



Universiteit
Leiden
The Netherlands

Emotional privacy: a philosophical investigation of emotion recognition technology

Prégent, A.P.

Citation

Prégent, A. P. (2026, September 2). *Emotional privacy: a philosophical investigation of emotion recognition technology*. Retrieved from <https://hdl.handle.net/1887/4309114>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/4309114>

Note: To cite this publication please use the final published version (if applicable).

Part Three:
The (Ethical) Analysis

5. Privacy and emotional life: Re-engineering the concept of emotional privacy

5.1. Introduction

The previous chapter argued that emotional expressions are probabilistically informative patterns of behaviour and physiological changes that provide perceivers with reliable, though defeasible, affective information. Emotional expressions thus sustain some of the core background conditions that make mutual understanding, interpretation, and warranted reliance possible in social life. They enable social orientation, decision-making, coordination, and trust, yet they remain situated, contested, and subject to deliberate modulation within context-sensitive practices of interpretation. It is this calibrated interplay between legibility and opacity that the prospect of fully functional ERTs threatens to reconfigure.

This chapter responds to that technomoral disruption by re-engineering the concept of emotional privacy. I argue that emotional privacy concerns the expressive ability of individuals to shape the legibility of affective information in interaction. This positions emotional privacy within the social perspective of (informational) privacy, according to which privacy is not primarily a retreat from others, but a constitutive condition of social life. I present emotional privacy as a subconcept of the informational privacy cluster, where the object of emotional privacy is (a type of) information. Its aim is thus not to introduce a new way of conceiving the broad concept of informational privacy, in the way that other accounts do (such as the Control-based approach, Access-based approach, and Contextual Integrity), but rather to offer a new conceptual tool within the informational privacy cluster in face of the new tech-induced conceptual gap that fully functional ERTs risk creating. After identifying the conditions under which emotional expressions can perform their social functions, I propose emotional privacy as a new conceptual tool to solve this problem.

Section 5.2 traces the conceptual evolution of privacy and shows how technological shifts repeatedly generate the need for conceptual adaptation. **Section 5.3** draws on the case of mental privacy to illustrate how subconcepts of privacy can be pro-actively added in anticipation of emerging technologies. Building on this precedent, I define two types of emotional privacy: actual emotional privacy and standing emotional privacy. While in everyday life we often lose actual emotional privacy, standing emotional privacy is typically retained. **Section 5.4** examines how emotional privacy operates within affective communication channels. It introduces

expressive conventions, affective leakages, and two types of privacy mechanisms – selective disclosure and concealment – as the dynamic tools through which individuals negotiate affective legibility. **Section 5.5** turns to the value of emotional privacy and shows that these privacy mechanisms are functionally significant at three levels of social interaction: they preserve self-presentation at the individual level; sustain role plurality and norm conformity at the interpersonal level; and contribute to social regulation and cohesion at the societal level.

5.2. Privacy

The previous chapter argued that emotional expressions (EEs) are probabilistically informative yet defeasible signals that sustain core conditions of social life. Building on that account, this section examines the relationship between privacy and emotional life, particularly in technosocial contexts. It traces the historical and conceptual evolution of privacy, showing how shifting social norms and successive technological innovations have repeatedly challenged and reshaped our understanding of both emotional life and privacy. By showing that privacy is not a static concept but one that evolves in response to technological and cultural changes, the section highlights why (inner) emotional life, once thought inaccessible, is now at risk of being exploited by ERTs.

In what follows, I first examine two relations that will help me foreground the concept of emotional privacy in light of new technological progress, namely the relation between privacy and emotional life and between privacy and technology. I then focus on the recent developments in (informational) privacy scholarship.

5.2.1. Privacy and emotional life

Although privacy was only recognised as a human right in 1948, when the United Nations enshrined it in the *Universal Declaration of Human Rights*, its concrete expression has taken various forms over the years. Particularly its related subconcepts, such as spatial privacy, proprietary privacy, bodily privacy, associational privacy, decisional privacy, and more recently, informational privacy, have been understood, recognised, and used differently over time but also across different cultures (Koops et al. 2016). For instance, whereas the bedroom is today generally regarded as a private space in French culture (spatial privacy or proprietary privacy), this was not the case for the monarchs, particularly Louis XIV, who turned waking up and going to bed into public rituals for his courtesans and subjects. Another notable example is the evolution of bodily privacy, which was shaped by the global eugenics movement and later

deconstructed and reshaped by feminist movements in the twentieth century. From the marital-rape exemption to the forced and coerced sterilisation of vulnerable groups in many Western countries, such as Denmark (Bryant 2025), the USA (Peal 2004), and Canada (Stote 2015), the concept of bodily privacy has changed profoundly in the past century, with most Western countries currently recognising bodily privacy as the capacity to control access to and interference with one's body, protecting against unwanted intrusion in medical, technological, and social contexts.¹⁰⁴

Against this backdrop of profound shifts in the conceptualisation and collective understanding of privacy, emotional life stands out as a persistent constant. Even in highly remote cultures – such as the Javanese, the Ngadju Dayaks of Borneo, the Chocó Indians of Panama, or the !Kung Bushmen of Southwest Africa – where physical, spatial, and even bodily privacy have no equivalent to the conceptions endorsed by modern Western societies, at least a part of emotional life is nonetheless voluntarily and carefully concealed under finely tuned decorum (Altman 1975: 15-16, 42).¹⁰⁵

Despite this, scholarship on the privacy of emotions remains scarce. Perhaps partly because the concept of privacy did not need to be as prominent in academic discussions as it is today, and partly because emotions were somewhat *de facto* recognised as a deeply personal and inaccessible individual phenomenon to such an extent that the violation of this sphere was deemed impossible.¹⁰⁶

The advent of new intrusive forms of technologies is what brings privacy and emotion scholarships together. In Warren and Brandeis seminal paper on privacy (1890), we can observe that their concern about public image, rising from the advent of new technologies, led them to go as far as to argue for a protection of individuals' "thoughts, feelings and sensations" (Warren & Brandeis 1890: 195). In 1988, Anita Allen expressed concerns over the individuals' ability to protect their inner life against new forms of (techno-)intrusions, specifically noting the potential access to emotions. She argued that to say that a person has privacy is to say that "in some respect and to some extent, the person (or the person's mental state, or information about the

¹⁰⁴ This excludes some countries, such as the United States, which has recently seen major setbacks on this subject, particularly with regard to reproductive rights and the recent overturning of abortion laws (see the overturning of *Roe v. Wade* by *Dobbs v. Jackson Women's Health Organization* in 2022).

¹⁰⁵ See also Westin (1967: 13-18).

¹⁰⁶ They are sometimes treated as a form of qualia (see Chalmers 1996; Nagel 1974).

person) is beyond the range of others' five senses and any devices that can enhance, reveal, trace, or record human conduct, thought, belief, or emotion.” (Allen 1988: 15).¹⁰⁷

5.2.2. Privacy and technology

The rapid evolution and development of privacy scholarship in the last two centuries is mainly due to the development of new technologies that forced a reconceptualization of our conception of privacy. New technologies have consistently pushed against the boundaries of what counts as private. Each wave of innovation has unsettled prevailing social norms and revealed conceptual gaps in the understanding of privacy, forcing philosophers and jurists to re-articulate its meaning.¹⁰⁸

To reuse Warren and Brandeis's case, instant photography and the phenomenon of mass press that emerged from it were what provoked anxiety about exposing intimate life without consent, pushing for the adoption of a right 'to be let alone'.¹⁰⁹ By the mid-twentieth century, the spread of computers, databanks, and record-keeping systems shifted attention to the circulation and control of personal information. The most pressing threat was no longer photographers at the window but the invisible, large-scale capture of personal data. Concerns about government databanks and commercial profiling crystallised around the concept of informational privacy. These early concepts of informational privacy came again under strain at the turn of the twenty-first century with the rise of digital networks, online platforms, and ubiquitous surveillance. These technologies blurred the delineation of public and private spaces and cast doubt on the feasibility of both informational “control” and limited access so understood.

Every new technological wave brought with itself a need for conceptual adaptation with regard to privacy. Instant photography, x-ray devices, phone trackers, computer databanks, and today, digital networks and ubiquitous technology-based surveillance are reshaping our very understanding of what it even means to ‘have privacy’. As for the mental and emotional life, the

¹⁰⁷ Interestingly, she also notes how one's emotions and other mental states may be readily available to others through external states such as emotional expressions. She observes that: “[f]acial expressions, gestures, and body movements can give us away. Happiness and anger, approval and disapproval, sincerity and deception are often discernible by external inspection” (1988: 14). Hence, “the external self can reveal the internal self, [and thus] inaccessibility to persons' bodies and products of their bodies may be needed to keep the private self private.” (1988: 14).

¹⁰⁸ See Hopster & Löhr (2023) for a detailed argument that technologies create conceptual disruptions.

¹⁰⁹ In the nineteenth century, privacy was mainly understood as a protective boundary around the individual that secured autonomy, dignity, and freedom from intrusion (Mill [1859]1985; Westin 1967; Thomson 1975; Moore 2010). In this liberal tradition, the private realm functioned as the counterpart to the public, a space in which individuals could retreat, develop, and exercise self-determination (Mokrosinska 2018).

apprehensions expressed by Warren and Brandeis now find concrete echoes in contemporary forms of affective (techno-)weaponization, such as nudges (Wilkinson 2013; Thaler & Sunstein 2008, 2021), neurotechnology (Ienca 2021; Magee, Ienca, and Farahany 2024), and fully functional ERTs. Emotional life, once considered inaccessible, now risks being exploited and commodified through technologies that promise to render it transparent.

Now that I have established the links between both privacy and the emotional life, and privacy and technology, I now turn more particularly to this very concept of (informational) privacy, and establish what is the current state of affairs in terms of protection against new and invasive technologies.

5.2.3. (Informational) privacy scholarship

The deployment of technology that could process and keep detailed records of the masses in the mid-twentieth century shifted scholarly attention from the previously dominant notions of privacy that were proprietary and spatial privacy toward the emerging concept of informational privacy.¹¹⁰ To grasp what ‘informational privacy’ should entail, two general, but also competing notions were seen as plausible, namely, the control- and access-based views.

A prominent defender of the Control-based approach was Alan Westin, who defined privacy in *Privacy and Freedom* (1967)¹¹¹ as the individual’s ability to control the *disclosure* and *circulation* of personal data, in other words, the ability to determine “when, how, and to what extent information about them is communicated to others.” (1967: 7). On this view, privacy is an exercise of choice and agency, modelled on a liberal notion of individual autonomy. Informational privacy thus emerged as a response to a new socio-technical environment, where harms arose less from direct intrusion into one’s private (physical) sphere and more from the accumulation, storage, and re-use of informational traces. This now Classic Control account of privacy can be formalised as

¹¹⁰ See Roessler (2005) and Moore (2010) for modern counterarguments against this focus on informational privacy. Roessler and Moore defends the relevance of other types of privacy, particularly local and decisional privacy for Roessler and spatial and bodily privacy for Moore. Roessler, for instance, argues that “Something counts as private if one can oneself control the access to this ‘something’.” (Roessler 2005: 8). However, as I am focusing on affective communication channels and EEs, I will only gesture at these competing views and restrict my analysis to informational privacy.

¹¹¹ See also Moore (2003, 2010) and Fried (1968).

Classic Control: Privacy obtains when person A has the ability to give access to information p, which she possesses, to (any) person B.

The access view, by contrast, argued that (informational) privacy is not adequately protected by the notion of control and instead should be best viewed as being about the notion of ‘access’, understood in broad terms as the extent to which others can reach the individual. What I take as the now ‘Classic Access-based accounts’ thus emphasise the *passive condition of being shielded from (unwanted) access*. Ruth Gavison (1980), for instance, defined privacy as “our concern over our accessibility to others” (1980: 423). Jeffrey Reiman defined privacy as “the condition under which other people are deprived of access to either some information about you or some experience of you” (1976: 30). Consequently, the Classic Access account of privacy can be formalised as

Classic Access: Privacy obtains when information p, possessed by person A, is inaccessible to (any) person B.

Mainly due to their dialogue with each other where defenders from both sides restlessly pointed at the weaknesses of the other’s accounts, the Control and Access views have evolved from their respective classic forms. Hence, building on the classic form of the Control-based account, contemporary defenders, such as Mainz and Uhrenfeldt (2021) and Menges (2021, 2022) recognised that the Classic Control formulation was too imprecise on the *meaning* of control, leaving the general account weak and difficult to defend, hence pressing for a revision of the classic control formulation.¹¹² Taking up this challenge, Mainz and Uhrenfeldt (2021) propose to revise the control account along the idea of ‘Negative Control’.¹¹³ Negative Control is understood, as its name indicates, as the ability to prevent access. Contrary to the classic account of control, which present an active notion of control in which people may give or prevent access to information (presumptively at any time and under any conditions), Mainz and Uhrenfeldt

¹¹² The formulation of the Classic Control indeed leaves the door open to many cases where *a loss of privacy does not entail a loss of the ability to control the disclosure of information p*, and vice versa. For instance, as it is formulated, person A could “keep the ability to control” who will know about p, while person B could still correctly infer p from other epistemic sources. Conversely, it has also been argued that A “loses control” but not her privacy if, for instance, she attempts to give B access to her information p, and B *prevents* it. Following Mainz and Uhrenfeldt’s assessment for example, this does not intuitively lead to a *plausible* loss of A’s privacy.

¹¹³ They also advance the possibility of a “Republican Control” notion of privacy, in which person A’s privacy is lost simply in virtue of B’s ability to access her information. But they flag it as also problematic as this ultimately means that the mere invention of new surveillance technologies, without their use, diminishes A’s privacy, which also does not intuitively lead to a plausible violation of A’s (moral) right to privacy (Mainz & Uhrenfeldt 2021: 300). Hence, I will not consider this formulation.

notion of negative control presents control as a reactive notion which focuses on the ability to prevent access *once* someone does attempt to access. Following this notion,

Negative Control: Privacy obtains when, given an access attempt by (any) person B, person A possesses the capacity to prevent B from accessing personal information p about A.¹¹⁴

Menges (2021, 2022) recently defended another alternative on how to best conceive the notion of control, which he mainly grounds in the moral responsibility literature (2021: 395). He argues that control should be defined as

Source Control: Privacy obtains when person A is the *right kind* of source for information p's flow (if information p flows at all).

Following the Source Control account, the focus is not on the way person A would lose control over B's access to information p – what Menges relates to (effective) voluntary choice – but rather on the legitimacy of A as the *right kind* of source for information p's flow (viz., sourcehood). Hence, Menges considers other accounts to be concerned with different kinds of (effective) voluntary choice (viz., control) over the access of (any) person B to information p about person A, while the Source Control account is concerned with the legitimacy of the source (viz. the right kind) making the information flow. This view has several advantages over the Classic Control view, as it tracks back privacy loss to instances where someone else makes personal information about the agent flow without the agent's input as the source. Take the case where person B steals the diary of person A, but does not read it. In such cases, the Classic Control account struggles to explain why there seems to be no violation of privacy because even though there is clearly a loss of control on the part of person A, this loss did not result in an access to the information. In contrast, the Source Control account does not (yet) identify a privacy violation the moment person B steals the diary. This is because the focus is not on whether B has lost (effective) control over the information, but on B being the wrong kind of source for the information flow, if the information flows at all. As a result, person A,

¹¹⁴ Note that for reasons of clarity, I only refer to the (last) formulation of Mainz and Uhrenfeldt's control account, in which they say that person A's right to privacy is violated iff "Agent A has involuntarily lost Negative Control over the access to personal information P about agent A, *due to action(s) of agent B, of which B is responsible.*" (2021: 298, authors' original italics). I will return to this last formulation in the next chapter as it is intended for a normative account. Here I only want to highlight the diverse views and evolution of (informational) privacy.

presumptively the only right kind of source for this diary, has lost privacy over the diary's information *only* in the case that B does access it, since B is the wrong source of the flow.

Similarly, on the Access view, compelling alternatives to the classic formulation are to be found in Anita Allen's (1988),¹¹⁵ Tavani and Moor (2001), and Macnish's (2018) work. Allen and Tavani and Moor both propose a Restricted Access account to privacy. The restriction is mostly determined in relation to the context, the situation, and the actors present. On Allen's account, restricted access means that privacy is a condition in which access to persons and personal information is appropriately limited. On Allen's view, to say that privacy has been violated is to say that someone has been made accessible in contexts where *restricted access* is the very point of the practice.

In a similar vein, Tavani and Moor write that there are many established "zones" in which the protection of privacy is already reasonably expected and normatively protected.

The normative aspect of these private situations restricts access by individuals, groups, or governments. This restricted access expresses a right of protection. It prohibits intrusion and information processing by someone or something. [...] Although there is some conventionality in how we carve up zones of privacy in social situations including Internet use, overall the carving should produce a set of zones that offer sufficient protection of personal information. [...] In defining a private situation it is necessary to define who has access to what under which circumstances. (Tavani & Moor 2001: 7).

Tavani and Moor also raise the issue of the repurposing of the data gathered on individuals' personal information. For instance, does the fact that I have willingly granted access to my personal information *p* to person B, in circumstance *x*, also mean that now they can sell it and have it accessed by person C, in circumstance *y* (Tavani & Moor 2001: 6)? Below I have summarised the idea of Restricted Access accounts as

Restricted Access: Privacy obtains when only members of a specified set *R* have access to information *p* in the given situation, where *R* is fixed by the situation's role- and purpose-structure.

Another alternative to the Access view has recently been proposed by Kevin Macnish (2018) who insists that privacy is lost only when there is *actual* (epistemic) access. On his account, automated capture or background processing that no individual meaningfully understands does

¹¹⁵ Later, Anita Allen defines it as "limiting access to personal information and electronic communications" (2011: 196), thus slightly modifying the view she had previously defended.

not by itself reduce individuals' privacy. The loss occurs when a person who can understand, in fact accesses information *p* in a way that can be brought to bear in practice (Macnish 2018, 2020). Björn Lundgren (2020, 2021) and Peter Königs both endorse Macnish's requirement of actual (epistemic) access.¹¹⁶ Addressing more specifically the problem of government surveillance, Peter Königs (2022) builds on the Actual Access account of privacy within institutions, tying privacy loss to the point of human access, and in a second step, legitimating the loss within role-justified purposes.¹¹⁷ The reduction of privacy is bound to the actual (epistemic) access while the right to privacy is then assessed on whether the access was carried out by an authorised individual or not (Königs 2022). Hence the Actual Access account of privacy reads as

Actual Access: Privacy obtains when no person *B* with understanding has actually accessed *A*'s information *p*.

In parallel to the co-development of the Control- and Access-based approaches is notable work in what is best conceived as the social perspective of (informational) privacy. In response to new socio-technological threats, scholars such as Helen Nissenbaum (2004, 2010), Julie Cohen (2013), and Beate Roessler and Dorota Mokrosinska (2013, 2015) reconceived the value of informational privacy as a constitutive feature of social life, reshaping privacy's relevance in the face of new technologies. Leaving both the notion of control and access aside, Nissenbaum's Contextual Integrity (CI) framework, for instance, focuses instead on the appropriateness of the information flow. Among its strongest advantages is its fine-grained property, enabling it to explain, for instance, why "ten different people observing the contents of your shopping cart on ten different occasions is not the same, from a privacy perspective, as one party recording and storing it" (Nissenbaum 2019: 237, italics in the original).

Nissenbaum leaves the realm of privacy being seen as controlled or accessed by an individual, and redirects the attention towards the actual movement of information. CI starts by considering privacy as *appropriate* flow of information correlating violations of privacy as *inappropriate* flow of information (Nissenbaum 2010, 2018). This appropriateness is defined by a respect for the context's own integrity. This integrity is, in turn, itself determine through an analysis of the information flow in a given context. To be able to trace back the information

¹¹⁶ Lundgren argues for a pro tanto right (to privacy) against risks, meaning that this right can be infringed before actual access occurs in the case of serious attempts or engineered risk-impositions (Lundgren 2021).

¹¹⁷ In that sense, his view can be seen as a hybrid account between Restricted Access and Actual Access.

flow, CI focuses on the informational norms at play in a context. These informational norms are characterised by key parameters that are present in every context: the actors (senders, data subject, recipients), the transmission principles, and the attributes. These parameters respectively define the parties who are the subjects of the information as well as those who are sending and receiving it, the principles under which this information is transmitted, and the types of information at stake (Nissenbaum 2010: 142–143). Hence, flows of information are shaped by key parameters, and embedded in specific contexts. Libraries, health care system, voting stations, schools, airport security, etc., are all contexts.

Contexts are considered social domains with their own teleological values and purposes (Nissenbaum: 2010, 134). Hence, social contexts have entrenched social norms embedded in them (Nissenbaum 2010: 132-134, 146). These social norms help achieve the context's values, which take the form of goals, purposes, or ends. Prevailing values in the U.S. health care system, for instance, include alleviating physical suffering, curing illness, and promoting the health of individuals as well as collectives (Nissenbaum 2010: 134). Hence, informational norms and values are both in a co-constitutive relationship with their context, shaping it and being shaped by it (Nissenbaum 2010: 141, 180).¹¹⁸

On her account, thus, CI can be summarised as

Contextual Integrity: Privacy obtains when information p flows in a way that respects the prevailing entrenched informational norm(s) at play in a given context.¹¹⁹

If a new practice or new technology provokes a change in the flow of information, the practice is “flagged as violating entrenched informational norms, which constitute a *prima facie* violation of contextual integrity” (Nissenbaum 2010: 150). This departure from the entrenched context-relative informational norms at play in the context is assessed by drawing a “comparison [...] between [the] entrenched practice and the novel practice” (Nissenbaum 2010: 148-149).

Once the practice is flagged, it is not rejected *ipso facto* but needs to be assessed normatively. The normative assessment can be seen as the second part of the evaluation and mainly relies on

¹¹⁸ Part of this discussion is from Prégent (2025b).

¹¹⁹ I will take the moral augmentation part of the framework into account in the next chapter.

the values of the context. CI uses the values as evaluative tools to analyse the new technology. The values of a context are the result of a collective judgment that is, similar to the context-relative informational norms, also historically, culturally, and geographically shaped. Hence, a new technology that induces a departure from these entrenched norms is thus seen as an *a priori* threat, but can nonetheless be upheld against the entrenched practice if it can show that it respects and promotes the values of the context; thus maintaining the integrity of the context (Nissenbaum 2010: 186).

Let me briefly compare CI to the Source Control and Restricted Access accounts.¹²⁰ In the case of Menges' Source Control account of privacy, the account remains a control-account of privacy, only the notion of control is interpreted differently. This means that, on the Source Control account, to have privacy is to be the right kind of source of a data flow while “no receiver and no conscious agency is needed” (Menges 2022: 696).¹²¹ Hence the focus remains on the individual and her information p, contrary to CI, where the focus is on the informational flow itself and where receivers are of particular importance as they often end up being the reason of a ‘prima facie’ violation (e.g. in cases where a receiver changes, which means that a new direction is imposed on the original flow).¹²²

The Restricted Access account is also intrinsically different from CI. On the Restricted Access account, privacy is the condition of “restricted access”, which means that privacy is obtained when there is either a non-access or an authorised access. The focus is on who may look, hear, touch, smell, etc., which is dependent on role-justified permissions. On the Restricted Access account, privacy loss occurs when *unauthorised* parties gain access to information p, regardless of whether information p flows in a way that respects the context's integrity. This means that, while in cases where the flow is disrupted by a change in actor, CI and the Restricted Access account are likely to be similar in their output, in other cases, such as for instance, repurposing information p from health data to marketing data, they will likely diverge. The reason is that change of context is not in itself considered a valid ground for a loss of privacy following the

¹²⁰ For the Source Control account, see Menges (2022), footnote 2, for his full distinction between Source Control and CI.

¹²¹ Although Menges (2021, 2022) remains silent about what is entailed by the “right kind of source”, I assume that this is what he has in mind. This would mean that in the case of Source Control, a violation occurs when the authority of the source is impeded, bypassed, or defeated.

¹²² Following CI, privacy is lost (on a prima facie basis) when the information flow is altered, which may be identified through a change in actors, transmission principles, or attributes (Nissenbaum 2010: 142-143). The violation counts as a *proper* privacy violation once the moral analysis, evaluating the new information flow, fails to demonstrate that it upholds the (prevailing) values, ends, purposes, and goals of the context, in the same way – or in a better way – than the entrenched information flow did.

Restricted Access account. A loss of privacy would happen only when a person (or party), who has no authorised access, nonetheless gains access to information p . In cases where a parameter other than the ‘actors’ changes, only CI can flag the new flow.

Control-, Access-, and Contextual Integrity accounts are foundational touchstones for work on informational privacy. They shape how we think about and analyse flows of information, conditions of access, and appropriateness. Yet as mentioned by many scholars, the breadth of (technology-mediated) cases appears to resist any single, all-purpose definition (Roessler & DeCew 2023; Koops et al. 2016; Solove 2008) of informational privacy. This partly explains why, as noted in §5.2.2, conceptions of privacy are being repeatedly redefined at an ever-increasing pace. In the next section, I introduce the concept of emotional privacy as a complement to informational privacy in cases specifically tied to emotional life. I clarify the notion of tech-induced conceptual gaps and show how conceptual re-engineering can help address them. Building on DeCew’s cluster account of privacy (1997), which distinguishes bodily, spatial, and informational privacy, among others, I propose that emotional privacy is best understood as a subconcept of informational privacy. In this sense, I do not offer a competing theory of informational privacy, as Control-based, Access-based, and Contextual Integrity accounts do vis-à-vis one another. Rather, I articulate a new subconcept within informational privacy to better capture what is at stake, in privacy terms, when nonverbal emotional expressions are targeted by new and emergent technologies.

5.3. Emotional privacy

In this section, I propose to re-engineer the concept of emotional privacy by relying in particular on the case of mental privacy, which I take as an informative precedent to emotional privacy. Then, following the insights from the last chapters, I anchor emotional privacy in the social privacy perspective and define the concept.

I suggest *re-engineering* emotional privacy because in some ways “emotional privacy” is neither old nor new but at the same time it is also both an old and new concept. Philosophers have long recognised that emotional life is, at least in part, inaccessible to others, sometimes even treating emotions as qualia.¹²³ Yet recent scholars, including Menges and Weber-Guskar (2025), Akar (2024), Gremsl and Hödl (2022), and McStay (2018, 2020) argue that the exposure of the inner

¹²³ See Chalmers 1996; Nagel 1974.

affective life has become an urgent concern in the age of affective computing. The emotion recognition industry now being a fast-growing, multibillion-dollar market, with rising interest across public agencies, workplaces, education, and social media in deploying fully functional tools, only adds to this concern (Schwartz 2019; Miles 2024).¹²⁴ McStay (2020) was among the first to raise the alarm about the lack of protection surrounding (emotional) soft biometrics, identifying it as a significant regulatory gap. In response, Gremsl and Hödl (2022) and Menges and Weber-Guskar (2025) propose recognising information about emotions as ‘emotional data’ and integrating it into existing categories of personal or sensitive data, including the GDPR’s special categories. Akar (2024), while offering no explicit definition, highlights the need for a concept like “emotional privacy”. Upon closer examination, however, this emphasis on the *inner* affective life¹²⁵ overlooks a crucial dimension of emotional life and therefore fails to capture the constitutive social dimension of the concept, leaving it incomplete.

Focusing only on safeguarding inner affective experiences can allow us to focus on issues linked to the risks of certain inferences, which may be helpful in shedding light on important relations that emotions share with values, judgements, and beliefs, for instance, but it also makes us overlook the strong interplay between the inner and outer parts of emotional life (see Chapter 4). As a result, this focus blinds us to the important role emotional privacy plays in sustaining the core social functions of affective life. As much of the literature has so far focused on the former, I dedicate the remainder of this thesis to shedding light on the latter. To address this issue, I propose re-engineering the concept of emotional privacy in a way that accounts for its social dimension, as discussed in the preceding chapter.

In the remainder of this chapter, I thus go beyond current scholarship on the matter and (i) articulate what ‘emotional privacy’ is and (ii) specify precisely what it protects. This definition will help me to anchor the normative conditions of a (moral) right to emotional privacy in the next chapter. Chapter 6 will consequently formulate a proposition for a moral right to emotional privacy using the concept developed here and applying it to the case of ERTs to identify the normative conditions of its protection and distinguish wrongful violations of emotional privacy from ordinary mindreading.¹²⁶

¹²⁴ See Chapter 2.

¹²⁵ Although Akar’s analysis remains focused on the inner emotional experience, she also acknowledges the importance of social life in shaping that experience (Akar 2024).

¹²⁶ ‘Mindreading’ here is simply meant as the human ability to understand and attribute emotions and intentions to others from the interpretation of physiological and behavioural changes.

5.3.1. The case of mental privacy as a precedent to emotional privacy

We use concepts and conceptual systems to make sense of our complex reality. Concepts help us to put into neat categories different phenomena of life (Haslanger 2000; Cappelen 2018; Burgess & Plunkett 2013). Scholars working on conceptual engineering argue that, within our collective conceptual systems, *conceptual gaps* may be induced by new and emergent technologies (Hopster and Löhr 2023). “Tech-induced conceptual gaps emerge when we lack the conceptual resources to classify and conceptualise novel technologies, yielding conceptual uncertainty” (Hopster & Löhr 2023: 10). Such gaps may motivate new needs for *conceptual adaptation*, which is best understood as a response to a need for either (i) the introduction of a new concept, (ii) the elimination of an existing concept, or (iii) the preservation of a concept (Hopster & Löhr 2023: 8; see also Haslanger 2000; Cappelen 2018).

A canonical case of a tech-induced conceptual gap is the concept of “brain death”. “Brain-Death” is the result-concept of patients with severe brain injury whose brain activity is non-existent (or considered equivalent to it) while lung function is retained (indefinitely) through the invention of mechanical ventilation. The advent of mechanical ventilation created a conceptual gap in our life-or-death collective conceptual system where people were unsure about the classification of these patients: were they alive or were they dead? The ensuing uncertainty prompted the introduction and consolidation of the concept of “brain death” to respond to this new tech-induced need (Beecher 1968; Capron 1988).¹²⁷

Another illuminating example, which I take as a great precedent to emotional privacy, is the recent case of *mental privacy*. Control-based, Access-based, and, more recently, Contextual Integrity, have forced clear splits in the privacy scholars’ ideas on what it means to have informational privacy. But while the broader understanding of how we should best conceive it is still debated, the development of the subconcept of mental privacy shows that there is value in supplementing these broader conceptions with complementary subconcepts that are narrower and targeted to the exact problems that arise from particular tech-induced conceptual needs. Hence, I assume that mental privacy is, like emotional privacy, a subconcept of

¹²⁷ This situation notably provoked high moral uncertainty, and a rather obscure path towards potential moral obligations and other normative considerations that typically follow the creation of a new concept. In this case, this led to moral uncertainty regarding the moral obligations towards these patients (Bernat 2006; Truog & Miller 2014).

informational privacy. On this understanding, the object of mental privacy is a type of information pertaining to brain-related mental content.

The emergence of the concept of mental privacy has been triggered by advances in neurotechnology. The term gained traction in response to the growing ability of neurotechnologies to record, decode, and infer mental content from brain activity. These technologies include electroencephalography, functional magnetic resonance imaging, and brain-computer interfaces. Although their fully functional forms, capable of reliably “reading” thoughts, do not yet exist, their potential has already exposed a conceptual gap in privacy theory.

Traditional notions of informational and bodily privacy, which are typically designed to regulate access to external representations of a person, such as medical and financial records, location data, and to protect against physical intrusions on the body (for example, strip searches or non-consensual examinations), fail to capture the distinctive exposure created by neurotechnologies, where the information at stake is not cleanly separable from its source (as they are ongoing neural processing) and thus exemplifies what Ienca and Andorno (2017) call the *inception problem*, namely a potential direct access to mental content itself.

This perceived tech-induced need has led scholars to propose mental privacy as a new (sub)concept. Ienca and Andorno (2017) argued that the rise of neurotechnologies requires a “neurorights” framework including the right to mental privacy, defined as the right to keep one’s brain data and mental content protected from unauthorised access.¹²⁸ Similarly, Yuste et al. (2017) called for international recognition of mental privacy as a fundamental right to safeguard individuals from invasive brain-recording and decoding tools. Subsequent literature has refined the notion. For example, Lippert-Rasmussen (2017) linked mental privacy to epistemic access to neural representations, Ienca (2021) and Magee, Ienca, and Farahany (2024) expanded it to include inferences about any mental states drawn from brain-derived data, and Susser and Cabrera (2024) showed that the field remains divided over whether the target of protection is brain-related data, mental content, or the broader integrity of mental life.¹²⁹

¹²⁸ In their paper, they proposed four neurorights: cognitive liberty, mental privacy, mental integrity, and psychological continuity. Mental privacy, which is defined as the right to keep brain-related information from *unauthorised* or *actual access*, therefore directly expands on the access view.

¹²⁹ Note that it is also not fully accepted that mental privacy is best understood as a protection over *information* per se (or data), as some argue that the loss is best understood as something like a bodily violation (Barnard et al. 2024).

What unites these positions is the recognition that existing privacy frameworks were not designed for the epistemic intrusions posed by neurotechnology.¹³⁰ Control-based accounts focus on disclosure and consent, Access-based accounts attempt to pre-empt the information flow; and Contextual Integrity approaches assess the appropriateness of the information flow given the prevailing informational norms of a given context. None of these, however, can maintain a single definition of informational privacy in which neural activity or brain-signals are fully protected while other types of information are protected (conditionally) by pro tanto reason. This is so because mental privacy is typically conceived and defended as a new fundamental human right, whereas informational privacy is typically conceived as a pro tanto right.¹³¹ Mental privacy was therefore engineered as a complementary concept to respond to this new privacy need created by neurotechnology.

Conceptually, mental privacy operates as a protective condition concerning brain-derived information. This move has methodological significance. Mental privacy was not introduced because existing privacy laws had already failed, but because scholars anticipated that they would fail once fully functional neurotechnologies appeared. It represents an instance of proactive conceptual adaptation, creating conceptual resources ahead of empirical necessity to guide future normative reasoning and regulation.

This case provides a methodological precedent for the present project. It exemplifies how conceptual engineering can function as a philosophical counterpart to technological anticipation, ensuring, as Hopster and Löhr (2023) note, that our conceptual landscape evolves in step with emerging forms of human-technology interaction. Just as mental privacy was introduced to anticipate the implications of fully functional neurotechnologies, emotional privacy is introduced here to anticipate the implications of fully functional ERTs. In what follows, I thus make use of this methodological approach in the affective domain.

¹³⁰ Recent national and supranational discussions have concluded that the concept of mental privacy is needed for the proper protection of brain-related content (see for instance Council of Europe (2019)).

¹³¹ In brief, mental privacy's advocates defend the idea that there are never any "good" reasons that would outweigh a person's right to mental privacy (hence why it is often presented as a fundamental right) while there might be good reasons to outweigh someone's right to informational privacy (hence why it is usually conceptualised as a pro tanto right). Consider the workplace. As an employee, you might have a fundamental right to have your brain-related activity safeguarded from your employer or the company, but you do not hold such right regarding your salary or your name.

5.3.2. Anchoring emotional privacy in the social perspective

In the Information Age, we have been prompted to first worry about the information itself, which is in most cases the core of the problem: who loses information, who gains information, and who does what with the newly acquired information. Within this paradigm, we often forget that with the loss or illegitimate acquisition of data come other forms of harms. With emotional privacy, I attempt to show that the kind of losses that are invoked by the prospect of fully functional ERTs include unique kinds of harms, which are interconnected with the functions EEs serve in social life. These harms are to be added to the kind of harms that have already been identified, most of them relating to emotions in particular (Gremsl & Hödl 2022; Menges and Weber-Guskar 2025). To shed light on this, I turn to the social perspective of privacy.

Several scholars have argued that privacy serves interests extending beyond the individual, with works highlighting its constitutive role in specific relations to society in general (Fried 1968; Rachels 1975; Altman 1975; Cohen 2012, 2013; Roessler & Mokrosinska 2013).

Already in 1968, Charles Fried's account, though still situated within an individualistic, control-based framework, marked a clear departure from the classical liberal approach to privacy (viz. Mill [1859]/1985) by placing relationally dependent values such as love, friendship, and trust at the centre of the concept of privacy. On Fried's view, privacy cannot be reduced to a merely individual interest in freedom or autonomy conceived as a retreat from the social world. Rather, privacy constitutes the necessary foundation of social relations, for it enables the informational space on which love, friendship, and trust depend (Fried 1968: 205). The possibility of being human without these relationships, Fried insists, is hardly conceivable, since they form some of the most valuable dimensions of human life. Privacy thus safeguards what is most precious in our moral and social existence: it regulates the flow of information in social interactions, enabling individuals to determine when and to whom they reveal aspects of themselves, thereby protecting the very conditions under which intimacy, trust, and friendship can develop and endure.^{132, 133}

¹³² Note that while it is true that Fried defined privacy in terms of control-based approach (Fried 1968, 482), his insistence on the idea that individuals' having privacy is a condition for their being able to engage in interpersonal phenomena such as trust and love, opening the door to the social privacy perspective which later came to rely on his account (see Roessler & Mokrosinska 2013).

¹³³ The informational privacy perspective, especially its control-based proponents, had already incorporated some of the core elements of a conception of social privacy. Hence, similarly to Fried, Westin (1967), also described privacy as a *dynamic process* that regulates role demands and creates spaces of respite from them.

Whereas Fried establishes privacy as the condition that makes intimate relationships possible, James Rachels (1975) generalises this insight. He argues that privacy's importance cannot be limited to intimacy alone but extends to the entire web of social roles through which persons engage with others. In his essay *"Why Privacy is Important,"* Rachels argues that the significance of privacy cannot be grasped solely through the values of love, trust, and friendship as it requires an analysis of privacy's role in structuring the diversity of relationships that constitute human life. On Rachels' account, each relationship presupposes a distinct pattern of informational access. Privacy is the regulative condition that allows individuals to manage these patterns, controlling who knows what about them in ways appropriate to context: "our ability to control who has access to us and to information about us [shapes] our ability to create and maintain different sorts of social relationships with different people" (1975: 326). What one discloses to a physician about one's health differs from what one shares with an employer or confides in a close friend; without privacy, these distinctions collapse, and all relationships risk flattening into an undifferentiated field of exposure. A loss of privacy in this sense would entail the loss of the very possibility of maintaining the diversity of relationships that gives human life much of its richness. Rachels therefore argues that privacy is indispensable to the very organisation of social life as it protects the plurality of social roles.

A further development of the social perspective is found in Helen Nissenbaum's theory of Contextual Integrity (2004, 2010, 2018), which reframes privacy as the preservation of appropriate information flows within socially defined contexts. On this account, each domain of social life, such as healthcare, education or friendship, is structured by context-specific norms that govern what kinds of information may be shared, with whom, and under which transmission principles (Nissenbaum 2004, 2010). Nissenbaum's approach strengthens the insight, already present in Rachels's work, that privacy is constitutively tied to the maintenance of social roles. Yet it also shifts attention from individual preferences to the integrity of social practices themselves. What matters is whether the informational flows enable or undermine the proper functioning of the context. This structural focus is highly relevant for emotional privacy, since emotional expressions are embedded in normative expectations that regulate how affective information circulates within different social settings.

Fried's and Rachels' postulate – that the value of privacy lies primarily in the virtue of constituting social relations – has also been the stepping stones of other, more recent work, notably by Roessler and Mokrosinska (2013, 2015) and Mokrosinska (2018), who argue that

privacy should be conceived not merely as an individual or interpersonal value but also as a social and political one.¹³⁴ On their view, privacy's function is not exhausted by the management of interpersonal access to information as it also benefits the relations between individuals. Their argument (2013) pushes the social conception of privacy further by allowing privacy to protect elements *within* the relationship and *of* the relationship. In intimate relationships for example, privacy protects both the individual's ability to withdraw from forms of surveillance and intrusion that would erode their ability to form independent judgements and the capacity to "live as mature and responsible citizens in democratic societies" (2013: 775). At the same time, it also protects the social practice of intimacy itself, which "has to be protected from social and societal publicness, for the function and intention of these relationships can be realized only if their privacy is protected" (2013: 775). Privacy therefore serves both individual and collective interests; individually, it enables autonomy and the cultivation of personal relationships, while collectively it sustains social and political practices.¹³⁵ Roessler and Mokrosinska hence expand the concept of privacy from a condition of interpersonal relations to a structural requirement of social and political life (Roessler and Mokrosinska 2013; Mokrosinska 2018). While earlier relational accounts could still accommodate the early classical liberal view of the self, the social conception of privacy requires a decisive break from this atomistic view and rests instead on a social ontology of the self.

Early liberal understandings of privacy, grounded in the idea of a self that exists prior to and apart from social relations, can no longer capture the full spectrum of privacy violations that occur in contemporary, interconnected societies. As Michael Walzer noted, "[t]he individual who stands wholly outside institutions and relationships and enters into them only when he or she chooses and as he or she chooses: this individual does not exist and cannot exist in any conceivable social world" (Walzer 1984: 324). Whereas the early liberal view of privacy protected the individual's capacity to withdraw from social and political life, the social conception of privacy grows out of a non-atomistic understanding of identity (Mokrosinska 2018). On this view, identity is *constituted* through social roles, relationships, and contexts (Durkheim [1897]1952; Goffman 1959; Boyd 2008).¹³⁶ Individuals are therefore situated beings whose sense

¹³⁴ This argument does not only follow Fried and Rachels' work but also to some extent that of Gavison (1980), Reagan (1995), and Solove (2008).

¹³⁵ Mokrosinska (2014, 2018) pushes the argument in favour of privacy as a political value further, claiming that privacy preserves autonomy not as an individual good, but as a condition of democratic life. If identity and autonomy are inherently social, then privacy cannot be understood as disengagement from social life, but as its precondition.

¹³⁶ Sandel (1982), for instance, in a direct reply to Rawls' work on liberal justice, criticised the limits of liberalism by arguing for a redefinition of individual identity which should be conceived as deeply tied to social and communal

of self is relational, shaped by social norms, and influenced by the institutions and structures that organise collective life.

Social accounts of privacy take this social ontology as their foundation. The possibility and legitimacy of a social conception of privacy depend precisely on recognising that the self is socially constituted. If individuals are socially constructed, they can only exercise autonomy and develop their identity through their relationships with others rather than by withdrawing from society. Privacy's function, on this account, is to protect those relationships and the practices through which autonomy and identity are realised (Mokrosinska 2018).

Importantly, this shift does not lead social-privacy theorists to conclude that privacy does not protect individual interests. Rather, they maintain that privacy's role includes the protection of individual interests, but *extends* to the protection of social life itself. As Solove explains, privacy's value lies not only in shielding individuals but also in sustaining the activities and practices that make both individual and collective life possible. "Although privacy certainly protects individuals," he writes, "this does not mean that privacy is an individualistic right. Understanding privacy as having a social value – the benefits it confers on society by enhancing and protecting certain aspects of selfhood – does not oppose the individual to the community or value privacy solely in terms of its benefits to particular individuals. [...] The protection of privacy shields us from disruptions to activities important to both individuals and society" (Solove 2008: 92-93).

On this view, privacy is not only a personal safeguard but a constitutive condition of social life. It sustains the network of relationships upon which both individual flourishing and collective life depend. I propose to apply this rationale to the emotional life, as the emotional life is an essential part of what gives individuals a sense of self, while also being (at least partly) socially constituted. Accordingly, emotional privacy should be conceived within this broader social framework. Emotional privacy therefore is best understood as protecting individuals' expressive ability over affective information, ensuring that the *possibility* for expressive ability is preserved. Yet if we maintain this social perspective, and accept that emotional expressions and their causes are embedded in dynamic social processes, both influencing and being influenced by them, then emotional privacy cannot be regarded as something wholly possessed by individuals. Rather, it

contexts rather than independent from it. Feminists and communitarians have also criticised the classical liberal view of identity and respectively pushed for a reconceptualisation of one's identity as intrinsically social and shaped by personal and formal relationships (MacKenzie & Stoljar 2000), and have highlighted the crucial role of community, tradition, and shared values in shaping moral and political life (Etzioni 1999; Taylor 1992).

is a condition that is socially regulated, emerging from natural dynamics of affective life and from privacy mechanisms that individuals enact through their expressive capacities to modulate the flow of information about their affective states.

5.3.3. Defining emotional privacy

I have already mentioned that Control-based accounts focus on whether a person succeeds in regulating an information flow in a particular moment. Within this view, the control-based variants typically treat “control” as the central concept of informational privacy, where control is understood as the *successful regulation* of access to information p. They also usually (implicitly) assume that a person A possesses the ability required for such control. This assumption invites them to take for granted that A stands in a position, as a matter of general ability, to control access (Classic control), to block access attempts (Negative control), or to serve as the appropriate informational source (Source control). This is perfectly understandable, as outside the affective domain, it generally makes more sense to simply focus on the notion of control while implicitly assuming the ability to control. Hence, Control-based accounts operates at the control level, asking whether, in a particular moment, A succeeds or fails in regulating information p given some conditions. However, this distinction between “control” and the “ability to control” is what ultimately matters in the affective domain.¹³⁷

Building on the social functions of emotions and emotional expressions (see Chapters 3 and 4), and on recent work done in the social privacy perspective (§5.3.2), I argue that *emotional privacy* is best understood as a relational and social condition whose significance must be examined in terms of affective communication. I postulate that it is best to understand emotional privacy as operating at two interconnected levels: the actual level and the standing level. At the actual level, a person enjoys emotional privacy over a particular expressive event only insofar as she can *successfully exercise* expressive discretion in that moment. When a spontaneous EE is produced, for instance, the person loses actual emotional privacy with respect to that specific disclosure. Such losses are not only common in ordinary social life, they are also inevitable. Yet the person may retain

¹³⁷ Menges’ Source Control account is the one that comes closest to it since he distinguishes Source Control from “effective control”. However, to *be* the (right kind of) source of a given information flow is what ultimately matter for his account. Additionally, the problematic identified in his account is still found at the point of a given “actual” safeguard or lost of control over the information p’s flow (as the right kind of source of that information flow).

emotional privacy at the standing level, because she still possesses the *ability* to exercise expressive discretion. This refers to the ability of a person to use certain types of emotional expressions to disclose, conceal, reinterpret, or refute affective information. It is only when this expressive ability is weakened or defeated that emotional privacy collapses at the standing level. That normative concern will be addressed in the next chapter.

I therefore define emotional privacy as

Actual Emotional Privacy: Actual emotional privacy obtains when person A has expressive control over her affective information p.

Standing Emotional Privacy: Standing emotional privacy obtains when person A has the expressive ability to control her affective information.

This two-level structure reflects the nature of emotional life. Emotional expressions are either voluntary or involuntary. People do not enjoy full control over all their expressions, and they regularly experience affective leakages in ordinary social interactions. “Actual” losses of emotional privacy are therefore a common feature of everyday life. Yet, even when they lose (actual) control, they still retain the ability to control, as they can anticipate, manage, or refute meaning carried by EEs even when those actual losses occur.

By focusing on ‘control’ and the ‘ability to control’, emotional privacy naturally aligns with the control tradition of informational privacy in its broad sense. However, it does not align neatly with any of the Control-based variants seen above. In the affective domain, control cannot be exercised unless the person has access to a repertoire of controlled expressions and privacy mechanisms through which she can shape the affective informational environment (see following §5.4). This implies that actual control over affective information p may fail even though person A’s expressive ability is preserved, but actual control can never be retained if the expressive ability is lost. This is why emotional privacy departs from the traditional Control approaches. Emotional privacy, contrary to mainstream Control accounts, begins with the question of whether the person possesses the expressive ability that allows her to engage in expressive discretion at all.

Imagine a person sitting in a quiet corner of a café with a close friend, discussing a sensitive matter. Let's call him Daniel. Daniel chooses his words carefully, lowers his voice and keeps an eye on his surroundings. In these conditions, he retains control over who can hear the information *p* and when it becomes accessible. During the conversation he leans forward and whispers a particularly delicate remark. His friend is unsure about whether he heard him correctly and repeats Daniel's words aloud in a clearer, louder voice in order to confirm what he thinks he heard. A customer seated at a nearby table hears the repeated remark. In this scenario, Daniel loses (actual) control in the classic form, since he has not succeeded in regulating access to *p*. He loses negative control, since he was not in a position to prevent the other customer from hearing what he intended to keep private. He also loses source control in Menges' sense, since the information flow that reaches the third party does not originate from Daniel's communicative agency but from his friend's. Yet Daniel's "ability to control" remains intact. He can lower his voice again, move to another table, shift to coded phrasing or postpone the conversation. He has lost actual control over information *p*, but he retains the ability to regulate the flow of information through his speech in ordinary circumstances.

Now suppose that during another visit to the same café an employee slips a drug into Daniel's drink. As the drug takes effect, Daniel's cognitive and motor coordination begins to deteriorate. His speech becomes erratic and he begins revealing sensitive and personal information aloud. Several nearby customers hear the unintentionally loud disclosure. In this state, Daniel has lost actual control because he has lost the ability on which his actual control depends. His underlying ability to regulate his speech, to correct mistakes, or to rephrase information has been removed. Control is not simply unsuccessful in this instance. It has become unavailable to him altogether because the basic expressive resources (i.e. his expressive ability) that make control possible have been compromised by the drug's effects.

These two scenarios show the distinction between control (viz. actual emotional privacy) and the ability to control (viz. standing emotional privacy). People routinely lose actual control over their EEs in everyday social life. A blush, a wince, or a momentary tightening of the jaw can reveal more than a person intends, yet these episodes do not undermine the person's broader expressive ability. She can still reshape the interaction, reinterpret what has been seen or correct the impression through intentional expressions. By contrast, when expressive ability itself becomes unavailable, emotional privacy collapses at the standing level. This is structurally analogous to Daniel's second scenario. When a person cannot manage how affective

information is disclosed or interpreted, she lacks the ability that standing emotional privacy requires. Importantly, while expressive ability is a *necessary* (pre-)condition of actual emotional privacy, it is not sufficient. For instance, an emoter's own affective leaks can thwart their own (actual) control or perceiver's ordinary mindreading skills may also thwart that actual emotional privacy.

Chapter 6 will shed light on why this distinction becomes essential for understanding what is truly at stake on a normative level with emotional privacy and the use of fully functional ERTs, but in a nutshell, it will help us distinguish between legitimate and wrongful losses of emotional privacy.

In this section, I have briefly introduced the idea that people can manage their emotional expressions through some 'privacy mechanisms', but I have not yet defined them. I do so in the following section, where I aim at deepening this conceptualisation of emotional privacy through an analysis of its functions within affective communication channels.

5.4. Emotional privacy in affective communication channels

I have introduced emotional privacy as a subconcept of the informational privacy cluster and as aligned with the broad Control-based understanding of privacy. I have noted that Control accounts typically presuppose the ability to control (standing) while concentrating their analyses on particular instances of success or failure of control (actual). In the affective domain, by contrast, both actual and standing levels matter, even though the standing level is ultimately primary. The remainder of this chapter examines how actual emotional privacy is sustained through the functions of emotional privacy, understood as privacy mechanisms. The next chapter then explains why the standing level is normatively decisive, since the successful deployment of these mechanisms depends on it.

These functions of emotional privacy are best understood as enabling and regulating emotional expressions' functions. Whereas 'first-order' functions of emotional expressions concern what expressions do in interaction, the functions of emotional privacy concern the background conditions under which those first-order functions can succeed. In particular, emotional privacy sustains a workable space of calibrated opacity in affective interactions, maintaining neither full

secrecy nor full transparency. This space of calibrated opacity is where individuals can partially shape the legibility of affective information across settings and audiences.

Many things may cause informational opacity in affective communication channels. The obvious one is noise.¹³⁸ Others include expressive conventions, affective leakages, and privacy mechanisms. These may increase the affective *opacity* in a given communication channel, leading to a situation where the perceiver has only partial information or not enough information to successfully infer the meaning of the EEs. In turn, thus, affective *legibility* is seen as dependent on the degree of opacity present in affective communication channels.

Together with expressive conventions and affective leakages, privacy mechanisms create an interactive process that is best understood as a tense dialectic, where individuals play both the roles of emoters and perceivers in each interaction, negotiating what they reveal and reflecting on what they infer. Emoters may thus partially shape this legibility through the use of privacy mechanisms. Such mechanisms do not simply conceal or reveal; rather, they are strategic techniques regulating affective legibility; making emotional life more or less accessible. I distinguish two broad types: disclosure mechanisms, which typically decrease opacity by enabling selective access under chosen conditions, and concealment mechanisms, which typically increase opacity by limiting access to affective information.

5.4.1. Expressive conventions and affective leakages

Before turning to the privacy mechanisms that individuals may deploy to preserve emotional privacy in affective communication channels, I first overview two important background dynamics that put pressure on an emoters' own negotiation of emotional privacy within affective communication channels. First, expressive conventions need to be acknowledged as they often govern when and how emotional expressions should be produced. These conventions tell us when to smile and when to look angry, but correlatively they also guide us in anticipating when smiling and looking angry will be looked upon or disapproved. In that sense, they also participate in regulating the negotiation of emotional privacy by guiding people around what is expected but also what would be considered inappropriate.

¹³⁸ In this section, I ignore the role of noise to focus on the different types of functions. See Chapter 2, §2.2.3 for a discussion of noise in affective communication channels.

Second, while some expressions are intentionally produced to support interactional goals, others arise *spontaneously*. I have already discussed the latter in Chapter 4 (see §4.2.3), and will now refer to them as ‘affective leakages’. Affective leakages are micro, partial, or complete EEs that are not intentionally produced by the emoters, giving perceivers access to meaningful affective information that is not deliberately shared. Sudden physiological changes, such as blushing, pupil dilation, sweat, or shaking are classic examples of changes that produce affective leakages. Affective leakages may also occur through unintentional behavioural changes such as jumping or screaming out of fear, clenching one’s jaw out of anger, or even rolling one’s eyes. Some affective leakages are non-intentional in the sense that they simply cannot be controlled or suppressed by emoters, while others are simply “let out” in Green’s sense (2007), meaning that, even though they are typically classified as spontaneous, they could have been pre-empted by the emoters if they had wanted to. However, this presupposes that control over some types of affective leakages requires important “emotional labour”.¹³⁹ Following emotional labour theory,¹⁴⁰ it is possible that, given the ‘labour’ required to prevent affective leakages, in situations where they would be numerous, for instance, it is unlikely that an individual will have the emotional bandwidth to successfully pre-empt or suppress *all of them*, hence the second type may merge with the first one, transmitting meaningful information to perceivers without emoters’ deliberate intention.

This means that emoters’ management of their own affective legibility is always imperfect. Expressive conventions and affective leakages are background pressures shaping the negotiation of emotional privacy within affective communication channels. With this in mind, I now specify what privacy mechanisms are and how they function in affective communication. In a nutshell, although they rely on different strategies, they all enable and protect the negotiated and *in situ* interpretation of expressions through calibrated opacity.

5.4.2. Privacy mechanisms

The privacy mechanisms that I introduce below are inspired by Irwin Altman’s work (1975) on boundary-processes. Altman’s work is best categorised within the social perspective of privacy

¹³⁹ Emotional labour theories usually focus on the control that one can hope to have over one’s affective life through concepts such as “emotion regulation”, “internal management”, “emotional adjustment”, or “emotional contagion”, mostly studying the extent to which we possess that control. See Wilms, Lanwehr, & Kastenmüller (2020); Eisenberg (2000); Needham, Mastracci & Mangan (2017) for more on this particular point.

¹⁴⁰ An interesting insight from emotional labour theories is that the emotional labour required to control these leakages varies depending on the type of interaction, the context, and the interlocutors. These are all recognised as primary determinants of the degree of control one can hope to have over affective leakages (Xu & Fan 2023).

as he conceives privacy as a relational dynamic where privacy facilitates the management of social interactions by regulating desired and achieved levels of interaction.¹⁴¹ Altman's account, and the framework he developed from it, provides useful conceptual tools that can be reshaped into privacy mechanisms within affective communication channels. Altman argues that privacy needs to be conceived as:

[An] interpersonal boundary-control process, which paces and regulates interaction with others. Privacy regulation by persons and groups is somewhat like the shifting permeability of a cell membrane. Sometimes the person or group is receptive to outside inputs, and sometimes the person or group closes off contact with the outside environment (Altman 1975: 10).

On his account, understanding privacy as a dynamic concept is key to understand the regulation of interpersonal relationships and the self. Privacy boundaries enable us to mark a division between ourselves and the world. Altman's privacy approach does not only see privacy as a withdrawal process but as an interplay between withdrawing from social interaction and engaging in social interactions.¹⁴² His framework is conceived around an optimisation function between the desired level of contact with others and the achieved level of contact with others (Altman 1975). The desired level of contact is understood to be constantly changing over time, shaped by social contexts and types of relationships. Hence, the value of privacy is found in this role of interface between the self and the non-self. Privacy is achieved when the dialectic between the self and others – enabled by the boundary mechanisms (and their respect) – is obtained, meaning that the desired level matches the achieved level of contact. Altman's privacy-boundary account identifies specific social privacy-mechanisms that individuals, groups, and societies use as a way to express and achieve their desired level of interaction in different contexts and relationships (Altman 1975: 30). For instance, individuals can use their own body as a space-regulatory indicator during an interaction, communicating the level of desired contact they aspire to achieve. Additionally, companies, for instance, might design the office spaces in a way that shape the general work environment (i.e. flex-offices, open offices, insulated/separated offices).

I see the tense dialectic between expressive conventions, affective leakages, and emoters' privacy mechanisms in a similar way. I thus rely on Altman's core idea that individuals deploy certain

¹⁴¹ Altman's view of privacy itself builds both on Simmel and Westin, taking inspiration from Simmel's work on the "dialectic quality of social exchange" (Simmel 1950, in Altman 1975: 19) and Westin's analysis, composed of four types (solitude, intimacy, anonymity, and reserve) and functions of privacy (personal autonomy, emotional release, self-evaluation, limited and protected communication) (Westin 1967, in Altman 1975: 18). Moreover, the idea of a dynamic process regulating interaction is also found in other's work (see Smith et al. (1969), Schwartz (1968), McGinley (1959), Bates (1964), Jacobs (1961), Jourard (1966), Rapoport (1972)).

¹⁴² His framework also allows for feedback and readjustment during a given interaction (Wisniewski & Page 2022, 21).

types of privacy mechanisms to obtain the desired level of *contact* and repurpose it for the idea that individuals deploy privacy mechanisms within affective communication channels to obtain the desired level of *informational opacity*. Below I discuss the role of disclosure and concealment within the nonverbal emotional realm.

5.4.3. Disclosure privacy mechanisms

Disclosure is generally understood as sharing one's information with others. In turn, *self-disclosure* generally refers more specifically to the act of sharing one's own information, be it emotions, motivations, goals, thoughts, personal information, or past, present, and future experiences with others (Westin 1967; Derlaga & Berg 1987). Disclosure and self-disclosure are also usually seen as *intentional* revelations. But while in other contexts both disclosure and self-disclosure may be considered "full disclosure", in the realm of affects, this would be erroneous and misleading. When it concerns emotional life, disclosure is never "full", not only because the emotional experience is hardly shareable in its entirety, but mainly because individuals do not, and cannot, disclose every single emotional experience.

Consequently, from now on, I will refer to the privacy mechanism of disclosure as *selective disclosure*, to remain closer to reality. Selective disclosure thus enacts emotional privacy by making one's affective legibility conditional, contextual, and recipient-relative. It is *not* a loss of privacy but a mechanism through which privacy's broader social function may be performed, i.e. it partly governs the circulation of affective information so that interaction remains intelligible. Hence in that sense, spontaneous EEs, which are typically not controlled, are not part of selective disclosure as this mechanism is intentionally deployed.

Emotional privacy thus protects the conditions of affective legibility, and selective disclosure functions as one of its tools: it permits partial, targeted legibility for chosen perceivers, facilitating their interpretation.

5.4.4. Concealment privacy mechanisms

Concealing information is commonly regarded as *withholding* information from others. From this perspective, concealment mechanisms are often regarded with suspicion in debates about

privacy, trust, and honesty (Bok 1983).¹⁴³ Taken as such, it is easy to assume, then, that concealment simply amounts to deception. But while deception is commonly defined as inducing false beliefs, concealment in affective life typically functions differently. For one, individuals do not always aim to convince others of (affective) untruths, nor do they necessarily hide their “true” selves. Rather, they rely on this mechanism to present the self that best fits a given context or relationship. For another, despite the given fact that conveying emotions or social motives that are not currently felt or desired may sometimes mislead, it also, at least just as often, preserves social appropriateness, protects dignity, or sustains intimacy. Similarly to selective disclosure, then, the extent to which concealing one’s information is acceptable, permissible, or encouraged is thus often seen as dependent on the relationship the individual has with others (as the nature of the relationship usually models a significant part of the exchange of information (Rachels 1975; Fried 1968)) and on the context.

Within the concealment mechanisms, I distinguish two types. I refer to the first as “direct concealment”. This subcategory of concealment concerns privacy mechanisms that are used to *signal* to the other that the emoter does not wish to share inner affective states or social motives. Practical examples of these mechanisms include avoiding eye contact, neutral EEs, and increasing or maintaining physical distance. In this case, there is an overt attempt at signalling a desire to not share affective information.

Broadly conceived, this type of mechanism maintains, increases, or even decreases¹⁴⁴ the level of informational opacity presents in the affective communication channels. Moreover, it is not necessarily the case that maintaining or increasing informational opacity in the interactions indicates an unsuccessful attempt at using privacy mechanisms, as one might want to have this level of opacity.

The second subcategory of privacy concealment mechanisms I refer to as “indirect concealment”. These are cases where the emoter *does not signal* to the other that *she does not want*

¹⁴³ The “nothing-to-hide” argument against privacy, for instance, was at the centre of these debates. But as Sissela Bok rightly notes: “[...] while all deception requires secrecy, all secrecy is not meant to deceive” (Bok 1983: 7).

¹⁴⁴ Even though using concealment to decrease the level of informational opacity seem counter-intuitive, “direct concealment” may have this result, as the privacy mechanisms deployed are “out in the open”, meaning that perceivers may easily interpret that this person, for instance, “does not wish to share”. In that sense, it cannot decrease informational opacity regarding the causes of emotional expressions directly, but it may provide information on the willingness to share them (or not).

to *share* her inner affective state or social motive.¹⁴⁵ This is done by relying on the success of one's intentional emotional expressions to convey what the emoter judges as the best expression even though they do not sincerely experience them. Indirect concealment mechanisms can thus be used to follow conventional norms and express what is expected rather than what one feels, or to make clear a given stance, commitment, goal, or desire. Hence it might also be done to preserve a relationship or one's appearance. These are cases where one may *choose* to conceal their inner affective states and/or social motives and instead use an emotional expression that fits the context better. Typical cases would be to show the expression of sadness for someone's loss even though one does not feel sad, anger to support someone's side in a conflict even though one does not feel angry, or even happiness for someone's wanted pregnancy even though one does not feel happy.

These privacy mechanisms are functionally important as they regulate the inferences and predictions perceivers can make from emotional expressions, allowing individuals to manage the strategic ambiguity of their own emotional lives. In standard interactions, emotional privacy does not merely protect individuals from exposure but also secures the social and epistemic conditions under which affective legibility occurs. Emotional privacy's role is thus to safeguard emoters' ability to manage their affective legibility.

An important point is that all privacy mechanisms rely on the deployment of *intentional* EEs. I have already mentioned in Chapter 4 that *intentional* EEs serve important social functions. Empirical research has shown that EEs serve a broader range of interpersonal and regulatory purposes (Van Kleef 2016; Fischer & Manstead 2016; Kumar & Senthil 2021). Expressing love, compassion, and care, for instance, conveys affective commitment, fosters intimacy, and enables cooperation in close relationships (Gonzaga et al. 2001; Ellis & Malamuth 2004). Conversely, expressing anger can rectify social misunderstandings, pinpoint norm violations, clarify expectations, or challenge status arrangements (Lerner, Goldberg & Tetlock 1998; Averill 1982; Fischer & Roseman 2007).

Particular attention to intentional EEs is necessary to understand their role in the successful deployment of privacy mechanisms. While direct concealment may be achieved by at least allowing the emoters to remain neutral in their expressions, indirect concealment means that

¹⁴⁵ Classic examples of this type that have been largely discussed in the literature on emotional expression are the "survival" type of cases, where for instance one has *advantage* in showing anger (pretended emotion) rather than fear (true emotion) for the sake of their own survival (see Darwin [1872] 1965).

emoters are for instance able to express anger *without* experiencing it, which may, for instance, support a mother's education of her child by prompting him to reflect on his behaviour or help friends signal that they do not endorse a particular behaviour, even though in both cases the emoters are not actually angry. In selective disclosure, it allows people to emphasise and carefully craft the way they will express their emotional experience. The existence and the possibility to successfully deploy intentional EEs is what these privacy mechanisms are about.

Below I turn towards the value of emotional privacy, and overview three primary functions of the privacy mechanisms at three different levels of social interaction (individual, interpersonal, and societal).

5.5. The functional value of emotional privacy

In this section, I focus on the value of emotional privacy. I show that the privacy mechanisms are not just options at the emoters' disposition but are *integral* to the regulation of social life. They should be recognised as necessary preconditions to the enactment of expressive discretion because they are functionally essential. I argue that this is valuable on three different interaction levels, namely individual, interpersonal, and societal. At the individual level, they allow individuals to manage the self they want to present; at the interpersonal level, they enable the management of role plurality and ensure adherence to role-bound norms of expressiveness; and at the societal level, they allow *society itself* to regulate, ensuring cohesion. These different levels are not to be seen as mutually exclusive but rather as complementary, shedding light on different levels of impact.

5.5.1. Preserving self-presentation

At the individual level, philosophers and sociologists have often argued that privacy primarily helps individuals to conceal what might be called their “uncivilised inner life.” By this, they do not refer to something pathological,¹⁴⁶ but to the ordinary stream of impulses, desires, and emotions that fail to conform to established social norms. A persistent assumption is thus that society depends on such concealment. For instance, Simmel said that if one's “subjective, impulsive, intimate matters [...] were exposed, they would drive everybody into the insane asylum” (Simmel [1908]1950: 311–312). Similarly, in “Concealment and Exposure”, Nagel

¹⁴⁶ The concealment of pathological desires – such as those of paedophiles – cannot be assimilated into the same category. Here concealment may serve the individual, but it directly endangers others and erodes the very social goods that privacy is meant to protect.

posits that “there is much more going on inside us all the time than we are willing to express, and civilization would be impossible if we could all read each other’s minds” (Nagel 1998: 4).¹⁴⁷ It is difficult to deny that humans constantly experience emotions – envy, irritation, hostility, lust – that, if rendered fully transparent, would strip people of the possibility to decide what they choose to reveal, to whom, and when (Westin 1967). Emotional privacy ensures that the raw material of affective states and social motives (often disordered, fleeting, or socially misaligned) does not overwhelm the civility required for social life.

A second interconnected rationale is that concealment mechanisms enable moral reflection and value-commitment. As previously discussed in Chapter 3,¹⁴⁸ emotions play an essential role in individuals’ moral reasoning. Emotional privacy protects this role by giving individuals the space required to turn their attention inwards, to their inner life. At times, individuals experience emotions or have social motives that directly conflict with moral principles they hold dear. In that sense, the active suppression of such inner states can be morally admirable in itself, insofar as it demonstrates a prioritisation of enduring moral commitments. Emotional privacy thus allows people to work on themselves when they experience internal (moral) conflicts, enabling the development of important virtues. Concealing a fleeting attraction towards a friend’s partner, for instance, can develop into a morally appropriate act of self-regulation. Opacity in such a case preserves the friendship and the relationship, perhaps also forging a core value that will follow the individual for the rest of their life. By shielding “uncivilised” elements of our inner lives or interior dilemmas from immediate exposure, privacy enables individuals to reflect, recalibrate, and ultimately sustain their moral commitments within social and moral communities.

A third one is highlighted in the work of Erving Goffman (1959, 1967), who notes that individuals meticulously craft the self that they best see fit, with that ‘fittingness’ being modulated by the identity of the perceivers, the context, the time, etc. (Goffman 1959). On Goffman’s view, the purpose of self-presentation is to shape how others interpret the situation, so that they respond in ways that support one’s aims, roles, or identity in that moment.

These different types of interests are best recast as benefits that result from emotional privacy’s role in protecting expressive autonomy, which sustains the crucial ability to present the self that

¹⁴⁷ Mokrosinska (2014: 371) discusses a similar idea, citing also Murphy (1964), Merton (1964), and Altman (1975).

¹⁴⁸ See Chapter 3, §3.3.

one sees fit, regardless of the motives behind. Emotional privacy in affective communication channels is an ongoing, practice-like management of legibility and opacity. In this domain, emotional privacy safeguards those interests (hiding uncivilised inner life, enabling moral reflection and value commitments, and impression management)¹⁴⁹ not by guaranteeing secrecy, but by preserving a strategic background of partial opacity that keeps expressive meaning negotiable and answerable to context. Privacy mechanisms enact this role. Through calibrated concealment and disclosure, emoters adjust perceivers' access to affective information, exercising autonomy over their affective and social lives, and more specifically enabling individuals' self-presentation abilities.

This ultimately points to the relation between privacy and autonomy, and this relation is found at the heart of most theories of privacy. This relation, as regards to the ability to present the self that one best sees fit, is also discussed by several authors. On Fried's relational view of privacy (1968), for instance, respecting persons involves recognising their entitlement to govern access to themselves. On Reiman's account (1976), which relies on Goffman's work on self-presentation, privacy is seen as "an essential part of the complex social practice by means of which the social group recognizes and communicates to the individual that his existence is his own" (1976: 39). On Julie Cohen's view (2012), privacy "shelters" the situated boundary-work through which self-determination develops.

However, here I aim at stressing its significance within the emotional paradigm more particularly. Recontextualised to affective communication channels, self-presentation is about the maintenance of calibrated opacity that prevents others – humans or institutions – from rendering the person fully legible to others and thus affectively transparent. This preserves a person's capacity to shape how they are taken up by others across settings, retaining the ability to frame what their expressions count as *in situ*.

5.5.2. Managing social role plurality and ensuring conformity

At the interpersonal level, privacy mechanisms manage social role plurality and ensure conformity with established norms of expressiveness, often related to these roles. Most societies

¹⁴⁹ For the sake of conciseness, when referring to these different types of interests at the individual level (protecting uncivilised inner life, allowing moral reflection and value commitments, and enabling impression management), I will abridge them by saying that emotional privacy preserves "self-presentation".

prescribe specific codes of affective expressiveness that dictate how people should act and react in particular social contexts (Ekman & Friesen 1969; Matsumoto 1990; Matsumoto, Yoo, & Fontaine 2008). Such rules may call for restraint – such as suppressing irritation in professional settings – or, conversely, for expressiveness, as when service workers are expected to smile to project friendliness (Matsumoto & Kupperbusch 2001). Individuals often follow, consciously or unconsciously, the social norms of expressiveness as these rules usually sustain social practices and norms, ensuring social conformity and acceptance.

Conformity to these norms is not about hiding what one feels, but about performing a contextually appropriate self (Schwartz 1968). By maintaining partial opacity, privacy mechanisms let individuals align their emotional expressions with socially expected forms of conduct. This not only allows individuals to keep distinct role-sets (colleague, friend, client, stranger) from collapsing into one another, but it also enables the compartmentalisation and rapid role-switching without friction (Goffman 1959; Altman 1975; Stryker & Burke 2000). Relationship-specific disclosure mechanisms make this diversification actionable: who may see what, when, and in which setting (Westin 1967; Rachels 1975).

Hence it is generally accepted, and assumed, that people's behaviours, demeanour, and personality traits tend to differ from one social setting to another as people try to adjust (Goffman 1959; Matsumoto, Yoo, & Fontaine 2008). Goffman's notion of "face-work" illustrates this phenomenon by highlighting how individuals *perform* contextually appropriate selves to maintain the smooth functioning of interactions and to appear in the ways they prefer (Goffman 1967). Importantly, this entails that individuals spend lots of cognitive effort *adjusting* their EEs with the socially expected ones so as to ensure acceptance and integration (Needham, Mastracci & Mangan 2017). This emotional labour can be seen as the counterpart of affective leakages (Hochschild 1983).¹⁵⁰ Individuals present themselves differently across settings, tailoring demeanour and expression to sustain the smooth functioning of interactions (Schwartz 1968; Goffman 1959, 1967; Tcherkassof 2018). This is only possible through the strategic deployment of privacy mechanisms, which allow this role-management to exist, ensuring the preservation of social conformity within these roles.

¹⁵⁰ Hochschild's (1983) concept of emotional labour highlights how this effort often involves managing inner feelings so that outward expressions conform to institutional and cultural norms.

5.5.3. Enabling social regulation and social cohesion

Within affective communication channels, privacy mechanisms of selective disclosure and concealment serve society's interests by enabling social regulation and ensuring social cohesion (societal level). A stable and functional society requires both social regulation and social cohesion to prosper. By social regulation, I mean the ensemble of public norms, roles, and procedures that steer behaviour predictably, allocate responsibilities, and proportion sanctions so that collective goods can be produced and conflicts managed. By social cohesion I have in mind a Taylorian concept, in which there is a society-wide willingness to cooperate under disagreement, anchored in mutual trust, shared identification with common purposes, and routinised means of containing and repairing conflict (Taylor 1992). Society's interests are best served when both are jointly sustained, and I argue that at present, emotional expressions are integral to each.

As signals, they mark the salience of norms, calibrate proportional responses, and confer or withhold legitimacy on decisions and directives: indignation flags transgression and may justify censure; reassurance typically lowers social temperature and sustains compliance; remorse invites a moderated response rather than escalation (Bourgeois & Hess 2008; Hess & Fischer 2013).¹⁵¹ At the same time, shared emotional expressions underwrite social cohesion: mourning and celebration sustain common identification; apology and forgiveness rituals enable repair and re-entry after wrongdoing; (in the right cultures) everyday cues of warmth and restraint facilitate coordination and keep cooperation attractive. In short, emotional expressions do not merely reflect a society's order; they are key mechanisms through which that order is enacted and maintained.

However, the regulatory and cohesive force of EEs is conditional on how they are transmitted through affective communication channels.¹⁵² When affective life becomes fully legible, society bears a cost. Emotional mishaps are exposed, and the guiding functions of intentional expressions lose their force. As a result, regulation weakens, and social cohesion gradually erodes. Selective disclosure and concealment are the constitutive safeguards that secure the 'middle zone' society needs. Therefore, within affective communication channels, privacy mechanisms of selective disclosure and concealment serve society's interests precisely because they are constitutive of social regulation and cohesion on which healthy societies depend.

¹⁵¹ See Chapter 4, specifically §4.3.2 for more empirical work supporting these examples.

¹⁵² By this, I am referring to those socially regulated pathways that carry emotional signals across face, voice, posture, textual tone, ceremonial display, and even platform affordances.

Removal of these mechanisms would prevent EEs from performing their regulatory and cohesive work.

What these three levels of interaction highlight is how intentional emotional expressions are the important instruments of everyday coordination. Through them, people deploy privacy mechanisms to help regulate the interactive process and manage the degree of opacity in affective communication channels. This protects multiple core functions of EEs, at different interactive levels: at the individual level, it protects individuals' self-presentation, at the interpersonal level it sustains the differentiation of roles and ensures the possibility for norm conformity, and at the societal level it allows for social regulation and social cohesion. It is through this regulation, secured by emotional privacy, that the affordances needed for emotional expressions to fulfil their social functions are preserved. I will explore the potential consequences that would result from technosocial contexts in which these functions are not preserved in the next chapter.

5.6. Conclusion

This chapter has introduced and developed the concept of emotional privacy as a subconcept of informational privacy tailored to the affective domain. I argued that by protecting individuals' expressive ability, emotional privacy protects the interactive practice of emotional expression through which affective meaning is negotiated *in situ*.

I started by investigating how privacy concepts historically adapt to technological change. I showed that successive waves of innovation have repeatedly required refinement of existing privacy frameworks, prompting the introduction of new (sub)concepts where older formulations became insufficient. This established the methodological basis for treating privacy as a cluster concept open to conceptual expansion. Drawing on the case of mental privacy as an example of proactive conceptual adaptation in anticipation of emerging technologies, I defined emotional privacy as the condition under which individuals retain expressive ability over affective information, that is, when they can shape the degree of opacity in affective communication channels.

I then examined how emotional privacy functions within affective communication channels. I first identified two background pressures (i.e. expressive conventions and affective leakages) that shape how affective information becomes legible. I then introduced privacy mechanisms through which emotional privacy is enacted in practice: selective disclosure, which allows partial and context-sensitive revelation of affective states, and concealment, which regulates or redirects affective legibility. Together, these mechanisms sustain a space of calibrated opacity between full transparency and full opacity. Emotional privacy is not a barrier to social life but a condition of its intelligibility and stability. It constitutes the precondition that allows emotional expressions to perform their social functions by guiding interaction, signalling commitments, coordinating behaviour, and maintaining social order.

Finally, I showed that the value of these privacy mechanisms extends across three levels of social life. At the individual level, emotional privacy supports self-presentation by allowing individuals to shape how they appear to others. At the interpersonal level, it preserves the differentiation of social roles and expectations across contexts. At the societal level, it contributes to social regulation and cohesion, enabling affective communication to guide behaviour and maintain shared norms. The upshot is that emotional privacy is an integral part of the infrastructure of social life.

This provides the conceptual groundwork on which Chapter 6 relies on to articulate the moral conditions under which emotional privacy should be protected. It then uses these normative criteria to evaluate the prospective use of ERTs. The aim is to articulate standards by which we can judge when emotional privacy has been violated and to mark the limits that any acceptable technology must respect.