Technologies of the Self

Diachronic Privacy and Self-Presentational Autonomy

Helena Ward

2023

University College London, Department of Philosophy

A thesis submitted for the degree of Philosophical Studies MPhil Stud

I, Helena Molly Elizabeth Ward, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated.

1st September 2023

Abstract

This thesis is concerned with our interests in self-presentation: in having a measure of control over which aspects of our selves – our bodies, thoughts, how we style our appearances – we allow others to observe, and on what terms. I consider our self-presentational interests in relation to our right to privacy, privacy being one of the main ways in which we retain self-presentational control. And in light of technologies of the self, which have transformed the ways in which we construct our own identities and observe other people's. The first aim of this thesis is to vindicate a distinction between the *synchronic* and *diachronic* aspects of self-presentation. Sociological and ethical accounts of self-presentation adopt a Synchronic Approach, I argue that this approach is incomplete – our self-presentational interests also have a diachronic character. In short, it is not only important that we are able to exercise a measure of control over how we present ourselves to others at any one time, but also that we're able to exercise a measure of control over other people's access to our past self-presentations. These are our diachronic self-presentational interests (DSPIs). The second aim of this thesis is to show how technologies of the self undermine our DSPIs. While these technologies do not single-handedly thwart our DSPIs, they undermine our diachronic self-presentational autonomy by introducing social customs and individual incentives to act in ways that are negligent, short-sighted, hostile, or in various other ways insufficiently attentive to our DSPI-related responsibilities. I argue that the nature and extent of this threat has been underestimated because of the implicitly Synchronic Approach. Our DSPIs provide us with some grounds for a right to diachronic privacy – in keeping our childhood foibles, jejune opinions, and bygone bodies to ourselves.

Impact Statement

University College London requires each thesis to submit an 'Impact Statement' describing 'how the expertise, knowledge, analysis, discovery or insight presented in the thesis could be put to beneficial use'.

My central theoretical aim is to propose and defend a hitherto unexamined distinction between synchronic and diachronic self-presentational interests; by doing so, I fill a significant gap in the existing literature on self-presentation. While there has been ample work discussing the diachronic character of autonomy, the diachronic character of privacy and self-presentational autonomy has not been emphasised in moral and political debates about privacy. The fact that these connections between autonomy, diachronicity, and privacy, have been under-explored, means that this thesis opens up a range of new and illuminating ways of thinking about both self-presentational autonomy and privacy. Some of these new ways are described in this thesis; others will be explored in future works on online friendships, shaming, and the ethics of social networking services.

While written for an academic audience, the topics of this thesis bear on puzzles, concerns, and ways of life that will be familiar to many beyond academia. Being able to control which aspects of ourselves we reveal to or conceal from others is a morally important capacity, its import recognisable to anyone who cares about what other people think about the kind of person they are. Ordinary reflection reveals this to be a common vein: many of us have been embarrassed at revealing more than we ultimately would have liked and recognise aspects of our internal monologue we would rather keep to ourselves.

The chapters of this thesis which address privacy, address topics on which a great deal has been written. Public discourse has repeatedly and resoundingly warned us of the ills of social media, and privacy is frequently identified as one of the most pressing ethical factors that should make users of social media cautious. By exploring privacy's role in facilitating our diachronic self-presentational autonomy, I raise a new and important challenge to the role that social media plays in our online worlds. In doing so, I hope to elucidate and galvanise our pre-theoretical intuitions about privacy.

Acknowledgements

Thanks are due to my supervisor, Dr Robert Simpson, who provided me with invaluable support throughout the duration of this thesis. The contents of this thesis have benefitted considerably from iterations of feedback, as well as insightful supervisions from which many examples within this thesis emerged. I am in addition incredibly grateful for generous financial support from University College London which allowed me to undertake the MPhil, including a Jacobsen Scholarship and Keeling Bursary.

Table of Contents

Introduction

Percival drowsily tunes in to the sound of Radio Four. He lifts his head from his pillow to face his radio alarm clock. The time reads 6:30 AM. Percival thinks of how pleasant his morning routine is; he despises those jarring iPhone alarms. Now awake, Percival emerges from under the covers, ready to start another day. He adjusts the volume dial as the morning news commences, recalling a conversation he had yesterday; his roommate had derisively commented on his habit of keeping up with the world – isn't that the responsibility of a sophisticated adult? He stops his mind drifting. Percival will be presenting his research in a few hours – he needs to be looking his best. Scanning his minimalist wardrobe, he selects a modest trouser pant, thick horn-rimmed glasses and a loose fitted casual shirt, conscious not to appear overdressed. As he briskly walks to his destination, he recites his opening line.

The presentation is going well, no slip-ups so far, and the audience at least appears to be engaged. Percival suddenly notices his leg jittering. How long has it been doing that for? He pushes the offending foot firmly to the floor, attempting to regain an aura of self-assuredness, speaking slowly but assertively.

Regardless of how much you identify with Percival, certain aspects of his morning routine will stand out as common to everyone. We present ourselves in certain ways in the presence of others: we plan what to wear and what to say, we perform certain mannerisms with some audiences and not others, we share different sides of ourselves with strangers, lovers, and friends, and keep certain aspects of our inner lives to ourselves. Nearly everything we do carries some sort of social meaning which tells other people what we are like: the shoes that we wear, the utterances we make about ourselves and others, and the way we enter a room, all communicate and express things about ourselves. We reveal our tastes through our clothes and our cars, whilst our political views show up in the form of badges, bumper stickers, and utterances. In short, we have interests in *presenting ourselves to others*, in having a measure of control over which aspects of ourselves – whether our bodies, thoughts or how we style our appearances – we allow others to observe, and on what terms. These are our interests in *self-presentation*.

This thesis has two aims. The first has to do with the implicitly *synchronic* approach which is adopted by ethical and sociological accounts of self-presentation. I argue that this 'Synchronic Approach' is incomplete, it overlooks the *diachronic* character of our interests. While synchronic accounts examine self-presentation within a particular moment in time, diachronic accounts examine our self-presentational interests as they change across some extended duration of time.

In short, it is not only important that we're able to decide how we reveal ourselves to others at any one given moment, but also that we're able to exercise some control over the public construal of our *past* self-presentational choices. Percival ought to have some control over which of his thoughts, trousers, and theories he reveals to others at any one time, but he also has an interest in exercising some control over other people's access to these presentations ten years later, by the time he considers former opinions to be ignorant, ludicrous, or naïve.

The second aim of this thesis is to show how Social Networking Services (SNSs), or *technologies of the self*, pose a significant threat to our self-presentational interests.[1] I argue that the nature and extent of this threat has been underestimated because of an under-appreciation of the diachronic aspects of self-presentation. While technologies of the self prima facie enhance our self-presentational interests, attending to the diachronic dimensions of our interests reveals how SNSs have a thornier effect on our self-presentational interests than is first apparent. SNSs make it easier for us to act in ways that set back or threaten both our own and other people's self-presentational interests. Minimising the harms that are liable to arise on SNSs will require us to reinstate diachronic privacy and move away from the default of remembering online.

Before providing a chapter-by-chapter outline of how these aims will come together, I need to put three assumptions on the table. The first has to do with our right to privacy.

One of the underlying claims of this thesis is that privacy-related ethical concerns are partly about our interests in self-presentation. This view is not unprecedented; however, my defence of this claim depends on a particular conception of rights.

For the purposes of this essay, I understand privacy within an interest theory of rights. Interest theorists maintain that the function of a right is to further a right-holders' interests (Wenar 2023: 2.2). Perhaps the most prominent proponent of this position is Joseph Raz, who holds that 'X has a right if X can have rights, and, other things being equal, an aspect of X's well-being (his interest) is a sufficient reason for holding some other person(s) to be under duty' (1986: 166). Following Raz, I will be assuming that our right to privacy is grounded in interests. Interests are aspects of our wellbeing, which make certain actions impermissible – acting against them would be a violation of our right. We move away from this default by waiving our right to privacy, allowing individuals to act in ways that would have otherwise been impermissible. I take these claims to be intuitively

[1] A term coined by Foucault in his lecture 'Technologies of the Self' (1982) more recently employed by Barkardjieva and Gaden in relation to SNSs (2012).

plausible. As Andrei Marmor writes, 'any moral right, worthy of that name, must protect certain interests persons have' (2015: 4).[2]

The second assumption has to do with the relevance of privacy in public. One of the claims of this thesis is that we ought to have greater diachronic privacy online.[3] In this way, our right to diachronic privacy concerns information that, though once private, is now public. By making self-presentational choices via SNSs we open aspects of our private lives up for public consumption. But is information in the public realm a matter of privacy?

Privacy theorists have traditionally understood there to be a clear-cut distinction between that which is private, and that which is public. Privacy, on this picture, is concerned exclusively with the private domain. Judith Jarvis Thomson is one notable proponent of this position. She maintains that by making information public we waive our rights over that information (1975: 296-306). Although we *did* have a claim to privacy over the outfit we tried on in our homes, and opinions we deliberated in our inner worlds, by presenting ourselves to the world we waived our right to privacy, absolving other people of any obligations that were previously in play.

The second assumption I am making, contrary to this traditional distinction, is that there are privacy-related duties in public domains. Several authors have provided convincing arguments for this claim.[4] To demonstrate the plausibility of this assumption, I will draw on an approach embraced by James Wilson and Benedict Rumbold (2019: 19) and Adam Moore (2018: 342). These authors differentiate two kinds of duties engendered by our right to privacy: duties concerning access to information, and duties concerning the use or distribution of that information. By intentionally disclosing aspects of our private lives in public, we grant other people access to us. However, other people can still violate our right by using that information or distributing it without our consent. As Moore writes, 'we should be careful not to assume that when an individual consensually shares personal information, this is construed as relinquishing all downstream claims to control this information' (2018: 342).These authors suggest that, instead of taking right waiving to be an all or nothing affair, we should adopt a granular approach. While we may absolve or render defunct some privacy related duties when we make information public, this does not mean

[2] Véliz (2018: 6); DeCew (1997: 73) and Rachels (1975: 323) also seem to understand privacy in this way.

[3] One has diachronic privacy when one is able to keep their past self-presentations private.

[4] See e.g., Véliz (2018); Wilson and Rumbold (2019); Moore (2018); Stahl (2020) and Nissenbaum (1997).

there are no privacy related duties in play. I will say something similar about diachronic privacy later in this thesis.

The final assumption has to do with personal identity. Given that I am presenting a *diachronic* account of self-presentation, I will be considering our diachronic selves: our past, current, and future selves. One pertinent question which emerges in the wake of this is whether or not our diachronic selves constitute the same person. Where being the same person is a question of numerical identity – a relation which holds between one thing and itself. For the purpose of this thesis, I will assume that my past, present, and future selves constitute a single person. Because of this, I am assuming that we have *some* reason to care about our future selves.

While our everyday understanding of personhood takes continuous identity in human lives as given, the two dominant approaches to personal identity, namely the 'Biological Approach' and 'Psychological Approach' concur. I take the fact of this agreement as good reason to treat this third assumption as fairly uncontroversial. In order to motivate this, I will briefly outline these two approaches. The Biological Approach states that a person P will be numerically identical with a person Q if and only if P and Q are biologically continuous, where biological continuity has to do with the persistence of an organism (Olson 1997). This view equates persons with biological organisms, so as long as my body persists from A to B, my diachronic selves will be the same person. Rather than taking numerical identity to be a physical relation, the Psychological Approach takes personhood to depend on one's psychology. P will be numerically identical with Q if and only if P is uniquely psychologically continuous with Q (Parfit 1984), where psychological continuity is ordinarily cashed out in terms of memories, desires, or character. Most accounts which take a Psychological Approach are consistent with diachronic selves constituting a single person. This is because psychological continuity is a transitive relation. My 10-year-old self does not need to be psychologically continuous with my 30-year-old self, so long as there are strong psychological connections between my 10-year-old self and my 11-year-old-self, my 11-year-old self and my 12-year-old self, and so on. Personhood does not, contrary to popular belief, require the persistence of psychological traits.[5]

Now we have on the table three assumptions: (1) our right to privacy is grounded in interests, (2) there are privacy related ethical duties in public realms, and (3) all the versions of myself that exist across time are time segments of the same diachronically continuous entity, i.e., me.

[5] See Matheson and Khoury for a defence of this claim (2018: 212).

Turning now to the outline of this thesis. Chapter One will motivate our self-presentational interests. I argue that being able to control the aspects of our selves we reveal to and conceal from others is a morally important capacity – infringements of this capacity give us legitimate grounds for objecting to others. Chapter Two will address the matter of privacy: what is it? And why is our right to privacy valuable? Given the contested nature of our right to privacy, I will not seek to defend a particular conception of privacy. Instead, I seek to defend one of the central underlying claims of this thesis – that self-presentational interests form part of the normative basis of our right to privacy.

Chapter Three seeks to address the first aim of this thesis: to vindicate a distinction between the synchronic and diachronic aspects of self-presentation. Once I have isolated the Synchronic Approach, I will explain and defend the diachronic dimensions of self-presentation. I describe how we have non-instrumental justifications for these interests, and how not all infringements to our diachronic self-presentational interests (are all things considered) wrong. Finally, I argue that the Diachronic Account grounds a right to *diachronic* privacy.

Chapter Four will begin to address the second aim of this thesis. I argue that SNSs complicate and potentially undermine our self-presentational interests. While they give us greater synchronic self-presentational control, the utilisation of these tools simultaneously endangers our diachronic self-presentational autonomy. In the final chapter of this thesis, I examine what we ought to do in light of these harms. Part of this will involve examining our own role in facilitating infringements to our own interests – the sense in which we are complicit in the violation of our own privacy. The other side of this will involve examining other people's role, and the role of the technologies themselves. I argue that while social media users collude in these harms, we have good reasons to modify technologies of the self so that they don't make it so easy for us to undermine both our own, and other people's self-presentational interests. In the final section of this thesis, I explore some changes that may be helpful to this end.

I

The Presentational Self

> 'It is probably no mere historical accident that the word person, in its first meaning, is a mask. It is rather a recognition of the fact that everyone is always and everywhere, more or less playing a role… it is in these roles that we know each other; it is in these roles that we know ourselves'.
>
> (Park 1950: 249)

We have an interest in being able to manage the public construal of ourselves to others; to have some say in determining what other people think about the kind of person we are. These are our interests in self-presentation. This chapter has two central aims. Firstly, to defend the idea that it is in our interests to have a measure of self-presentational control. Secondly, to demarcate a kind of self-presentational choice which will be central in motivating the Diachronic Account.

In §1.1 I will describe what our interests in self-presentation are, and why having a measure of self-presentational control is important. In §1.2 I describe how we lack complete self-presentational control by exploring the non-voluntary aspects of self-presentation. In §1.3 I highlight various ways in which our self-presentational interests might be set back. In §1.4 I describe the notion of identity which is in play in discussions of self-presentational impediments. Finally, in §1.5 I distinguish self-presentational choices per se, from self-presentationally autonomous choices, thus isolating the autonomous presenter this thesis is concerned with.

1.1 The Presentational Self

The idea of self-presentation was most influentially elaborated by the sociologist Erving Goffman in his iconic book, *The Presentation of Self in Everyday Life* (1959). There, Goffman describes the repertoire of human behaviours driven by our need to present ourselves to others differently in front of different audiences, almost as if we're performing as actors on a stage (Ibid., 57). Goffman carves up the world into front-facing realms and back stages. When we are on stage, we work to present the best versions of ourselves. We put on our finest clothes and smiles, politely greeting passers-by, and holding back the insincerity of forced congratulations. It is only when we

step off stage, entering the safe space of a friend's home, or sit on an unobserved park bench that we may relax, let our hair down, slouch, and speak ill of our enemies. Here, 'we engage in all the little eccentricities we'd never dream of exhibiting in public' (Schlenker 1980: 33) and hide from the world the processes by which we come to present ourselves to others.

The social psychologist Barry Schlenker developed Goffman's early articulations of self-presentation. He describes how the aspects of ourselves we share with others carry social meanings which tell other people what we are like (1980: 6). Our publicly enacted choices, the clothes that we wear, the way that we enter a room, and the utterances we make about ourselves and others, all communicate and express things about ourselves. It is on the basis of our presentational self that other people form an idea of who we are. We make evaluations like this all the time: is this person nice? Are they genuine? Are they fun, or boring? Conservative, or liberal?

Our interests in self-presentation are our interests in shaping the impressions other people form of us. By exercising control over the aspects of ourselves we share with others, we influence how other people perceive our personality, intentions, values, friends, possessions, abilities, and behaviours (Schlenker 1980: 6). This is called *impression management*. Impression management is both positive, and negative. As Juha Räikkä notes, while we are often conscious to moderate which aspects of our inner life we share with others, very often we are conscious *not* give a particular impression of ourselves (2017: 224). Influencing how other people perceive us is thus a vital part of establishing both who we are and who we are not (Douglas 2016: 202).

Having a measure of control over which aspects of ourselves we allow others to observe is important for a number of reasons. These reasons are in some sense coextensive with, and constitutive of our interests in self-presentation.

One of the most cited reasons why self-presentational control is important, is that it enables us to maintain different kinds of social roles and relationships with others.[6] These roles and relationships are valuable both instrumentally, as with doctor-patient relationships and valuable per se, as with friendships and significant others. As James Rachels has argued, many of the roles and relationships we stand in require us emphasize certain aspects of ourselves and in some cases conceal things about ourselves entirely (1975: 326). Relationships involve a conception of the degree of information that would be appropriate to share – a joint understanding of what it would be to under-share or overshare. While it is appropriate for me to reveal my darkest thoughts to

[6] See e.g., Mead (1934: 142), James (1990: 294), Fried (1970: 484), Schlenker (1980: 36) and Cocking and van den Hoven (2018: 61); Cohen (2002); Gerstein (1978); Owens (2022: 225); Nagel (1998: 20) Marmor (2015: 8); and Shiffrin (2014: 10).

best friend or therapist, the same is not true of my Uber driver. In this way, expectations concerning self-presentation partly define what relationships are – that I am related to her as a best friend and them as a stranger, is partly constituted by the fact I'd tell her things I wouldn't tell a stranger. This concealment need not be understood as dishonest, nor deceptive. Rather, our relationships function as a sort of division of labour, where loading one kind of relationship with a burden of another can be inappropriate or disconcerting (Véliz 2018: 7). Without having a measure of control over which aspects of ourselves we reveal, these different kinds of relationships would be much more difficult to create and sustain.

Having a measure of self-presentational control also enables us to maintain relationships on an even keel. Thomas Nagel describes how conventions of concealment and reticence allow us to present ourselves for appropriate interactions with other people, avoiding unnecessary conflicts with others (1998: 4). Such conventions give us space to bicker about our antiquated relatives and irksome colleagues behind their backs. As David Velleman notes, social harmony is often best promoted by 'assuming an amiable expression' and keeping one's 'true sentiments to oneself' (2009: 69).

A final reason why our self-presentational interests are important, as Marmor notes, is that being able to control how, when, and to whom we publicly reveal ourselves to others means being able to control when and if we subject ourselves to social scrutiny (2015: 9).

Sometimes, we welcome the judgement of others, whether appraisal or criticism. However, there are often parts of ourselves we want to safeguard from social scrutiny. We may not want to expose ourselves because we are embarrassed or ashamed, because we haven't quite figured out what we think, or simply because it is not the kind of thing we want to share. It is not always possible to choose, of course, but sometimes our choice to put ourselves on display is significant. As Marmor notes, a clear example of this is the decision of gay people to make public their sexual preferences. This can unfortunately be a momentous decision, because it immediately subjects the individual to public scrutiny (Ibid., 10). Though less weighty, many of our day-to-day decisions involve the same processes: we want the space to write drafts on our laptops before pressing send, to try on outfits before we step outside of our homes, and to sing in the shower without being judged for our dissonant tone. Having self-presentational control grants us the space to decide whether or not to invite aspects of ourselves for social scrutiny. Without it, we would exist on stage at all times.

Self-presentational choices involve a whole range of choices concerning the ways in which we navigate and mould our social environment. We depend on one another to construct a reliable

picture of the world, of each other, and of ourselves. The fulfilment of these foundational social values requires a capacity for discretionary self-revelation. We need to have some control over who has access to what, and when, not because we have anything to hide (although that may be true too), but because self-presentation enables us to create the social life that we want to have (Marmor 2015: 10).

1.2 Giving and Giving Off

While we ought to have some control over the aspects of ourselves that we reveal to and conceal from others, we cannot command the impressions other people form of us. If we could, then much more of Percival's social circles would see him as the charming, intellectual, and hilarious person he tries to be. In short, there are limits to self-presentational control – we can shape and influence the impressions other people form, but impression management is never fully within our control.

Not all of our public presentations are the result of purposeful and controlled actions. My shoe size, height, psychological dispositions, and facial expressions are aspects of my public self I cannot easily change. Other aspects of our inner lives reveal themselves despite our best efforts to conceal them: our embarrassment radiates from our flushed cheeks, and nervousness exudes from our body language. As Mark Leary notes, 'virtually every aspect of our behaviour provides information from which other people can draw inferences about us' (1996: 16). Information is in this sense *inferentially fertile* – we can infer information about individuals on the basis of how they present in public (Manson and O'Neill 2012: 104).

Even self-presentations which are the result of highly controlled actions give off more than we intend. As we begrudgingly congratulate our enemies, or greet our ex's new lover, the insincerity and forced nature of our interaction is often obvious. This is because the presentations we intend tend not to replicate the presentations they seek to mimic (Cocking 2008: 129) – we can tell the difference between a Duchenne Smile and a smile through gritted teeth.[7] These involuntary forms of self-revelation have the potential to undermine our self-presentational interests when they conflict with the image of our selves we wish to portray.[8]

[7] The kind of authentic smile that reaches your eyes, making the corners wrinkle up with crow's feet.

[8] Cocking and van den Hoven (2018: 77-88); Véliz (2022: 38); and Marmor (2015: 7) all recognise this limitation.

Goffman's distinction between the expressions that we give and give off is useful here. While we consciously 'give' information to others through verbal and non-verbal symbols, in doing so we often unwittingly 'give off' certain signs and expressions which extend beyond the locutions we intend to portray (1959: 14). Because of this, the information that we intend to convey frequently exceeds our capacity to manage it.

These limits to self-presentation illustrate two aspects of our presentational self: (1) *self-presentational choices* – the parts of ourselves we intentionally give, and (2) *self-revelations* – the aspects of ourselves we involuntarily give off. Our interests in self-presentation recognise that we cannot fully command our presentational self. They are ethically legitimate interests in having a measure of control over our public appearances, but not complete control – *complete* control is not only unachievable, but undesirable.

Marmor speaks at length of the balancing act between our right to exercise a certain level of control over our self-presentations, and the fact we can end up with too much control (2015; 2021). Complete presentational control is undesirable because it compromises authenticity.

The authenticity Marmor has in mind is not a mild one: most of us want to appear slightly better (more beautiful or interesting) than we actually are, and there is little wrong with that. What does seem wrong, is getting other people to believe you are something quite different from what you really are (2022: 3). There is something intuitively unsavoury about another person to have grossly inaccurate views about the kind of person you are. An unkind person is bad enough in being unkind, but inducing other people to believe they are goodhearted is a form of deceit we may rightly condemn (Ibid.). Excessive self-presentational control risks blurring the distinction between truth and fiction – a problematic result in a world with a default moral presumption of truthfulness (Shiffrin 2014: 11).

No one has a legitimate ethical interest in exercising complete control over their self-presentations: complete control is both unachievable and undesirable. But we should be able to exercise some control over the ways in which we present ourselves to others – to have some say in determining what other people think about the kind of person we are.[9]

[9] I follow Marmor in leaving this goldilocks level of self-presentational control vague. Marmor uses the deliberately opaque notion of 'reasonableness' to capture this – the idea is that people are likely to have differing views and attitudes about how much control they want with respect to different aspects of their lives (2015: 15). I will borrow this notion of responsibility, although I take most of the infringements in this thesis to be clear instances of lacking self-presentational control.

1.3 A Survey of Impediments

We ordinarily have in mind an image of ourselves we want to project. Goffman calls this our 'face' (1955: 213). This is our publicly acceptable self: the self we want to share with others. We are not always able to maintain face. Sometimes, we are prevented from exercising a reasonable measure of control over our self-presentations. When this happens, we may have grounds for objecting to others on the basis of our self-presentational interests.

There are a range of ways in which our self-presentational interests may be set back. In this section, I identify six features which are useful in characterising various kinds of impediments: Willingness (§1.3.1), Accuracy (§1.3.2), Valence (§1.3.3), Involvement (§1.3.4), Voluntariness (§1.3.5) and Awareness (§1.3.6).

1.3.1 Willingness

The impressions that other people form of us may be formed on the basis of aspects of ourselves we wanted to share, or aspects of ourselves we wanted to hold in private.

Many self-presentational infringements involve other people accessing parts of ourselves we wanted to keep hidden. When someone reads our diary, or distributes information we shared in confidence, they undermine our self-presentational interests because they expose us to the world in ways we do not wish to be exposed. There are often many grounds for objecting to such disclosures: whether breaking norms of confidentiality, or the sensitivity of the information in question. However, there are examples in which the wrongness of the exposure is clearly an infringement of our self-presentational interests.

Suppose my best friend knows I once saved someone from drowning. I have not explicitly asked them to keep it a secret, but I have told them I'd rather other people didn't know, as I feel embarrassed about how it puts me in the spotlight. Now suppose the two of us meet up with a third person, and in order to brag about how wonderful I am, my friend tells the third person about my heroic deed. While they had the best intentions and did not break any confidentiality demands, I have grounds for objecting to them because they ran roughshod over the Willingness dimension of self-presentation.

1.3.2 Accuracy

Other people can undermine our self-presentational interests by revealing information which is accurate, or information which is not.

If I were quadrilingual, then I might not want other people to see me as such. I may have grown tired of the incessant 'how do you say, "[insert phrase]" in [insert language]' and I feel embarrassed when I am asked to translate words with no direct equivalent. Because of this, I try my best to conceal this part of myself in my interactions with others, pretending to be your typical monolingual Englishman. If someone were to reveal to others that I can indeed speak Italian, Mandarin, English, and German, then they would undermine my self-presentational control. I would lose control over the terms by which my linguistic capabilities were knowable to others, but the information on the basis of which other people form an impression of me is accurate.

There are other examples where our self-presentational control is impeded by other people revealing information about us which is not at all accurate. When someone spreads a rumour, makes wildly speculative claims about my inner life, or accuses me of a crime I did not commit, then other people may form an impression of me on the basis of information that has no bearing on the kind of person I am. Such cases are referred to in privacy tort as *false light,* and include cases where information is taken out of context, cases where opinions or utterances are falsely attributed to individuals, and false imputations of criminality (Morvan 2018: 85). Such fictitious impressions threaten our self-presentational interests by inducing other people to think we are a kind of person we are not.

1.3.3 Valence

Both positive and negative construals of us can undermine our self-presentational interests.

Valence has to do with how other people construe us. When my friend bragged about me saving someone from drowning, they did so in an entirely positive light – they wanted and potentially succeeded in getting the third party to think well of me. But the fact that the construal of me was positive does not preclude its being injurious to my interests.

Other self-presentational infringements involve the disclosure of information which reflects badly on individuals. If I were to reveal Percival's infidelity or knowingly accuse him of misdemeanours he did not commit, then I could induce other people to form unfavourable impressions of him.

1.3.4 Involvement

Sometimes, other people undermine our self-presentational control even though what is disclosed is congruent with our public self-image, and we are willing to share it, because they deny us the opportunity to present it ourselves.

The involvement condition only applies to aspects of ourselves that we wanted to share with others. An example by David Shoemaker illustrates this dimension: 'suppose I've just won a big award and I can't wait to tell my wife about it, but you call her up and tell her before I can. You've prevented me from presenting the news in the way I wanted to, and despite the fact that I'm still proud of my accomplishment, you've tarnished it' (2010: 13). Unlike the drowning case, the wrongness of the infringement does not concern my unwillingness to share the news, but that *I* wanted to be the one to share it. As Shoemaker puts the point – 'when it comes to the exposure of ourselves, most of us prefer that our selves be the ones doing the exposing' (Ibid; 14). By preventing us from participating in our own disclosure, violations of the involvement clause deny us a key aspect of self-determination.

1.3.5 Voluntariness

Our self-presentational interests may be undermined by self-revelations – the less controlled, and non-voluntary aspects of self-presentation.

Suppose my ex walks into the room. I try to keep it together, but my inner life breaks through the surface: I am flustering, visibly anxious, surprised, and uneasy. There's a sense in which my self-presentational control is already very minimal. I have, against by best efforts, put myself on display. However, other people can still violate my self-presentational interests with respect to this revelation by making my reaction obvious to people who had not noticed, or telling people who were not present how I had acted. In order to respect my self-presentational interests, other people ought to turn a blind eye to these involuntary revelations, practising what Goffman calls *civic inattention* – a kind of 'dimming of lights' (1963: 84).

1.3.6 Awareness

Sometimes we know when other people form impressions of us. In other situations, we may be unaware that our self-presentational interests are at stake.

Part of what explains why data mining – the gleaning and analysis of personal information – is wrong, is that it undermines our self-presentational interests without us knowing.

There are two reasons why it is difficult to know what parts of ourselves we are revealing online. Firstly, because the information we give off online surpasses the information that we give. When we surf the web or scroll through Twitter, we leave a weighty data trail behind us, unknowingly revealing information about our tastes, health, preferences, and friendships. Secondly, because information may be gleaned from the information we give. While we are aware of the relatively mundane information we share on the internet, the deep learning models which underpin social medias analytics clusters and segments individuals into more complex and higher dimensional spaces than humans are capable of, creating meticulous profiles of 'what we are like'.[10] Helen Nissenbaum calls this 'the problem of aggregation' (1998: 587).

Daniel Susser (2016) and Shoemaker (2010) both critique the practice of data mining on the basis of our self-presentational interests. Data mining is problematic, Shoemaker writes, because it denies us 'the opportunity to participate in the process by which we are perceived and understood' (2010: 233). Most of us do not know that we are being monitored and evaluated, and without knowing who has information about us, and what the information is, we cannot positively influence the conception that other people have of us. Here, 'the boundaries between ourselves and others are no longer ours to draw' (Susser 2016: 235).

There are a number of reasons why we might lack self-presentational control. In some cases, other people expose us to the world in ways we do not wish to be exposed, in other cases, we are prevented from exposing ourselves in the way that we want. We may object to the actions of others when we lack a reasonable measure of control over the terms by which we present ourselves to others – when our self-presentational interests are set back.

1.4 The Self of Self-Presentation

Many self-presentational impediments involve the exposure of aspects of ourselves we wanted to hold in private – parts of ourselves we were not willing to disclose. In this way, the Willingness dimension is central to many self-presentational infringements. However, when I think about all of the properties attributable to me, I only have interests in controlling other people's

[10] See Tavani 1991 for a discussion of this.

access to a certain subset of them. If someone were to go around saying 'Helena is British, short and has brown hair' I might be slightly bemused, but I wouldn't take this to deny me self-presentational control. If, on the other hand, someone revealed my medical history, embarrassing phobias, or quirks, then I might. What we need, in the wake of this, is a notion of the self which can help to make sense of this – a notion of the self that we have interests in controlling the presentation of.[11]

One notion of the self that is clearly irrelevant here is the self of numerical identity: the metaphysical glue which holds a person together across time. Instead, we are looking for the answer to the more existential question – who am I really? Maria Schectman calls this the characterisation question of identity. It asks which features make us the kind of person we are (1996: 73). Think of it this way: there are various properties which may be attributed to each of us. I, for example, have ten fingers, a protruding chin, and brown eyes. I also have beliefs, values, emotional dispositions, and desires. But not all of these properties play a role in understanding me; only a certain subset of all the properties attributable to me may rightly be called mine (Shoemaker 2010: 7).[12] When we ask the question of 'who are we, really?' there are some facts about us that are more relevant than others. These properties are central to our personality, character, and view of ourselves, such that 'understanding what these facts are is crucial in some important way to understanding [us]' (Copp 2002: 355).

I do not have the space to argue that we ought to conceive of self-identity in this or that way, however, I will present what I take to be a plausible candidate. David Copp argues that a property P is part of our self-identity when the following conditions hold (2002: 371):

(1) We believe we instantiate P.
(2) Our belief that we instantiate P grounds emotions of self-esteem.[13]

---

[11] Another pertinent difference between these kinds of properties is that other people can see me for what I am – other people do not have to overcome the opacity of my mind to know that I have brown hair. However, it remains true that I am indifferent to the exposure of these aspects of myself – this section seeks to capture the relationship between these kinds of properties and our tendencies to care.

[12] We can think of the notion of understanding here in contrast to mere knowledge. If you know a bunch of mundane facts about me, but no facts that are part of my self-identity, then you have mere knowledge of me rather than understanding.

[13] Shoemaker defends Copp's account as an apt in the context of our self-presentational interests, I borrow both his usage and defence of Copp's account (2010: 7-11).

Self-esteem has to do with how we value and perceive ourselves. If we have low self-esteem over an aspect of our self-identity, then we are likely to feel negative feelings towards it, of guilt, shame, regret, and embarrassment. If we have high self-esteem over an aspect of our self-identity, then we are likely to feel positive feelings of pride and confidence (Ibid., 370). These self-conscious emotions pertain to the self, as Robbins and Parlavecchio write, 'the self stands before itself as if the self were an exterior event, object or person' (2006: 326).

In Copp's account, the properties which make up our self-identity are the parts of ourselves which ground emotions of self-esteem – they are the parts of ourselves we are proud of or embarrassed about. This connection between the self and self-conscious emotions is well rehearsed – if we did not think that the property P was rightly ours, then it is hard to see what reason we would have to feel shame, embarrassment, or pride over it.[14]

This conception of self-identity elucidates the willingness dimension – how there are parts of us we want to keep to ourselves, and others we want to subject for social scrutiny. The parts of our self-identity which ground negative emotions of self-esteem – the things we are embarrassed of and ashamed about – are usually aspects of our self-identity we wish to conceal from rather than reveal to others. The parts of our self-identity which ground positive emotions of self-esteem – the things we take pride in and draw confidence from – are parts of ourselves that most of us want to reveal to rather than conceal from others. In this way, the self that we are willing to present to others forms only a subset of our self-identity.

We can take Copp's theory as a sketch of the kind of theory that can help us to explain what notion of the self is relevant in the context of self-presentational infringements. We have interests in exercising a measure of control over the presentation of our self-identity; to conceal or reveal aspects of ourselves which form a relevant part of the picture we have of ourselves. Especially those parts of our self-identity we are not willing to expose. My fears, infatuations, and achievements may form part of my self-identity, while more mundane facts about me don't. We have interests in controlling the former but not the latter.

When other people expose us to the world in ways we do not wish to be perceived, we lose control over the terms by which the grounds of our self-esteem are known by others. While there are a range of ways that we could specify improper identity access, a plausible way to think about this is that access is unauthorised when it involves the exposure of esteem-grounding properties.

[14] See Lewis for an exploration of the self in self-conscious emotions (1997).

1.5 Self-Presentational Autonomy

So far, I have drawn a distinction between two aspects of the presentational self: (1) self-revelations – the aspects of ourselves we involuntarily give off, and (2) self-presentational choices – the parts of ourselves we intentionally give. The purpose of this section is to identify a third aspect of the presentational self – to mark a difference between self-presentational choices per se, and self-presentationally autonomous choices.

Before we can describe which of our choices are self-presentationally autonomous, we need to have in mind some conception of what autonomy is. Broadly speaking, autonomy is the ability to govern oneself – the capacity for self-determination or self-governance. As adult human beings we are capable of deciding how we want to live our lives, to determine what we value, and to live in accordance with those values. While specifying the precise bounds of autonomy is a much-contested task, we only need to appeal to two relatively uncontroversial aspects of it. Firstly, there is the question of whether or not it is your deliberation and choice shaping your life, as opposed to some other persons choice, or brute circumstance. As Raz describes it, autonomy is an 'ideal of self-creation' in which agents are 'part authors of their lives' (1986: 370).[15] Secondly, there is the question of whether your deliberations and choices are responsive to reasons or values, as opposed to being arbitrary, capricious, neurotic, or in some other way insufficiently responsive to the kinds of things that endow our choices with value. Autonomous decisions are the kinds of decisions we endorse upon reflection – choices that express our values and wants.

While autonomy concerns being able to craft our lives in accordance with our values – the governance of our self, *self-presentational* autonomy has to do with crafting our self-presentational choices in accordance with our conception of a publicly acceptable self – the governance of our presentational self. This ability to curate our presentational self is not simply being able to decide which aspects of ourselves to reveal or conceal from others but deciding to reveal or conceal aspects of ourselves we have exercised a measure of control in shaping.

Consider my decision to reveal to you that I have three brothers or that I'm British. These are self-presentational choices, and I may take these facts about myself to be an important part of you understanding who I am. But these aspects of myself are not meaningfully within my control – I did not decide to be this way rather than another. Because of this, they are not self-presentationally autonomous choices.

[15] Similar notions of autonomy are found in Colburn (2016: 61) and Hurka (1993: 148).

To the contrary, my decision to share with you my political tendencies, or my decision to leave the house with my face and hair styled in a certain way *is* a self-presentationally autonomous choice. My beliefs and appearances are aspects of my self-identity I exercise a degree of control in shaping – we craft them in ways we do not craft our nationality, or siblings. In contrary to other parts of myself, I am involved in a sort of project of passion with respect to my beliefs and outfits, I figure them out as I go and discover what feels authentic to me.

One important caveat to note is that we can stand in stronger custodial relationships over some aspects of our selves than others. On the one hand, we have aspects of ourselves we are clearly custodian of: the way that we style our hair and the clothes that we wear fall into this category. As do our beliefs, opinions, and hot takes. These are paradigmatic examples of self-presentationally autonomous choices. On the other hand, we have peripheral cases. Take my height as an example. While I cannot choose my height, there are some things I can do to control how my height is presented to others – I can wear clothes that elongate my limbs, wear platform shoes, or have reconstructive surgery. I stand in some sort of custodial relation with respect to my height, but a relatively weak one. The same is true of our accents (which we can change via speech therapy), nose shapes (which we can have surgery to alter) and our physique (which we can influence through exercise). While I may feel proud that I have worked to alleviate my hollow and dreary cadence, embarrassed about my failure to keep in shape, or confident in my platform shoes, these aspects of our self-presentations are not as straightforwardly amenable to volitional control as the clothes that I wear, or the opinions I disclose.

Our self-presentationally autonomous choices form a subset of our self-presentational choices: self-presentationally autonomous choices are self-presentational choices in which we reveal aspects of ourselves we have exercised a measure of control in shaping. With these distinctions in mind, we can posit three aspects of the presentational self:

(1) **Self-Revelations:** These are the aspects of our presentational self over which we have little control. How I reveal to you the shape of my body, my anxiety, embarrassment or dominating personality traits without meaning to.

(2) **Self-Presentational Choices we are not custodian of:** These are decisions to reveal or conceal aspects of our identity we cannot change, for example, my decision to reveal to you that I have three brothers, my sex, or where I was born.

(3) **Self-Presentationally Autonomous Choices:** These are our decisions to reveal to or conceal aspects of ourselves we have exercised a measure of control in shaping. These include both *paradigmatic cases* such as our opinions, outfit choices and hair styles, and *peripheral cases* such as one's height, voice, or weight.

Note that the fact that I *can* exercise a degree of control in shaping some aspect of my appearance doesn't mean that it *is* a self-presentationally autonomous choice. My hair, for example, is something that I could cut, curate, and dye. But if I do not do any of these things, then there is a sense in which I do not *decide* to reveal to you my natural hair colour and type when I step outside. This part of my appearance seems to fall under the not-fully voluntary aspects of self-presentation. However, when I leave the house having dyed my hair a lurid red, or after styling it, I make a self-presentationally autonomous choice. In this way – at least with respect to the peripheral cases – it is only when we have exercised self-presentational autonomy over an aspect of our inner world or appearance, in shaping, curating, and changing it, that it becomes a self-presentationally autonomous choice.

While most accounts of self-presentation mark a distinction between the more and less purposeful aspects of self-presentation, no one – to my knowledge – has differentiated these different kinds of self-presentational choices. I take this distinction to be pertinent. It seems important to recognise the sense in which we shape and curate our opinions and beliefs in ways we do not shape our nationality, head shape and psychological dispositions. This thesis will largely be concerned with the aspects of our self-identity we are custodian of, in this way, my discussion is confined in a way other accounts of self-presentation are not.

SUMMARY

This chapter had two overarching aims. Firstly, to identify and explain the moral reasons which constitute our interests in self-presentation. We ought to have some control over when, how, and to whom we reveal aspects of our self-identity – those parts of ourselves which ground emotions of self-esteem. Our self-presentational interests are set back whenever we lose control over the terms by which the grounds of our self-esteem are knowable to others. When this happens, our interests in self-presentation give us legitimate grounds for objecting to others.

The second aim of this chapter was to demarcate a kind of self-presentational choice which will be central to the Diachronic Account – our self-presentationally autonomous choices. These are the parts of our presentational self we exercise a degree in shaping. Without giving the game away,

noticing the sense in which we craft certain aspects of our self-identity is central to understanding the diachronic dimensions of our self-presentational interests.

II

What is Privacy?

Privacy is something that is held by individuals. It can be lost, and institutions, situations and individuals can undermine, intrude upon, and invade it. But what exactly is privacy? For a term so widely used and understood, it's surprisingly difficult to explain the meaning, scope, and value of our right to privacy. This chapter seeks to examine privacy by considering some of the leading attempts to characterise it. The main aim of this chapter is to defend one of the central claims underlying this thesis: that self-presentational interests form some part of the normative foundations of our right to privacy.

In §2.1 I will outline three key central categorisations in privacy theories and motivate a non-sceptical view of privacy. In §2.2 I will outline some of the core attempts to classify our right to privacy. Finally, in §2.3 I defend the tight connection between privacy and self-presentation by drawing on the interest theory of rights I outlined in the introduction.

2.1 Key Distinctions

Several important distinctions arise in discussions of what privacy is: descriptive and normative accounts, relative and non-relative accounts, and reductionist and non-reductionist accounts. I will review each in turn.

Descriptive, or non-normative accounts, describe the conditions under which privacy obtains or is lost. They tell us when we have privacy – what is in fact protected as private. On the other hand, we have normative accounts, these use moral obligations or claims to explicate our right to privacy; they defend the value of our right to privacy, and the extent to which it ought to be protected (Moore 2008: 413). I will explore some of the central descriptive accounts of privacy in the next section. The way that our self-presentational interests relate to privacy (which I will defend in §2.3) is a normative claim – our self-presentational interests explain why we have a right to privacy.

Giving a definition of where privacy obtains, and why privacy is valuable is difficult – at least in part – because there are significant variations across time and between cultures on what

counts as an individual's 'private affairs' (Schafer 2011: 13). The second important question to consider in theorising about privacy is whether or not privacy is relative. However, as Ferdinand Schoeman notes, there are two important aspects to this question. Firstly, the question of whether or not privacy is relative to cultural or temporal differences, and secondly, the question of whether there are any aspects of our lives which are conventionally private (1984: 200). On the first of these questions, there is wide agreement. Most privacy theories recognise that cultures differ in their ideas of what privacy is, how we obtain it and how we protect it.[16] There is much less agreement on the second of these questions. On the one hand, we have theorists such as Hannah Arendt, who argues that the boundaries between the private and public are set. Arendt argues that there's an underlying social ontology under which certain aspects of our lives are more likely to be regarded as private. On this picture, what counts as an individual's private affairs is fixed (1998). On the other hand, we have theorists who think that the distinction between the private and the public is entirely conventional; that the boundaries between private and non-private affairs 'call for constant reinterpretation, are always open to dispute, and will never be fixed for good' (Roessler 2009: 708).

The final important contention among privacy theorists is between reductionists and non-reductionists. While we ordinarily speak of '*the* right to privacy' as if there's a single underlying interest which underpins our right, there are numerous contexts in which it may be claimed that we have a right to privacy: whether we're confiding in a friend, disclosing sensitive medical information, or reading a book in the privacy of our home. This has led several theorists to question the unity of our right to privacy – 'to doubt that there is truly a common thread' (Schafer 2011: 12).

Reductionists are critical of privacy – they think that there is nothing useful about the concept of privacy, and nothing distinctive about the interests that privacy protects. Non-reductionists, or coherentists, defend the value and coherence of the concept of privacy – they maintain that there is something 'fundamental, integrated, and distinctive' about the concerns which typically fall under the rubric of privacy issues (Shoeman 1984: 5). Shoeman's analysis of these positions is helpful. He distinguishes two components – the rejection of which characterises the reductionist account. Firstly, the *coherence thesis* – this is the claim that there is something common to the various situations in which we have a claim to privacy. Secondly, the *distinctive thesis* – which claims that

[16] Westin (1967); Rachels (1975) and Moore (2003) outline ways in which cultural differences shape conceptions of privacy.

privacy claims are explained or defended in terms of principles, rights or interests which are distinctive to privacy (1984: 5).

Thomson is one notable reductionist. In her seminal paper *The Right to Privacy,* Thomson claims that there is no such thing as a distinctive right to privacy (1975: 306). Instead, the concept of privacy is a placeholder for a cluster of rights: namely, property rights over our possessions and bodies. These rights are the true constituents of privacy, meaning that we can explain what is wrong about cases we would typically describe as privacy violations, in terms of these other rights *with no conceptual loss* (Ibid., 332). As Frederick Davis remarks – privacy violations are 'in reality, a complex of more fundamental wrongs' (1959: 20).[17] It's not that these non-derivative interests are illegitimate, however, there is nothing useful, distinctive, or illuminating about the concept of privacy itself.

Thomson seems right to recognise that we use the concept of privacy to denote a plurality of ethical concerns. When we say that something safeguards, or otherwise impedes our right to privacy, we may be referring to a range of interests that are advantaged or disadvantaged by some practice, policy, or person, and it is unclear whether or not these adjacent interests possess a shared normative foundation beyond being beneficial to our welfare. However, although the supposed unity of these interests may be called into question, several authors have provided convincing grounds against Thomson's scepticism.[18] In analysing Thomson's critique of privacy, I will leave to one side the question of coherence – instead focusing on her rejection of the distinctive thesis.[19]

It is worth recognising that Thomson's account is at the very least partially correct. There are cases in which other people accessing our possessions constitutes a privacy violation – if someone reads my diary or examines my body without my permission, my privacy is clearly at stake (Inness 2018: 33). However, privacy concerns are not reducible to infringements of property rights. While property rights over our possessions and bodies *can* explain privacy infringements, it is not true that privacy rights *are* property rights.

---

[17] Prosser (1960) and Bork (1990) are also privacy sceptics.

[18] Lever argues that Thomson overlooks how our interests in confidentiality are central to understanding privacy (2012: 49). Scanlon argues similarly to Inness that rights over our possessions fails to exhaust privacy concerns – 'rights of ownership over objects do not play the primary role Thomson assigns to them' (1975: 319).

[19] As I will explain in the following section, I think that self-presentational interests are common to most privacy violations. However, this does not strictly speaking mean advocating the coherence thesis, because privacy violations may occur in the absence of self-presentational infringements.

As Julia Inness argues, having a justified claim of ownership over X is neither a necessary nor sufficient condition for having a justified privacy claim over X (2018: 33). Consider sending a love letter to your betrothed. By sending the letter, you relinquish possession of it – it no longer belongs to you. However, it seems that your right to privacy would be violated were it to be distributed to others without your consent. Here, your right to privacy adheres to the object even though your property rights have been relinquished. This demonstrates that privacy claims may attach to objects in the absence of property claims – ownership is not necessary for privacy (Ibid., 33). On the other hand, having a justified claim to possess an item is not sufficient for having a privacy claim over it. Say that there are two items on my desk – a bottle, and a diary. Both of these items are mine, but the fact of their being mine does not necessarily justify my claim to privacy over them. Someone examining my diary would violate my privacy, but the same is not true of my bottle. This demonstrates that only certain sorts of property rights over objects are relevant to privacy – mere ownership is not sufficient for privacy (Ibid., 34).

As Rachels points out, understanding privacy in terms of property rights requires us to dissociate privacy infringements from ideas of intimacy (1975: 333). Part of the reason why it's wrong for other people to access our diaries, letters, medications, and bodies is that they are ours, but ownership fails to exhaust our privacy claims over these things. Rather, it is the fact that these things are the kinds of things we do not ordinarily want everyone to access – that they are the kinds of things we want to keep private – which grounds our right to privacy over them. Ideas of access and intimacy reveal how there are distinctive liberties, goods, and ills which are illuminated by talk of a right to privacy – aspects of privacy which cannot be easily accommodated by talk of other rights.

Despite the multifariousness of what the concept of privacy denotes, most theorists take the view that privacy is a meaningful and valuable concept, as Annabelle Lever writes 'talk of privacy… is not pointless or simply confused' (2012: 57). In the following section, I outline some central, non-reductionist, and descriptive accounts of privacy.

2.2 Accounts of a Right to Privacy

In discussing what a right to privacy is, it seems fitting to begin with Samuel Warren and Louis Brandeis, who famously describe privacy as 'the right to be left alone'. Their article *The Right to Privacy* marks the beginning of a systematic discussion of the concept of privacy (1890). Warren and Brandeis focus on invasions of privacy brought about by the public dissemination of

information concerning a person's private life, thus laying the foundation for the now dominant view of privacy as informational control.

Accounts of informational privacy maintain that privacy concerns having control over our personal information.[20] These accounts start from the idea that there are facts about us which we do not want other people to know, or at least certain people in certain contexts to know (Westin 1967: 7). Personal information such as one's salary, health conditions, sexual orientation, or weight would fall into this sphere (Parent 1983: 270). Accounts of informational control can be understood as setting epistemic boundaries; defining the aspects of us that other people should know and the aspects they should not.

While there is almost universal acceptance of the value of informational privacy, certain privacy violations seem to exist which cannot be accommodated by talk of information control. The most frequently cited cases are based on modified versions of Sartre's 'Peeping Tom'. Say Tom peers through the keyhole and observes his former lover. He may gain no new information about him: he knows the exact layout of his room, his routine, and the features of his unclad figure. Despite the inapplicability of information, Tom seems to violate his former lovers right to privacy – he had a right not to be observed. Another counterexample is given by Moore, who imagines a sleepwalker wandering into someone's room and proceeding to pet their sleeping head (2008: 417). This seems to be a blatant violation of privacy, but it would not be a privacy infringement under the informational view.

Some have argued this is not an issue. William Parent argues that these intrusions are irrelevant to privacy; that they are best understood as cases of trespass or harassment (1983: 282). However, many others take such cases to demonstrate that privacy is not just concerned with information. I am inclined to agree. Examples such as Moore's show us that ordinarily language ascribes a more capacious range of meanings to the concept of privacy than the informational view of privacy will allow. And while Arthur Schafer is right to note that our definition of privacy may not precisely match our ordinary usage of the term (2011: 13), our lexical definition of privacy ought to fit our intuitions as to when our right to privacy has been violated. These intuitions are liable to vary, however, I suspect there will be near universal agreement (certainly among ordinary users of the concept 'privacy') that the above cases are privacy infringements. Because of this, Parent seems

[20] Prosser (1955); Fried (1970); Parent (1983); and Westin (1967) all support some version of the control theory.

unjustified in artificially policing the meaning of privacy. Other people can violate our privacy in the absence of information, by smelling us, observing us, touching us, or listening to us sing.

This hunch is accommodated by accounts of privacy as access, which define privacy in terms of being able to negotiate physical or sensorial boundaries. Such views take privacy to be a domain or territory in which we are free from other people's access (Parker 1974; Allen 1988). Carissa Véliz describes this territory as *the autotopus* – a 'metaphorical sensorial personal space', 'the kind of space that people in a society commonly would not want anyone, other than him/herself (and perhaps a very limited number of other people chosen by him/her) to access' (2019: 5). When other people see, smell, hear, or touch us in this zone, they infringe our right to privacy (Ibid.). Véliz gives an account of privacy which combines accounts of privacy as access with accounts of informational privacy.[21] On this picture, an agent will have privacy when they are able to control other people's access to them: whether that access is informational, physical, or sensorial.

Nissenbaum is one notable theorist who stands outside of this standard framing. She takes privacy not to be about who has what information, but whether information flows appropriately. Different information norms will hold in different domains, so what counts as appropriate information flow will depend on the context of the situation, the social norms in play, and the kind of information in question (2010: 4). In this way, Nissenbaum ties adequate protection for privacy to specific contexts (2004: 101). Our privacy is violated when information is shared in inappropriate contexts. For example, when our therapist discloses our conversations on Twitter, or our medical records find themselves in the hands of insurance companies.

I suspect that each of these accounts reveals something important about privacy. We seem to have privacy when we are epistemically, physically, and sensorially inaccessible, when we are free from the prying minds, unbidden ears, eyes, and touch of others. Privacy also restrains the flow of information from one context to the next.

How do these conceptions of privacy relate to the interest-based understanding of rights I outlined in the introduction? By defining privacy as the spaces in which we are epistemically or sensorially inaccessible, or as obtaining when the information we share flows appropriately, we are giving a descriptive account of privacy – an account of where privacy obtains. We can hold any of these conceptions of what privacy is while maintaining that privacy is grounded in interests that we have; that these spaces and information flow conventions are valuable *because* they protect various interests that we have. Understanding our rights in terms of interests explains the function

[21] As do Moore (2003: 215) and Gavison (1980: 423).

of privacy – it explains why we value our right. And just as with other concepts, such as liberty and justice, we can describe our right to privacy in both descriptive and normative terms.

Véliz is one example of a privacy theorist who conceptualises privacy in this way. She maintains that sensorial and informational conceptions of privacy are unified under the heading of 'a right to privacy' because of the various interests that privacy protects (2018: 6). We have privacy when we are epistemically and sensorially inaccessible, and privacy functions to protect the various interests we have in not being accessed by others: our interests in not having other people knowing intimate details about our private lives, our interests in not being judged or ridiculed by others, and our interests in not being watched or heard (Ibid.).

2.3 Privacy and Self-Presentation

In this section I argue that our interests in self-presentation partly ground our right to privacy. I defend two central claims. Let's start with the first.

(1) One of the key functions of privacy is to protect our self-presentational interests. Our interests in self-presentation thus form some part of the normative basis of our right to privacy.

This claim is not unprecedented. A multitude of authors recognise the deep-seated connection between privacy and self-presentation.[22] Marmor is one prominent proponent of this position. He argues that there is 'a general right to privacy grounded in people's interest in having a reasonable measure of control over the ways in which they can present themselves (and what is theirs) to others' (2015: 4).[23] Our interests in self-presentation explain why we have a right to privacy; privacy is valuable because it makes self-presentation possible.

[22] See e.g., Gross (1971: 341); Gavison (1984: 336); Rachels (1975); Fried (1970); Nissenbaum (1998: 592); Posner (1991: 349); Nagel (1998: 20); Hadjin (2018); Cocking and van den Hoven (2018); Shoemaker (2010); Simpson (2020); Owens (2022); Eliot and Soifer (2022: 3); Shiffrin (2014: 144); Véliz (2018: 6); Marmor (2015; 2019); Räikkä (2016: 223); O'Callaghan (2016: 162) and Austin (2014: 178, 182). There are other privacy theorists, which, while they do not speak of self-presentation per se, consider privacy in relation to identity. Parker (2002), Michelfelder (2001) and Reiman (1976) explore how privacy allows us to determine and develop our conception of our self, and our conceptions of who we are and who we ought to be.

[23] There's an important sense in which 'what is ours' is central to the presentation of our selves. We may hide and organise particular objects in our room before a guest enters, wear particular jewels, or polish our cars. In doing so we are able to exercise some measure of control in 'determining the information an audience is able to acquire' (Goffman 1975: 98).

Consider how self-presentation is entangled with access. What you think of me depends on your access to me. You form impressions of me based on what you know, what you've seen, smelt, or heard. It is by setting limits on your epistemic and sensorial access to me (i.e., by having privacy) that I am able to exercise a measure of control over the impressions you form of me. As David Eliot and Eldon Soifer observe, self-presentational choices ordinarily require us to be in a position of epistemic privilege (2022: 3). And although it is sometimes true that our friends 'know us better than we know ourselves', most of the time we know ourselves better than other people do. I am the only person with direct access to my mind, and this epistemic privilege allows me to decide which of my infatuations, beliefs, and values I reveal. It is only through communicating to others that we are able to overcome the opacity of one another's minds (Shiffrin 2014: 9).

Self-presentational autonomy also requires spaces in which we are sensorially inaccessible; spaces in which we can figure out our views and deliberate our appearances before subjecting ourselves for social scrutiny.[24] Mane Hadjin calls these 'drafting spaces' (2009: 67). In order to make decisions about how and when we present ourselves to others, we need to have control over who – if anyone – gains access to us (Ibid.). This kind of control would be very difficult without drafting spaces. As Robert Simpson notes, drafting spaces do not have to be completely solitary – a confidential conversation between two friends is still in an important sense a drafting space (2020: 7). In such cases, we are permitted 'to explore unpublic feelings in something other than solitude' (Nagel 1998: 20). Here, rather than inaccessibility, it is conventions of confidentiality and limits on information flow which enable and safeguard our self-presentational interests.

It is very difficult to have a reasonable measure of control over which aspects of ourselves we reveal to and conceal from others in a world without privacy such as Thomas More's Utopia (2002). A central reason why we value privacy is because it makes self-presentation possible.

The second claim is a corollary of the first. Given that one of the main functions of privacy is to look after our self-presentational interests:

(2) We often lack a reasonable measure of self-presentational autonomy because we lack privacy. In this way, the wrongness of various types of privacy violations can be explained in terms of infringements of our self-presentational interests.

We can illustrate claim (2) by examining self-presentational infringements. Consider the following case:

[24] Cogen (2012) writes extensively on the value of privacy in preserving spaces for free moral and cultural play.

> BAD NEWS. You are walking out of the office and reach into your pocket to retrieve your phone. Glancing at your recent notifications, you lay your eyes on some very upsetting news. As you read on, you feel your stomach plummet towards your feet, and the message begins to blur. Better not, you think, rubbing your eyes with your sleeve. No need to make a spectacle of yourself. Besides, you might bump into colleagues this close to work. You swallow down the lump in your throat, and take the next left, recalling an unfrequented garden nearby. Finally finding yourself alone, your body crumples onto the nearest bench. You burst into a harrowing sob.
>
> Little do you know, the cobbled wall beside you is in fact a two-way mirror. Your colleagues are stood behind it and watch you as you clutch your abdomen and wheeze, gasping in between inarticulate grunts.

In Bad News, your self-presentational interests are clearly at stake. You have been denied a reasonable measure of control over which aspects of your inner life you revealed to others. The vital point for our purposes, though, is that you lack self-presentational control because you lack privacy. And part of what explains what is wrong with this privacy infringement is your lack of self-presentational control.

Marmor's analysis of self-presentational infringements is helpful here. Unless we can predict the flow of information about ourselves and the likely consequences of our actions, we cannot make informed choices about how and what we present to others (2015: 13). In short, self-presentational control requires us to have some idea of who we are exposing ourselves to. Now, things in public do not always happen as we expect. When we have conversations on park benches, someone we know might happen to overhear the content of our conversation, but this is arguably bad luck rather than a violation of privacy – some level of unpredictability is to be expected in the public realm (Ibid., 12). However, in Bad News, other people have manipulated, without adequate justification, the environment in ways that have significantly diminished your self-presentational control. Had you known they were keenly observing you, you would not have exposed yourself in the way that you did – you would have tried your best to hold it all together until you found the appropriate setting.

Another example which illustrates (2) is found in Simpson's discussion of drafting spaces:

> DATE. Suppose you are getting ready for a date. You're halfway through the process of styling your hair and face in the way that you want, and you're yet to settle on an outfit, when… knock knock! Your date is, to your dismay, overly punctual. What's worse, against your perfectly proper request to wait downstairs, they insist on perching at the end of your bed and watching you get ready.

You would most likely feel intruded on at this point (2020: 6). This is because you want other people to see the 'impressive *result* of you deciding how to present yourself', not the 'deeply revelatory *process* of you deciding how to present yourself' (Ibid., 7). In this case, you have been denied a drafting space in which to deliberate your self-presentational choices – you lack self-presentational autonomy because you lack privacy. Crucially, the self-presentational infringement partly explains why the privacy violation is wrong. The same is true of the self-presentational infringements we considered in §1.3. You lacked self-presentational autonomy when your wealth, infidelity, heroic act, and award was disclosed against your wishes, because you lacked privacy.

The idea behind (2) is this: many of us have the intuition that infringing upon people's self-presentational interests are wrong, and this intuition has some explanatory power in making sense of the judgements we have about the wrongness of various types of privacy violations. In order to safeguard ourselves from self-presentational infringements, we need to safeguard our right to privacy.

2.3.1 Interrelatedness

While many privacy theorists endorse the tight connection between self-presentation and privacy, there is disagreement as to how closely they are related. On the one hand, we have Marmor, who argues that self-presentation is the *only* interest underpinning our right (2015: 4).[25] On the other hand, Rachels argues that although self-presentation is 'one of the most important reasons why we value privacy', 'there are multiple interests which may be protected by guaranteeing people's privacy' (1975: 323). My own view is closer to Rachel's than Marmor's. I take our interests in self-presentation to be among the plurality of interests underpinning our right to privacy – they

[25] It is unclear whether or not Marmor has since changed his mind. In a more recent work, Marmor speaks of the '*main* interest protected by the right to privacy' (2021: 3 *emphases added*).

give us one, but not the only reason to uphold privacy.[26] Our self-presentational interests cannot be the *only* interests which ground privacy because there are (a) cases where privacy violations occur in the absence of self-presentational infringements, and (b) cases where our self-presentational interests are compromised but not our privacy. Both Räikkä (2016: 223) and Véliz (2022: 30) defend this position.

To give an example from Véliz. Suppose you become of interest to an Intelligence Agency – they follow you around, listening in to your conversations, but have no plans to disseminate any of the information they gather, you are only of interest to them because they are interested in some other person with which you are acquainted (2022: 34). It seems clear that your privacy has been violated, and yet your self-presentational interests do not appear to be at stake.[27] This is an example of (a). Another more commonplace example is how we voluntarily disclose aspects of our private lives. Here, we lose privacy (by waiving certain privacy related duties) yet retain self-presentational control.

There are also cases where (b) our self-presentational interests are infringed but not our privacy. Consider how certain schools ban piercings and 'unnatural' hair colours. These regulations prevent students from presenting themselves in certain ways; their self-presentational autonomy is thwarted, and yet our ordinary, intuitive ways of applying the concept of privacy would not interpret these regulations as infringements of privacy.

Claims (a) and (b) demonstrate that while privacy often supports self-presentation, we can have self-presentational control without privacy. And while privacy infringements often involve self-presentational infringements, not all privacy violations are injurious to our self-presentational interests. Self-presentational interests do not exhaust privacy infringements.

This is not to deny that self-presentational interests are appurtenant in discussions of privacy. I take privacy's protection of our self-presentational interests to be one of the main reasons why we value privacy, it's just that there may be other interests which ground privacy or otherwise explain why privacy infringements are wrong.

It may be helpful here to distinguish between central cases and peripheral cases. As Moore notes, given the range of contexts in which it may be claimed that we have or lack a right to privacy, there may be core features of the central cases of privacy which are not present in the outlying cases

[26] Notice that the interest theory of rights I am assuming leaves open whether or not our right to privacy furthers one or several interests.

[27] In this case, the loss of privacy may reveal the function of privacy in protecting us against the discomfit of being watched and heard.

(2008: 211). Aristotle distinguishes central, and peripheral cases in his discussions of friendship. Where, in attempting to define friendship one will be unable 'to do justice to all the phenomena of friendship; since one definition will not suit all' (EE VII 2: 1236a-16-31). The same may be said of our right to privacy – many violations of our right to privacy involve infringements of our self-presentational interests, but this does not mean that self-presentational interests are present in every privacy violation.

SUMMARY

One prominent view of privacy argues that privacy has to do with access. We have privacy when we have control over other people's access to us, whether that access is epistemic or sensorial. We also have privacy when the information we grant access to in one context does not lead to us inadvertently granting access to ourselves in other contexts – when the information we share flows appropriately. These are descriptive claims about privacy, they tell us where privacy obtains.

The central aim of this section was to defend a normative claim concerning the function of privacy – to argue that one of the key functions of privacy is to protect our self-presentational interests. The claims I will go on to make in this thesis do not depend on any particular descriptive account of privacy, they are compatible with any account which can accommodate the normative claim I have defended in this section. For ease of exposition, I will construe privacy in terms of sensorial and epistemic inaccessibility.

III

Diachronic Self-Presentational Autonomy

This chapter will address the first aim of this thesis: to vindicate a distinction between the *synchronic* and *diachronic* aspects of self-presentation. Most accounts of self-presentation are implicitly synchronic; they evaluate our interests at specific moments in time. I take this approach to be incomplete - our self-presentational interests also have a diachronic character. In short, it should matter to us – and does matter to us, upon appropriate reflection – that we retain some control over how and when other people observe our past self-presentations, whether an unabashed emo phase or an old political stance. These are diachronic self-presentational interests (or DSPIs).

In §3.1 I explore the Synchronic Approach and begin to motivate the diachronic turn by examining the notion of diachronicity in discussions of decision making, rational conflict resolution and theories of autonomy. In §3.2 I set out the Diachronic Account by exploring four key features of our DSPIs. In §3.3, I explore an objection to the Diachronic Account: do our DSPIs legitimise evading accountability for our past mistakes? I explore how sometimes, DSPI-infringements are all things considered permissible. Finally, in §3.4 I explore the relation between our DSPIs and diachronic privacy.

3.1 The Synchronic and Diachronic

The Synchronic Approach examines self-presentation in terms of how we present ourselves to others at any one time; whether our ability to present ourselves as being a certain way within a 'particular social interaction' (Austin 2014: 178, 182) or revealing some subset of information about ourselves in *an* act or *a* situation. Goffman describes self-presentation as the kinds of things we may or may not do while sustaining a performance before an individual (Goffman 1963: xi).[28] His distinction between 'on stage', and 'off stage' exemplifies an episodic perception of self-presentation, speaking as if we hop in and out of distinct episodes of impression

[28] See also (1975: 33); (Ibid., 34); (Ibid; 43) and (1959: 32) (Ibid., 43) where Goffman emphasizes our presentational efforts within a specific moment or encounter.

management (1959: 57). Schlenker also seems to embrace the Synchronic Approach, speaking of impression management within 'individual actions in real or imagined scenarios' (1980: 6). More recent contributions to the literature examine individual self-presentational choices online, where we re-frame texts and carefully select photos (Marmor 2021: 4; Cocking and van den Hoven 2018: 60).

Given that the Synchronic Approach examines how we alter our performances and expressions in particular circumstances, their analysis of self-presentation is often qualified. They consider how we present ourselves as industrious *in the workplace* or how our jovial side emerges *when with children*.[29] Examining our self-presentations within the context of environments or relationships is clearly pertinent. However, the diachronic character of self-presentation seems indispensable in understanding how we present ourselves within these environments and relationships. This is because our presentations within these contexts shift over time. The way that I present myself to my colleagues on the first day will be different to the way I present myself on the hundredth day, as I begin to loosen up after initial periods of reservation. The same is true of the self-presentational choices we make before our children; they transform as both parties in the relationship mature. The kind of humour we disclose to our children adjusts as that child goes from a toddler who sweetly giggles at slapstick humour, to a teenager who sardonically smiles at one-liners with political undertones.

Our relations with others are not static. They mature and decline in ebbs and flows: our acquaintances develop into close friends, our close friends become lovers, and our lovers become distant memories. As these relationships shift, so too do the aspects of ourselves we share with one another. When we begin to get to know another, there are aspects of ourselves we are not yet ready to disclose. We may feel exposed, vulnerable, and embarrassed as they come to know such parts of ourselves. As time passes, we often find mutual comfort, ceasing to find these exposures intolerable.

These shifts in our self-presentations seem inextricable from the dynamic and prolonged nature of relationships and social roles; understanding our self-presentational interests, even within the context of individual role or relationships, requires us not only to consider our self-presentational choices at any one time, but to recognise the sense in which our self-presentational choices adjust as time passes. This is the *diachronic* character of self-presentational autonomy.

[29] See Rachels (1975: 326); Schlenker (1980: 5, 36); James (1990: 293); Cocking and van den Hoven (2018: 61); Mead (1934: 142).

The diachronicity of self-presentation is also observable in the temporally extended process of impression management. When I attempt to induce another person to form a particular impression of me – to think me charming and warming – I do so over an extended duration of time. Others do not ordinarily think we are trustworthy or credible on the basis of one trustworthy or credible interaction. Instead, such qualities are earnt over time – we have to be consistent in our performances. Many of the expectations tied to trustworthiness, such as one's ability to keep a secret or not break a promise, may only be evinced over time. Saying this is not to deny that other people *can* form impressions of us in a single interaction, i.e., in a synchronic self-presentation, but to acknowledge that the process of impression management is temporally extended. If I were in an inexplicably hot-tempered mood and responded snappily to your query, you may quite reasonably form an unfavourable image of me on the basis of this interaction. However, it seems relevant to consider ongoing episodes: did you ever come to understand why I had been especially irascible? Did this lead you to see me differently?

My claim is not that the Synchronic Approach is fundamentally misguided – clearly, it is important to evaluate our self-presentational interests at particular moments in time. However, we also ought to consider how our public self-identity extends and shifts – how our self-presentational choices change with the passage of time.

Others have acknowledged the extensional nature of impression management. Susser writes that 'the images of us that others carry in their minds are not static images; they change constantly in light of new revelations… shaping those images is not a one-time task' (2016: 5). This claim echoes Nagel, who remarks that 'managing what appears on the surface… is the constant work of human life' (1998: 5). Goffman also seems to recognise the importance of consistency within impression management. Stressing the gravity of 'the initial interaction in an extended series of interactions involving the same participants' (1959: 16). In order to be authoritative, a teacher must be so from the very first interaction with their students – one must maintain face from start to finish (Ibid.). While these various authors recognise the temporally extended nature of our self-presentational efforts, no one (to my knowledge) has treated the diachronicity of self-presentation as a focal point for their analysis of the concept, nor with respect to the ethical concerns attached to it. This chapter seeks to address this gap.

3.1.1 Broader Discussions of Diachronicity

In developing the Diachronic Account, I will be utilising insights from parallel philosophical work on autonomy, the self, and future-directed decision making. These accounts utilise the notion of diachronicity to capture our tendency to change over time. L. A. Paul speaks at length on diachronic decision making involving *transformative experiences* which fundamentally alter our preference structure by shifting our epistemic or mental states (2014: 6). Paul seeks to understand how we can rationally have preferences transcending transformative experiences; how we can rationally hold a preference at t1 to be in some later state t2, when the preferences we hold at t2 are inconsistent with those at t1. Agnes Callard also seeks to shed light on an aspect of diachronicity in human decision making, taking aspirational conflict resolutions as her focus. Aspiration is the diachronic process by which an agent effects change on their ethical or aesthetic point of view – it's the rational process by which we work to care, value or desire something new (2018: 4). I might aspire to be the kind of person who enjoys wine. I try my best to notice the subtle notes of oak, cherry, and blackcurrant as I tentatively sip what tastes to me as vinegary as the last. I do so, not because I enjoy wine, but because I aspire to be the kind of person who does. Callard asks how we can make sense of my aspiring to enjoy wine, despite lacking the cultivated aesthetic qualities to tell the difference between a rich buttery oaky wine and an aromatic fruity one. Both Callard and Paul seek to explain the distinctive rationality at stake when we strive to do something which radically alters our outlook – it is the diachronic features of decision making which require explanation.

Some accounts of autonomy also utilise this notion of diachronicity. There are broadly two different kinds of autonomy being theorised: (1) autonomy as a property of persons, which is most plausibly theorised diachronically, and (2) autonomy as a property of desires or preferences, which can plausibly be theorised in either diachronic, or synchronic terms.

Starting with (1), most sophisticated theories of autonomous personhood recognise that autonomy is diachronic ideal, an ideal whose realisation depends on how one's life progresses. Autonomy, on this view, is a global rather than local concept – it is a property of an individual's entire life (Pettigrew 2022: 206). Gerald Dworkin writes that autonomy 'evaluates a whole way of living one's life and can only be assessed over extended portions of a person's life' (1981: 16). There are a number of reasons why diachronicity is central to this kind of autonomy. We, the objects of self-governance, are temporally extended entities, so whether or not we are autonomous is not simply a function of our way of living at a particular point in time. Instead, it is a function of our long-term values, projects, commitments, and cares. While we have desires and values that express

themselves episodically, the grounds of our self-esteem ranges across the sums of episodes and moments. Whether or not I hold myself in high self-regard isn't a function of how well my values and intentions are realised at any one moment in time, but rather in how they are realised in the sum of all of those moments. Take a spirited desire to nurture close friendships. One may both be pleased with oneself in an overarching sense that this desire has fore fronted one's conduct, and yet be episodically disappointed that one currently holds what they take to be a less valuable pursuit in higher regard.

Turning now to (2). Accounts which take autonomy to be a property of desires conceptualise autonomy locally rather than globally – autonomy is a property of persons at a particular time, including the choices they make at that time (Pettigrew 2022: 206). While it seems implausible for accounts of autonomy as a property of persons to overlook the diachronic character of autonomy, not all accounts of kind (2) are diachronic: they may plausibly consider whether or not some desire or preference is autonomous relative to a particular moment in time. However, in *Autonomy and Personal History,* John Christman argues that accounts of autonomy as a property of desires should take a diachronic approach (1991). Christman argues against 'Coherentist Accounts' of autonomy – who argue that an agent acts autonomously only if their motivation for acting is harmonious with the mental state or action guiding it (Shew and Garchar 2020: 31). We need to know not only whether or not an agent identifies with their desires at one moment in time, but whether or not they identify with the process by which they came to form that intention – i.e., the diachronic character of desire formation (1991: 1-2).

I will be employing these general ideas of diachronicity to reconsider our self-presentational interests from the perspective of its temporal extendedness; analysing self-presentation in light of our history and future, the sense in which we develop our identities, and in view of our tendency to change our minds and take different directions.

We can see now why distinguishing between the different aspects of our presentational self – i.e., between our self-revelations, self-presentationally autonomous choices and the self-presentational choices we are not custodian of – is central to the Diachronic Account. There are parts of our public appearances that we shape, curate, and change. It is recognition of these shifting parts that motivates the diachronic turn.

3.2 Past Self-Presentations

Our self-identity will alter over the course of our lives. I am in many ways different to the person I was ten, even three years ago. This underlying change of who we are and who we take ourselves to be coincides with changes in self-presentation: parts of self-identity that we did not want to share, become parts of ourselves that we do, and parts of ourselves that we once shared, become parts of ourselves we wish to hold in private. The Synchronic Approach overlooks these diachronic features of self-presentation. Specifically, it overlooks our interests in exercising control over other people's access to our *past* self-presentational choices; those parts of our self-identity we no longer wish to share.

On the one hand, part of our interests concern how we present ourselves to others at any one time: we want to appear quick-witted on a first date, maintain an aura of professionalism on our first day at work and reveal only a subset of our self-identity when meeting our in-laws. These are our *synchronic* self-presentational interests. On the other hand, we have interests in having a measure of control over how we present our past selves to others: the beliefs we used to hold and clothes we used to wear. This isn't just about how we're perceived at a particular moment in time, but in how we're perceived now, on the basis of our past self-presentations. This clearly concerns our interests in self-presentation, yet it seems importantly different from the synchronic side of our interests.

Before explicating the Diachronic Approach, it's worth acknowledging the somewhat pseudo-paradoxical character of our ways of conceiving of these identity shifts. Someone who has undergone a gender transition, may see themselves as an entirely different person to the person they were pre-transition, describing their past self as 'not the real me'. But they might also want to say things such as 'I always knew that the gender I was assigned wasn't my real gender' which suggests that the pre-transition person who knew this was them. The point of the example is that our ways of describing and understanding our own (and others) diachronic selves is somewhat elusive. On the one hand, an alteration of one's identity plausibly entails that two different identities exist. But the fact that it is *their* identity that is altering, suggests that one entity endures the change. I have assumed that all the versions of ourselves that exist across time are time segments of the same diachronically continuous entity, so it's not that these shifts prevent positing interests which traverse our diachronic selves. Rather, it is to recognise the somewhat elusive character of the claims I will go on to make in this chapter, which will become pertinent as a result of the shift to the diachronic approach.

In order to vindicate the diachronic character of our interests, I will outline cases in which our current self-presentational efforts are undermined by the self-presentational choices we made in the past. We have DSPIs because our past self-presentations can impede our present ability to exercise a reasonable measure of self-presentational control.

Each of the following cases have the following features. Firstly, the ways in which these individuals present themselves has changed considerably; their past self-presentations differ from or conflict with their current self-presentational choices. I will refer to these conflicting presentations as *prior self-presentations*. Secondly, their prior self-presentations ground negative emotions of self-esteem – the individuals are ashamed of, embarrassed about, or regret their past presentations.

RHINOPLASTY. Nora was born with an unfortunately long and rather prominent nose. It was positively crooked, and perpetually runny. For much of her childhood, Nora's nose caused a great deal of torment. She despised looking at herself and welcomed its reconstruction as soon as the opportunity arose. Nora's nose is now delicate, dainty, and perfectly upturned. She is finally happy with the way that she looks, but dreads photos of her past resurfacing.

GENDER DYSPHORIA. After a lengthy period of social and medical transitioning, Morgan finally feels comfortable in her own skin. She has found the hair, make-up, clothes, and voice that feels authentic to her. The thought of other people seeing photos of the way she used to look is deeply disturbing. Morgan wants other people to see her as she is now, not on the basis of a perceived gender she does not want to project.

AMENDED ATTITUDE. Isaac used to detest non-human animal testing, believing it to be inherently cruel. Looking back, Isaac considers his teenage opinions to be misinformed and naïve. He had overlooked several (now obvious) countervailing moral reasons and would be humiliated if anyone read his callow and outdated views.

STRUGGLE. Kurt is deeply embarrassed of the way he used to look. He was in poor physical shape, and the clothes he wore were scruffy and stained. Since then, Kurt has put in considerable time and effort into his appearance. He is proud of how far he has come

but wants the space to move on from his past. To avoid prying eyes he meticulously reviews family photo albums to eradicate evidence of him at the time.

ALTERED APPEARANCE. Elma used to be an 'Emo Girl'. She constantly listened to melodically aggressive hardcore punk, wore striped socks, studded belts, and fingerless gloves. Now an adult, Elma's style preferences have drastically changed – she is deeply mortified by her emo-phase and tries her best to conceal her past from others.

For each of these cases, we can imagine someone happening across evidence of these prior self-presentations. They scroll down Twitter and find Nora's crooked nose, find a crumpled-up photo of Morgan pre-transition, and discover Isaac's article in a university archive. This evidence of their prior self-presentations threatens to set back an interest that the protagonist has in each of these examples, because it gives other people the means to observe them on the basis of their bygone selves. It seems that these individuals would be wronged were someone to distribute this evidence against their wishes. To do so would be to expose them to the world in ways they do not wish to be perceived. This would clearly undermine their self-presentational interests. But it would frustrate them in a specific, and as I have argued, underexplored way. Their ability to mediate their public self-presentation is being compromised because of other people's access to their prior and outdated selves; presentations which no longer conform with their conception of a publicly acceptable self.

It seems that these individuals ought to be able to conceal evidence of their past from others, whether the beliefs they're embarrassed about, the appearances they're ashamed off, or the aspects of their past that never felt like them. In short, they ought to have a measure of control over the terms by which their past is knowable to others, and this gives them grounds for objecting to the unbidden exposure or distribution of their prior self-presentations. The intuition that we ought to have some control is captured by our DSPIs. All this is to say that we seem to be justified in asserting DSPIs, not because of some deeper explanation or ideal, but because we intuitively feel the pull of the interest when we're thinking about these sorts of cases. If you share the intuitions I have, then your intuitions bespeak your recognition of our DSPIs.

Notice that these exposures may differ in their ramifications. Some DSPI-infringements may result in a mild embarrassment, others a deeply traumatic violation of our interests, in some cases permanently shaping the impressions other people form of us. How severe the infringement is

may be a function of both the magnitude of the exposure: how deeply the exposure penetrates into one's social circles, and the extent to which the prior self-presentations ground negative emotions of self-esteem. But these will be differences in degree rather than kind – any exposure of our prior self-presentational choices which exposes us in ways we do not wish to be perceived is injurious to our diachronic self-presentational interests.

Analysing these cases in light of the features I outlined in §1.3 will prove fruitful:

Accuracy. Each of these exposures involve a construal of an individual which was accurate – the evidence of their past is not fabricated.

Willingness. DSPI-infringements occur against the Willingness dimension. In each case, the agent's prior self-presentational choices are observed and or distributed against their wishes.

Valence. DSPI-infringements do not necessarily involve a negative construal of one's past – someone might distribute a 'before-and-after' photo of Morgan in order to celebrate trans women, and yet undermine Morgan's DSPIs despite portraying her positively.

Voluntariness. Our DSPIs may be undermined by both the aspects of ourselves that we decided to share with others (beliefs, opinions, hairstyles), and the aspects of ourselves that we couldn't avoid sharing (our nose shape, gender, physical features). Because of this, our DSPIs advocate a measure of control over both our prior self-presentational choices, and prior self-revelations. Notice that although our prior self-presentations may not have been self-presentationally autonomous at the time, one must exercise self-presentational autonomy over their appearance in order for it to be a prior self-presentational choice. A prior self-presentation is by definition a past self-presentation which conflicts with our current self-presentational choices, because of this, our DSPIs will be undermined by a prior non-voluntary self-presentation only if we have exercised a degree of control in shaping it.[30] For example, by having surgery to change a facial feature, or having speech therapy to change our tone of voice.

Involvement. The involvement dimension concerns our participation in the process by which we are perceived and understood. By obtaining and or distributing another person's prior self-presentations, we position these aspects of their past into their public presentation without

[30] Strictly speaking, conflicting presentations may occur even if both the past and present appearances are non-voluntary. There are ample examples of changes appearances which occur at no deliberate effort of individuals. Changes that come about through hormonal changes in puberty: skin conditions such as acne which appear in ebbs and flows, visible health conditions which may diminish over time. Such individuals may desire to conceal images of themselves when they were younger even if they were unable to change it themselves.

their permission. This denies them the opportunity (with respect to this particular presentation) to define who they are in the present.

Let us take stock. DSPI-infringements involve accurate representations of our past, they expose aspects of our past we are not willing to share, they may construe us positively or negatively, they may involve both past self-presentationally autonomous choices (outfits and opinions) and past self-revelations (gender, nose-shape), and they frustrate our involvement in the process of constructing our public self-image. Our diachronic self-presentational autonomy is undermined when our ability to shape our public presentations are undermined by prior self-presentations i.e., ways that we used to present ourselves which differ from or conflict with the kinds of self-presentational choices we make now. By undermining our DSPIs, other people deny us a key aspect of self-determination. While we can never fully control the public construal of ourselves, by others, on our own terms, DSPI-infringements wrest us away from having a reasonable measure of control over the terms by which our past is knowable to others.

I will proceed by analysing three pertinent features of our DSPIs. In §3.2.1 I explore prior-self-presentations. In §3.2.2, I explore how prior self-presentations ground emotions of self-esteem. Finally, in §3.2.3, I explore why DSPI-infringements involve our self-presentational, rather than non-self-presentational past.

3.2.1 Modified Selves

By marking a distinction between our current and former self-presentational choices, I am not marking a distinction between the self-presentational choices I am making this instant, and all of those I made over an hour ago. Rather, I am differentiating between different *kinds* of self-presentations. In this way, the notion of a prior self-presentation has built into it a divergence, or conflict with the self-presentational choices we make now. This section will explore the sense in which our self-presentationally autonomous choices shift over time.

We all have some conception of who we are as a person. We have certain beliefs, preferences, personality traits and feelings, and self-presentational control allows us to decide which of these different parts of ourselves we express to others. Over time, these aspects of ourselves change – the things we prefer, value, desire, and believe shift over the course of our lives. It is because of these shifts that our past self-presentations differ from those present.

This uncoupling of our former and current self-presentational choices is exemplified by our tendency to dissociate from our past self-presentations:

> 'At the age of twenty I sent to the Academy of Lyons various writings which I subsequently withdrew. When I read them, I found that their author deserved to be whipped. What ridiculous things I said and how annoyed I would be if they were preserved' - Napoleon Bonaparte (Markham 1963: 11).
>
> 'Very little about the Joseph of a year ago pleases me. I cannot help laughing at him, at some of his traits and sayings' – Saul Bellow's *Dangling Man* (1994: 21).

You may have experienced a similar sense of recoil reading an old diary or looking through teenage photo albums. We blush recalling the unabashed goth attire we once thought becoming and feel regret recalling the flagrant jokes we once thought witty. We tend look back on our past with an odd mixture of both familiarity and foreignness, knowing that it was in an important sense us, and yet finding it absurd, embarrassing, comical, or shameful, precisely because we would not present ourselves in those ways now. That we divorce our current self-presentational choices from those prior bears witness to the diachronicity of self-presentation; the ways in which we make different self-presentational choices over the course of our lives. It is because of these shifts that we have interests in exercising control over other people's access to our past.

In his analysis of privacy, Hadjin uses an example of drafting letters which is illustrative of this feature of our DSPIs. When I write a letter, there may be various drafts that I write before settling on the final version. As I make changes, I revise what I say, developing my thoughts until I say to myself "Yes, this is it! This is the letter I want to send out!". For Hadjin, the example is supposed to illustrate the boundary between figuring out what we want to say in drafting spaces, and the public space in which we actively present ourselves for judgement (2009: 66). For my purposes, it's a useful way to illustrate the diachronicity of our self-presentational choices, and the boundary between those past and present. For just like drafts of letters, there are various views that we've had in the past, as we've been going through life, figuring out what we think, who we are, what kind of self-presentations we want to give off. These prior beliefs, attitudes, and appearances are no longer congruous with our perception of a satisfying public version of ourselves. They are drafts – old versions of letters we no longer wish to share.

3.2.2 Self-Identity, Past and Present

In Chapter One, I described how the aspects of our self-identity that we want to present to others forms only a subset of our self-identity. DSPI impediments involve the exposure of prior self-presentational choices which ground negative emotions of self-esteem *at the time of the infringement.* These prior self-presentations are part of our self-identity, but they are not aspects of our self-identity we want to share.

Given that we can be proud of or embarrassed about our past, our self-identity is made up of both properties we currently instantiate, and properties we no longer instantiate. I used to have the belief that animal testing was wrong, that black eyeliner was epic. I now think that animal testing is sometimes permissible, that black eyeliner is tragic. Our conceptions of ourselves in the past and in the present are thus relevant to the question of 'who are you, really?'.

Some DSPI-infringements can be explained by shifts in the subset of our self-identity we want to present. Parts of our self-identity which were once congruent with our conception of a publicly acceptable self – aspects of ourselves we were willing to share – become parts of ourselves we want to hold in private. Such shifts can be explained by changes in the grounds of our self-esteem. In Amended Attitude and Altered Appearance, the prior self-presentations previously grounded positive emotions of self-esteem – the agents were proud to utter their opinions and show off their appearances at the time of making them. However, with the passage of time these same opinions and appearances have come to ground negative emotions of self-esteem. These shifts from grounding positive to negative emotions of self-esteem are common in cases where our prior self-presentations are straightforwardly subject to volitional control; the paradigmatic cases of self-presentationally autonomous choices I discussed in Chapter Two.

On the other hand, there are cases in which our prior self-presentations were never congruous with our idea of a publicly acceptable self; cases where our past presentations never grounded positive emotions of self-esteem. This is true of Rhinoplasty, and Gender Dysmorphia. Nora felt embarrassed about her nose then, and she feels embarrassed about that same nose now. Morgan was ashamed of her appearance then and she feels ashamed of that same appearance now. While these parts of their appearances formed part of their presentational self, they never fell into the subset of their self-identity they wanted to share. Prior self-presentations which grounded negative emotions of self-esteem *at the time,* tend to be self-presentations that were not straightforwardly amenable to volitional control (i.e., self-revelations rather than self-presentationally autonomous choices) – the peripheral cases I discussed in Chapter Two.

When other people undermine our DSPIs, they expose us to the world in ways we do not wish to be perceived- they reveal aspects of our self-identity we wanted to hold in private. Our DSPIs give us grounds for objecting to others because they wrest us away from having a reasonable measure of control over the public presentation of our past.

3.2.3 Non-Self-Presentational Past

The cases I have examined thus far involved our current self-presentational autonomy being undermined by our self-presentational past (whether self-revelations or self-presentationally autonomous choices). But our current self-presentational autonomy can also be undermined by our *non*-self-presentational past. Consider the following case:

> DIARY. You are spending Christmas in your childhood home. You pop on The Pogues (a classic) and begin the monotonous task of peeling potatoes. All of a sudden you hear raucous laughter next door. Curious, you meander into the neighbouring room – your entire family is stood around your father who is reading a passage from a tattered red book. You recognise it immediately – your diary. They are laughing at your expense.

In this case, your self-presentational interests are not set back by a prior self-presentation. A diary is, by convention, something private – the thoughts and feelings recorded within it are not aspects of yourself you once disclosed to others: intentionally or otherwise. In Diary, your self-presentational interests are set back by your past, but not your presentational past.

There is a distinction then, between impediments to self-presentation which arise because of (1) other people observing our self-presentational past (a letter, an outfit we wore in public, a speech we gave at our friend's wedding) and those which arise because of (2) other people observing our non-self-presentational past (a letter we never sent, a diary entry of thoughts we never disclosed, a photo we took but never shared, an outfit we tried on but never left the house in).

Cases such as Diary are pertinent in the context of both self-presentational control and privacy, but DSPI-infringements involve our self-presentational past, rather than our non-self-presentational past. This is because self-presentational infringements of kind (1) are diachronic in a way that setbacks of kind (2) are not. To illustrate this, let's compare Diary with a second case: Blog. In Blog, the agent expresses these same views, but publicly. For simplicity, stipulate that in

each case the aspect of one's past is an opinion, and that it is the agent's family who observes, exposes, or critiques it. The central difference between Diary and Blog is that Diary involves non-self-presentational past, and Blog involves self-presentational past.

Call t1 the time at which the agent expressed their opinion (whether written in a diary or blog) and t2 the time at which their self-presentational autonomy was setback. We can then consider two variations of these cases. In 'Old', time t1 was 10 years ago. In 'Recent' time t1 was yesterday. These temporal deviations modify the cases in the following ways:

> OLD. With so much time passed, the agents no longer agree with their opinion at t1.
>
> RECENT. With so little time passed, the agents still agree with the opinions they expressed at t1.

Along these two dimensions, we have four variations: Recent Diary, Old Diary, Recent Blog, and Old Blog. Each of these cases are alike in Awareness, Accuracy, and Involvement, but they differ in Willingness.

Let's start by comparing Recent Diary with Old Diary. In both variants of this case, the agent was never willing to disclose the opinion that had been exposed. A diary, whether ten years old or ten weeks old, is a form of self-expression intended for their eyes only. The fact that they were never willing to disclose these opinions means that your family never ought to have read them. This is true whether the diary was written yesterday, or ten years ago. So, introducing diachronic features to the case (i.e., a period of time, along with a change in attitude) is immaterial to the wrongness of the exposure.[31]

But introducing diachronic features to Blog does make a difference; it alters our intuitions on the wrongness of the exposure. While it seems apt for the agent's family to critique the agents' opinions from yesterday (at most mean spirited), it seems wrong (clearly worse) for the agent's family to dig up and critique the agents' opinions of ten years ago. Both Recent Blog and Old Blog involve a consciously public self-presentational choice: an active decision to disclose one's opinions to the

[31] We could imagine a variant of this case in which one's past is so far-fetched and detached from one's self-identity that one does not care at all that their family are reading it. I would still take this to be wrong on account of infringing their right to privacy, but one may question whether this really impedes our self-presentational interests. To avoid this complication, we can assume that the diary entry from 10 years ago still grounds emotions of self-esteem.

world, so the difference in wrongness cannot bear on that. Instead, it bears on the diachronic features of the case. The notion of Willingness is pertinent here. In Recent Blog, the agent is still willing to disclose the opinions in question. In Old Blog the agent is not. While at t1 they held these views and expressed them willingly, at t2 they no longer hold and no longer want to express them. This shift in willingness is relevant to the wrong in question – the wrongness of the exposure bears on the diachronic features of the case.

While there are cases in which our self-presentational interests are undermined by our non-self-presentational past, these cases are not diachronic in the sense I am interested in. Because of this, our DSPIs concern our presentational past, not our non-self-presentational past.

Our DSPIs are our interests in being able to exercise a measure of control over other people's access to our past self-presentations. These interests are infringed whenever our self-presentational interests are compromised by other people observing our prior self-presentations without our permission, where these prior self-presentations ground negative emotions of esteem and are different from our current self-presentations (in that we have exercised self-presentational autonomy in changing them). When other people undermine our DSPIs, they wrest us away from having a reasonable measure of control over the terms by which our past is knowable to others. This gives us grounds for objecting to exposures on the basis of our right to privacy.

3.3 Skeletons in the Closet

Exposing other people's prior self-presentational choices seems wrong – it sets back an interest we have in controlling the terms by which the grounds of our self-esteem are knowable to others. But what about cases where the prior self-presentational choice is itself wrong?

Many of us publicly express views we later come to regret, we endorsed values we now find abhorrent, and wore disrespectful costumes we are now ashamed of. By advocating a legitimate ethical interest in exercising a measure of control over the public presentation of our past, our DSPIs appear to problematically excuse individuals from morally reprehensible past self-presentations. They seem to grant individuals a right to hide from their past mistakes – to keep their flaws and moral transgressions out of sight. Are the individuals pictured smiling at a lynching, the people who wore blackface for World Book Day, or the individuals who repeatedly made deplorably racist remarks *wronged* when someone exposes their prior self-presentational

transgressions against their wishes? The wrongness of the DSPI-infringement seems dubious. To the contrary, it seems good that we're able to expose blameworthy prior self-presentations, whether to hold individuals accountable, express value judgements or deter would-be wrongdoers.

This section seeks to clarify when DSPI-infringements involving prior mistakes are all things considered permissible. I defend the following claim: it is not *always* wrong to expose prior self-presentational choices contrary to an individual's preferences. Sometimes, exposing the skeletons in the closet is a good thing. However, it is still sometimes wrong to expose other people's prior transgressions, and the intuitive wrongness of these exposures are not overridden by countervailing intuitions about the appropriateness of holding individuals accountable.

Whether or not a DSPI-infringement is wrong seems to hinge on whether or not we think the self-presenter is blameworthy for their prior transgression: when a prior self-presentation is not blameworthy, a DSPI-infringement will be wrong, and when a prior self-presentation is blameworthy, we are liable to doubt the wrongness of the DSPI-infringement. In analysing our self-presentational mistakes, I will utilise Andrew Khoury and Benjamin Matheson's distinction between diachronic and synchronic blameworthiness. Where synchronic blameworthiness is the extent to which an agent is blameworthy at the time of committing a morally objectionable act – call that time t1, and diachronic blameworthiness is the extent to which an agent is blameworthy for an earlier action at some later time – call this time t2 (2018: 206).

While an agent who commits a morally objectionable act is liable to blame in principle, there may be exculpatory factors which make blaming inappropriate in practice. In particular, there may be exculpatory factors which make it such that, even though blame was appropriate at t1 (they are synchronically blameworthy at the time of acting), blame is inappropriate at t2 (they are not diachronically blameworthy at the time of the DSPI-infringement). My analysis of the relation between blameworthiness and the wrongness of DSPI-infringements will be somewhat speculative, in seeking to suggest connections rather than establish them.

Exploration of the following cases reveals that the extent to which an agent ought to be held accountable for their prior actions can reduce as time passes. It is this dearth of diachronic blameworthiness at the time of the exposure which explains why some DSPI-infringements are wrong even if they involve prior transgressions. In each of these cases, the agent committed a moral infraction at some unspecified time t1.

> The agent, A was caught in the grip of a visible addiction.
> The agent, B wore a culturally insensitive fancy dress costume.

The agent, C made a racist utterance.

The agent, D smiled at a lynching.[32]

Let's assume that all of these agents satisfy the conditions for synchronic blameworthiness: they are blameworthy at t1. And stipulate that, at t2, evidence of their past self-presentational choices emerges, unbidden. For each DSPI-infringement, the following conditions hold: the past self-presentational choice is incongruous with their current self-image, the prior self-presentation grounds negative emotions of self-esteem, and the exposure occurs against their will but in their knowledge. The central difference between these cases is the severity of the moral infraction – A being a relatively minor moral infraction, D being a serious moral transgression.

To determine whether or not the agents would be wronged by these DSPI-infringements, we need to consider whether the agents are diachronically blameworthy at t2 for their publicly enacted choice at t1. Compare agent A to agent D. While it seems unfair to unearth a photo of A's previous addiction against their wishes, I feel little sympathy towards D. In this way the nature and severity of the bad act being dredged up from the past seems to affect our blameworthy intuitions, and in turn, the wrongness of the exposure. Assigning blame to each of these agents, however, will require us to fill out a few more details: how much time has passed since the mistake and the moment of its present revelation? Does the agent have the same psychological characteristics that led them to perform the prior self-presentational choice? In what sense do these self-presentations ground emotions of self-esteem? Analysing variations of these cases demonstrates that exculpatory factors may come into play as time passes.

Scenario One: B_1 was 14 when they wore the offensive costume, they are now 30. They feel deep remorse for their past mistake, and have changed a lot since then, so much so, that they'd never dream of wearing such disgraceful attire today.

Scenario Two: B_2 was 30 when they wore the offensive costume, they are now 32. While they have come to regret their choice of outfit, they still find it slightly funny. They have changed as much as one can change in two years, and it's not clear whether or not they'd do the same again.

[32] Lynching photos which were distributed in the late 19th and early 20th century frequently captured white people smiling as willing participants of the 'spectacle' (Salahu-Din 2017: 10).

I have the intuition that a DSPI-infringement would be wrong in Scenario 1, but not in Scenario 2. I also have the intuition that while B_1 is *not* diachronically blameworthy, B_2 is.

Where a moral infraction is relatively minor, but little time has passed, an agent does not feel particularly bad about their past, and they hold many of the beliefs and motives which led them to perform the act in the past, it seems legitimate to expose their past. On the other hand, if a moral infraction is relatively minor, a long time has passed, an agent feels deeply negative emotions of self-esteem towards their past, and they no longer hold the beliefs and motives which led them to perform the act in the past, then it seems wrong to expose their past. Crucially, the wrongness of the DSPI-infringement in Scenario One is not overridden by countervailing intuitions about the appropriateness of holding them accountable for their past. This is because there are exculpatory factors which make it the case that even though blame was appropriate at t1 (they are synchronically blameworthy), blame is inappropriate at t2 (they are not diachronically blameworthy).

Scenario One reveals that there are cases where we ought to be able to manage the public presentation of our past mistakes. However, this does not preclude cases in which DSPI-infringements are permissible. As with Scenario Two, calling B_2 out on their past mistakes might be a good thing.

The situation is different when considering serious moral infractions. Consider agent D, who willingly and happily attended a lynching. Given the severity of the moral transgression, I suspect that there are relatively few circumstances in which exposing this past public appearance would constitute an objectionable DSPI-infringement. This is because there are relatively few exculpatory factors under which D would *not* be diachronically blameworthy: even if a long time has passed, and D feels terrible remorse. Perhaps this can be explained by a doubt that D has really changed – there are surely traces of the motivations and character traits that led D to find amusement in such a morally despicable act.

Other cases may be less clear. If someone made a racist remark when they were 17, and 3 years have passed, are their DSPIs infringed when this self-presentational remark is exposed in the present? The ambivalence we feel may be explained in part by our uncertainty of the relevant psychological change – it is difficult to know whether they have really turned over a new leaf.

These cases reveal how our intuitions about the wrongness of DSPI-infringements are liable to vary. I suspect something like the following will hold:

Variable	An Agent will be Synchronically and Diachronically Blameworthy:	An Agent will be Synchronically but not Diachronically Blameworthy:
Severity of the moral transgression:	Serious	Relatively Minor
Age at which agent committed the moral transgression:	Mature	Immature
How the moral transgression grounds Self-Esteem:	Mild Embarrassment	Deep Remorse
How their character, beliefs and motives have changed:	Largely the Same	Significantly

If we are diachronically blameworthy for our past mistakes, then exposing our past may be permissible. If we are synchronically but not diachronically blameworthy for our past mistakes, then our DSPIs legitimize exercising a measure of control over the terms by which our past

mistakes are knowable to others. Revealing other people's past mistakes is not *always* wrong. Instead, it is sometimes wrong to expose other people's past mistakes.

While it is pro tanto wrong to infringe other people's DSPIs, the wrongness of a DSPI-infringement may be outweighed by the right-making features of the act in question – the DSPI-infringement may be permissible all things considered. When I expose someone who attended a lynching, or someone who still finds the racist comment they made 'funny' then I infringe their DSPIs, but my actions may be all things considered justified. In such cases, the normative presupposition of the wrongness of the DSPI-infringement is suspended because the infringement serves other valuable purposes whose achievement depends on the exposure. Valuable purposes such as holding agents accountable or reaffirming norms against certain actions in the present. If we allow other moral considerations to factor into our moral evaluations, then we can allow DSPI-infringements to come apart from moral impermissibility – we may encroach on an individual's DSPIs without it being, all things considered, wrong.

DSPIs do not advocate evading accountability for our past, but they do recognise that the extent to which we ought to be blamed for our past may diminish as time passes. When this happens, we ought to be granted the breathing room to unhitch ourselves from our past mistakes. I do not take this to be advocating an ignorant future, but one which acknowledges that we often come to realise our flaws and learn from them with the passage of time.

3.4 Diachronic Privacy

In Chapter Two, I explored the relation between our self-presentational interests and right to privacy. I defended two claims: (1) one of the key functions of privacy is to protect our self-presentational interests, and (as a corollary of the first) (2) we often lack self-presentational control because we lack privacy. In this section, I explore the relation between our DSPIs, and diachronic privacy, which we have when our prior self-presentations are private.

The diachronic counterparts to the claims above are as follows:

(1*) One of the key functions of diachronic privacy is to protect our DSPIs.

(2*) DSPI-infringements often occur because we lack diachronic privacy.

Starting with (1*), it should matter to us – and does matter to us, upon appropriate reflection – that we retain some control over how other people observe our prior self-presentations.

Diachronic privacy is one of the main ways we can ensure we have diachronic self-presentational autonomy. Where having diachronic self-presentational autonomy is having a reasonable measure of control over the ways in which other people gain access to our past self-presentations We ought to value diachronic privacy because it protects our DSPIs. In this way our DSPIs provide us with some normative basis for a right to diachronic privacy. As I argued in §3.3, it's not that we should *always* be able to control other people's access to our prior self-presentations – sometimes the skeletons in our closet ought to come out. However, we ought to have *some* control over people's access to our past – we should have *some* diachronic privacy.

Turning now to (2*), we often lack diachronic self-presentational autonomy because we lack diachronic privacy – it is other people's unbidden access to evidence of our prior self-presentations that threatens harm. However, there are two quite different ways in which evidence of our prior self-presentational choices may be unearthed: evidence of our prior self-presentational choices may be public (as with Instagram or Twitter posts) or private (as with photo albums in our lofts). Given that I will consider our self-presentational interests in online public spaces, this difference is pertinent.

When someone observes or distributes evidence of our prior self-presentational choices in a private domain, they clearly violate our right to privacy. But when someone observes a prior self-presentational choice on Instagram, I am in some sense complicit in the violation of my diachronic privacy. By making a self-presentational choice online, we grant other people access, and it's unclear whether or not our *privacy* is violated when other people observe these self-presentational choices later in time. It's not that we shouldn't have diachronic privacy, but that we don't: we have waived certain access-related duties over the information by posting it on a public domain. This demonstrates how, similarly to the claims I made in Chapter Two, our DSPIs can come apart from diachronic privacy – our DSPIs may be infringed even if our diachronic privacy is not.

This is not to say that no diachronic privacy infringements occur on the basis of information in public realms. As I indicated in the introduction, the fact that we have granted other people access to us does not mean that there are no privacy-related duties in play. It may be that someone violates both our DSPIs and our diachronic privacy by distributing our online prior self-presentational choices without our permission. Distributing our prior self-presentations increases their accessibility – it means that people who may otherwise not have had access to them, now do.

These duties engendered by our right to privacy seem to emerge as time passes. Consider the Blog case we discussed earlier. While your family may not violate any privacy related duties by reposting a blog post you published yesterday, I have the intuition that they might with another you

published ten-years-ago. As the grounds of our self-esteem shift over time, so too does the function of diachronic privacy in allowing us to retain control over the terms by which these grounds are knowable to others. While you may have waived your right not to have these opinions accessed by others, duties concerning distribution may be reinstated over time. To echo Moore, we should be careful not to assume that presenting X at time t1 means relinquishing our diachronic privacy over X at every time t2 prior to t1 (2018: 342).

Differences aside, protecting our DSPIs from either public or private evidence requires us to retain diachronic privacy over them: either by preventing access to them, or preventing their distribution.

3.4.1 The Right to be Forgotten.

There are existing privacy norms which treat the past as a focal point. One prominent example of which is the Right to be Forgotten (RTBF). The RTBF centres around the idea that we ought to allow 'legally available, truthful information about a person to sediment, without being constantly rehearsed' (Floridi 2015: 163). Discussions of this right emerged on the basis of individuals wanting to determine their lives autonomously 'without being perpetually or periodically stigmatized as a consequence of a specific action performed in the past' (Mantelero 2013: 213). This sentiment – that we ought to alleviate individuals from their past and outdated selves – is shared with our DSPIs.

The EU has begun to incorporate iterations of the RTBF into regulations. One example of which is that EU citizens are now able to submit requests to remove search engine results which appear when an individual searches their name, 'where that information is inadequate, irrelevant, or no longer relevant, or excessive in relation to those purposes and in light of the time that has elapsed' (CJEU 2014: 4). These are called *delinking requests*.[33] The intuitions behind our DSPIs can help to explain why we have interests in removing information about our past. Our lack of privacy impedes our self-presentational interests: it allows other people to understand us on the basis of our prior and no longer relevant selves.

[33] Another more recent example is Article 17 of the GDPR – the 'right to erasure' which allows individuals to ask organisations to delete their personal data where: they no longer consent to processing of it, there are significant errors, or the data is held unnecessarily (Wolford 2018: 4). While the personal data that organisations hold of us can threaten our self-presentational interests (c.f., the data mining objection in Chapter One), the right to erasure does not address publicly accessible information, because of this, I will focus on delinking requests.

Delinking requests can be understood as encouraging diachronic self-presentational autonomy online. However, there are important differences between the RTBF and our DSPIs. While the RTBF is restricted to the online realm, in tackling information appearing in search results, our DSPIs are not. They also concern evidence of our past resurfacing offline, such as photo albums, and memory recall.[34] Secondly, while our DSPIs are concerned with our self-presentational past, the RTBF tackles our non-self-presentational past, such as our criminal history and debt records. This non-self-presentational past may impede our self-presentational interests, but it does not impede our DSPIs as I have conceived of them. Although our DSPIs and right to diachronic privacy capture some of the issues that are at stake when we speak of a RTBF, they are not one and the same.

SUMMARY

Most accounts of self-presentation are implicitly synchronic; they evaluate our interests in light of particular scenarios, interactions, and performances. I have argued that this approach is incomplete. Self-presentational interests aren't just about controlling our present interactions with others, they also concern having a measure of control over the terms by which our prior self-presentations are accessible to others. These are our diachronic self-presentational interests. The Diachronic Account takes as a focal point our self-presentationally autonomous choices – the parts of our presentational self we change over the course of our lives.

Not all DSPI-infringements are all things considered wrong – when the aspect of our past which is being dredged up is a moral transgression, the wrongness of the infringement may be outweighed if the infringement serves other valuable purposes which depend on the exposure, such as holding individuals accountable. However, we have legitimate ethical interests in exercising some control over our prior self-presentations, even some of our prior mistakes. These interests provide us with some grounds for a right to diachronic privacy – in keeping our childhood foibles, jejune opinions, and bygone bodies to ourselves.

---

[34] DSPI-infringements ordinarily involve the resurfacing of physical or digital evidence of our past, but our DSPIs may also be set back by testimonial evidence: where someone recalls what we used to be like, or something we said or did. These kinds of DSPI-infringements are important, and yet the extent to which it is reasonable for us to have control over this evidence of our past is dubious – clearly, we cannot ask other people to forget what they know about us. But there may still be DSPI-related responsibilities in play: it seems reasonable for us to ask other people not to tell anyone about our past, relying on norms of confidentiality and expectations of trust. However, such control is importantly different from the kind of control we can expect over photos or letters. I will return to this distinction in Chapter Four.

IV

Expressing Ourselves Online

At the outset of this thesis, I set out two central aims. The first was to argue that the Synchronic Approach is incomplete and to vindicate the diachronic character of self-presentational autonomy. The second, which I will now begin to address, is to show how Social Networking Services (SNSs) threaten our self-presentational interests. Several authors have argued that SNSs enhance self-presentational control. I take this to be a mistake. By attending to the diachronic dimensions of self-presentation, I show how technologies of the self complicate and potentially undermine our self-presentational interests. Avoiding these threats will require greater diachronic privacy online.

In §4.1 I explore how SNSs prima facie enhance self-presentation. In §4.2 I argue that this seeming advancement of our self-presentational interests' rests on its implicitly Synchronic Approach and explore how technologies of the self facilitate DSPI-infringements.

4.1 SNSs and Synchronic Self-Presentational Control

In the early 2000's there was a substantive reconfiguration of the internet. What has been dubbed 'Web 2.0' transformed the internet from a collection of largely static, institutional, and commercial webpages to a highly interactive and dynamic environment (Vallor 2016: 161).[35] The distinctive feature of Web 2.0 is its focus on user generated content: the videos, photos, stories, reviews, comments, and opinions that make up Web 2.0 are created by *us*. SNSs such as Facebook, Myspace and LinkedIn give even individuals with modest computer skills platforms on which to maintain, perform, and negotiate their public self-presentations. This shift has completely transformed how we initiate and maintain nearly every social interaction and role we have with others, whether seller to buyer, or friend to friend. Technologies of the self are especially pertinent in discussions of self-presentation because they modify how we self-present, and how we encounter other people's self-presentations.

[35] A term coined by Darcy DiNucci (Aced 2013: 6).

Several authors have recently argued that technologies of the self enhance our self-presentational interests. How SNSs give us greater choice and thus control over which aspects of our selves we reveal to others (Marmor 2021: 4).[36] Technologies of the self grant us more time to deliberate our self-presentational choices, enabling us to 'offer one another stylized versions of ourselves' (Borgmann 1992: 92). In contrary to face-to-face interactions, the space between 'being' and 'appearing' is potentially vast; we can alter our appearances with filters, retaking a photo over and over until 'our face exudes precisely the calm mastery of life we wish to project' (Miller and White 2021: 12). The careful construction of our online profiles allows us to conceal the dull and unsightly aspects of our lives, to craft our online personas 'on our own terms' (Cocking and van den Hoven 2018: 50).

In face-to-face encounters, our self-presentational control is often curbed by involuntary self-revelations which conflict with the impressions we intend to give. SNSs grant us communicative spaces in which we can avoid these uncooperative aspects of ourselves obstructing our intended presentations. They erase our tone of voice, facial and bodily expressions, wiping out the insincerity of bitter congratulations and concealing our blushed cheeks. In this way, they give us greater self-presentational control than in ordinary nonvirtual contexts (Cocking 2008: 137).Marmor describes such possibilities of self-construction as 'endless' (2021: 10). We can fabricate much of our online identities, convincing other people of public personas which have very little to do with reality (Ibid., 4). We're able to 'develop new and exciting selves' as Hubert Dreyfus puts it (2004: 75). Our ability to forge our online identities shows that, if anything, we're given 'too much control over the aspects of ourselves we present to others' (Marmor 2021: 4).

Such stances align with the famous meme about online anonymity 'on the internet, nobody knows you're a dog' – the idea being that, unlike our offline lives, we can be whoever we want online.[37]

[36] This position is commonly held in Information Studies see e.g., Rosenberg (2009: 1).

[37] Originally published in *The New Yorker,* Steiner (1993).

"On the Internet, nobody knows you're a dog."

Under Marmor's conception of privacy, SNSs give us greater self-presentational control *and* greater privacy. Given that Marmor takes self-presentation to be the only interest underpinning our right to privacy, the conclusions that he draws on the basis of our self-presentational interests concern privacy as a whole (2015: 4). The trouble with social media is not that it compromises our right to privacy, but that it gives us *too* much privacy (2021: 2). The solution, it seems, would be to break down privacy online, setting limits on our capacity to self-present. It's not that Marmor takes SNSs to pose zero risks to self-presentation – he recognises that these platforms can be used to deliberately undermine users' privacy by exposing private information online. Cases of doxing – for example – are clear privacy violations which undermine self-presentational interests (2021: 7). However, on the whole, Marmor takes social media to be 'generally conducive to privacy – often too much so' (Ibid., 1).

SNSs enable us to command our public personas in ways that offline spaces do not permit. The literature is largely unanimous on this point. But this increased self-presentational control isn't necessarily a good thing. For Marmor, excessive control compromises authenticity (2021: 2). For Albert Borgmann, hyperrealities subvert our organic social realities; we present ourselves for

'amorous or convivial entertainment' (1992: 92) and when we are forced to emerge out of the 'insubstantial and disconnected glamour' we are left 'resentful and deflated' (Ibid., 96). For Dean Cocking, online platforms seriously limit the scope for our developing intimate friendships by enshrouding involuntary self-revelations (2008: 125). While these authors highlight significant issues with excessive self-presentational control, the self-presentational affordances that social media offer us have a more ambiguous effect on our self-presentational interests than is initially apparent. I suspect that the nature and extent of this threat has been underestimated because of the implicitly synchronic approach.

Before I explore how technologies of the self jeopardize our diachronic interests, it is worth noting that there are reasons to doubt the extent to which SNSs enhance our synchronic self-presentational control. As Véliz notes, while we cannot give off involuntary facial expressions via text-based communication, there are other forms of involuntary disclosures in online spaces; when I open or respond to a message too quickly, I may give off to others that I am bored or lack will power and when I am active late at night I may reveal my restlessness and inability to sleep (2022: 38). There are ways in which we continue to be susceptible to involuntary disclosures online, of saying more than we'd ultimately want to say. In addition, though Marmor is right to note that social media gives us the opportunity to fabricate our online identities, SNSs more often than not anchor online identities to real and embodied selves. While we often distort aspects of our self-presentations, encouraging others to view us through rose tinted glasses, our public personas by and large have more than a grain of truth in them.

The main obstruction to synchronic self-presentational control is that we cannot predict who we are presenting ourselves to. When we decide to present aspects of our private lives for public consumption, we don't ordinarily present ourselves to everyone and anyone. Instead, a self-presentation is public to a particular individual, context, location, or group of individuals. When we write a letter, or have a conversation face-to-face, we ordinarily know who we're disclosing ourselves to. I can see who I am at the pub with, showing a photo to, or addressing a letter to, and a quick glance around can be an effective method of ensuring that what I say goes no further than I intend. This is not true of online spaces. The internet is a massively scaled environment in which people who do not know each other interact (Fyre 2022: 144). We face what Rima Basu calls *context collapse* – the blurring or merging of multiple contexts into one (2022: 484).

Context collapse distorts our ability to maintain differing levels of intimacy with different audiences, and although this messy collision of family, friends and co-workers can be somewhat managed using tools offered by the site, where we can decide to direct posts to specific sub-

networks that we define (Vallor 2022: 3.2), aspects of ourselves that we share in one context can easily pervade into other contexts. While technologies of the self give us more self-presentational control in some respects, they give us much less control over who we reveal ourselves to than traditional methods of communication.

We should thus be careful not to overstate the extent to which SNSs heighten our self-presentational interests, even under the Synchronic Approach.

4.2 Digital Archives

Expressing ourselves online complicates, and potentially undermines our diachronic self-presentational interests. I will examine threats to our DSPIs according to four dimensions: accessibility (§4.2.1), distribution (§4.2.2), deniability (§4.2.3), and online shaming (§4.2.4). Once we bear in mind the diachronic character self-presentation, it turns out that we have less self-presentational autonomy and less privacy.

4.2.1 Accessibility

Before we started expressing ourselves online, our past was preserved in private: in letters and photo albums, and while these archives could be used to undermine our self-presentational interests, it was relatively difficult for others to do so. The opportunities we have in accessing the letter you sent your ex 10 years ago, or your friend-of-a-friend's photo albums, are comparatively slim to the opportunities we have in finding Facebook or Twitter posts. Given the curtailments on access, these records of our past pose a relatively small threat to our DSPIs.

Online archives have transformed the ease at which other people are able to access our prior self-presentations. Formerly covered in dust, the things we once said, wore, and did are now neatly organised and accessible via the smartphone in our pockets; other people can instantly, autonomously, and effortlessly access us.

We have never remembered as much as we remember today, both as individuals and as societies (Véliz 2020: 145). For most of history, keeping comprehensive records of our past required huge amounts of effort: paper was expensive, and we needed space to store it, writing demanded time and dedication, and even with considerable effort only a tiny fraction of our individual experiences could be preserved. Forms of external record keeping were used with the utmost care, preserved

for the most precious and impressive moments (Mayer-Schönberger 2009: 37). Today's teenagers have no such constraints. Advances in technology have 'fundamentally altered what information can be remembered, how it is remembered, and at what cost' (Ibid., 52). We have gone from having to select what to remember, to having to select what to forget (Véliz 2020: 145), and this shift has made it harder for us to control the terms by which the grounds of our self-esteem are accessible.

It's not that we have no capacity to prevent other people from surveying our past, but the process of preventing access is much more laborious and less effective than it used to be. Unlike physical photos, which we can burn or stash away in bedside drawers, removing digital archives involves spending hours trudging through our online profiles and deleting individual posts. And even this cannot guarantee our liberation from the gaze of others, for other people have the tools to monitor, record, and duplicate our self-presentational choices without our knowledge. Though we still have some control, we have a lot less control than we used to.

Part of what makes both self-presentational control and privacy valuable is their role in enabling us to decide if and when we present ourselves for public scrutiny. However, the practice of archiving, and the ease at which other people can access our past, means that what we say now may be judged by others long into the future. As Borgman had predicted, no users of SNSs are cut off from the instruction of being seen and judged (1992).

By expressing ourselves online, we knowingly grant other people access to us at the time, but we also grant other people access to those self-presentational choices long after we have made them, by the time such choices ground feelings of embarrassment, regret, and shame. Our self-presentational choices cast a far longer shadow into the future than they would have had we shared them in traditional settings. The fact that our self-presentational choices often end up grounding negative emotions of self-esteem – becoming aspects of ourselves we want to hold in private, means that this default of archiving our self-presentational choices is problematic. By giving other people unmediated access to our past, SNSs allow other people to understand us in ways that ignore the sense in which our character, values and thinking have evolved, stripping away any humanising or otherwise mitigating context.

In *Discipline and Punish,* Michel Foucault describes the panopticon prison originally developed by Jeremy Bentham. In the panopticon there is a central tower from which each prison cell is 'constantly visible' (2008: 5). The inmate must 'never know whether he is being looked at any one moment; but he must be sure that he may always be so' (Ibid., 7). Within this panopticon individuals are monitored, surveyed, and recorded. And while it is ordinarily described in terms of

disciplining individuals in space, SNSs have a similar panoptic architecture, enabling any one internet user to survey and scrutinise another across time – a *temporal* panopticon.

The permanency of our online self-presentations, and consequently increased accessibility, means that we have less privacy and less diachronic self-presentational autonomy. 'For those of us over the age of 30, give or take, consider how often we have thought: "thank goodness social media, with its terrible elephantine memory, didn't exist when I was a teenager" (Simpson 2022: 8).

4.2.2 Distribution

The introduction of SNSs has also changed the nature and extent of DSPI-infringements. All interested internet users have the tools not only to unearth our prior self-presentational choices, but platforms on which to effortlessly distribute them to potentially monumental audiences. DSPI-infringements are thus more severe than they used to be.

Before Web 2.0, DSPI-infringements typically occurred in constrained environments. If someone came across an embarrassing photo of us, it was relatively hard for them to broadcast it to a wide audience. They could still impede our DSPIs by exposing us to others in ways that we did not wish to be perceived, but it was relatively difficult to do so outside of particular social contexts or social relations. This isn't to say that threats were negligible: someone could have photocopied the valentine's card you made them and distributed it fairly widely within your social circles without your permission, but it would at least be relatively arduous for them to do so. This lack of pervasiveness is advantageous to our self-presentational interests, as it gives us the option to regain face elsewhere – to start over in terms of the knowledge that other people have of us.

This process of withdrawing oneself is much harder with online infringements. Online exposures penetrate deep into our social circles, potentially impacting every social relation we have. This is what Marie Franks calls *virtual captivity,* a term used to describe how the effects of online exposures can 'manifest anywhere, to anyone, at any time' (2012: 682). Our prior self-presentations may reach not only our friends and friends-of-friends, but complete strangers – placing us in an almost 'unbearably transparent state of exposure to other people's gaze' (Simpson 2020: 8). SNSs give all of us the tools to shape the perception of other people throughout and beyond their social circles, because of this, they advance more severe threats to our interests than was previously possible.

Distributing our prior self-presentations infringes our diachronic privacy in two ways. Firstly, by making information easier to access. Other people can find our prior self-presentational choices with little snooping. Secondly, by making our past available to people who may not have had access to our profiles in the first place, whether because they do not know we exist, or because their access is blocked under privacy settings offered by the site.

When we make self-presentational choices online, we do so in a particular sort of context – namely, at a particular time. But the temporal panopticon we are in means that these choices can easily be accessed and distributed in different contexts – namely, at different times. We experience a temporal version of Basu's situational context collapse. As I explained in Chapter Two, privacy violations (which involve our self-presentational interests being thwarted) often occur when an environment is manipulated in ways that significantly diminish an agent's ability to control which aspects of themselves they reveal (Marmor 2015: 25). One kind of case which is relevant here, is when we know that by doing X we reveal Y to A. If we have good reason to assume that this is the case, then we can choose whether or not to X. Our right to privacy would be violated if someone were to manipulate the environment such that by doing X we reveal Y not only to A, but also to B et al., (Ibid., 14).

Our decisions to express ourselves online resemble this kind of case, but temporally. The distribution of our prior self-presentational choices means that by doing X we not only reveal Y to A at t1, but also to A (and B et al. if it is distributed) at t2. Although we have some reason to assume that the presentational choices we make online could be exposed at some later date, the online environment is extraordinarily unpredictable. We cannot predict if, when, or to whom our prior self-presentational choices will be unearthed for social scrutiny. And we need a measure of predictability to have a reasonable measure of self-presentational control. Utilising technologies of the self thus wrests us away from control over the temporal context in which we construe ourselves to others.

It's not that everyone will be worse off as a result of using these technologies. Some savvy and motivated users may figure out how enjoy the benefits of these technologies whilst shielding themselves against potential harms. Others may have the good fortune of not having been the kinds of teenagers to overshare, so that when they're older, they have a lot less to be embarrassed about. Moreover, people's differing degrees of personal change mean that the harms that can be inflicted as a result of DSPI-infringements are liable to vary. However, the relative number of losers, and the relative weight of their losses will be of greater ethical significance than whatever benefits the winners can get.

4.2.3 Deniability

Digital copies of our prior self-presentations also grant us less breathing room to deny that we actually did, said, or thought something.

Consider how DSPI-infringements utilising digital records of our prior self-presentational choices are different from oral transmissions of prior self-presentations – of 'so-and-so said this before'. When I entrust sensitive information to someone who later shares it with a third party without my permission, the danger of transmission is somehow mitigated by their capacity to share it. There's still a threat of harm, but the harm seems different in kind to digital sharing of information. When information transmission is oral, we seem to have more opportunity to pretend it was a big misunderstanding. Even if the gossiping penetrates deep into our social circles, there's still this potential for us to explain it away through a sort of fallacious Chinese Whispers.

Conversely, when we think about information transmission in the digital age, we are faced with information that is perceived as undeniably precise and objective, not prone to the shortcomings of human remembering, or distortion of gossiping. What we once said, advocated, or rejected is presented as a perfect replica of the original – no crumbled edges, or faded words to demonstrate the pass of time. It is 'information from which time has been eliminated' (Mayer-Schönberger 2009: 124). So, while our lives as embodied humans continue to change, digital representations shackle us to our prior self-presentational choices.

Does promoting plausible deniability wrongfully defend people's right to be devious, or deceptive? I suspect not. As Seana Shiffrin has argued, while we should have a 'default moral presumption of truthfulness in communication' there are circumstances in which being insincere is reasonable and justified (2014: 11). The normative presumption of truthfulness may be suspended in contexts which serve other valuable purposes; valuable purposes which depend on the suspension of the truthfulness presumption (Ibid.). Deceptiveness, so Shriffin argues, may be permissible within these justified suspended contexts (Ibid., 7). Being able to deceive other people of what you once said, advocated, or did, may be permissible (even desirable) within a justified suspended context, because it serves other valuable purposes. Deceptiveness may be necessary to avoid invasions of privacy – in shielding ourselves from invasive and inappropriate inquiries (Ibid., 149). Moreover, the function of communication is not always to advance propositions to be taken as true, sometimes, our speech functions to uphold conventions of etiquette (Ibid., 152), or to strengthen relationships that would be made fraught with too much truth (Ibid., 151). When we describe our past deceptively, our communication may be serving one of these other valuable functions.

In short, deceptive utterances can occur in justified suspended contexts, when this happens, we may be deceptive and not wrongfully so because the presumption of sincerity does not operate.

4.2.4 Online Shaming

Allowing people that you trust to access you can be a good thing – they can note your tastes and likes to find you the perfect gift or use your date of birth to organise a surprise party. But not everyone that accesses you will have your interests in mind. A multitude of authors have noted the harrowing culture of exposure which has come to dominate our online worlds, a culture which has been enabled by the ease at which we are able to access and distribute information online. Pertinent and oft cited examples of this are cases of shaming and doxing.

While doxing ordinarily pertains to information which is privately held, some cases of doxing involve the distribution of publicly available information. David Douglas describes 'delegitimization cases' of doxing, in which information about an individual is revealed with the intent of undermining a subject's credibility, reputation, and/or character, whether to embarrass, humiliate, shame, or simply expose the other (2016: 204). Some instances of online shaming and doxing can thus be understood as deliberate attempts to undermine self-presentational interests.

When a self-presenter is exposed but not diachronically blameworthy, the impermissibility of doxing may be explained (at least in part) in terms of DSPI-infringements. When an agent is exposed and diachronically blameworthy, the DSPI-infringement may be legitimate, as I noted in Chapter Three, however, it may be questioned whether SNSs are the best way to execute these exposures. As Paul Billingham and Tom Parr argue, users often fail to harness the appropriate measure of blameworthiness in the online world, frequently causing harms which are disproportionate to the moral transgression in question (2020).[38] Even granting that some DSPI-infringements are permissible, online exposures may involve both our irreproachable self-presentations (e.g., prior poor health and fitness, pre-rhinoplasty photo), and our synchronically but not diachronically blameworthy self-presentations. As such, SNSs pose considerable threats to our self-presentational interests.

[38] See Thomason (2021) for a discussion of permissible shaming practices, and Frye (2021) for a compelling argument as to why shaming practices cannot achieve the end of promoting cooperative behaviour when they are conducted online.

The internet is often applauded when compared to traditional media, for allowing bidirectional communication – empowering individuals to express their views and selves to others. However, in doing so, the internet has simultaneously unleashed a power that we have failed to constrain, or even exercise judiciously. In providing the means for us to publicise ourselves and each other in ways previously impossible, SNSs have made all too easy harms to both ourselves and others. It's not that SNSs have fundamentally shifted the nature of our activities – humans have always had deep interests in enabling the commemoration of the otherwise easily forgettable past, and societies and communities have encouraged the public shaming of one another since the pillories.[39] But technologies of the self have made shaming and remembering easy. They have democratised DSPI-infringements by making inflicting harms on one another collective and excessive – achievable on a whim, barely lifting a finger. Every image, word and click is collected, made available for others to scrutinise, and potentially torn apart in a cathartic act of public shaming. By flatlining the threshold for harm, we end up – on the whole – worse off.

SUMMARY

SNSs prima facie enhance our self-presentational interests; they give us tools to carefully curate and deliberate our public personas in ways that offline spaces do not permit. However, attending to the diachronic dimensions of self-presentation reveals how SNSs have a thornier effect on our self-presentational interests than is first apparent. SNSs give other people the means to analyse our prior self-presentations autonomously and effortlessly without our knowledge. They also give other people the means to distribute them more pervasively than previously possible. They put us more at risk of our past re-emerging unbidden, and in this sense, we have less control over the aspects of ourselves we reveal to others.

Giving users greater control over their identity sharing practices intuitively increases self-presentational control, however, contrary to appearances, giving users greater control actually leads to decreased self-presentational control and privacy. Preventing DSPI-impediments will require us to restore diachronic privacy online.

[39] See Frye for an overview of historical artefacts used in public shaming (2021: 133).

V

Diachronic Privacy for Diachronic Selves

> 'Perhaps after enough time has passed, the intrusion will be muted by distance, but with people whose lives have overlapped with ours, there is something excruciating about all this exposure, something wrong with our now having access to Bertrand Russell's desperate love letters, Wittgenstein's agonized expressions of self-hatred, Einstein's material difficulties' (Nagel 1998: 22)

There's an important sense in which we are complicit in the violation of our own privacy. When I decide to wear a particular outfit, Tweet my beliefs or share a photo on Facebook, I decide to present aspects of myself for social scrutiny. If these self-presentational choices end up undermining our DSPIs in the future, a plausible response might be: 'well, you shouldn't have done that in the first place'. Even if we have since changed our minds, at the time of making the (now prior) self-presentational choice, we purposefully revealed ourselves to the world. This response seems particularly cogent in the context of SNS, and a multitude of authors have argued that we ought to be careful not to unveil too much of ourselves online.

The aim of this chapter is to determine what sorts of norms we need to adopt in order to insulate us from the harms that are liable to arise on SNS. I examine three relevant players: (1) Me: should I adjust my conduct to mitigate the risk of being subject to future DSPI-infringements? (2) Others: is the onus on other people to refrain from unearthing and exposing our past self-presentational choices? And (3) Technologies of the self: should we modify SNSs to alleviate the threat of DSPI-infringements? I argue that, while we ought to adjust our online conduct to better protect our self-presentational interests, we still have good reason to modify technologies of the self.

One important question to consider when examining our self-presentational control online is whether or not there are important differences between online and offline forms of self-presentation: we have an ethically legitimate interest in controlling what we reveal to and conceal from others, but are online self-presentational choices forms of self-presentational autonomy we want to protect? One might argue that having control over our *online* presentations is not as important. In one sense I am inclined to agree – there is much to be said about the deficiencies of

self-presentations online.[40] However, this does not preclude their deserving protection. The affiliation of our online and offline worlds means that online self-presentational infringements often compromise our self-presentational interests offline, and we can simultaneously discourage exercising self-presentational autonomy online while advocating the preservation of our future DSPIs when we do. For the purposes of this chapter, I will be putting to one side the question of whether or not we ought to refrain from posting content or communicating with others via SNSs. Instead, I will consider what we ought to do about DSPI-infringements given that many of us are active participants of Web 2.0 technologies.[41]

In §5.1 I explore the phenomenon of undermining our future selves. The following three chapters examine what we ought to do in light of our role: in §5.2 I explore our tendency to over-share online, in §5.3 I examine the ease at which we undermine our future selves, and in §5.4 I consider the hyper-cautious self-presenter. In §5.5 I turn to the role of other people in facilitating DSPI-infringements. Finally, in §5.6 I suggest a way of modifying technologies of the self which would curtail DSPI-infringements by empowering users to retain diachronic privacy online.

5.1 Undermining Ourselves

When our DSPIs are infringed, our current self-presentational autonomy is set back because of a self-presentational choice we made in the past. While other people facilitate the infringement by unearthing our past and exposing it, there's an important sense in which we participate in the violation of our own interests; had we not made the self-presentational choice at t1, then our DSPIs would not be setback at t2. Our decisions to air our thoughts on Facebook or disclose photos on Instagram are in this sense *diachronically self-regarding*. They are self-regarding since they involve us making a decision and our relationship with ourselves. And they are diachronic since they involve our current and future selves (Viganò 2022: 1). Sometimes, our self-presentational choices end up making things worse for our future selves, we retrospectively assess our conduct as deficient and come to regret the choices we made.

This complicated form of diachronic self-sabotage has been noted in discussions of diachronic autonomy. Joel Feinberg's classic paper highlights how decisions made on behalf of children can undermine their future autonomy. He argues that children have anticipatory autonomy rights

[40] See Vallor (2016); Vallor (2022); Cocking (2008).

[41] See Simpson (2022) for a discussion of the goods and ills of quitting social media.

which cannot be exercised 'until later when they are more fully formed and capable' (1980: 76). Education systems which render a child fit for one way of living are objectionable on account of this right; we ought not foreclose irrevocably a child's future options (Ibid., 82).

Much of the literature on adjusting present conduct for the sake of future options focuses on children, but pertinent questions arise in the context of self-governing agents too. From the perspective of the individual in the present, should it matter to them whether or not they frustrate their future self-presentational autonomy? Even once we have decided that we have responsibilities to our future selves, there are further questions surrounding how we ought to accommodate our future interests when we make diachronically self-regarding decisions. Should we act so as to reduce the risk of regretting our decisions in the future? Should our current selves have greater weight in decision making? If some of our diachronic selves have more sway than others, then why?

Several authors of have begun to address these complex questions. Eleonora Viganò considers the nature of the moral relationship holding between our current and future selves. She has developed a prudential moral theory to guide our present self when making diachronically self-regarding decisions (2022). In response to the problems that Paul's work on transformative experiences raises for theorising about rational choice, Richard Pettigrew considers how we ought to weigh our current interests relative to our future interests within standard decision theory. He argues that we should use an aggregate utility function and incorporate the utility function of our present self along with the utility functions of our future selves (2022: 205). These substantial and complex questions are well beyond the scope of this thesis. However, drawing on the phenomenon of undermining our future selves will be useful in analysing our own role in facilitating DSPI-impediments, for it introduces the option of adjusting our own conduct to avoid DSPI-infringements in the future.

5.2 Speaking Out Loud on Social Media

If infringements of our DSPIs are harmful, then shouldn't we refrain from making choices which make us susceptible to DSPI-infringements; to moderate the kinds of self-presentations we make online? Various authors have rightly noted our tendency to overshare online. The psychopathology of our social milieu is, as Franco Berardi observes, antithetical to that described by Sigmund Freud. Rather than repressing memories, today's psychopathology is characterised by

'excess of visibility' or 'hyper-expressivity' (2007: 179) – we voluntarily strip our souls bare for the ear and eye of anyone who passes by.[42]

This inclination to express our private lives publicly is both unpalatable and puzzling. While the degree of control that we have over our online profiles is blurred and often deceptive, the expectations we have concerning privacy are seemingly contradictory and paradoxical. We 'treat social media as if it is a private conversation – a suitable drafting space for our identities – when it is functionally anything but' (Simpson 2022: 1). Perhaps internet users should accept liability for this diachronic self-sabotage. By shifting away from hyper-expressivity towards reticence, we could avoid a large proportion of DSPI-infringements. As Simpson eloquently puts the point: 'we do not need to air our thoughts on social media, and we have good reasons for finding other ways to scratch the psychological itches that nudge us in this direction. Thinking aloud with others is important, but social media isn't a private drafting space for doing so, however much it may feel that way' (2020: 8).

I am sympathetic to this line of thought. We ought to move away from the current culture of hyper-expressivity online. However, saying this does not imply that modifications to SNSs are gratuitous, or that other people's role in DSPI-infringements are extraneous to the harms in question. While norms of reticence would lessen our susceptibility to DSPI-infringements, they do not alleviate the need for auxiliary protection. This is because (1) there are features of diachronic self-regarding decisions which make the task of looking after our future selves virtually futile (§5.3), and (2) we sometimes have good reason to act in ways that put us at risk of future DSPI-impediments (§5.4).

5.3 Considering Our Future Selves

There are three features of diachronic self-regarding decisions, which complicate the project of adjusting our present conduct for the sake of our future selves. My aim here is not to defend an account under which we morally ought to take our future self-presentational autonomy into account. Instead, I aim to show why refraining from making choices which make us susceptible to DSPI-infringements is not as simple as 'well, you shouldn't have done that'.

[42] Bakardjieva and Gaden (2012); Simpson (2020); Véliz (2022) and Nussbaum (2007) highlight this propensity.

(1) The Epistemic Deficit

The Epistemic Deficit concerns our lack of knowledge about our future selves. We do not know what values, beliefs, and preferences we will hold in the future. Because of this, we do not know whether our future selves will agree with our current self-presentational choices.

While we can predict some things about our future selves. There are occasions throughout the course of our lives where we will face significant and unforeseeable changes. When one becomes a parent, moves to a new city, falls in love, or is diagnosed with a terminal illness they may undergo *transformative experiences*. Experiences which 'change our point of view, and by extension, our personal preferences, and perhaps even the kind of person we are or at least take ourselves to be' (Paul 2014: 16). We cannot know exactly how we will feel in the future about the self-presentational choices we make now, because of this, we cannot invoke our future preferences to guide the decisions we make now.

(2) The Involvement Imbalance

The Involvement Imbalance has to do with the fact that our future selves do not take part in the self-presentational choice, despite it primarily affecting them. My decision to get a face tattoo, for example, is a decision that I make at time t1, and it is a decision which will affect every future diachronic self that comes after me. Those future selves may not have made the same decision that I did; and yet it becomes an aspect of their presentational self.

Part of why it is difficult to accommodate our future interests when we make diachronically self-regarding decisions, is because our current selves have to decide on behalf of our future selves. I have greater decisional power than my future self purely on the basis of my temporal positioning, and this asymmetric power imbalance means that almost anything I do puts my future interests at risk.

(3) A Preference for the Present

P. F. Strawson marks a distinction between 'episodic people' whose sense of self is concentrated in the present and 'diachronic people' who understand themselves as temporally extended, integrating their present sense of self within the broader story of their lives (2004: 431). Acting so as to alleviate future DSPI-infringements would require us to see ourselves as diachronic. However, as Viganò has recently emphasised, there are a multitude of neurobehavioral studies

suggesting that we are predisposed to favour our present selves in decision making – we see our future selves as distinct from 'us' (2022: 23). Of particular interest for our purposes is how 'the brain areas activated in mentalizing about one's future selves are similar to those activated for mentalizing about other people' and different from those activated when we think about our present selves (Ibid., 45). So, rather than understanding ourselves as temporally extended, we view ourselves in line with Strawson's episodic agent – an agent who takes care of themselves in a period which is shorter than the life of which they are a part. Because of this, we often lack the foresight required to adjust our conduct for the sake of our future selves.

These features of diachronic decision-making reveal the ease at which we are able to undermine ourselves. While norms of reticence would help, moderating our online activity would not completely alleviate future DSPI-impediments. It is impossible to know for certain how the grounds of our self-esteem will shift, what we will end up regretting, being embarrassed about, or ashamed of. Refraining from making choices which make us susceptible to DSPI-infringements is thus not as simple as 'well, you shouldn't have done that'.

If I am right to think we have interests in diachronic self-presentational autonomy, and that DSPI-infringements threaten harm, then we ought to have some kinds of protections against avoiding these kinds of setbacks to our interests – especially conventions and environments which make the already difficult task of avoiding DSPI-infringements more difficult. As you might have guessed, SNSs do exactly this: they facilitate temporally asymmetric power imbalances by giving us tools which radically transfer power to our present selves to haunt ourselves in the future. Under previous communicative frameworks, it was much harder for us to inadvertently undermine our own interests, that SNSs make setbacks to our future interests relatively easy suggests that these technologies should be much more carefully regulated.

5.4 Hyper-Conscious Performers

The badness of DSPI-infringements only entails that we ought to avoid actions which put us at risk of future DSPI impediments if we accept something like the following: we always ought to act in ways that minimise impediments to our future interests. Call this Anti-Recklessness. This section argues that we ought to reject Anti-Recklessness. We sometimes have good reason to act in ways that put us at risk of future DSPI impediments, because of this, we're going to need norms to insulate us from the dangers that are liable to arise.

Given the epistemic deficit we are in with regard to our future selves, adjusting our conduct for the sake of our future interests is not as simple as avoiding *certain* self-presentational choices; any exercise of synchronic self-presentational autonomy runs the risk of compromising our future interests. Abiding by Anti-Recklessness would thus amount to living self-presentationally cautious lives, transforming both our online and offline worlds. Notice that abiding by Anti-Recklessness does not mean considering ourselves as either an episodic or diachronic person under Strawson's distinction, but some sort of 'future episodic' agent. Instead of guiding our actions on the basis of our current selves, or considering ourselves as extensional from the present, our sense of self is concentrated in the future – our present actions are guided by, and for the sake of our future selves. Consider the following case:

> PLAIN JANE. Jane values her diachronic self-presentational autonomy. She recognises that her future interests might be undermined by the self-presentational choices she makes now. Because of this, Jane endeavours to eliminate as many future risks to her DSPIs as possible. Jane starts by eradicating every record of her past, she burns her childhood photos and deletes her social media accounts. In her day-to-day, Jane strives to avoid the observation or recording of her self-presentational choices by evading the gaze of others. When Jane cannot avoid being in some state of exposure, she is conscious to appear as plain as possible. She avoids wearing expressive, garish, or atypical clothes, and confines her utterances to small talk. She avoids revealing herself to others, and melts into the shadows of obscurity.

Plain Jane reveals why we ought to reject Anti-Recklessness: abiding by it encourages an unhealthy self-presentational mindset – it thwarts various values which in some sense constitute our interests in self-presentation.

By encouraging us to remain in the safety of drafting spaces, Anti-Recklessness deprives our social worlds. Self-presentational control is valuable because it enables us to maintain differing levels of intimacy with others – to distinguish best friends and lovers from in-laws and colleagues. Part of what defines these intimate relationships is sharing what makes us vulnerable – giving other people the power to hurt us and trusting them not to take advantage of their epistemic privilege. These kinds of trusting relationships seem pivotal to our social lives, but they would be very difficult to maintain in a world which encourages evading others.

A future episodic person lives in fear that what they say now will be judged by individuals observing them now and far into the future. Anti-Recklessness thus encourages mindsets of mistrust and neuroticism – it makes it better to melt into the shadows of obscurity than put oneself on display. This is clearly an undesirable result. In order to be an object of particularised respect and appreciation, other people need to know us for who we are. But they cannot do this if we refrain from exposing our inner worlds. By fostering artificially stunted communicative exchanges, abiding by Anti-Recklessness also snubs out a whole hoard of other benefits that come with free speech.

Although being hyper-future-risk-averse would largely minimise our risk of future DSPI-impediments, it would simultaneously threaten a variety of other ethically significant values: values in having genuine and trusting relationships, of having stable inner worlds, and of free speech. While there is some sense in fussing over our superficial, everyday self-presentational choices, it doesn't make sense to obsess over them. We shouldn't be neurotic about disclosing our appearances, beliefs, and selves. We should wear the outfits we want to wear, say the things we believe, and experiment with different hairstyles without agonising over our self-presentational choices and their consequences. In short, it seems natural and permissible to engage in behaviours that put us at risk of future DSPI impediments. Because of this, it cannot be true that we ought to live by Anti-Recklessness.

While it seems intuitive that the badness of DSPI-infringements gives us some reason to refrain from acting in ways that make us susceptible to DSPI-infringements, the badness of DSPI-infringements only entails that we ought to avoid self-presentational choices that put us at risk of future DSPI impediments – to live like Plain Jane – if we accept Anti-Recklessness. I have argued that we should not. And if we ought not live hyper-cautious, hyper-future-risk averse lives, then we're going to need to establish norms to help insulate us from the dangers liable to arise.

Turning now to changes in Jane's online worlds. Avoiding DSPI-infringements online would amount to abstaining from using these technologies. Given that everything we do online puts us at risk of future impediments to our interests, we ought to disclose nothing at all. As I have indicated, I'm not sure this would be a bad thing. While SNSs have created opportunities for us to build relationships with others and fruitful avenues for exploring our ideas, there is a growing list of reasons why we ought to refrain from using SNSs: addiction, lack of data protection, and disinformation, to name just a few.[43] However, even if leaving social media would be a good thing, the nature of the situation we are in is not. By recognising our DSPIs, and the sense in which

[43] Dwivedi, et al (2016) provide a comprehensive overview of the goods and ills of social media.

technologies of the self threaten harm, we are placed in a somewhat tragic conflict: either we act so as to reduce future DSPI-infringements but remove ourselves from communicative channels which have come (for better or for worse) to form a centrepiece of our social worlds, or we utilise these tools but simultaneously put our future interests at stake.

While we ought to be more conscious of the ease at which we divulge our private lives online, the badness of DSPI-infringements only entails that we ought to abstain from SNSs if we adopt some sort of principle of incessant Anti-Recklessness. As I have argued, this is a way of living we ought to reject. In light of this, we're going to need safeguards to alleviate the risks these technologies are liable to pose.

5.5 Undermining Others

Turning now to the role of other people in facilitating DSPI-infringements, all of the cases I have considered thus far involved the active participation of other people trudging through our archives. It is other people who find unbecoming photos of us, dig up our naïve blog posts, and recall things we once said. And it is other people who distribute these prior self-presentational choices contrary to our wishes. Our DSPIs are at stake because other people are able to carry out these infringements without penalty. While, as I argued in Chapter Three, some DSPI-infringements may be permissible, I will focus here on cases where the exposure is wrong (either because the prior self-presentational choice is beyond reproach, or because it is synchronically but not diachronically blameworthy).

One interesting question that arises in the wake of other people's role, is why people feel the need to expose one another: vindictiveness? A good laugh? Or the main currency which seems to animate the moral economy of SNSs: popularity. This is one of the core themes in Cocking and van den Hoven's book, *Evil Online*. They describe the moral fog of the internet which has led us to be blind to the morally salient features of our actions, features that we are otherwise quite capable of recognising (2018: 84). Online worlds distance us from the people we are harming, they block our view of the aftermath of our activities and distort our ways of engaging with one another (Ibid., 87). As Roger Crisp notes, 'the moral authorities that provide us with guidance are no longer our parents, our teachers, ordinary role models, but our internet friends' (2018: 1). This kind of authority legitimises 'it's just a laugh' as an appropriate defence.

Cocking and van den Hoven seem right to acknowledge that many of the conventions, laws, and settings of our traditional worlds have been removed, minimized, or altered within online spaces

(2018: 72). Alleviating impediments to our self-presentational interests, among others, will require us to establish the requisite norms and practices, and emerge out of the moral fog of our online worlds. Part of this will involve cultivating moral understanding in online spaces – an undertaking that Shannon Vallor has recently explored in developing a framework for practicing *technomoral virtues* (2016). Some of these norms ought to encourage consideration of other people's self-presentational interests, others should be concerned with protecting our own. I suspect that norms of reticence and empathy are a good place to start.

There is something subtle about the ethical predicament we have found ourselves in. I have noted a number of ways in which all of us have a responsibility, to ourselves and to others, to try to act in ways that don't needlessly set back or threaten our DSPIs. These responsibilities have not been brought into existence by SNSs. But SNSs have made it easier for us to act in ways that set back or threaten our DSPIs. And they have created social affordances that, though they do not single-handedly thwart our self-presentational autonomy, undermine that autonomy by introducing social customs and individual incentives to act in ways that are negligent, short-sighted, hostile, or in various other ways insufficiently attentive to our DSPI-related responsibilities. Therefore, even granting that these are problems that the users of SNSs collude in, we have good reason to try to change technologies of the self so that they don't make it so easy for us to set back both our own and other people's DSPIs. In the final section of this chapter, I will explore some changes that may be helpful to this end.

5.6 Reintroducing Forgetting Online

In order to mitigate the risk of future DSPI-infringements, we need to empower users to protect their diachronic privacy online. Here's a rough principle I think we should abide by: we should (a) give users more control over how the self-presentational choices they post online are preserved into the future, and (b) give other people less access to our past (and consequently less capacity to distribute it). This section will explore the viability of 'expiry dates' in achieving these two aims.

There are a range of reasonable views concerning how best to achieve normative change, whether by government regulation, industry self-regulation or some sort of hybrid view. My discussion will leave the question of implementation to one side. My aim here is not to provide a definitive policy or roadmap to the implementation of policy. Rather, I seek to suggest the kinds of modifications

that would enable us to eradicate many threats to our DSPIs. Expiry dates, so I will argue, enable users to exercise their synchronic self-presentational autonomy without concurrently jeopardising their diachronic self-presentational interests in the future.

Practices of indefinite storage have become standard for SNSs. But the default shouldn't be to hold onto data indefinitely – otherwise it is all too easy for us to undermine both ourselves and others. Dismantling the temporal panopticon of Web 2.0 will require us to move away from the default of remembering – to restore diachronic privacy online.

One way of affording social media users' greater diachronic privacy online would be to require user contributions to set expiry dates – dates at which digital storage systems would automatically delete the self-presentational choices we make online.[44] This requirement could be built into the platform design in a simple, user-friendly way, and invite readers to reflect (at least momentarily) on how long they wish their self-presentation to be accessible. For example, before publishing a blog post, sharing a photo, or responding to a Tweet, we could be asked: 'How long would you like this post be accessible for?' The idea is that we should set expiry dates with some thought but minimal clicks. One option would be to set relative, rather than specific expiry dates (e.g., a week, month, or year from today). SNSs could also give users the option to set a window in which they would be able to revisit their expiry dates. For example, the prompt 'would you like to revisit this expiry date?' with a yes or no response could be given. Expiry dates have the benefit of being relatively uncomplex to implement, with limited technical modifications.

One potential concern is that users might avoid setting proximate expiry dates on the basis of wanting to be able to access their archives. We could address this by giving users an option to retain their posts in an archive only accessible to them. Instead of the self-presentation being discarded, it would simply be made inaccessible to everyone other than the author. Notice the distinction here between information availability, and information accessibility. In the analogue world, the availability of information (i.e., that the information is there) is ordinarily coupled with the accessibility of information (i.e., that the information that is there is known to be there). If I have a physical photo that I don't want other people to see, the best way to block other people's access to it is to destroy it – it is only accessible if it is available. However, the availability of information online is completely detached from its accessibility – we find information on the basis

[44] Mayer-Schönberger (2009: 171) and Véliz (2020: 147) both discuss the viability of expiry dates. Signal, an app which is trying to provide a more private and data-secure alternative to text-based communication services such as WhatsApp, also provides expiry dates as an option.

of its access, rather than its availability (Floridi 2015: 166). If other people cannot access it, then it cannot threaten our DSPIs.

Expiry dates would empower users to avoid DSPI-infringements in two ways. First, they would make the currently arduous process of deleting old posts simple; they would build into the infrastructure of sharing a self-presentational choice, a time cap. Secondly, they afford users greater diachronic privacy by limiting other people's access to and consequently distribution of our prior self-presentational choices. They give us greater control over the terms by which the grounds of our self-esteem are knowable to others.

While expiry dates would significantly reduce risks to our future self-presentational interests, they would not completely eradicate the risk of DSPI impediments. This is because other people would still be able to record anything we post online: a simple screenshot allows other people to retain our self-presentational choices until after the expiration date has passed. One option here would be to make the expiry date 'stick' to photos or posts, such that when it is shared, copied, or screenshotted, the expiry date would be copied along with it, much like other meta-information such as ratings, date of creation or location. Once the expiry date is reached, the deletion would form a sort of domino effect whereby photos and posts would not only be removed from the authors profile but from anyone else's devices who had replicated it.

Although this would increase self-presentational control, it's unclear whether this domino deletion would be desirable. This is because there are cases where we may want to overrule an expiry date and retain duplicates despite an individual's preferences concerning information retention. Cases in which an individual has committed a serious moral wrong may fall into this category. But the question also arises for artists who want their work to be transitory, or authors who want to modify their works: should they be able to remove older versions of their artworks? While an individual ought to be able to remove their prior self-presentations from their own profiles, a domino effect arguably grants individual's excessive self-presentational control.[45]

If we care about self-presentational autonomy, then it's hard to see what grounds we could have for *not* wanting expiry dates to be part of the platform design. Users would pay an incredibly small cost (namely, a brief reflection, decision, and selection) in exchange for far greater control over

[45] George Lucas, unsatisfied with the original Star Wars trilogy re-edited the text in ways that significantly altered the characters and narrative. Should we consider this as an expression of his self-presentational autonomy as the creator of that fiction? Or would this be more akin to cultural vandalism?

the grounds of their self-esteem.[46] Although expiry dates would not completely protect users from DSPI-infringements, they grant users much greater protection than they currently have.

As well as making it easier for people to keep their past self-presentations private, expiry date policies and tools help sensitise people to their own DSPIs and the ways in which their self-presentational choices may jeopardise those interests. The very fact of being prompted to consider how long you want this image, or block of text to be viewable by others puts us in the habit of thinking about our self-presentations as things that take on a temporally extended life of their own.

Perpetual data storage didn't emerge haphazardly – the individuals behind SNSs actively discourage reticence; they push everyone to express more than is necessary under the guise of serving friendship, communication, and public debate. If users share less personal content, then platforms such as Facebook modify the platforms to encourage more sharing (Hoffman 2016: 4). 'Share everything you can, is the message. Tell us who you are, how you feel… tell the world what you think about other people' (Véliz 2020: 109). The more information they have, the more they have of you stored. The more complex network of understanding they have of you, the more profit they make. Their business models depend on excessive and extensional disclosure.

SNSs shouldn't make it this hard for us to enjoy privacy. However, we live in an unideal world, one in which technologies of the self do not have our best interests as their priority. Expiry dates make it harder (even if just a little bit) for these companies to wring sellable data out of their users: they set limits on how long data can be stored (and thus used) and encourage us to think twice before speaking out loud online. By moving away from perpetual disclosures, expiry dates call for a remodelling of SNSs which places users in the driving seat – to move closer to a world where technologies of the self further the self, rather than inhibit it.

SUMMARY

The proliferation of SNSs has completely transformed the contours of privacy. We disclose aspects of our private lives online with little rumination; baring aspects of our selves for the ears and eyes of anyone who happens to be passing by. Alleviating some of the concerns I

[46] Though each individual decision may be simple, the cumulative effect of setting an expiry date for each individual post may not be an 'easy' task. SNSs users, especially those who post frequently, may opt for whatever happens to be the fastest route to posting. This illuminates how the use of nudges may be central – for example, there could be a default for all posts (say, 6 months) at the end of which a post will be archived, and can either be manually un-archived, or permanently deleted. Thank you to Micheal Garnett, for this discussion.

have outlined in this thesis will require us to distance ourselves from the current culture of oversharing and clear the moral fog of the internet that leads us to be blind to the morally salient features of our actions.

Though social media users collude in these sorts of harms, SNSs have created social affordances which have introduced social customs and individual incentives to acts in ways that are insufficiently inattentive to our DSPI-related responsibilities. They thwart our diachronic self-presentational autonomy by making it easy for us to undermine both ourselves and others. Because of this, we have good reason to move away from the current guises of SNSs. As I have argued, expiry dates could be a good place to start.

To Conclude

Our self-presentational choices vary over the course of our lives: we remodel our wardrobes and modify and embellish our facial features – we change our beliefs and revise our opinions. This is the diachronic character of self-presentational autonomy. The central claim I have made in this thesis is that we ought to have some control over the public construal of our past self-presentational choices, by the time we consider our former opinions and bodies discordant with our current selves. When other people observe and or distribute our prior self-presentational choices without our permission, we lose control over the terms by which the grounds of our self-esteem are knowable to others – we are presented to the world in ways we do not wish to be perceived. Prior to SNSs these interests were relatively shielded – there were relatively few ways in which other people could set-back or otherwise undermine our DSPIs. But technologies of the self have democratised these sorts of harms – we have less diachronic self-presentational autonomy and less privacy.

Forgetting has (for most aspects of self-presentation, in most societies) been very easy. If individuals or groups wanted to commemorate aspects of their lives, then huge amounts of effort were needed: paper was expensive, and we needed space to store it, writing demanded time and dedication. Modern technology has fundamentally altered the ease at which we are able to remember; we have found ourselves in a world in which forgetting has become the exception and remembering the default. A world in which our callow views and childhood foibles are available for others to observe and scrutinise. Given the harms liable to arise in this temporal panopticon,

we ought to start thinking carefully about what we can do to make this world less precarious – to make forgetting just a little easier than it is to remember.

References

Aced, C. (2013) Web 2.0: The Origin of the Word that has changed the way we understand public relations, in *The Open University of California* pp.1-12.

Allen, A. L. (2016) Protecting One's Own Privacy in a Big Data Economy, in *Harvard Law Review* 130(2) pp.71-80.

Allen, A., 1988, *Uneasy Access: Privacy for Women in a Free Society*, Totowa, N.J.: Rowman and Littlefield.

Altman, I. (1975) *The Environment and Social Behaviour: Privacy, Personal Space, Territory, Crowding* Brooks/Cole, Monterey.

Arendt, H (1998) *The Human Condition* The University of Chicago Press.

Arendt, H. (1963) *Eichmann in Jerusalem: A Report On The Banality Of Evil.* Introduction by Amos Elon. Penguin Books.

Aristotle, (1984) *The Complete Works of Aristotle, vol. II*, ed. Jonathan Barnes. Princeton University Press.

Austin, L. M. (2014) Enough About Me: Why Privacy is About Power, not Consent (or Harm), in *A World Without Privacy*. United States: Cambridge University Press. pp.131–189.

Bakardjieva, M. & Gaden, G. (2012) Web 2.0 Technologies of the Self. In *Philosophy & Technology*. 25(3) pp.399–413.

Basu, R. (2022) The Importance of Forgetting. In Episteme 19(4) pp.471–490.

Baumann, H. (2008) Reconsidering Relational Autonomy: Personal Autonomy for Socially Embedded and Temporally Extended Selves, in *Analyse & Kritik*. 30(2), pp.445–468.

Bellow, S. (1984) *Dangling man*. London: Secker & Warburg.

Berardi, F. (2007) The Pathologies of hyper-expression: Discomfort and Repression, Translated by Arianna Bove, in Art and Police (6) available at: https://transversal.at/transversal/1007/berardi-aka-bifo/en (Accessed: 1st August 2023).

Bernal, P. (2018) *The Internet, Warts, and All: Free Speech, Privacy and Truth* Cambridge University Press.

Billingham, P. & Parr, T. (2020) Online Public Shaming: Virtues and Vices, in *the Journal of Social Philosophy*. 51(3) pp.371–390.

Borgmann, A. 1992, *Crossing the Postmodern Divide*, Chicago: University of Chicago Press.

Borgmann, A., 1984, *Technology and the Character of Contemporary Life*, Chicago: University of Chicago Press.

Bork, R. (1990) *The Tempting of America: The Political Seduction of the Law*, London: Sinclair-Stevenson.

Bowles, C. (2018) *Future Ethics*, Now Next Press.

Boyd, D. (2007) Why Youth (Heart) Social Network Sites: The Role of Networked Publics in Teenage Social Life, in *the MacArthur Foundation Series on Digital Learning – Youth, Identity, and Digital Media* (ed. David Buckingham). Cambridge, MA: MIT Press.

Buss, S. and Westlund, A. (2018) Personal Autonomy, in The Stanford Encyclopaedia of Philosophy (Spring Edition), Edward N. Zalta (ed.), URL = https://plato.stanford.edu/archives/spr2018/entries/personal-autonomy/. (Accessed: 2nd January 2023).

Callard, A. (2018). *Aspiration: The Agency of Becoming*. Oxford: Oxford University Press.

Christman, J. (1991) Autonomy and Personal History, in *the Canadian Journal of Philosophy*. [Online] 21(1), pp.1–24.

CJEU, Court of Justice of the European Union (2014) An internet search engine operator is responsible for the processing that it carries out of personal data which appear on web pages published by third parties, Press Release 70(4) Luxembourg Available at: https://curia.europa.eu/jcms/upload/docs/application/pdf/2014-05/cp140070en.pdf (Accessed: 5th May 2023).

Cocking, D. & van den Hoven, J. (2018) *Evil online,* Hoboken: Wiley.

Cocking, D. (2008) 'Plural Selves and Relational Identity: Intimacy and Privacy Online', in *Information Technology and Moral Philosophy*. [Online]. United Kingdom: Cambridge University Press. pp. 123–141.

Cohen, J. (2002) *Regulating Intimacy: A New Legal Paradigm*, Princeton: Princeton University Press.

Cohen, J. (2013). What is Privacy For, in Harvard Law Review, 126(7) pp.1904–1933 Available at: http://www.jstor.org/stable/23415061 (Accessed: 11th November 2022).

Copp, D. (2002) Social Unity and the Identity of Persons, in *The Journal of Political Philosophy* 10(4) pp.365–391.

Crisp, R. (2018) Evil Online and The Moral Fog. Available at: http://blog.practicalethics.ox.ac.uk/2018/10/evil-online-and-the-moral-fog/ (Accessed: 02nd August 2023).

Davis, F. (1959) What Do We Mean by 'Right to Privacy' in Philosophy and Public Affairs 4 pp.1-24.

Dean R. and Parlavecchio, H. (2006) The Unwanted Exposure of the Self: A Phenomenological Study of Embarrassment, in *The Humanistic psychologist*. 34(4) pp.321–345.

DeCew, J. (1997) *In Pursuit of Privacy: Law, Ethics, and the Rise of Technology*, Ithaca: Cornell University Press.

DeCew, J. (2018) Privacy, in The Stanford Encyclopaedia of Philosophy (Spring Edition), Edward N. Zalta (ed.), Available at: https://plato.stanford.edu/archives/spr2018/entries/privacy (Accessed: 10th October 2022).

Desmond, H. (2022) Reclaiming Care and Privacy in the Age of Social Media, in *Royal Institute of Philosophy Supplement*, Cambridge University Press 92, pp.9245-9266.

Douglas, D. (2016) Doxing: a conceptual analysis, in *Ethics and Information Technology*. 18(3), pp.199–210.

Dreyfus, H. L. (2002) Anonymity versus Commitment: the dangers of education on the Internet, in *Educational Philosophy and Theory* 34(4), pp.369–378.

Dwivedi, Y. K. et al. (2016) *Social Media: The Good, the Bad, and the Ugly 15th IFIP WG 6.11 Conference on e-Business, e-Services, and e-Society, I3E 2016, Swansea, UK, September 13–15, 2016, Proceedings*. 1st ed. Springer International Publishing.

Dworkin, G. (1981) The Concept of Autonomy, in *Grazer philosophische Studien*. 12(1), pp.203–213.

Elliott, D. & Soifer, E. (2022) AI Technologies, Privacy, and Security, in *Frontiers in Artificial Intelligence*, (5)82673, pp.1-8.

Feinberg, J. (1980) The Child's Right to an Open Future, Chapter Three in *Freedom and Fulfilment,* Princeton University Press pp.76-96.

Floridi, L. (2015) 'The Right to Be Forgotten': a Philosophical View, in *Jahrbuch für Recht und Ethik Annual Review of Law and Ethics*, 23, pp.163–171.

Foucault, M. (1982) Technologies of the Self. In L. G. Martin, H. Gutman, and P.Hutton (Eds.), *Technologies of the self: A seminar with Michel Foucault* Amherst: University of Massachusetts Press, pp.16-50.

Foucault, M. (2008) Panopticism, from Discipline and punish: The Birth of The Prison, in *Race/Ethnicity: Multidisciplinary Global Contexts* 2(1), pp.1-12.

Frankfurt, H. G. (1971) Freedom of the Will and the Concept of a Person, in *The Journal of philosophy* 68(1), pp.5–20.

Franks, M. (2012) Sexual Harassment 2.0, in *Maryland Law Review* 71(3) pp.665-704.

Fried, C. (1970) *An Anatomy of Values: Problems of Personal and Social Choice*, Cambridge University Press.

Frye, H. (2021) The Technology of Public Shaming, in *Social Philosophy & Policy*. 38(2), pp.128–145.

Gavison, R. (1980) Privacy and the Limits of Law, in *The Yale law Journal*. 89(3), pp.421–471.

Gerstein, R. 1978, Intimacy and Privacy, in *Ethics*, 89(1), pp.76–81.

Goffman, E. (1955) On Face Work: An Analysis of Ritual Elements in Social Interaction, in *Psychiatry* 18(3) pp.213-231.

Goffman, E. (1959) *The presentation of self in everyday life*. New York, N.Y: Anchor Books.

Goffman, E. (1963) *Behaviour in Public Spaces*. The Free Press.

Gorzeman, L. & Korenhoff, P. (2017) Escaping the Panopticon Over Time, in *Philosophy & Technology*. 30(1), pp.73–92.

Gross, H. (1971) Privacy and Autonomy, Chapter 9 in *Privacy and Personality,* Routledge pp.169-181.

Hadjin, M. (2018) Privacy and Responsibility, *in Core Concepts and Contemporary Issues in Privacy*, Springer Publishers pp.63-78.

Hoerl, C. & McCormack, T. (2016) Making Decisions about the Future: Regret and the Cognitive Function of Episodic Memory, Chapter 12 in *Seeing the Future*. [Online]. New York: Oxford University Press, pp.241-266.

Hoffman, A. (2016) Facebook is Worried About Users Sharing Less – but it only has itself to blame, Guardian. Available at: https://www.theguardian.com (Accessed 24th August 2023)

Holstein, J. and Gubrium, J. (2000) *The Self We Live By Narrative Identity in a Postmodern World*, Oxford University Press.

Inness, J. C. (1992) *Privacy, Intimacy, and Isolation*. New York: Oxford University Press.

James, W. (1890) *The principles of psychology, Vol I.* New York, NY: Henry Holt and Co, Inc.

Johnson, J. L. (1989) Privacy and the Judgment of Others, in *The Journal of Value Inquiry*. [Online] 23(2), pp.157–168.

Jones, M. (2012) You are What Google Says You are: The Right to Be Forgotten and Information Stewardship, in *International Review of Information Ethics*, 17 Available at SSRN: https://ssrn.com/abstract=2154353 (Accessed: 3rd March 2023).

Khoury, A. C. & Matheson, B. (2018) Is Blameworthiness Forever? In *The Journal of the American Philosophical Association*. 4(2), pp.204–224.

Leary, M. (1996) *Self-Presentation: Impression Management and Interpersonal Behaviour*. Boulder, CO: Westview Press.

Lever, A. (2012) Privacy, Private Property, and Collective Property. In *The good society* - PEGS (Organization), 21(1), pp.47–60.

Lewis, M. (1997) The Self in Self-Conscious Emotions, in *Annals of the New York Academy of Sciences*, 818(1), pp.119–142.

Manson, N. C. & O'Neill, O. (2007) *Rethinking Informed Consent in Bioethics*, Cambridge University Press

Mantelero, A. (2013). The EU Proposal for a General Data Protection Regulation and the roots of the 'right to be forgotten, in *Computer Law & Security Review*, 29(3), pp.229–235.

Markham, F. (1963) *Napoleon / Felix Markham*. London: Weidenfeld & Nicolson.

Marmor, A. (2015) What Is the Right to Privacy? in *Philosophy & public affairs*, 43(1), pp.3–26.

Marmor, A. (2021) Privacy in Social Media, in *The Oxford Handbook of Digital Ethics*. Oxford University Press. pp.1-16

Marwick, A. E. & Boyd, D. (2011) I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience, in *New Media & Society*. 13(1), pp.114–133.

Mayer-Schönberger, V. (2009) *Delete: The Virtue of Forgetting in the Digital Age*, Princeton University Press.

McCullagh, K. (2008) Blogging: Self-Presentation and Privacy, *in Information & Communications Technology Law,* 17(1), pp.3–23.

Mead, G. H. et al. (1934) *Mind, self & society from the standpoint of a social behaviorist / George H. Mead ; edited, with introduction, by Charles W. Morris.* Chicago, Ill: The University of Chicago Press.

Michelfelder, D. P. (2001). The Moral Value of Informational Privacy in Cyberspace, in *Ethics and Information Technology,* 3(2) pp.129–135.

Miller, M. & White, B., (2021) The Warped Self, Aeon. Available at: https://aeon.co/essays/social-media-and-the-neuroscience-of-predictive-processing (Accessed 2nd August 2023).

Moore, A. (2008) Defining Privacy, in *The Journal of social philosophy*, 39(3), pp.411–428.

Moore, A. D. (1998) Intangible Property: Privacy, Power, and Information Control, in *American Philosophical Quarterly* (Oxford). 35(3), pp.365–378.

Moore, A. D. (2003) Privacy: Its Meaning and Value, in *American philosophical quarterly* (Oxford). 40(3), pp.215–227.

Moore, A. D. (2018) Privacy, Interests, and Inalienable Rights, in *Moral Philosophy and Politics*. [Online] 5 (2), 327–355.

Moore, Adam D. (2016) Waiving Privacy Rights, The Brookings Review. Available at: https://www.researchgate.net/publication/283420850_Waiving_Privacy_Rights (Accessed: 2nd January 2023).

More, T. et al. (2002) *Thomas More : Utopia / edited by George M. Logan and Robert M. Adams*. Rev. ed. Cambridge: Cambridge University Press.

Morris, G. H. M. A. C. W. (2009) *Mind, Self, and Society: From the Standpoint of a Social Behaviourist*. University of Chicago Press.

Munch, L. A. (2020) The Right to Privacy, Control Over Self-Presentation, and Subsequent Harm, in *The Journal of applied philosophy*, 37(1), pp.141–154.

Nagel, T. (1998) Concealment and Exposure, in *Philosophy & public affairs*. 27(1), pp.3-30.

Negley, G. (1966) Philosophical Views on the Value of Privacy, in *Law and Contemporary Problems*. 31(2), pp.319–325.

Newman, A. L. (2015) What the 'right to be forgotten' means for privacy in a digital age, in *Science* (American Association for the Advancement of Science), 347(6221), pp.507–508.

Nissenbaum, H (2010) *Privacy in Context: Technology, Policy and the integrity of Social Life*, Stanford Law Books.

Nissenbaum, H. (1997) Toward an Approach to Privacy in Public: Challenges of Information Technology, in *Ethics & Behaviour*, 7(3), pp.207–219.

Nissenbaum, H. (1998) Protecting Privacy in an Information Age: The Problem of Privacy in Public, in *Law and philosophy*. 17(5/6), pp.559–596.

Nissenbaum, H. (2004) Privacy as Contextual Integrity. *Washington Law Review*, 79(1), pp.119–158.

Nussbaum, E. (2007) Say Everything, in *New York (1968)*. 40(5), pp.24–35.

Nussbaum, M. (2004) *Hiding from Humanity: Disgust, Shame, and the Law* Princeton, NJ: Princeton University Press.

O'Callaghan, P. (2018) Privacy and Responsibility, *in Core Concepts and Contemporary Issues in Privacy*, Springer Publishers pp.63-78

Olson, E. T. (1997) *The Human Animal: Personal Identity Without Psychology* New York: Oxford University Press.

Owens, D. (2022) *Bound by Convention: Obligation and Social Rules / David Owens*. First edition. Oxford: Oxford University Press.

Parent, W. A. (1983) Privacy, Morality, and the Law, in *Philosophy & public affairs*, 12(4), pp.269–288.

Parfit, D. (1984) *Reasons and persons / by Derek Parfit*. Oxford: Clarendon.

Park, R. E., (1950) *Race and Culture*, Free Press.

Parker, L. S. (2002) Information(Al) Matters: Bioethics and The Boundaries Of The Public And The Private, in *Social Philosophy & Policy*, 19(2), pp.83–112.

Parker, R. (1974) A Definition of Privacy, in *Rutgers Law Review* 275 pp.283-295

Paul, L. A. (2014) *Transformative Experience*. Oxford: Oxford University Press.

Pettigrew, R. (2022) Autonomy for Changing Selves, Chapter 18 in *The Routledge Handbook of Autonomy,* Routledge pp.199-209.

Prosser, W. L. (1960) Privacy, in *California law review*. 48(3), pp.383–423.

Rachels, J. (1975) Why Privacy is Important, in *Philosophy & Public Affairs*. 4(4), pp.323–333.

Räikkä, J. (2017) Privacy and Self-Presentation, in *Res Publica* (Liverpool, England) 23(2), pp.213–226.

Raz, J. (1986) *The Morality of Freedom* Oxford: Clarendon Press.

Reiman, J. H. (1976) Privacy, Intimacy, and Personhood, in *Philosophy & Public Affairs*. 6(1), pp.26–44.

Roessler, B. (2008) New Ways of Thinking about Privacy, Chapter 38 in *The Oxford Handbook of Political Theory*. Oxford University Press.

Rosenberg, J. (2009) Online Impression Management: Personality Traits and Concern for Secondary Goals as Predictors of Self-Presentation Tactics on Facebook, A Thesis for the Degree of Master of Arts**,** Kent State University. Available at: file:///C:/Users/helen/Downloads/kent1259197291.pdf (Accessed: 8th August 2023).

Salahu-Din, T. (2017) The Evidence of Things Unsaid: The Lynching of Matthew Wiliams, The National Museum of African American History & Culture, Smithsonian Open Access. Available at: https://nmaahc.si.edu/explore/stories/evidence-things-unsaid (Accessed: 30th August 2023).

Scanlon, T. (1975) Thomson on Privacy, in *Philosophy & Public Affairs*. 4(4), pp.315–322.

Schafer, A. (2011) Privacy: A Philosophical Overview, Chapter 1 in *Aspects of Privacy Law*, Department of Philosophy, University of Manitoba.

Schectman, M. (1996) *The Constitution of Selves*. Cornell University.

Schlenker, B. (1980) *Impression Management: The Self-Concept, Social Identity, and Interpersonal Relations*. Monterey, CA: Brooks/Cole.

Schoeman, F. (1984) Privacy: Philosophical Dimensions, in *American Philosophical Quarterly*, 21(3), pp.199–213.

Shew, M. & Garchar, K. (2020) *Philosophy for Girls: An Invitation to the Life of Thought*. Oxford University Press.

Shiffrin, S. V. (2014) *Speech Matters: on lying, morality, and the law*, Princeton: Princeton University Press

Shoemaker, D. W. (2010) Self-exposure and Exposure of the Self: Informational Privacy and the Presentation of Identity, in *Ethics and Information Technology*, 12(1), pp.3–15.

Simpson, R. (2020) Speaking Freely, Thinking Aloud, and Embarrassing Yourself on Social Media, in *ABC Religion and Ethics,* pp.1-10 Available at: https://www.academia.edu/42969627 (Accessed: 2nd February 2023)

Simpson, R. (2021) The Ethics of Quitting Social Media, in Carissa Véliz (ed.), *The Oxford Handbook of Digital Ethics* Available at: https://doi.org/10.1093/oxfordhb/9780198857815.013.34 (Accessed: 10th August 2023)

Smart, B. (1976) Diachronous and Synchronous Selves, in the *Canadian Journal of Philosophy,* 6(1), pp.13–33.

Spurgin, E. (2019) Why the Duty to Self-Censor Requires Social-Media Users to Maintain Their Own Privacy, in *Res Publica* (Liverpool, England), 25(1), pp.1–19.

Stahl, T. (2020) Privacy in Public: A Democratic Defence, in *Moral philosophy and politics*, 7(1), pp.73–96.

Steiner, P. (1993) On the Internet, Nobody Knows You're a Dog, Published in The New Yorker, archived by The University of North Carolina, Available at: https://web.archive.org/web/20051029045942/http://www.unc.edu/depts/jomc/academics/dri/idog.html (Accessed: 16th November 2022).

Strawson, G. (2004) Against Narrativity, in *Ratio (Oxford)*, 17(4), pp.428–452.

Susser, D. (2016). Information Privacy and Social Self-Authorship, *Techné: Research in Philosophy and Technology* 20 (3) pp.216-239.

Tavani, H. T. (2007) Philosophical Theories of Privacy: Implications for an Adequate Online Privacy Policy, in Meta philosophy, 38, pp.1-22.

Tavani, H.T. (1999) Informational Privacy, Data Mining, and the Internet, in *Ethics and Information Technology* (1), pp.137–145.

Thompson, Judith Jarvis (1990) *The Realm of Rights*, Harvard University Press.

Thomson, J. (1975) The Right to Privacy, in *Philosophy & Public Affairs* 4(4), pp.295-314

Vallor, S. (2016) *Technology and the Virtues: a Philosophical Guide to a Future Worth Wanting.* New York, NY: Oxford University Press.

Vallor, S. (2022) Social Networking and Ethics in *The Stanford Encyclopaedia of Philosophy*, Available at: https://plato.stanford.edu/archives/fall2022/entries/ethics-social-networking (Accessed: 5th June 2023).

Véliz, C (2020) *Privacy is Power: Why and How You Should Take Back Control of Your Data*, Transworld Publishers.

Véliz, C. (2018) In the Privacy of Our Streets, in *Surveillance, Privacy, and Public Space*, edited by Bryce C. Newell, Tjerk Timan and Bert-Jaap Koops, London: Routledge, pp.16-32.

Véliz, C. (2019) The Internet and Privacy, in *Ethics and The Contemporary World,* Abingdon: Routledge.

Véliz, C. (2022) Self-Presentation and Privacy Online, in *the Journal of Practical Ethics* 9(2) pp.30-43.

Velleman, J. D. (2009) *How We Get Along.* Cambridge: Cambridge University Press.

Viganò, E. (2022) *Moral Choices for Our Future Selves: An Empirical Theory of Prudential Perception and a Moral Theory of Prudence*, Routledge Focus.

Warren, S. and Brandeis, L. (1890) The Right to Privacy, in *Harvard Law Review,* 4(5) pp.193–220.

Wenar, L. (2023) Rights, in *The Stanford Encyclopaedia of Philosophy* (Spring Edition), Edward N. Zalta & Uri Nodelman (eds.), Available at: https://plato.stanford.edu/archives/spr2023/entries/rights/ (Accessed 3rd March 2023).

Westin, A. (1968) *Privacy and Freedom,* The Association of the Bar of the City of New York.

Wilson, J. & Rumbold, B. (2019) Privacy Rights and Public Information, in *The Journal of Political Philosophy*, Volume 27(1), pp.3-25.

Wolford, B. (2018) Everything you need to know about the "Right to be forgotten" Available at: https://gdpr.eu/right-to-be-forgotten/ (Accessed: 21st August 2023).