. In short, run the contextual analysis and then let other theories about the value of privacy do the rest of the work. However, the second part (of the third option) is to point out that context is not just a necessary component of the analysis in the sense of providing useful background information; it drives our understanding of what is genuinely important for the case at hand. Context, and in particular the values (goals, purposes, ends), is part of the actual prioritization of what is important in a given case involving privacy and informs deliberation about cases using theories about privacy.

---

[96] Nissenbaum, 2010, p. 161
[97] Nissenbaum, 2010, p. 165
[98] Nissenbaum, 2010, 165

Returning to our school administrator example would be helpful. In this example, the idea would be to employ concepts such as the ones linked to privacy by the more general normative theories, such as respect for autonomy and freedom. But it would not be just a matter of applying respect for autonomy in the more general sense, such as how I have described it in earlier chapters, but in a way that relies on the context to arrive at the prescriptive recommendation. It is worth quoting Nissenbaum's example at length here.

> [C]onsider whether a local company should have access to student records for hiring purposes on the grounds that this benefits the company as well as the students. A critic resisting the move need not dispute the general on balance benefits, but may raise concerns about the effects on internal purposes of education. Intellectual experimentation might be inhibited as the practical orientation of the company punctures the relative safety of a school to try out ideas. Further, there is the potential that the school might adapt its curriculum to pander to the needs of the local business at the expense of other educational goals.[99]

Nissenbaum continues, but the basic idea is clear enough. The goals associated with education as a context are driving the considerations here. Education should provide for things like the opportunity for intellectual experimentation, the opportunity to make mistakes without fear of significant consequences down the road, etc. Allowing a private company access to academic performance information could very well hinder these goals, and thereby have an impact on the autonomy of the students affected.

---

[99] Nissenbaum, 2010, p. 171

But notice, this is not a direct application of a global theory about autonomy. It is definitely an expression of concern for students' autonomous capacity and liberty, but one that is fleshed out according to the context of education. Now, I think it possible, and likely to be productive, that one could integrate a particular conception of autonomy into considerations like these. I will return to this point towards the end of this section.

Contextual Integrity and Walzer's Spheres of Justice

As Nissenbaum describes it, the motivation for how context is motivating the normative component comes from Michael Walzer's notion of spheres of justice. The core idea here is that social life is "made up of autonomous spheres defined by their ideologies and social goods."[100] These spheres are not collectively governed by a comprehensive moral principle, but instead are ruled by criteria distinct to each sphere. It is this idea of the normative independence among various areas of social life that Nissenbaum intends to motivate the importance of context. As in the example above, the goals of education, insofar as education is a separate sphere or context, are the dominant considerations for determining how to apply more general normative principles such as respect for autonomy.

At this point I would like to make a few comments about the general structure of the normative component of contextual integrity, as well as how it may or may not connect with my discussion of the value of privacy up to this point. First, there is a methodological observation. What Nissenbaum seems to be doing here is applying concepts from political philosophy to inform and motivate her understanding of the value of privacy. This is very much in line with other privacy scholars who relate their theories about privacy's value to particular concepts in political philosophy, and (on a methodological level) is similar to the approach taken by Anita Allen. [101]

---

[100] Nissenbaum, 2010, p. 167

In her work *Unpopular Privacy* Allen evaluates cases that involve the question of when paternalism is justified in the sense of forcing privacy on citizens in a liberal democracy. In order to place her thoughts on the subject in a broader context she identifies her underlying political view as "comprehensive deontic liberalism." This she defines as the idea that here is a comprehensive moral view underlying or driving the structure and goals of political institutions (and is ultimately the basis for justifying those institutions). And furthermore this moral view is deontological (as opposed to consequentialist), and based on "toleration, religious diversity, and gender and race equality."[102] Hence, whether a particular instance of "privacy paternalism" is justified will in part rely on what would constitute legitimate paternalism given her understanding of comprehensive liberalism. In other words, her view of comprehensive liberalism will inform her understanding of what qualifies as a legitimate reason for enforcing privacy regulations, specifically those that will by definition go against the will of citizens.

Of course, the big difference between Allen and Nissenbaum in this respect is the nature of the political theories in question. Allen relies on comprehensive deontic liberalism, whereas Nissenbaum relies on conservatism and Walzer's spheres of justice. Since it is beyond the scope of this dissertation to evaluate or criticize these concepts independently of their connection with privacy, my analysis is limited to two things. First, I can look into the extent to which Nissenbaum's commitments to particular political theories actually influence the normative results of the contextual integrity approach to the value of privacy. And this would be mostly just a matter of understanding the theory of privacy itself more thoroughly. Second and pending the results of the first point, I can try to analyze the extent to which her making assumptions regarding conservatism and Walzer's spheres of justice might potentially limit the ways in which my approach to the value of privacy might integrate with contextual integrity. I will say more about this later, but the worry is that putting an emphasis on the independence of different "spheres" or areas of social interaction has the potential of limiting the applicability and significance of more comprehensive moral concerns such as respect for autonomy.

---

[101] Allen, 2011

[102] Allen, 2011, p. 16

Starting with the first point, let us consider the element of conservatism, and then the application of Walzer's spheres of justice. Based on what Nissenbaum has presented, I think that the element of conservatism is mostly negligible in contextual integrity. Yes, the process is to flag all changes as prima facie violations of contextual integrity; however, the remainder (and majority) of the analysis is based on other normative concepts that do not contain conservative elements. That is, it is the combination of theories linking privacy's value to certain moral ideals (such as autonomy) and Walzer's spheres of justice that really drive the normative component. So in this respect, I think that the conservatism ends up being more a matter of thoroughness and an insistence on evaluating as many things as possible; and hence, Nissenbaum may be slightly overstating the role of the conservative intuition in contextual integrity.

Next let us consider Walzer's spheres of justice. Importing this notion into contextual integrity would seem to create a significant limitation on the role of normative principles, which in turn may not be consistent with how I see the relational conception of autonomy connecting with privacy. This is because, as I interpret it, context (or the sphere in question), determines whether a normative principle is applicable in the first place. However, whether or not this is the case depends on how exactly we interpret the normative independence of spheres on Walzer's view, and the extent to which Nissenbaum is actually importing that part of his view into contextual integrity.

So one avenue of analysis would be to investigate these two points further. In lieu of that, I will (for the most part) just make the following tentative claim regarding the second issue above (about how my theory might integrate with Nissenbaum's). There is enough flexibility both in how we might interpret Meyer's and Mackenzie's conceptions of autonomy as well as how we might interpret the values latent in various contexts such that there is a reasonable consistency in applying those relational conceptions of autonomy to privacy cases via contextual integrity.

Let me unpack that last sentence a bit. My concern here is the extent to which my view is consistent with Nissenbaum's, at least in the sense that relational conceptions of autonomy can be brought to bear on particular contexts in the manner she describes via contextual integrity. Depending on how we interpret the normative independence of contexts (given her application of Walzer's notion of spheres of justice) it might be the case that there is an inherent inconsistency between applying a normative concept to multiple contexts. This is because there are conceivable scenarios wherein the nature of the context is such that the substantive elements of the particular conception of autonomy (or perhaps the notion of respect for autonomy itself) just do not apply.

This might seem like an odd worry, but I can use part of Nissenbaum's discussion of context to illuminate it somewhat. Nissenbaum makes the distinction between "contextual values versus the values of contexts."[103] The distinction here is between the values contained in contexts, and what society deems to be valuable about the context itself. On her view, and this is again largely based on Walzer, it is entirely possible for a context itself to be "jettisoned" as she puts it, because it does not in some way cohere with the larger collection of contexts and their associated values.

The question then becomes, to what extent might we be committed to Walzer's view of how to evaluate contexts as a whole and/or their associated values, instead of the more straightforward notion of weighing some value against a normative principle such as respect for autonomy? Going back to our example of the school administrator and the context of education, it seems like a safe bet that we would be able to make sense of how respect for autonomy plays a role in evaluating contextually based informational norms. This could be due in part to the nature of the context itself, wherein it is easier to think of education as being highly compatible with respect for autonomy. However, are there other cases where this compatibility is in question? And if so, are we then obligated to view the legitimacy of the values present in that context (or the context itself) as a matter of consistency with society's chosen collection of contexts and values, or can we appeal to the normative principle that we were trying to apply (in this case respect for relational autonomy)?

Consider an example. Analogous to our school administrator considering whether to implement a new student record system, suppose a prison warden is considering implementing a new surveillance system. Suppose that this system is capable of tracking prisoners' whereabouts at all times anywhere in the jail, is capable of reading heart rate and body temperature via some kind of telemetry, and via infrared it is able represent on a monitor a rough outline of the prisoner’s body thereby giving prison guards a sense of what activity the prisoner is engaged in. In short, this would be close to Bentham’s panopticon.

---

[103] Allen, 2011 p. 180

The context would be correctional justice, and some of the values of that context would include prisoner and facility worker safety, as well as rehabilitation. It is clear that this new system would be flagged as a violation of contextual integrity because the transmission principles would be changing drastically. Regarding the values in the context, some of the questions would involve facility worker safety (prison guards, medical personnel, etc.), the balance of prisoner safety relative to rehabilitation, and how we understand rehabilitation. So if the new surveillance system promotes these values more effectively than the current system (depending on how those questions are answered), then the change in transmission principles is allowed for and no violation of contextual integrity is deemed to have occurred.

Remember that one of the steps in the decision heuristic is to apply the more general theories of the value of privacy, such as those that link privacy to autonomy. With that, I might apply Mackenzie’s weak normative competence view here and determine that the experience of constant surveillance disrespects the autonomy of prisoners in virtue of diminishing self-esteem. Furthermore, I could determine that constant surveillance diminishes self-trust which in turn undermines the development of a healthier nonviolent pathology. Finally, I might conclude that the development of a non-violent pathology is a necessary part of what we take rehabilitation to be.

At this point, we would have to return to the relative weight of the values in the context to determine whether the instance of disrespect for autonomy has any real weight or relevance here. If the context is interpreted in such a way that rehabilitation is deemed secondary to worker safety, then it would seem that this instance of violating respect for autonomy has no traction despite its connection to one of the stated values of the context itself.

Furthermore, this result would appear to be consistent with Nissenbaum's adoption of Walzer's view. Normative principles are not supposed to play the role of being normatively informative independent of contextual values. Instead, in Walzer's view, contexts are evaluated according how they sit with other values and contexts, so the question in this context would have more to do with the importance of rehabilitation itself, not whether autonomy has been respected.

To be clear, I have only picked out one consideration in this debate, and it is obvious that there is a great deal more that could be said here about the relative weight of worker safety, prisoner safety (and how that might relate to autonomy even), and rehabilitation. The main point I am trying to get across is that, because of her adoption of the notion of spheres of justice and contextually driven values, it seems possible that there are cases where the normative principles we are supposed to be applying, according to the contextual integrity decision heuristic, may not end up having any normative force at all.

So there is a possible tension between my approach to the value of privacy and Nissenbaum's contextual integrity, resulting from her application of Walzer's spheres of justice. As stated previously, I think that in the vast majority of cases, something as fundamental as respect for privacy will be genuinely compatible, and that cases like the prison example above offer enough room for debate that, once the dust settles, the normative principle in question will have been given at least due consideration even in cases where it does not do any of the actual heavy lifting. I say this because it is important to keep in mind that things like respect for autonomy are only one part of the equation. Hence, I find cases as these acceptable so long as there is the opportunity for extensive deliberation about the relative merits of the normative principles in question, instead of their being set aside strictly for no other reason than because the nature of the contextual values are such that the normative principles in question are not considered at all. If there ever is an instance of the latter kind of case, then I think that there would be genuine reason to question whether my approach is consistent with Nissenbaum's contextual integrity given her adoption of Walzer's spheres of justice.

Conclusion

In this chapter I covered two different aspects of my approach to the value of privacy. First, I discussed the non-normative aspects of the notions constitutive of relational conceptions of autonomy, and how they relate to privacy. My claim here is that privacy can be seen as either promoting or hindering both the skills described by Meyers and the core ideas behind Mackenzie's account such as self-esteem.

Second, I addressed the idea of integrating my theory of privacy with Nissenbaum's contextual integrity. While her theory explicitly leaves room for approaches such as mine in order to incorporate the sorts of normative elements typically associated with privacy's value, it was necessary to address the extent to which this integration really works given her adoption of Walzer's spheres of justice. While I think that there is some degree of compatibility between the two, it will be important to gauge the extent to which values latent in contexts end up being dominant over the normative principles imported by theories such as mine.

The result of this analysis is what I hope to be a more robust approach to analyzing cases involving information privacy, and possibly other forms by extension. Nissenbaum’s contextual integrity has given us a detailed method for analyzing the various components of cases involving privacy, as well as a way to import normative content in order to provide actual prescriptive recommendations. That said, it will be helpful to go through a more elaborate analysis of a particular example involving information technology and associated practices. So in Chapter 6, I will apply the approach combining relational autonomy and Nissenbaum’s contextual integrity to Facebook.

CHAPTER 6

APPLYING THE RELATIONAL APPROACH TO PRIVACY

Introduction

In this chapter I will discuss two examples of privacy issues that I hope will bring together the ideas discussed thus far in this dissertation. Before describing the two examples, a brief summary would be helpful. In Chapter 2, I discussed two different kinds of approaches to the value of privacy. Functionalist accounts focused on the goods enabled by privacy, such as personhood and relationship creation and enhancement. Autonomy-based accounts understand the value of privacy as one kind of respect for autonomy. The goal is to try to develop an understanding of the value of privacy that incorporates the intuitions that drive both of these approaches.

In Chapter 3 I introduced the key idea behind achieving this goal, which is the notion of relational autonomy. Relational conceptions hold the idea of embeddedness as conceptually necessary and do not include individualistic ideals of autonomy whether they are ideals of agency, ideals of autonomous decision making or ideals of respect for autonomy. The claim is that understanding autonomy-based accounts of the value of privacy according to a relational conception allows us to understand respect for autonomy according to the goods described by functionalist accounts

I also described, in Chapter 3, Mackenzie's and Meyers' views of respect for autonomy. Mackenzie's conception of autonomy is described as a weak normative competence view. The core idea is that autonomy is largely constituted by the ideas of self-esteem, self-trust, and self-respect. To respect autonomy in her view is to promote those characteristics in agents. The same sort of idea holds true for Meyers. Instead of self-esteem, she identifies agentic skills that are constitutive of autonomy. Respecting autonomy for her entails promoting those skills.

To tie these ideas together, respect for privacy as a kind of respect for relational autonomy is both an acknowledgment of the goods that make privacy important and respect for *how* we achieve them. That is, to respect someone's privacy is to respect their capacity to pursue goods that are inherently social whether that capacity is conceived of as being constituted by self-esteem or agentic skills. To respect this capacity, we must understand it first as being social in nature, which is the insight that relational conceptions of respect for autonomy provide. In this way we respect agents' ability to exercise autonomous capacity where that capacity is understood as the ability to formulate conceptions of interpersonal goods. Furthermore, we respect the fact that privacy can be a necessary part of how those goods are realized.

Recall from Chapter 4 that I developed a test for understanding whether a particular action, policy, law, etc. constitutes disrespect for autonomy. The test goes as follows: Were I to commit a particular action towards an agent, is there any conceivable competency or skill associated with autonomy that would be diminished on the part of the agent as a result? Simply put, we want to know if taking a particular action has a diminishing effect on an agent's autonomy, where that autonomy is conceived as according to either Meyers' or Mackenzie's conceptions. In cases involving privacy, we are more often concerned about disrespecting autonomy through diminishing skills or competencies than we are promoting them such as in the case of Mackenzie's hypothetical hospice patient Mrs. H. We are preserving the opportunity for individuals to manage the access of others and to manage information about themselves. We provide them the opportunity to manage interactions with others, and the extent to which other individuals or society as a whole intrudes upon their lives. The way in which these interactions and (possible) intrusions play out has a formative and/or sustaining effect on agents' autonomy. Therefore denying agents the opportunity to manage them is tantamount to disrespecting autonomy.

While I take this connection between privacy and relational autonomy to be a genuine one, it needs a second component if it will ever do the work of making any practical recommendations about real life situations involving privacy. In Chapter 2 I discussed the requirement that any complete account of the value of privacy will need to offer some sort of prescriptive recommendation in this practical sense. If we look at what we have so far, it is an emphasis on social factors. These factors are highly dependent upon context. In each conceivable situation involving privacy, what we take to be the social context will be made up of things like informational norms, canonical roles played by various actors, and the values embedded in those contexts. This is where Nissenbaum's notion of contextual privacy comes into play.

At a very minimum, I claim it is necessary to go through the steps laid out by the basic version of her heuristic. That is we need to analyze the constituent pieces of the context and understand how they in turn relate to the normative content presented by relational conceptions of respect for autonomy. In other words, how do things like informational norms, values, and transmission principles connect to things like agentic skills and/or self-esteem? If we understand these connections, we understand how privacy has value in that context and at the same time how to respect the privacy of the individuals involved. Furthermore, we do so in a way that leaves us more likely to make an actual concrete recommendation for how best to respect privacy.

In this chapter I will discuss two examples of how I see this dynamic playing out; but before describing those examples it will be helpful to lay out what they should accomplish or explain. The relational approach to privacy in part claims that privacy can be a means to the development and/or preservation of autonomous capacity. Examples should therefore show *how* this can play out. They should show how privacy as a state of affairs constitutes respect or disrespect of autonomous capacity through the analysis of the context. Keep in mind that these are not wholesale relationships being described here. In other words, I will not embark on an attempt at condemning a particular practice or recommending it wholesale. In both of the cases that will be discussed, there are aspects that can be understood as respectful of privacy and practices that are not. Nor will I attempt to offer analyses of either example that are exhaustive in virtue of dissecting every conceivable aspect of them. The goal of this chapter is to elucidate the relational approach to the value of privacy and demonstrate how that approach goes about analyzing cases involving privacy, and of course the practical recommendations that fall out of that analysis.

The first example is the well-known social network Facebook, and the second one is known as lifelogging. Lifelogging is generally the practice of recording many aspects of a person’s life through the use of multiple technologies such as video and audio recording, data storage of different kinds, as well as the frequent retrieval of that information for various purposes.

As an example, Facebook is meant to point out some of the ways in which every day use of social networks can constitute a significant potential violation of privacy, and some ways in which I take concerns about Facebook and privacy to be overblown. As I will discuss in more detail later, much of the regular use of the site really is not the scary sort of “privacy creep” that people sometimes think of it as. Furthermore, I take this result to be a strength of the relational approach to privacy in that it can, in some cases, demonstrate how privacy concerns are overstated and why. As I will argue, the single biggest issue with Facebook is not so much the potential invasion of privacy it represents, but the lack of material disclosure via informed consent. Users must be made more aware of what use of the site and its associated services actually entails in relational terms so that they can manage their experiences in more effective ways.

The second example of lifelogging I will use here as an extension of the same sorts of phenomena represented by social networks such as Facebook. If we extrapolate the basic idea behind lifelogging as well as online social networks, it is not difficult to see that this is the general direction in which the personal use of technology is headed, and in the not so distant future. The various technologies used to record and share human behavior and thoughts (opinions, instances of self-expression, etc.) are generally coalescing into a more integrated platform and set of services. Hence, lifelogging offers a good example in the sense that it will push contextual integrity to handle privacy issues that cover a broader “context,” in the sense that Nissenbaum uses the term. And it will also show some of the need for identifying privacy’s value in a global sense, in other words in a sense that cuts across multiple contexts.

What I will try to show is that more extensive phenomena such as lifelogging demonstrate a limitation of Nissenbaum’s contextual integrity approach. These phenomena are extensive in the sense that they offer individuals the ability to give away larger amounts of data, a wider variety of data, and disseminate that data to a wide variety of people. As Nissenbaum points out, they therefore do not function within one context, but instead operate as a medium through which different contexts play out. However, in cases like these, we have a need to evaluate the practice itself on a normative level, yet cannot necessarily avail ourselves of context-based values or informational norms. Therefore, there is still a need for analysis of privacy issues outside of the context-based approach offered by Nissenbaum.

What will be the upshot of this analysis of Facebook and lifelogging? The cluster of technologies and practices constituted by online social networks and lifelogging has a great deal of potential to influence the autonomy of individuals. Whether we understand that influence according to Meyers’ agentic skills or Mackenzie’s emphasis on self-esteem and similar values, the importance of privacy is best understood through the relational lens. Maintaining privacy while using these technologies and engaging in these practices is one part of managing the extent to which they affect our sense of self, the ways in which we are able to understand our personal conception of the good, and ultimately who we are and who we can become.

Facebook

Facebook (www.facebook.com) is what is often referred to as a social networking site. The idea of the site is that users should be able to connect with the people currently in their lives, people from their past (schoolmates, former coworkers, etc.), as well as identify future or possible connections. Facebook is, in short, intended to facilitate social interactions. Mark Zuckerberg, the founder of Facebook, often uses the term “openness” to describe the goal and mentality of the site. This of course can be antithetical to privacy hence there is a definite challenge in promoting increased accessibility between over a billion users while at the same time respecting the privacy rights of those users.

Users of the site can create a free account by providing an email address and basic personal information including name, gender, and date of birth; however, the site requires that users be at least thirteen years of age in order to sign up. Once registered, the site creates a homepage for the user. This page contains the personal information that the user has decided to share (this can include things like their hometown, schools attended and names of significant others and family members to name a few), as well as any “updates” or comments that the user has decided to post on that homepage. In addition, users may post links to other sites or pages, photos, videos, join groups based on common interests, and chat with other users who are currently logged on. Users can also invite other users to be “Friends” on the site. A user's friends on Facebook are typically granted greater access to that user's profile, such as posts, photos, contact information, etc. In addition to increased access, Facebook typically automates the posting of a user's comments to their friends' news feeds. This way, it is relatively convenient to keep up with the events in their friends' lives.

As should be clear from this description, users are able to generate and make public a great deal of personal information, both about themselves and other people. Facebook has established a privacy policy that allows users to select certain options or settings that determine how this information will be handled. In what follows I will provide a brief overview of Facebook's current data use policy. The first thing to note is that Facebook has a general “Data Use Policy.” This describes the company’s overall policies about how data is managed within the site itself, and about what gets shared with third parties and why. Policies concerning data management within the site govern the company’s privately managed information technology infrastructure. Parts of this data use policy would be more properly thought of as a confidentiality policy; however, the word “confidentiality” appears nowhere in Facebook’s data use policy. I will return to this issue at a later point.

Second, Facebook has particular “Privacy Settings” that are configurable by each user. These settings are quite intuitively divided up into different categories, including: “who can see my stuff” and “who can look me up.” Each of these is further divided into categories described according to the site functionality that they manage. Many of these settings are organized according to the degrees of separation that a potential viewer has to the user. So “friends” are immediately connected to the user (and must be selected or agreed to by the user to receive that status), friends of friends are second degree connections, and so on outwards until they get to the category of everyone. Users are also able to set up specific lists according to how they want to subdivide their friends (such as school, family, work, etc.).

Facebook also offers a collection of settings that address third party applications (apps) that run on the site; and gives users the ability to block certain apps, or limit their access to user information. Facebook states that it is not responsible for data after it has been transferred to third parties. It is the responsibility of the end user to make herself aware of any additional privacy policies that may apply. However, Facebook frequently requires third parties with whom it partners to meet contractually based standards of data governance. For example, Facebook has what it refers to as an “instant personalization program.” This program provides partner sites the ability to search a user's friends list in order to customize that user's experience when they are visiting the third party site. Bing.com (a search engine) and RottenTomatoes.com (a movie review site) are examples listed in the description of instant personalization on Facebook's site; however, at the time of this writing both require the user to login in order to personalize the experience on the site.

Contextual Integrity and Facebook

Having laid out the very basics of Facebook’s functionality and its privacy and data use policies, I turn now to the kind of analysis that Nissenbaum employs with contextual integrity. Recall that the basic process behind the decision heuristic is first to identify the context (or contexts), the constituent parts of that context (values, actors, transmission principles, etc.), and then to determine what changes are taking place or are constituted by the practice in question. Nissenbaum offers some thoughts on Facebook specifically, which can get us started in this process.[104] As mentioned, the first goal is to make sense of what the context is in the case of Facebook. Nissenbaum considers the notion that there really is no context given the nascence of online social networks. They are, after all, a relatively new phenomenon and hence it may be the case that there simply are no “canonical” activities or roles. She rejects this idea in favor of viewing social networks as media through which we interact, using telecommunications as an example.

---

[104] Nissenbaum, 2010, p. 222

> In a similar vein, one might conceive of the telephone system not as constituting a distinctive context, but as a medium for interactions occurring within diverse distinctive contexts, such as family, workplace, and medical.[105]

So social networks such as Facebook are not the context itself, but are the media through which we interact when we interact with family members, schoolmates, coworkers, etc. Social networking sites facilitate interaction across multiple contexts such as these simultaneously. Individuals may make disclosures that are governed by contextual components (values, transmission principles, etc.) that are part of pre-existing contexts, but that occur in a single medium. For example, a user may lament what an “epic win” last night’s party was, and intend that message to be seen by her peers and understood in that particular context.

Further, Nissenbaum claims that this notion of social networking sites functioning as medium instead of context explains why people are often surprised by the results of certain disclosures that they make online. For example, it has become fairly common for employers to do searches of candidates for their open positions as part of the interview process. This practice has resulted in employers discovering what we might call “unprofessional” posts, photos, comments, etc. by those candidates, and in turn the candidates being rejected by the employers. Many times, candidates are in fact surprised that employers would conduct such searches, and feel wronged that the information was used as part of the evaluation process.

---

[105] Nissenbaum, 2010, p. 223

Nissenbaum’s analysis in cases such as these is that the individuals in question (i.e., the candidates) treated certain disclosures, say having their picture taken while obviously intoxicated at a party, as having occurred in a particular context. This context in turn is one in which its values made the nature of the photo acceptable (i.e., it was shared among friends). Furthermore, they see the search conducted by the employer as a surprise because the assumption, however incorrect, was that the photo would stay in that context and thereby be judged accordingly. In short, people simply fail to realize the reach of social networks, and the extent to which they constitute a medium that reaches far across many of the contexts in which they interact.

Thus far, the upshot of our contextual integrity analysis is that Facebook functions as a communication medium across multiple contexts, and that this fact can explain (when not understood properly) why some individuals experience surprise and/or feel treated unfairly when a particular disclosure migrates from one context to the next. The next step is the normative component.

What we want to look at here are the possible ways in which Facebook can impact the actors, transmission principles, and attributes of the particular context in question. Clearly, Facebook and social networks like it offer a change to the transmission principles. The opportunity for and nature of disclosure seems quite different than previous modes of communication. And this change has the potential to alter the way in which those certain relationships may play out.

> In social contexts of friendships and family, there is a morally relevant difference between knowing information is available with some effort, having it specifically shared with you by your friend, and having it scroll unavoidably across your screen. Bonds of trust, crucial to the myriad other duties and obligations of

> kinship and friendship, are one of many values supported by norms of information flow.[106]

In other words, how information flows within relationships, whether of friendship, family or otherwise, impacts how those relationships play out; and sudden and drastic changes need to be evaluated as potentially damaging as well as potentially helpful.

> Understanding the ways that norms of information flows relate to values, ends, and purposes of social contexts is crucial to judgments of whether novel flows are acceptable, and if not, constitute reasons for resisting change and weighing in favor of entrenched norms.[107]

We need to understand the impact of changing transmission principles through the lens of how the management of relationships will be affected. And as Nissenbaum points out, in order to draw substantive conclusions, we would “need to elaborate and demonstrate key dependencies.”[108]

I take this last point to mean that more would need to be said about the nature of those relationships themselves, how they affect people, and how people seek to manage them. In other words, the contextual integrity approach has done its initial job of analyzing what is really changing relative to the contexts involved. And it has given reason to be wary of the changes constituted by social networking sites like Facebook. It is at this point that it would be beneficial to bring in one of the “stratospheric” or normative theories such as the relational approach to the value of privacy. Something like my approach, which ties in relational factors, would seem well suited to fleshing out how and why we should really be concerned about the changes in transmission principles in this case. This will be the focus of the next section.

---

[106] Nissenbaum, 2010, p. 228
[107] Nissenbaum, 2010, p. 228
[108] Nissenbaum, 2010, p. 228

Facebook and the Relational Approach to Respect for Privacy

Meyers' and Mackenzie's conceptions of relational autonomy occupy a central space in the continuum of relational conceptions. Insofar as substantive constraints are concerned, both of these approaches are more constraining than strictly procedural conceptions such as Christman's, and less so than Oshana's. Primarily for that reason I will be considering how each of these views would apply to the case of Facebook's data use privacy.

I will start with Meyers' approach. What we are concerned about here is respect for the skills that generally make up autonomous capacity. These skills include introspective skills (having a sense of how well one understands oneself), communication skills, memory skills, imaginative skills (the ability to see oneself otherwise and thereby develop goals and aspirations), and others. At first glance, it may seem that viewing the average person's experiences on Facebook through this lens is essentially overstating the significance of it. After all, your typical user simply posts random things when she feels the desire to do so, peruses the posts and links on her friends' timelines, and not a whole lot else ever comes of it. I think that to a certain extent this is true. People use social networks like Facebook to varying degrees, and in some cases, the impact of the technology on their autonomy is negligible or even nonexistent.

However, because of the scale of Facebook (i.e., the number of users and the fact that they are distributed globally) as well as the wide range of possible kinds of interactions one can have on the site, there is a great deal of potential (and plenty of actual) experiences that can be very formative of an individual's agentic skills. To get at this dynamic in more detail, I will return to the example of Carol Gilligan, discussed previously in Chapter 4. What if Facebook had been around during the time that Carol Gilligan had been working her way through graduate school? It is not too hard to see how having access to social networking sites such as Facebook could have been highly beneficial for someone in her situation, as she may have been able to find a greater base of support, and a larger forum for her ideas. This would have reinforced her ability to envision alternatives to the picture being presented to her by her male professors, and could have exposed her to even more possibilities that she would not have considered otherwise. We see real life examples of this every day in the form of groups focused on particular interests and/or social causes.

Then again, this all depends on the nature of the ideas to which she is exposed, and on the nature of the interaction she has with other individuals on the site. Social networking sites increase the scope, and in some cases, the intensity of human interaction, for better or worse. In this way we can see the potential impact on agentic skills generally speaking. However, in some cases the impact on agentic skills does not involve crossing contexts, in the sense discussed in the previous section. In those cases, what we have is the impact of social networks on autonomy in a more general sense not directly involving privacy. In these cases our interactions with others through social networks may impact us on the level of agentic skills, but not in a way that constitutes a violation of privacy. However, there are other cases where individuals disseminate information in one context, and do so with the understanding (or hope) that this information will somehow remain in that context. In many of those cases, when that information crosses over into a different context, harm can be done and there can be a perception of a violation of privacy.

The lessons to be drawn here are first that there is a great deal of potential for Facebook to impact agentic skills. Second, this can often occur as a result of misguided expectations about privacy and how users genuinely expect information to be managed. Hence the relational approach to respect for privacy suggests that we understand the transmission principles constituted by the use of social networks according to their potential impact on agentic skills. Furthermore, it suggests that there is a strong obligation to inform users of transmission principles in light of their potential impact on agentic skills. In other words users should be made aware of how information may flow from party to party. This should be done in such a way that users are encouraged to think of "party" by the nature of their relationships with the individuals who make up "party" or group category on Facebook. At the time of this writing, Facebook offers no such guidance.

There is a similar point to be made when we apply Mackenzie's conception of respect for autonomy. Social networking sites have a great deal of potential for impacting such things as self-esteem, self-trust and self-respect much in the same way that they can impact agentic skills. There has been a great deal of work done by psychologists about how the use of social networks can impact self-esteem. A recent study by Gonzales and Hancock found that the opportunities for the sort of selective self-disclosure offered by Facebook actually enhances self-esteem in that it gives individuals the opportunity to review and select positive aspects of themselves.[109] There have also been studies attempting to understand the links between such things as self-esteem and the number of friends individuals tend to have on Facebook ("compensatory friending" and some people do compensate by adding more friends).[110]

---

[109] Gonzales, Amy L. and Jeffrey T Hancock. Cyberpsychology, Behavior, and Social Networking. January/February 2011, 14(1-2): 79-83. doi:10.1089/cyber.2009.0411.

Yet another study by Swedish researchers identified a negative correlation between Facebook usage and self-esteem. With greater use came lower self-esteem, and more notably in female subjects than male.[111] Studies such as these become an important part of how we would want to understand respect for privacy in social networks, if we take things like self-esteem, self-trust, and self-respect to be a necessary component of how we understand autonomy.

The logic goes something like this. There is an obligation to respect the self-esteem of agents. Privacy plays a part in how self-esteem is impacted through interaction and how it is managed by individuals. Individuals should be made aware of how social networks like Facebook operate, and they should be made aware of how information is likely to flow within those networks. Again to use Nissenbaum's terminology, the transmission principles should be explained to users in such a way that is cognizant of the potential impacts on self-esteem. Users should be made aware of what constraints the site places on the flow of information from party to party; and of course these constraints should be honored. Hence what the relational approach tells us is that we have an obligation to understand how the preservation or loss of privacy on social networking sites can affect self-esteem. Furthermore, that particular relationship is what should inform whether we take a particular policy or feature of Facebook to be respectful of privacy. Clearly there is an empirical (i.e., psychological and sociological) component to this. For that reason there is a need for a lot more research about how the relationship between social networking and self-esteem actually works. But again, the point here is to understand the direction in which the relational approach is taking us.

---

[110] Lee, Jong-Eun Roselyn, David Clark Moore, Eun-A Park, Sung Gwan Park, "Who wants to be 'friend-rich'? Social compensatory friending on Facebook and the moderating role of public self-consciousness" Computers in Human Behavior Volume 28, Issue 3, May 2012, Pages 1036–1043

[111] See http://psychcentral.com/blog/archives/2012/10/04/facebook-happiness-and-self-esteem/

Regarding Meyers' and Mackenzie's approaches taken together (i.e., the relational approach), the lesson is that we should base our concerns about privacy and social networks first on the content of the conceptions of autonomy themselves and second on how the flow of information works relative to the context. We get a different list of concerns depending on which conception of autonomy we go with. However, the core idea is that agents should be encouraged to think through the context in which they will be using Facebook as a medium. Second, they should be encouraged to think proactively about the consequences of disseminating certain types of information in light of the fact that there is a strong likelihood that such information will cross over into other contexts.

Lifelogging and the Relational Approach to Respect for Privacy

Social networking sites like Facebook are one part of a larger trend driven by advances in information technology. Devices like mobile phones, personal computers, wearable cameras and cameras embedded in phones, etc. all contribute to the ability of individuals to record their surroundings and experiences. Add the Internet and they are then able to share those recordings with many people instantaneously and retrieve them at will. The practice of consistently recording one's experiences and aspects of one's state (also including such things as blood pressure, heart rate, etc.) is generally known as lifelogging. The idea is to keep a sort of constant and comprehensive journal that can then be shared with others and/or viewed by the individual for reflection and insight into past experiences.

As mentioned in the introduction of this chapter, I take the general practice of lifelogging to be indicative of where the personal use of information technology is headed, primarily because various services and hardware are being integrated and converging on an ever increasing ability to offer lifelogging capabilities. For this reason, lifelogging makes a reasonable choice as a second example through which to consider the relational approach to the value of privacy.

Anita Allen discusses lifelogging at length in her work *Unpopular Privacy*, although her focus is on the question of whether there are reasonable and fair forms of privacy related to lifelogging that the government should enforce on citizens. My concern here is not about possible paternalism, but some of her insights into the issue bear mention nonetheless. On the one hand, there are certain appealing things about lifelogging becoming a common practice. There are many sorts of crimes that may not be committed due to the deterrence of surveillance. People can greatly benefit from the opportunity to review and reflect on past periods of life, much in the same way people currently do with journals. Inheriting a lifelog from a deceased relative could provide the opportunity to heal and remember in better and more effective ways.[112]

On the other hand, there are genuine worries about the extent to which individuals are giving away information about themselves. Allen summarizes the concerns by stating that “the potential would be great for incivility, emotional blackmail, exploitation, prosecution, and social control surrounding lifelog creation, content, and accessibility.”[113]

---

[112] Allen (2011) p. 170

[113] Ibid.

These worries are not too hard to envision. Think of the practice of electronically recording and storing every activity every day as becoming ubiquitous. Then, think of how unscrupulous companies, individuals bent on self-promotion, criminals and the like could conceivably make use of such information. The possibilities are endless, in the worse kind of way.

With pros and cons on the table, I would first like to consider the application of contextual integrity to lifelogging. Initially we could make a claim similar to the case of Facebook: as a general practice, lifelogging is more the medium than the context. This would seem to be true to some extent. Our concerns about privacy in the case of lifelogging would appear to have more to do with the particular ways in which it was put to use, and particular instances of how some recording either fell into the wrong hands or was taken in some invasive way.

I think that there are some insights to be gained from making this move; however, it is not clear how contextual integrity is supposed to help us talk about lifelogging as a general practice. It seems that, insofar as Nissenbaum’s approach is driven by some version of independent spheres of justice, making any sort of evaluative claim about the practice would be fraught at best. Whether or not some instance is a prima facie violation of entrenched norms and therefore something to worry about simply depends on how it is used in which context.

For that reason, it would be helpful to consider the practice from the “stratosphere” as it were and evaluate it along relational lines. I think many of the same observations that were made regarding the Facebook example apply here at least insofar as individuals have the opportunity to disseminate a wide variety of information about themselves to a wide variety of people. It will not be long before we can all upload the last five minutes of our lives to a large network like Facebook, regardless of whether we remembered to record it with the video camera in our phone. The potential impacts on our ability to envision alternatives for ourselves could be greatly impacted for better or worse, as well as the potential impact on our self-esteem.

But in the cases of both Facebook and lifelogging, there is an underlying worry about how they cause us to live in an environment where the stakes are constantly so high. With both sets of technologies, there is a significant increase in the potential for both good and bad effects. Regardless whether we think of these as agentic skills or self-esteem, self-trust, and self-respect, there seems to be a strong sense that we would be under constant bombardment of situations wherein the consequences are very significant. There would be an omnipresent possibility of social disaster. In this sense, being afforded privacy controls means being afforded the opportunity to interact on a scale and at a pace that is far less stressful and much more manageable. We need opportunities for rest and respite from being under this constant pressure.

Notice that context does not play a role in this concern. It is a global quality of the practice that is at issue here. So while I think the contextual integrity leg of the analysis would be important, being able to evaluate lifelogging independently of context is crucial as well. In fact, the kind of worry that I sketched out in the previous paragraphs may not even come up in a strictly contextually based approach. It is a concern that essentially supervenes on multiple contexts.

This leads us back to the question brought up during the initial discussion of contextual integrity in Chapter 5, which is the primacy of the spheres of justice versus comprehensive normative theories. It might very well be that certain contexts speak in favor of lifelogging, and that others are a mixed bag as it were. If this turns out to be the case, do we then consider the practice as not being a violation of contextual integrity and therefore allowable regardless of concerns brought up by respect for autonomy?

As mentioned in the previous section on Facebook, there is a great deal of work in psychology that seeks to understand the connection between use of social networks and self-esteem. There also many cases of individuals and groups of individuals using social networks to relate over particular issues but who do so relying on contextual integrity. So there are clearly empirical issues in play here incorporating psychological and sociological elements. The benefit of the relational approach is that we know where to focus our concern, and that we have more specific questions to ask. Furthermore, it allows us to understand the benefit of going beyond the need for informational control than what is offered by more elaborate privacy settings. The relational approach highlights the importance of disclosing to users the potential dangers of using these services. It directs us to encourage them to think through the experiences that they hope to have on a relational or interpersonal level. To round out my discussion of social networks and lifelogging, I turn now to confidentiality and obligations associated with third parties.

Confidentiality, Facebook, and Lifelogging

At this point I would like to turn to an issue that was initially discussed in Chapter 2, confidentiality. Recall that Francis made the distinction about how information is gathered initially (and whether we would consider that invasive or not) as opposed to what is done with that information. Furthermore, the claim was made that it is important to keep these two processes (and their associated moral evaluations) separate for the sake of properly understanding certain policies or practices. For example, in cases of healthcare policy we do not want worries about the minor invasiveness of a pin prick or mouth swab to deter us from the possible larger public health benefits of gathering the information in the first place.

In the cases of Facebook and lifelogging however, it is possible for this distinction to work in a different way. The experiences of using Facebook, or of initially recording some episode of our lives, are not the sorts of things that we would consider invasive in and of themselves. But we do not want that fact to cause us to lose sight of the importance of the confidentiality of the information that they generate. As mentioned above, Facebook's data use policy establishes guidelines for how the company will handle the information users generate through use of the site. And again, this information can vary greatly in its personal significance to the user.

So the first thing to note here is that Facebook, and any hypothetical lifelogging service, have definite obligations towards their customers centering on confidentiality. Furthermore, the nature of those obligations turns on how we understand both the significance of the information to the individual(s), as well as an obligation to inform them of how the information will be managed. Second, as per my modification of Francis' point, we should not let the typically innocuous nature of the initial gathering of that information lull us into a diminished sense of the importance of that information.

Before evaluating these in more detail, I would like to appeal to some of the work done by Chris Hoofnagle. In some recent surveys, Hoofnagle inquired as to the general public's level of understanding of the online privacy and user tracking policies set by private companies.[114] He has also investigated general sentiment regarding those policies, specifically what individuals would like to see as part of the Do Not Track initiative (DNT).[115] Primarily a set of guidelines put forth by the Federal Trade Commission (FTC), DNT provides guidelines for providing consumers the ability to opt-out of advertisers tracking of consumers' Internet activity.

Hoofnagle surveyed a large number of Internet users first as to whether they were aware of DNT generally, and second what they would want something like DNT to do. The results are pretty telling. The majority of those surveyed (87%) were not aware of DNT at all, and 60% expressed an interest in DNT forcing companies not to track their activities online. Only 14% of the respondents expressed an interest in DNT requiring companies not to tailor online ads specifically to their Internet history or activity. A general prohibition against tracking is closest (of the survey options provided to respondents) to the version of DNT suggested by the FTC. The version prohibiting tailoring of ads was closest to what was suggested by the Interactive Advertising Bureau (IAB), which is a nonprofit group supported by the largest companies involved in online advertising.

Hoofnagle also investigated how well informed users are of existing regulations regarding online advertising. For example, respondents were asked whether companies are required to get an individual's permission before tracking their activities across the

---

[114] Hoofnagle, Chris Jay, Urban, Jennifer M. and Li, Su, "Privacy and Modern Advertising: Most US Internet Users Want 'Do Not Track' to Stop Collection of Data about their Online Activities" (October 8, 2012). Amsterdam Privacy Conference, 2012. Available at SSRN: http://ssrn.com/abstract=2152135

[115] This was at one point a bill (not enacted) sponsored by Sen. Rockefeller from West Virginia in 2011. The bill was referred to committee and died there. However, the actual content of the DNT policy came from staff at the FTC. Hoofnagle discusses additional versions, such as the one created by the Interactive Advertising Bureau (IAB).

Internet; 48% responded incorrectly by saying that this was true, and 19% responded "Don't Know/Refused." And finally, responding to the statement "When you use the Internet to learn about medical conditions, advertisers are not allowed to track you in order to target advertisements," 22% said "Yes," 36% False, and 41% "Don't Know/Refused."

There are a number of different conclusions one might try to draw from Hoofnagle's research, and there is of course room for debate here. However, I think the most straightforward conclusions are that users of the Internet are generally misinformed about existing regulations; and they have interests regarding regulations that are more in line with aggressive regulation than what has been proposed by the online marketing industry itself. Hence, there is a disconnect between consumer perception, consumer preference, and actual corporate behavior when it comes to online advertising practices.

Going back to data use and confidentiality in social networks and lifelogging, it seems clear that this sort of disconnect described by Hoofnagle constitutes some significant challenges. Internet users are a) generally ignorant of what is allowable and what is not, and b) have preferences about what should be allowable that seem very far apart from reality. If we were to add to this current state of affairs a significant increase in the extent to which technologies track our activities and hence the information we give up, then the net result would be very far removed from any kind of respect for autonomy, much less confidentiality. I will return to this point at the end of this section.

Another layer of complexity regarding confidentiality has to do with personally identifiable information (PII), nonpersonally identifiable information (non-PII) and "big data." Just as it sounds, PII is information about an individual that is detailed and/or comprehensive enough to allow a third party to identify exactly who that individual is. Non-PII would be information about a person that is not comprehensive enough to allow such identification to be made. Big data typically refers to very large amounts of data sets (exabytes, or one billion gigabytes) and the data that is generated in virtue of analyzing those data sets. Over the last several years, the technology has gradually become available to analyze those large data sets very rapidly and to provide corporate decision makers visualizations of those data.

There is a growing trend in business to make use of big data. Partly as a result of this trend, pretty much every trackable bit of data possible will be gleaned from the activities of users online, including which sites they surf, in some cases the details of their interactions with those sites, and “trend data” such as how often they engage in those activities. The same sorts of data would be conceivable in the case of lifelogging as well. What sorts of files are uploaded and when, any detectable content of those files (such as video, audio, text, etc.) and the like would all be crunchable by big data operations. On a practical level, this information is seen as incredibly valuable because of the speed at which it is gathered, and the extent to which it can generate actionable insights.

Up to this point we have two main concerns. First, there is essentially no limit nowadays as to the amount and kind of data that can conceivably be harnessed by any number of actors (and as a result of this, the potential breaches of confidentiality are limitless as well). Second, as Hoofnagle’s work shows us, consumers are largely ignorant of regulations and have expectations that differ extensively from actual practice. So where does that leave us? Starting with large sets of non-PII, there is the potential for a much more rapid kind of feedback loop between consumers and producers of pretty much any conceivable product or service. For example, the online video provider Netflix based its decision to produce the show *House of Cards* (starring Kevin Spacey) on big data. Essentially, they tracked several different trends such as which actors and directors were popular, what kinds of shows tended to be watched and when, etc. Based on that data, they produced *House of Cards* which has turned out to be the most watched TV series ever on Netflix.[116]

Assuming that actions and decisions like this ultimately result from non-PII data, it would seem that there are not very significant confidentiality concerns. I base this claim on the how I defined non-PII earlier. There is no way to discern the identity of the individuals from whom it has been aggregated. I say this with the understanding that there are shades of grey here; however, I will not be able to address that complication in this dissertation.[117] There is the related worry, although not ultimately a matter of privacy or confidentiality, of how we may be influencing ourselves to a scary degree (and the amount of control this may give certain parties). However, in cases involving PII, the potential confidentiality disasters are staggering. Large amounts of very intimate data could conceivably be handled in any sort of way, such as being sold, put on display inadvertently, used for a wide variety of immoral purposes, and so on. So the trend is that users are giving away more and more data about themselves, of ever increasing variety, they are less and less aware of third parties' obligations about that data, and what they *hope* will happen is getting farther away from what *is* happening.

---

[116] See "The ''Big Data' Revolution: How Number Crunchers Can Predict Our Lives" http://www.npr.org/2013/03/07/173176488/the-big-data-revolution-how-number-crunchers-can-predict-our-lives

[117] I do not mean to downplay the importance nor the complexity of the distinction between PII and non-PII. There is a great deal of work being done on how those categories should be understood, especially with the advent of cloud computing. Currently, it is not at all clear what really constitutes a personal record. Is it a single file about an individual, or bits of information spread across different platforms and services that can be recalled because of how they are linked? Even though I am bracketing this issue, I hope that how I have presented the distinction between PII and non-PII speaks to the importance of determining a hard line between the two.

The upshot of all of this is that it is important for the parties responsible for social networks and lifelogging services to inform users of what will become of the information that they generate. This is because users must be able to evaluate what the use of these services means in a way that connects to competencies and agentic skills, and they cannot do so without a proper understanding of how both PII and non-PII will be handled.

Conclusion

Whether the issue is confidentiality or privacy, the net result of social networks, lifelogging, and big data are that individuals are constantly pressed to represent themselves not just publicly by being seen or heard, but also by being evaluated and influenced. In relational terms, our agentic skills and our sense of self-esteem are constantly being manipulated and tested in this kind of environment. I take the core ideas behind relational conceptions of autonomy to speak to this directly, and essentially to say there is a great deal of potential danger here. Furthermore, I take the central insight behind the relational approach to be that respect for privacy and confidentiality should be understood by how use of these technologies impacts competencies and agentic skills. If we understand the impact on autonomy, we understand how the opportunity to manage our interactions becomes important whether those interactions involve privacy directly or not.

CHAPTER 7

CONCLUSION

The goal of this dissertation has been to a) construct a view of the value of privacy that incorporates functionalist intuitions and autonomy-based intuitions in such a way that b) allows us to understand the value of privacy as being commensurable with other moral considerations. As previously mentioned, what we want is a view of the value of privacy that is complete in the sense that it addresses both of the primary views (functionalist and autonomy-based) within the existing literature. But we also want to do this in a way that allows us to make ready use of that account when it comes time to tackle real life applied issues.

I consider the relational approach that I have outlined here to be, at the very least, a good first step in this direction. Privacy is an inherently relational concept. Whatever form of inaccessibility or separation it takes, privacy only makes sense in the context of our relationships with others, as well as in the context of our relationship with the rest of society and the cultural values of that society. Hence, both the goods (and harms) and the autonomous preferences that we would associate with privacy are inextricably bound to

relational factors. Making sense of privacy without relational or social concepts is a bit like trying to make sense of silence without sound.

It follows from this that respecting privacy as a form of respecting autonomy only makes sense if the underlying notion of autonomy captures, or is able to make sense of, those relational factors. Furthermore, by following an approach that is both relational and autonomy-based, we are able to avail ourselves of existing conceptual apparatuses within the literature on autonomy, such as autonomy's connections with political notions, which provide additional insights into the various ways in which privacy can and should be respected.

While I take the relational approach to the value of autonomy to be effective as a global theory about privacy's value, it still needs to be augmented with sensitivity to context. This is the benefit I see from incorporating Nissenbaum's work. Our expectations about privacy vary from one situation to the next. Our sense of rights, obligations, and even our hopes about what privacy should do for us are in large part driven by details and existing conventions. Nissenbaum's contextual integrity provides a very effective way to capture those concerns, and does so in a way that is highly compatible with the relational approach.

As stated in the introduction of this dissertation, cases involving privacy naturally draw in concerns about other moral issues. We could ask a lot of questions about the case of Gen. Petraeus; and some of those would center on the value of privacy itself, while others would incorporate additional concerns such as how individual liberty stacks up to public welfare. It is my contention that we have little chance of resolving questions about privacy's value unless we figure out a robust and systematic way of connecting privacy to those sorts of concerns. I consider the relational approach to be a good first step in this direction because of its use not just of autonomy, but of a particular conception of autonomy. Relational autonomy is well discussed in its own right. Philosophers such as Christman, Benson, and Stoljar have done a great deal of nuanced work considering how the idea of autonomy plays out in other areas. It is this kind of work that strengthens the relational approach to the value of privacy because making these sorts of connections lessens the extent to which privacy must be treated as an incommensurable good or right.

What I take these points to speak to, at least insofar as a theoretical view of the value of privacy is concerned, is the importance of first starting with a more thorough and comprehensive understanding of the value of privacy and then connecting that understanding to other considerations in such a way that allows us to explore the full importance of privacy's value without losing anything in translation. There is no need to treat privacy as a riddle wrapped in a conundrum as it were; its value is quite accessible.

SELECTED BIBLIOGRAPY

Alderman, Ellen, and Caroline Kennedy. *The Right to Privacy*. New York: Knopf, 1995.

Allen, Anita L. *Uneasy Access: Privacy for Women in a Free Society*. Totowa, NJ: Rowman & Littlefield, 1988.

Andrews, Lori B. *I Know Who You Are and I Saw What You Did: Social Networks and the Death of Privacy*. New York: Free, 2012.

Barvosa-Carter, Edwina. "Mestiza Autonomy as Relational Autonomy: Ambivalence & the Social Character of Free Will." *Journal of Political Philosophy* 15.1 (2007): 1-21.

Beardsley, Elizabeth. "Privacy: Autonomy and Selective Disclosure." *Nomos* 13th ser. (1971).

Beauchamp, Tom L., and James F. Childress. *Principles of Biomedical Ethics*. Oxford: Oxford UP, 2001.

Benn, Stanley. "Privacy, Freedom and Respect for Persons." *Nomos* 13th ser. (1971).

Benson, Paul. "Feminist Second Thoughts About Free Agency." *Hypatia* 5.3 (1990): 47-64.

Benson, Paul. "Free Agency and Self-Worth." *The Journal of Philosophy* 91.12 (1994): 650-68.

Benson, Paul “Feminist Intuitions and the Normative Substance of Autonomy” in *Personal Autonomy.* Taylor, 2005.

Christman, John. "Constructing the Inner Citadel: Recent Work on the Concept of Autonomy." *Ethics* 99.1 (1988): 109.

Christman, John. “Autonomy and Personal History.” *Canadian Journal of Philosophy* 21 (1991): 1-24

Christman, John. "Relational Autonomy, Liberal Individualism, and the Social Constitution of Selves." *Philosophical Studies* 117.1/2 (2004): 143-64.

Dept. of Justic. "The Dru Sjodin National Sex Offender Website." *United States Department of Justice National Sex Offender Public Website*. Department of Justice, 05 Jan. 2013.

Dworkin, Gerald. "Autonomy." *A Companion to Contemporary Political Philosophy*. Oxford, UK: Blackwell, 1993.

Facebook. "Sign Up." *Welcome to Facebook*. Facebook, 4 Feb. 2004. Web. 05 Jan. 2013.

Gavison, Ruth. "Privacy and the Limits of Law." *The Yale Law Journal* 89.3 (1980): 421-71.

Gonzales, Amy L. and Jeffrey T. Hancock. Cyberpsychology, Behavior, and Social Networking. January/February 2011, 14(1-2): 79-83. doi:10.1089/cyber.2009.0411.

Google. "Google." *Google*. Google Inc., 28 June 2011. Web. 05 Jan. 2013.

Hoofnagle, Chris Jay, Urban, Jennifer M. and Li, Su. Privacy and Modern Advertising: Most US Internet Users Want 'Do Not Track' to Stop Collection of Data about their Online Activities (October 8, 2012). Amsterdam Privacy Conference, 2012.

Jordan, Judith V. "A Relational Perspective." *Women's Growth in Diversity: More Writings from the Stone Center*. New York: Guilford, 1997.

Kukathas, Chandran. "Liberty." *A Companion to Contemporary Political Philosophy*. Oxford, UK: Blackwell, 1993.

Lane, Frederick S. *American Privacy: The 400-year History of Our Most Contested Right*. Boston, MA: Beacon, 2009.

Lee, Jong-Eun Roselyn, David Clark Moore, Eun-A Park, Sung Gwan Park. “Who wants to be “friend-rich”? Social compensatory friending on Facebook and the moderating role of public self-consciousness.” *Computers in Human Behavior* Volume 28, Issue 3, May 2012, p. 1036–1043.

LinkedIn. *LinkedIn*. LinkedIn, 5 May 2003. Web. 5 Jan. 2013.

Mackenzie, Catriona. "Relational Autonomy, Normative Authority and Perfectionism." *Journal of Social Philosophy* 39.4 (2008): 512-33.

Nissenbaum, Helen Fay. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford, CA: Stanford Law, 2010.

Oshana, Marina A. L. "Personal Autonomy and Society." *Journal of Social Philosophy* 29.1 (1998): 81-102.

Peterson, Eric "Branded" *Salt Lake City Weekly* April 26, 2007 Vol. 23 No. 49

Premack Sandler, Eleana. "Promoting Hope, Preventing Suicide." *Privacy and Mental Illness, Part II*. Psychology Today, 27 July 2011. Web. 05 Jan. 2013.

Robb, Christina. *This Changes Everything: The Relational Revolution in Psychology*. New York: Farrar, Straus and Giroux, 2006.

Solove, Daniel J. *Nothing to Hide: The False Tradeoff between Privacy and Security*. New Haven, CN Yale UP, 2011.

Solove, Daniel J. *Understanding Privacy*. Cambridge, MA: Harvard UP, 2008.

Stallman, Richard. "On Hacking - Richard Stallman." *On Hacking - Richard Stallman*. Richard Stallman, 1 Jan. 2002. Web. 05 Jan. 2013.

Stoljar, Natalie, and Catriona Mackenzie. *Relational Autonomy: Feminist Essays on Autonomy, Agency, and the Social Self*. New York: Oxford UP, 2000.

Stoljar, Natalie. "Autonomy and the Feminist Intuition" *Relational Autonomy* Stoljar & Mackenzie ed.'s Oxford Univ. Press 2000.

Taylor, J.S. (ed.), 2005, *Personal Autonomy: New Essays on Personal Autonomy and its Role in Contemporary Moral Philosophy*, Cambridge: Cambridge University Press.

Tietjens Meyers, Diana. "The Rush to Motherhood: Prenatal Discourse and Women's Autonomy." *Signs* 26.3 (2001): 735-73.

Warren, and Brandeis. "The Right to Privacy." *Harvard Law Review* 4.5 (1890).

Westin, Alan F. *Privacy and Freedom*. New York: Atheneum, 1967.

Westlund, Andrea C. "Rethinking Relational Autonomy." *Hypatia* 24.4 (2009): 26-49.

Young, Robert. "The Value of Autonomy." *Philosophical Quarterly* 32.126 (1982): 35-44. Web.