universität wien

# MASTERARBEIT / MASTER'S THESIS

Titel der Masterarbeit / Title of the Master's Thesis

„Taking Diaries and Analyzing Every Word. Why NSA and GCHQ Practices Invade our Privacy"


verfasst von / submitted by

Sebastian Sattlecker BA


angestrebter akademischer Grad / in partial fulfilment of the requirements for the degree of

Master of Arts (MA)

Wien, 2022 / Vienna 2022

| | |
|---|---|
| Studienkennzahl lt. Studienblatt / degree programme code as it appears on the student record sheet: | A 066 642 |
| Studienrichtung lt. Studienblatt / degree programme as it appears on the student record sheet: | Masterstudium Philosophy and Economics |
| Betreut von / Supervisor: | Univ.-Prof. Dr. Hans Bernhard Schmid |

## Acknowledgements

I want to thank my supervisor Univ.-Prof. Dr. Hans Bernhard Schmid, for his guidance in the process of writing this thesis, his much-valued feedback and encouragement, and his interesting courses.

I also want to express my gratitude to Dr. Justin Leduc for taking the time to sit down with me and test my understanding of the economic models I refer to in this work, and for spurring me to closely adhere to my economic mind when thinking about privacy.

I want to thank Ass.-Prof. Dr. Leonhard Menges for his guidance over the years, for countless discussions, his clear and stimulating feedback, and for taking an interest in where my research is leading me.

Thank you to those who listened to me, argued, and thought about privacy with me; who helped me clear my mind and change perspective on the matter; and sometimes just distract me from the work in moments of despair.

And thank you to my parents for enabling me to study, and to take my time in writing this thesis.

## Table of Contents

## Introduction

Concerns about privacy have been rising at least since computing technology has become more and more easily available to governments, private companies, and, finally, even to private households and individuals. While the legal debate around privacy has started way before these developments, concerns over the possibilities digital information processing provides has fueled the debate over decades. In 2013 the revelations of Edward Snowden have drawn a lot of public attention to how governing bodies in the US and UK have been putting up large scale surveillance systems, monitoring most of the world's internet traffic (MacAskill et al. 2013). More recently, the business models adopted by social media platform providers like Meta (formerly Facebook) have revealed the potential damage they can do to democratic societies and the way public discourse functions. The connected practices of collecting information about virtually all activities of individuals online in order to feed them tailored ads, news reports, posts of other users, and in that way nudging or even manipulating their behavior and beliefs raise worries that have been around in the philosophical literature for quite some time (e.g. Foucault 1995).

As governments and NGOs try to find ways of dealing with such issues, finding a clear definition of privacy is an important step in developing a clear understanding of what is going wrong.

In the first section, I am therefore going to provide an overview over some important conceptual distinctions being made in the existing literature. This will help find some orientation in the vast literature available on the topic, as well as narrow down the part of the debate that will provide the context for my arguments in the latter sections. Most importantly, the focus on informational privacy rather than spatial or decisional privacy will be adopted throughout the following sections.

In the second section, I am going to present two kinds of cases which are employed to test our intuitions regarding two opponent accounts of privacy. On the one hand, limited-access accounts are assumed to be able to deal with these cases rightly (Lundgren 2020; Macnish 2018). On the other hand, Menges (2020b) sets out to defend control accounts against the arguments made by Macnish and Lundgren, and claims understanding privacy as source control can meet the challenges of the two kinds of examples.

As I am going to argue, however, there are some challenges the source control account has to meet. Some minimal conditions of knowledge or adequate reasons have to be imposed on these accounts in order to maintain their plausibility. Using novel variations of exemplary cases

found in the literature, I am going to argue that contextual features matter a lot more than the recent discussing seems to admit.

Next, the third section presents a shift in focus. More practically connected to cases of state surveillance, I am going to argue that the kind of access relevant to privacy can be gained by machines and software just as well as by individual human agents. While Menges and Macnish share the assumption that only human access can diminish privacy, I am going to draw on theories of extended cognitive agency to argue that, in a context where machines are deeply integrated with humans to aid and undertake their tasks, machines are part of an extended agent. As such, their accessing certain information is sufficient to be speaking of access in the sense relevant to privacy. Government agencies and companies are therefore accessing personal information, even if none of their human agents ever access it directly.

Finally, in the last section I am going to assess a claim about the value of privacy as control. Namely, Menges (2020a) claims privacy as control protects the bargaining power of individuals. When bargaining power is understood as the ability to raise one's profits from trade, economic analysis clearly stands against this kind of reasoning. When information relevant to a trade is not or cannot be shared by one side of the transaction a possible consequence is less profits being realized for both parties. That is, the informed subject, not being able to share their information, cannot realize all possible profits, although source control is clearly maintained.

## 1. Conceptual Distinctions

In this part I am going to introduce some of the main conceptual problems and opposing positions held by privacy theorists, as well as some of the examples they employ for making their intuitions clear.

While in the legal debate the interest in privacy can be said to have peaked in the 1970s, with many scholars beginning to understand possible implications of data collection by companies in the form of customer records and the like, the earliest conception of privacy in its modern form dates back to the 1890 essay by attorneys Samuel D. Warren and Luis D. Brandeis, in which they argue for understanding privacy as the "right of the individual to be let alone" (ibid, p.205) and in this sense as "part of the more general right to the immunity of the person" (ibid, p.207). While this definition was consequently endorsed by a number of legal scholars (see e.g. Parent 1983, 341f.), it is far from uncontroversial to understand privacy merely as a right to being let alone.

In this section, I am going to give a brief overview of the conceptual oppositions in the debate and discuss different ways of understanding privacy. In order to not get lost in the vast array of literature while still providing a sufficiently broad picture of the relevant debates, I am going to loosely follow Inness' approach of dividing the conceptual debates into those about the *function* of privacy, its *content*, and finally debates about the *source of the inherent value* of privacy (Inness 1992, 5ff.). On a first remark on this approach it is important to keep in mind that not all conceptions of privacy attach inherent value to it, and not all approaches attribute a unique function to privacy. One such approach is directly addressed by Inness. The target of her criticism, Judith J. Thompson (1975), is proposing a conception that aims at reducing privacy to other rights such as property rights, and the right to one's person.

Because the literature is so vast, and because the three strands of the debate are sometimes closely intertwined, Inness' categorization will not be applicable to all debates about privacy. Rather, it serves to provide a rough overview over the debates around privacy, and to set the scene for the arguments made in the following sections.

### 1.1 Function

As Inness notes, the function of privacy can in general be grouped into two different conceptions. On the one hand, privacy might be understood as a kind of limiting of access of others to whatever the exact content of privacy is. On the other hand, it is also understood as a

form of control an agent exercises over whatever the exact content of privacy is (Inness 1992, 23) – e.g. over their private space, private information or their own person . This conceptual division is also at the heart of the lately revived conceptual debate on privacy, which will be introduced in the next chapter.

### 1.1.1 Restriction Based Accounts

One side of the debate around the function of privacy claims that privacy functions through separating individuals from each other. Privacy thus understood is a state in which an individual enjoys secrecy about personal information, separation from others, or limited access of others to one's personal space, personal information, one's person or whatever else the exact content of privacy might be. Important proponents of this view include Judith Thompson (1975), William Parent (1983b), Anita Allen (1988), Ruth Gavison (1980), and more recently Kevin Macnish (2018) and Björn Lundgren (2020).[1] In this understanding of privacy for example, whenever someone else gains access to personal information or my personal living space, my privacy will be diminished regardless of whether I allowed this or not. The main function of accounts in this category is thus a certain form of limitation of what others have access to – in cases where we take the content of privacy to be access to personal information, then access to personal information must be limited for a person to enjoy privacy. Whenever this limitation is suspended, the individual loses privacy.

While all of the above share the basic conviction that the function of privacy is some form of restriction of access, spelling out the function takes very different forms. First of all Daniel Solove notes that "the right to being let alone" (as defended by Warren and Brandeis) and enjoying solitude can both be seen as part of limited-access accounts, but limited-access accounts extend beyond solitude and the right to being let alone (2008, 19). Rather, in a much wider account of privacy, Gavison claims the extent to which we enjoy privacy is directly related to "the extent to which we are known to others, the extent to which others have physical access to us, and the extent to which we are the subject of others' attention" (1980, 423).

However, limited-access accounts surely face some problems. For example, focusing on restricted or limited access renders all willful acts of sharing personal information with others a loss of privacy. At least to some authors this is not intuitively plausible (e.g. Inness 1992,

[1] For a more extensive list of authors adhering to this account see for example Inness (1992, 19) or Macnish (2018, 420).

Menges 2020b). These authors claim that allowing others to access personal information can be privacy preserving if the individual has a certain form of control over who access to them or information about them. The debate between proponents of limited-access accounts and those understanding privacy as some form of control will be at the heart of section 2. Next in this section I am going to take a closer look at understanding privacy as some form of control over personal or intimate matters.

### 1.1.2 Control based accounts

On the other side of the debate around the function of privacy we find accounts that hold the relevant function to lie in some form of control over personal information, access to personal information, one's personal space, one's intimate decisions, or whatever else the content of privacy is. This position is defended for example by Inness (1992), Beate Rössler (2001), and more recently Menges (2020a; 2020b; forthcoming), as well as Mainz and Uhrenfeldt (2021a; 2021b). The latter authors hold that one enjoys privacy just in case one has the right kind of control over personal information and over another individual's gaining access to this information. In this way, sharing personal information with a friend does not necessarily diminish privacy. However, losing control will be taken as a loss of privacy under some of these accounts, even in cases where no one actually accessed the relevant information in the end.

In chapter 2 I am going to discuss recent attempts at formulating complete accounts of the control view in more detail. For now it is sufficient to note that, depending on how we define control accounts of privacy, actual access to personal information might or might not be necessary for an intrusion of privacy.

## 1.2 Content

The debate about the content of privacy is concerned with defining what kinds of things we must limit access to, or what we must have control over in order to enjoy privacy. Oftentimes individual strands of the many debates around privacy explicitly focus only on one kind of content, without assuming that this is the only possible content or necessarily the most important content of privacy.

Inness makes a number of important distinctions with regard to the content of privacy (1992, Chapter 2). For one thing, we need to discern informational privacy from decisional privacy.

While the former conception is concerned with information about ourselves including access to our voice, images of us etc, the latter is concerned with granting the individual a certain autonomy or freedom in making her own decisions. Important cases of this include, for example, the decision to use contraceptives or abort a pregnancy.

In addition to this distinction, it is important to note that when thinking of privacy in terms of limited access we can also distinguish a spatial meaning of privacy. One can access private information just as well as one can access a bedroom. However, while accessing the bedroom might be problematic because private information is gained by doing so, it might also be problematic simply in virtue of entering a space that is considered private. Personal space can therefore be considered as a distinct kind of content of privacy.

The scope of either of these contents can vary significantly on different accounts of privacy, which drastically influences their plausibility. For example, it might be that all information about, decisions of, and access to an individual is relevant to privacy, and that therefore any access to information that is about a specific person constitutes a violation of informational privacy. However, if this were so, then even learning someone's name, or seeing them walk in public space might be violating their privacy. Since this would surely have absurd consequences, it seems more likely that the scope of the realm protected by privacy is at least somewhat limited. Privacy might be limited to *intimate* decisions, or information (Inness 1992), or *personal* information (e.g. Parent 1983). In the recent debate providing the focal point of chapter 2 the content of privacy is assumed to lie in personal information only (compare Menges 2020a, 32).

Clearly, this requires us to define what we understand by *intimate* or *personal*, information, access, and decisions. Some argue that these understandings are clearly subject to cultural differences, and therefore the scope of privacy will be different in differing contexts (e.g. Roessler 2001, 33ff.). In what follows, I am not going to discuss the definitions of these classifiers in any detail.

The discussion that forms the background for my arguments in section 2 is explicitly limiting the content of privacy to personal information. As the authors in the debate, I take our everyday view of what information is *personal* to provide a sufficiently clear guideline for analyzing the examples that are going to be discussed.

### 1.3 Reductionism and Contextualism

Radically opposed to the discussion that aims at pinning down one particular function of privacy, some authors argue that privacy cannot be limited to one function, can be reduced to other, more fundamental goods and rights, or varies significantly across various social and cultural contexts. Importantly, the above distinctions are not always relevant to this debate, since, for example, the context dependence of privacy might concern function, content and value of privacy equally.

For example, while many control accounts take the relevant content of privacy to be personal information (compare Solove 2008, 24 f.), authors like Judith DeCew include matters of accessibility and personal expression into their account (1997, 76 f.). However, DeCew does not seem to be holding control as the only function relevant for privacy. Rather, depending on what is at stake, DeCew claims that different functions of privacy can be observed. When personal information is at stake, control is the relevant function; when (physical) accessibility is at stake, limitation of access is the relevant function; and finally, when expression of one's person is at stake, control over intimate decisions is the relevant function (DeCew 1986, 75 ff.).

We therefore have to consider a third category of accounts of the function of privacy here, which does not take privacy to have a unique function or a unique kind of content at all.

When looking at accounts that do not ascribe a unique function to privacy we can distinguish between two kinds. On the one hand, accounts like that of DeCew (1997) take the function of privacy to vary, depending on what is at stake. Similarly, Solove (2008 Chapter 6) and Nissenbaum (2010 Part III) take privacy's function to vary depending on context. That is, depending on what is at stake, who is involved, and which method of agency are involved, privacy has different functions. Nissenbaum, for example, claims that privacy can neither be determined as a right to secrecy nor as a right to control of personal information, but as a "right to *appropriate* flow of information" (2010, 127). How appropriate flow can be understood then depends on the context in which the flow takes place.

On the other hand, authors such a Thompson (1975) claim that the reason privacy does not have a unique function is that ultimately privacy can be reduced to other rights, and therefore the relevant functions must be found by looking at how those rights function. For her, the right to privacy is not a distinct right at all, but can be reduced to property rights and rights over the person (1975, 306). Reducing the function of privacy to that of other rights also implies that the content of privacy will also be reduced to the content of these other rights.

Even though they both deny privacy having a unique function, these two kinds of

accounts have very different implications. Contextualist approaches merely claim that the function (as well as content and value) might vary significantly in different contexts. However, proponents of these accounts often accept that privacy is still a concept of its own that is worth defining. Reductionist approaches, on the other hand, not only hold that the right to privacy can be reduced to other rights, but they deny those rights share any common foundation (Scanlon 1975, 315). Therefore, reductionist approaches urge us to focus our attention on these other rights rather than privacy.

Now, let me briefly lay out and answer to some of the reasons for holding a skeptical view regarding privacy that one can make out in the existing literature. For example, Thompson claims that "it is possible to explain in the case of every right in the [right to privacy] cluster how come we have it without ever once mentioning the right to privacy" (1975, 313). That is to say, in all kinds of cases we can imagine, where we intuitively think the right to privacy is violated, we can actually find another right that is more fundamental than privacy, and which can explain the rights violation occurring in that case. Whenever we want to explain why privacy has been violated in a specific case, the explanation will refer to some other right. As Inness notes, in Thompson's view the right to privacy is reducible to the two clusters of property rights, and rights over the person (1992, 29).

One problem for such a view is that it is possible to find examples in which property rights and respectively rights over the person come apart from privacy rights. Inness argues that property claims come apart from privacy claims in cases where someone violates a property claim, but the object and nature of the violation do not concern privacy in any way. For example, a friend taking my pen without asking violates my property right over the pen without violating any right to privacy (Inness 1992, 33f)[2]. Moreover, rights over the person come apart from privacy rights in cases where another person is not actually gaining access to the person, but there merely exists a threat of access occurring. For example, consider a case where a peeping Tom attempts to take a glance at someone through their closed but slightly transparent curtains, and the person successfully hides at the last moment (ibid. 1992, 34). In this case, since no access to the image of their body has been gained, no right over the person has been violated, while Inness claims that privacy has been violated since the person had to take what she calls 'emergency maneuvers' (1992, 51). As Menges argues, however, this argument presupposes a

[2] A similar thought is articulated by Marmor in his example where a dry cleaner lends your jacket to another customer (2015, 6). While this action violates your property rights, it does not violate your privacy. He also focuses on the fact that a lack of permission to use someone's property does not necessarily constitute a violation of the (in his case) *right* to privacy.

specific understanding of control, where control is essentially a choice among different options (Menges forthcoming, 5ff). As we are going to see later, this view of control has some flaws of its own.

Another line of argument against a reductionist understanding of privacy is formulated by T.A. Scanlon, who argues that just because property rights and rights over the person might be protecting the same interests as privacy rights, it does not follow that the former are more fundamental than the latter. In order for the reduction of the right to privacy to other rights to be truly helpful in an investigation of privacy, it would be necessary that the interests at the basis of these other rights were different from each other. However, to him "it seems more likely […] that rights of ownership and rights of the person as Thomson understands them will turn out to be based in part on [the interests underlying] norms of privacy" (Scanlon 1975, 322).

Finally, an important reason for not drawing this conclusion is of a practical nature. Thomson attempts to simplify the discussion about privacy by reducing it to other rights. However, a unified concept of privacy would spare us the need to consult other rights whenever we are unsure about some relevant privacy norm. At the very least it is not clear that a reduction of the right to privacy to other rights is actually making things easier.

With this in mind, I want to highlight another distinction that has become important in the recent literature. Namely that between the *right* to privacy hand, and the *concept* of privacy.

### 1.4 The Right to Privacy

As some scholars have pointed out, there is yet another distinction sometimes overlooked in the literature. Naturally, many legal scholars will be concerned with discussing the *right to privacy* and its implications. On the other hand, however, the recent philosophical discussion has tried to avoid referring to the right to privacy and rather restricted the discussion to the object of this right, i.e. the *concept of privacy* (Lundgren 2020; Menges 2020a)[3]. This shifts the focus to discussing a state (i.e. privacy) that individuals can be in, as opposed to a right to being in this state they might or might not possess. Surely, privacy is the object of a right to privacy. However, focusing on the concept allows us to avoid confusing normative with descriptive arguments in the discussion.

Ultimately a full discussion of the topic would of course involve some clarification of whether we should have a right to be in a state of privacy, what this right involves and how it

[3] In doing so they follow earlier scholars like Parent (1983a), (1988), or more recently Nissenbaum (2010).

can be secured. However, clarifying these questions sure requires us to first get a clear idea of what it is that we have a right to when talking of a right to privacy. To do so, authors in the recent discussion usually focus on conceptual questions rather than axiological or normative ones (e.g. Menges 2020b, 368).

Only after having settled the conceptual question to a reasonable degree can we meaningfully start to ask ourselves whether we should have a right to privacy, or whether enjoying privacy is valuable. What we should avoid, according to the recent debate I am referring to here, is to start defining privacy by appealing to arguments involving normative claims. Whether or not privacy is valuable can only be answered after we have defined it, and we should not ex ante hold the requirement that privacy (however we define it) be valuable.

However, of course privacy and the right to privacy are closely related and often claims about the understanding of the concept of privacy will have implications for our understanding of the right to privacy. For example Lundgren claims that "if the *right to privacy* should be conceptualized as a control account of privacy, then so must the concept of privacy and vice versa" (2020, 167 fn.3).[4]

Importantly in this discussion the focus is not on why we would or would not want to have privacy , but on getting a clear idea of what our intuitions concerning privacy are. Based on these intuitions the authors try to conceptualize a coherent account of privacy.

A central feature of such an account of privacy is its extensional adequacy, i.e. whether it enables us to identify just and all of the cases we intuitively take to exemplify a violation of privacy as such violations (compare Solove 2008, chap. 2; Menges forthcoming, 11ff). In the next section I am therefore going to present some of the most important cases from the literature analyzing the concept of privacy, and present the recent arguments around those cases. Using these seminal examples I am going to focus on laying out what proponents of limited-access accounts take to be the most pressing problems for control accounts of privacy. After this, I am going to consider variations of these exemplary cases in order to point out some pressing problems for the proposed understanding of control.

[4] However, as Mainz and Uhrenfeldt point out, it is not made clear whether Lundgren here takes defining the right to privacy in terms of control to mean the *right* itself being defined in terms of control, or takes it as a right over *privacy* where privacy is defined in terms of control (Mainz and Uhrenfeldt 2021b, 6).

## 2. The Concept of Privacy

In this section I am going to focus on the current debate around defining the *concept* of privacy and present some of the exemplary cases employed to test the intuitive plausibility of limited-access accounts and control accounts of privacy.

An important restriction of the content of privacy I am adopting here is to focus on *personal information* only. In this I am following Macnish (2018), Menges (2020a; 2020b), and Lundgren (2020), who restrict their analysis to personal information throughout the works mentioned here.[5]

Moreover, as mentioned above, the arguments are not aimed at answering normative questions like whether we should have a right to privacy, whether privacy is valuable and what this value might be grounded in. Rather, the arguments in the recent discussion I am focusing on are assumed to be neutral towards axiological questions (Menges 2020b, 2), or explicitly aimed at establishing the neutrality of the concept of privacy. For example, against Inness' claim that privacy is indeed a valuable concept (1992, 44), Macnish holds that there are cases in which giving up privacy is valuable to individuals, and privacy itself is not inherently valuable (2018, 421). On a first note, we can hold against this argument that just because giving up privacy can mean gaining something valuable there is no strong reason for us to believe that privacy is not valuable too. For example, I can give up a relationship with a good friend in order to find more time for other relationships, but just because I gain something valuable by finding more time for other relationships does not mean I do not also lose something valuable by giving up the friendship. Similarly, if establishing intimacy means giving up some amount of privacy, then the mere fact that we gain something valuable by doing so does not mean we do not also lose something valuable by giving up privacy. Goods can be weighed against each other, and it might just be that sometimes intimacy is more valuable than privacy. Since establishing the value of privacy would be too big of an issue for the scope of this work, I am very much sympathetic to Menges' approach of at least initially leaving aside axiological questions.

In this section I am going to present two standard cases of the existing literature, and conceptual arguments that flow from these examples. On the one hand there are *voluntary divulgence cases*, and on the other hand there are *threatened loss cases*. These cases are central to the debate between the defendants of limited-access accounts and control accounts of

[5] However, that does not mean that all the arguments depend on assuming this restriction. Compare Lundgren on this issue: "I will follow [the limitation to informational privacy] even if none of my arguments depend on restricting the discussion to concern only information" (2021, 671).

privacy, and capturing our intuitive responses to these cases is assumed to be a minimum requirement for any account of privacy. While Lundgren and Macnish argue that control accounts are not able to accommodate these intuitive responses, Menges develops a novel account of privacy in terms of source control, which he argues can make sense of our intuitions in both kinds of cases.

While the source control account can initially answer the worries articulated by Macnish and Lundgren, I am going to argue that it is far from clear what source control is. I am going to develop variations of the voluntary divulgence case that force proponents of source control accounts to clarify what the exact criteria for having source control are. My focus will be to highlight pressing problems for the most obvious way of defining source control.

### 2.1 Voluntary Divulgence Cases

As mentioned above, when investigating accounts of the concept of privacy, a standard procedure is to test the account in question against the intuitions we have about certain kinds of cases. While some of the cases we encounter in the literature are based on real world examples, it is not a necessary criterion for the cases we employ in our reasoning to be based on real events. Rather, it is sufficient for the cases to be similar to cases one might encounter in the real world. Their purpose is to test our intuitive reactions concerning privacy in situations one might plausibly expect to encounter in the real world. While I am not going to elaborate on what conditions must be fulfilled for an exemplary case to be plausible in this way, I am assuming that the reader and her experience of the social world we experience every day has equipped her with the capability to decide for herself whether the fictional cases employed in the following are plausible and sufficiently similar to real world cases.

Let me now turn to introducing the first type of case which I am going to investigate in this section, the so-called *voluntary divulgence* cases. A first version of these cases has been introduced by Parent and has since been referred to, sometimes in slightly modified form. The general idea behind voluntary divulgence cases is a situation in which an individual "freely divulges everything, no matter how intimate the fact, about himself" (Parent 1983, 344), in a state where they can be assumed to be in control of what they are doing. So while control over the information that is shared is supposedly given, Parent (1983) and Lundgren (2020) claim the privacy of the individual is nonetheless violated in these cases. That is to say, voluntary divulgence cases are supposed to show that control accounts are too broad, i.e. they compel us to take cases where privacy is actually lost as cases where it is not.

Parent notes that voluntary divulgence should be taken as at least sometimes diminishing privacy, since "any adequate definition of privacy must allow for the possibility that persons can exhibit a lack of respect for their own" (1983b, 344). In this sense, an individual can exhibit some sort of control over her personal information, but in doing so might be diminishing her own privacy, if she is giving away her personal information all too readily without proper consideration or simply for the wrong reasons. "Concern for privacy may dictate the protecting of people against their own consent [or other forms of control they can exercise]" (McCloskey 1980, 23).

Note that this line of argument now can be taken to flow from axiological assumptions rather than only conceptual ones. Parent and McCloskey claim that because individuals sometimes act against their own best interest in cases where they have control, we should not take control to be a sufficient condition for maintaining privacy (Parent 1983b, 344). Since voluntary divulgence is one way in which they can act against their own best interest, an adequate account of privacy as control should take those cases as cases in which privacy might be diminished. However, since control accounts (as Parent and McCloskey understand them) take voluntary divulgence as compatible with privacy, even when voluntary divulgence has bad consequences, control accounts are claimed to be wrong.

Clearly, it is true that sometimes people fail to act in their own best interest and act in ways that harm them and the goals or desires they have. More relevantly for the conceptual debate, individuals can clearly fail to maintain privacy with their voluntary actions. These two claims, however, must be kept apart. Maintaining privacy can mean a person is not acting in their own best interest, just as one can lose privacy and still benefit from the consequences of this loss.[6]

The conceptual debate I am focusing on at this point does not evaluate the consequences of voluntarily sharing information, but rather is only concerned with whether voluntarily sharing information can diminish privacy.

I take it that sharing is *voluntary* if and only if the individual has full control over her actions. Here, my understanding of voluntariness departs from at least that of Menges, who claims that even in cases in which individuals reveal facts under the influence of drugs or similar conditions, they can nonetheless act voluntarily (forthcoming, 20). I adopt a stronger notion of

[6] An interesting question still is whether one can deliberately give up privacy in Menges' understanding. Since deliberation seems to ensure a sufficient degree of (source) control, this does not initially seem to be the case. However, restricting the source control account in a meaningful way might open up this possibility and thereby accommodate the worry expressed by Parent. I.e. if only certain desires are seen as privacy protecting sources of a flow of information, one might lose source control, but still have important other desires satisfied as a consequence of such a flow.

voluntariness, which implies what Menges calls "full control".

At this point it is important to highlight that there are radically opposed intuitions about voluntary divulgence in the literature. While for authors such as Macnish (2018), voluntarily sharing personal information is always diminishing privacy, Menges claims that this intuition is not sufficiently clear, and the question of whether privacy is diminished by sharing personal information remains to be solved (Menges 2020b, 389).[7] I share Menges' intuition that voluntarily sharing personal information need not necessarily diminish privacy. However, I take it that even under full control, individuals can diminish their privacy. That is, I also accept Parent's condition that even under full control an individual "can exhibit a lack of respect for their own [privacy]" (1983b, 344).

The following passages present an attempt to resolve this conflict of intuitions. By arguing that even full control over ones actions only protects privacy in some cases, I will try to do justice to both kinds of intuitions about voluntary divulgence. My conclusion will be that full control over sharing information only protects privacy if adequate concern was given to privacy related issues and reasons in the decision. For example, sharing personal information for purely economic reasons will diminish privacy, even if the individual had full control over doing so.

To assess whether control accounts can make sense of the two seemingly conflicting intuitions, I am going to compare two notions of control that are prominently featured in the recent debate around privacy. On the one hand there is the standard understanding of *leeway* control, which holds control being given if there are alternative courses of action available to the individual. On the other hand, there is Menges' novel account of *source control*, which is intended to avoid some of the problems leeway control faces with respect to the cases presented in the next section. In Menges' understanding, both of these accounts "can be defended against the objection from Voluntary Divulgence Cases" (forthcoming, 23). In the following I am going to assess this claim.

---

[7] There is some discussion about whether philosophers are strongly relying on intuitions towards cases like the ones we encounter in the existing literature on privacy, and what role such intuitions should be playing in philosophical investigations (for a general overview see Pust 2019). Some philosophers have contested that intuitions really play a central role in contemporary analytic philosophy (Cappelen 2012), while others have defended the claim (Chalmers 2014). While I am not going into this discussion here, I take the side of those who argue that philosophers rely on intuitions, and that this can be perfectly fine (e.g. Goldman 2015).

Moreover, since evidence suggests trained philosopher's intuitions about philosophically interesting cases are not significantly different from those of laypersons (Knobe and Nichols 2017, sec. 3.2), and that therefore employing methods of experimental philosophy might be useful in deciding between these conflicting intuitions, or at least showing why they arise. This can be helpful in resolving some of the controversies in the debate between advocates of privacy as control and privacy as limited access.

Assessing the plausibility of this claim requires us to pay close attention to what kind of control individuals can be exhibiting in in cases of voluntary divulgence.

On the one hand, leeway control accounts assume control is given if an individual could have just as well chosen not to act in the way they have. That is to say, they had a significant amount of choice and could have acted otherwise.

A general problem with such an understanding of control is, however, that assessing whether a person has meaningful control over an action intuitively requires at least some kind of explanation over and above that fact that they could have acted otherwise. It seems plausible to assume that even if several options are available, sometimes things go wrong on the side of the individual. That is they might not understand the type of situation they are in and what their counterpart is capable of capturing (compare Menges 2020b, 389), they might be completely drunk, under the influence of some other drugs (compare Menges 2020b, 388), be in an emotionally upset state and therefore not acting clear-headedly, they have some kind of mental disorder, and so on. The point here is that many of these possible assumptions listed here would question that the person acting is in a state of control, and whether having the choice to act otherwise is really the account of control we should adopt here[8]. Proponents of control accounts can accept the intuition that in many such cases individuals diminish their own privacy because their conditions might call into question their ability to act otherwise. Conversely, however, there are cases in which agents could have acted otherwise, but it is still questionable whether they maintain privacy in doing so. For example, when the options available are not covering a meaningful spectrum of alternatives, having a choice might be irrelevant to maintaining privacy.

Let me briefly recapture what this means for the claim that leeway control can answer to the voluntary divulgence objection. Answering this objection requires the relevant account of control to ensure that fully exhibiting leeway control in an act of sharing personal information with others does not diminish one's privacy.

One specification of leeway control that might avoid this problem is the following. As

[8] Most of the examples listed here seem to be compatible with the individual having had some alternative. However, the truth drug example Menges (2020b, 388) gives does not grant the individual any kind of leeway control. Still, even in this example, privacy might be preserved if the individual chose to take this drug knowing it will lead her to freely divulge personal information. Again, the intuition that truth drugs cannot be compatible with privacy as control seems to come from the intuition that it might have bad consequences to take truth drugs. However, this is not necessarily so, and it seems to me the intuition rests on axiological assumptions, which Menges wants to avoid.

Menges mentions[9] the options available in a certain situation must be such that they guarantee the individual's autonomy (forthcoming, 17). However, autonomy includes many more dimensions than just privacy, and it is far from clear that individual autonomy is only guaranteed when an individual shows sufficient care for her own privacy. Conversely, assume the options available to an individual guarantee leeway control in an autonomy ensuring way. Still, if the individual shows a lack of care for her own privacy, she might nonetheless diminish her privacy. That is to say, the kind of autonomy that needs to be secured in order to secure the kind of leeway control sufficient for maintaining privacy is not the only possible kind of autonomy an individual might enjoy. Therefore, leeway control can be given and still privacy might be diminished.

However, if the right kind of autonomy over personal information is secured, then the individual might enjoy privacy even when voluntarily divulging information. Still, what exactly the right kind of autonomy is needs to be further developed.

Privacy as leeway control can therefore make sense of the idea that privacy need not be diminished when individuals voluntarily divulge personal information, and at the same time individuals might exhibit (full) leeway control and still diminish their privacy by doing so. This is so because only a certain kind of leeway control ensures privacy, namely leeway control that ensures the agents autonomy with respect to personal information. If the agent does not care about privacy at all, she might enjoy some other kind of leeway control even if acting with this control diminishes her privacy. Leeway control in voluntarily divulging information therefore only sometimes ensures privacy.

Defining the leeway control account such that the relevant kind of leeway control always ensures privacy will involve finding out which kind autonomy over what kinds of matter the individual must possess to ensure privacy. Doing so would, however, exceed the scope of this work. For now it is sufficient to see that by allowing only some kinds of leeway control to be ensuring privacy, we can make sense of why sometimes even full control is not sufficient to maintain privacy, while still sometimes control ensures privacy.

On the other hand there is the account of source control, which holds that only an individual's having source control over her personal information ensures privacy. Since this account is rather novel, I am going to at this point introduce the source control account as Menges understands it.

In formulating his account of source control Menges draws on an argument originating from

[9] Note that Menges calls leeway control "control as choice" in the cited work.

the debate around the concept of responsibility. In this debate, the conception of leeway control as a necessary condition for responsibility is called into question by so called Frankfurt cases (compare McKenna and Coates 2020, sec. 3.2) – named after Harry Frankfurt, who first introduced this kind of case and the argument around it (1969). Very briefly, in Frankfurt cases an individual B has installed some mechanism that intervenes if person A chooses to take action 1 rather than 2 and forces the person to take action 2 instead. For example, suppose I am in a situation where I can decide to either harm a person or refrain from doing so. If I choose not to harm the person, a chip that B has implanted in my brain sends a signal and makes me harm the person anyway. In this case it is hard to see how I can be said to be in control of what I am doing, and it is not plausible that I should be held responsible for harming the person. One conclusion from this could be that the possibility to act otherwise, i.e. leeway control, is a necessary condition for being held responsible. Analogously, not having the possibility to act otherwise in a voluntary divulgence case would be necessarily diminishing privacy.

Continuing with the example, however, in case I decide to harm the person, the fact that no alternative course of action was open to me does not imply that I am not responsible for harming the other person. Rather, I can be said to nonetheless have been exhibiting an important sense of control and therefore I can be held responsible. The conclusion Frankfurt draws from this is that I am responsible just in case the decision to act in a certain way, and the consecutive action were connected to my deepest wishes and desires, or whatever else is considered to constitute my true self, in the right way (see McKenna and Coates 2020).

Building upon this argument, Menges concludes that "an agent's having privacy with regard to a piece of information consists in the agent's being such that if the information flows to others, then the agent is the right kind of source of this information flow" (2020b, 393). Where *being the right kind of source* means something like the individual's choice being grounded in her desires, deepest wishes, or simply being "grounded in the agent and not someone or something else" (Menges forthcoming, 9). Throughout his arguments, for the sake of simplicity, Menges adopts a Frankfurtian understanding that being the right kind of source of an information flow, where that means having the right kind of first- and second-order desires about a particular flow of information occurring. That is, an individual has to desire the flow to happen, and have the desire to desire the flow to happen. I am going to loosely stick to this, but mostly just talk of desires in order to not complicate the formulations in the following sections too much.

The account of privacy as source control appreciates that, if personal information flows to other individuals by my own choice, privacy is diminished if the flow is not based on the right

wishes and desires. If no information flows, privacy is maintained in any case. However, my claim is that privacy might also be diminished if the relevant wishes and desires do not feature any privacy related reasons for my acting the way I do. That is, in order to account for the intuition that privacy might be diminished even when an individual has full control, it is important to see that having desires that are connected to my action in the right way is not sufficient. Rather, the desires themselves must exhibit some relevant concern for privacy-related issues. That is to say, having source control is a broader notion than having the right kind of source control to ensure privacy.

The account of privacy as source control can make sense of the fact that individuals can voluntarily divulge personal information and not thereby diminish their privacy. That is so because having source control does not initially seem to guarantee the individual also enjoys privacy, since this kind of control does not necessarily feature a regard for privacy or privacy-related wishes and desires. If an individual shares information for purely economic reasons, they might exhibit source control, but might nonetheless be exhibiting a lack of regard for their own privacy. Likewise, if they are ignorant of some fact that might be considered relevant with regard to privacy, simply having desired to act in the way she does will not plausibly ensure privacy.

For example, one can buy shoes on a digital shopping platform as a consequence of desires that ensure source control. These desires can be purely economically motivated and not be based on any considerations regarding private information, digital surveillance and so on. Assume the individual acts only on economic reasons because she does not know the company is collecting information about her, tracking her use of the site and so on. That also means that changing any facts about being tracked on the website, having customer information collected and so on would not change the behavior of the individual. My intuition towards such a case is that the kind of control the individual exerts over her personal information is not enough to ensure privacy.

Coming back to the analogy with the theory of responsibility and free will, we can note that in general lack of knowledge about the situation one is in might exempt an individual from being held responsible (compare Shoemaker 2015). Even if the individual can be attributed the right kind of qualities to be seen as a responsible agent in such cases, the fact that they were lacking important information about the situation exempts them from being held responsible.

Likewise, and individual might generally have the right kind of control over what they are doing, but because they are lacking certain information, we cannot say that they are maintaining privacy by acting the way they do.

Again, I take it to be quite plausible to allow the possibility of full control not ensuring privacy for a general account of source control. However, since we want to have an account of control that necessarily ensures privacy, it will be important to spell out more clearly, which kinds of wishes and desires are to be seen as possible sources for a privacy ensuring kind of source control, and which other conditions should be fulfilled.

That is, the relevant kind of source control cannot be assumed to arise from just any kinds of deep wishes and desires. Otherwise, even cases that are intuitively diminishing privacy would be compatible with maintaining privacy, if only individuals have some deep wishes and desires about this flow of information occurring. I am going to address this worry again in later sections of this chapter. As we are going to see, clarification and restriction of which kinds of desires establish source control is necessary to maintain its plausibility. I take it that only a specific kind of source control will maintain privacy when personal information flows.

Now let me briefly consider a possible answer to the above argument. Proponents of source control or leeway control accounts of privacy might argue that a lack of regard for privacy will result in an individual having less than full control, and therefore the whole argument just would collapse. Likewise, the kind of control that does not ensure autonomy over personal information, they could argue, is simply not full control in the relevant sense.

As an answer to this, first, it is important to see that privacy or considerations relevant to privacy are only one kind of reason individuals might have for acting the way they do, and especially for sharing information. They might choose to ignore privacy considerations when doing so, and still have a multitude of other highly important reasons for sharing their personal information. An account of *full* control, however, that requires individuals to not be disregarding any of the possible dimensions of a certain choice is clearly asking too much of individual agents. We can have full control of what we are doing and still be disregarding certain reasons for acting.

Of course, it matters whether this lack of regard results from a more or less deliberate choice, or from ignorance, mental conditions, or drug use for example. Still, in some cases individuals can be perfectly aware certain reasons exist, and their mental state would allow them to deliberate about these reasons, but they choose not to. That is, as mentioned above, control accounts should be leaving open the possibility that people disregard certain reasons to act while still being taken to have full control when doing so.

For example, when an individual shares personal information for solely economic reasons, we do not usually have reason to doubt they can be in full control of doing so, but rather we have reason to doubt they are considering their own privacy at all. The kinds of desires and

beliefs source control can be based on are not limited to those that ensure privacy, but rather I take it that privacy will only be assured by a very specific kind of source control that flows from the right kinds of beliefs and desires. Quite plausibly, these beliefs and desires have to be about privacy or closely connected to privacy-relevant considerations.

Spelling out which kinds of considerations are relevant to privacy, and which kind of autonomy must be ensured to maintain privacy, is clearly up to proponents of these accounts. Defining them will, however, require us to consider at least some contextual features of the situation in which information flows, and therefore push the debate in the direction of Helen Nissenbaum's (2010) conception of *contextual integrity*.

Similarly, the relevant kind of autonomy for ensuring privacy under a leeway account of control could be seen as closely connected to or directly based on contextual integrity with respect to personal information. In order to assess whether the relevant kind of autonomy is maintained clearly requires us to consider a multitude of facts about the situation the individual is in, her own role, goals etc. Nissenbaum's account provides a lot of useful insights on this and would provide an obvious lead in going about this task.

In conclusion, while control accounts are able to answer the objection raised by voluntary divulgence cases, specifying which kind of control ensures privacy is crucial to defending both leeway and source control accounts. Importantly, the relevant kind of control must be either flowing from the right kinds of desires and wishes of the individual on a source control account, or ensuring the right kind of (informational) autonomy on a leeway control account. Individuals can be in *full* control of a certain action without these conditions being fulfilled, and therefore even full control does not necessarily ensure privacy. That is, full control is clearly necessary for privacy being maintained, but since full control can be based on reasons completely irrelevant to privacy, only a specific kind of full control will be relevant for answering the question whether and individual enjoys privacy or not.

With this being said, in the next part, I am going to turn to the second kind of standard case from the literature.

### 2.2 Threatened Loss Cases

The second type of cases often referred to in testing our intuitions about privacy are so called *threatened loss* cases. In these cases, effective control over a certain piece of information is lost, while no information flows to the person controlling the information. Intuitively, if

information does not flow, informational privacy cannot be diminished. However, control accounts defining control as leeway control imply a loss of privacy in these cases. In other words, threatened loss cases show "that [leeway] control accounts are too broad" (Lundgren 2020, 170).

In this section, I am going to lay out why threatened loss cases pose a challenge to understanding privacy as control, as well as explain how privacy as source control can deal with such cases. Finally, I am going to present some novel variations of threatened loss cases that expose some of the challenges the source control account faces.

To get a clearer idea of what is going on in threatened loss cases let me start by considering the following example:

> "[…] Imagine that I leave my diary on a table in a coffee shop and return to that shop 30 minutes later to retrieve it. When I enter the shop I see a stranger with my diary on her table […] I panic for a moment, but on seeing me the stranger smiles and hands me the book. She explains that she has not opened it, but saw me leave without it and collected it to await my return. She knows how intimate her own diary is, so she respected my privacy and kept it shut, as well as making sure that no one else would be able to read it" (Macnish 2018, 420)

In this case, Macnish acknowledges that we might have strong feelings about the loss of control that occurred, but we should not take privacy to be violated. Control accounts are wrong in taking privacy to be violated in cases where (leeway) control over private information is lost, as long as the information is not actually accessed by another individual.

Another example of this kind can be found in the following passage:

> "[A] peeping Tom looks in a window, but does not see the intended object because she has concealed herself under the bed to avoid his gaze" (Inness 1992, 50).

While Inness claims that in this case privacy has been violated because the individual had to take "emergency measures" (1992, 51) to avoid being seen, authors like Macnish, Menges and Lundgren would claim that in such a case privacy has not actually been violated, since no information flow has occurred.

A final example of threatened loss cases has been introduced by Thomson:

> "If my neighbor invents an X-ray device which enables him to look through walls, then I should imagine I thereby lose control over who can look at me: going home and closing the doors no longer suffices to prevent others from doing so. But my right to privacy is not violated until my neighbor actually does train the device on the wall of my house. It is the actual looking that violates it, not the acquisition of power to look" (1975, n. 1).

Again, we have a strong intuition that what the neighbor is doing is wrong. However, so the argument goes, as long as the neighbor doesn't use the device, no information flows and privacy is not violated.

This intuition poses a problem for control accounts like that of Inness (1992), since those accounts are based on the idea of leeway control or control as choice rather than source control. Essentially, the worry is that while effective choice over what happens to the information in the diary or who gets to see me naked in my home at what times is gone, I still enjoy privacy contrary to what this version of the control account would imply.

The source control account clearly fares much better in correctly addressing cases like the diary case, since it too takes information flow to be a necessary condition for privacy loss. Since information does not actually flow in any of the cases presented here, the source control account rightly concludes that privacy is not diminished.

The dilemma Lundgren (2020) sees for control accounts seems to be resolved at least for the moment.

However, in the next sections, I am going to raise some critical objections to the source control accounts, that show what needs to be clarified in order to make the source control account more convincing.

### 2.3 Source Control and its Challenges

On a short note, let me recall that the source control account provides a clear response to the worries raised by voluntary divulgence cases and threatened loss cases. As such it deals with the intuitive responses in the way demanded by Lundgren (2020) and Macnish (2018), who both argue that (leeway) control accounts fail to capture our intuitive responses to these cases and are therefore implausible.

While I believe the source control account can be successful in addressing the worries expressed in the existing literature, I believe that there are cases that challenge the plausibility

of a conception of privacy as source control. In the next two sections I am going to present such challenging cases and lay out their implications for privacy as source control. After that, in section 2.4 I am going to present some worries for defining the kind of sourcehood relevant for source control in terms of higher-order desires. This definition, although presenting only one possible way of spelling out source control more explicitly, is based on Frankfurt's compatibilist theory. For this reason, it initially seems to be the most intuitive way to go in defining source control. However, as we are going to see, defining source control in terms of higher-order desires faces some serious problems.

### 2.3.1 Underspecified Desires

In order to spell out the source control account more clearly, Menges refers to a variation of the standard cases, involving a couple revealing intimate facts about their relationship (2020b, 394; originally Lundgren 2020, 171). In the case that is interesting for us here, the couple has the desires to reveal intimate facts, but forgets to turn on the web-camera that is supposed to get the flow of information going. However, some software bug turns on the camera nonetheless and the information flows. In this case the source control is supposed to be violated, since the couple's desires are not causal for the information flow. As Menges himself admits, however, this intuition is not totally clear (2020, 394), and in the following I am going to clarify what is going on in the example and some similar but slightly different cases.

One reason for the intuition that privacy might not be violated is that when accidentally revealing personal information one wanted to share anyway one does not lose anything valuable. However, as we have seen above, the question of whether something valuable is lost should be separated from the question of whether privacy has been lost (Menges 2020b, 394 fn. 4). I.e. axiological arguments are not relevant at this point in the discussion.

Still, there is an interesting case to be made from this example, that might help spell out the source control account more explicitly. Specifically, the following case will call into question whether the problem in the web-camera example is merely about the couple not having the right kind of desire. As I am going to argue, in the context of the example it is simply not possible for the couple to hold a desire that would make the flow of information maintain privacy, even if the couple's actions were to be seen as distantly causal for the flow occurring. That is so because the desire we would have to ascribe to the couple would be utterly underspecified, and therefore source control cannot be grounded on that desire. Let me first turn to the variation of the web-camera example.

We can imagine that the couple knows the bug sometimes turns on their camera. While they forget to turn it on in the very moment they desire to be filmed, they are aware that they might be filmed at any time anyway, know what being filmed by a web-camera implies, and have the desire that this be so.[10]

In this case it seems the general desire to be filmed at any time, and the desire that this be so, would compel proponents of the source control account to hold that privacy is not diminished in case the camera is turned on and personal information flows. After all, the couple could have gotten rid of the camera that films at random times, but chose not to, i.e. their agency is distantly causal for their being filmed, and they have the relevant desires.[11]

This conclusion, however, seems implausible. Not knowing the specific content of what is being recorded and ultimately transmitted to another (and possibly multiple other) individual(s) is clearly problematic for the claim that nonetheless privacy is being maintained in moments where information flows.

Moreover, the utterly unspecified nature of the desires we would ascribe to the couple is problematic for the source control accounts in two ways.

First, it is not clear whether unspecified desires, such as the desire to be filmed at any time, can be ascribed to individuals at all. That is, because it is not clear when this desire will be fulfilled – after being filmed one time, being filmed multiple times, after being filmed in a specific interval that can be characterized as random, etc. – it is hard to see how an individual can hold such a desire. Moreover, in the web-camera setting the number of possible flows of information is incredibly vast. Holding an unspecified desire about such a situation implies a vast number of beliefs about how this desire might actually be realized. It seems rather implausible that the couple in the example would be able to anticipate more than a small fraction of the possible routes their personal information might take. It is not that this might be problematic for Menges' claim that in the initial web-camera example the couple has source control. However, in the software bug example, the problem gets amplified by the couple's lack of knowledge about the time they are being filmed.

Second, and closely connected to the first point, one can easily see how there will be some moments in which the couple desires not to be filmed. When they are the ones turning on the camera, they control the time frame in which they act as they desire in front of the active camera.

[10] As Menges does, I am assuming source control to be defined in terms of first and second order desires for the moment (2020b, 393).

[11] If the couple were not in some sense causal for the filming occurring, source control accounts would rightly conclude that the couple's privacy is diminished. (Many thanks to Leonhard Menges for pointing this out to me.)

If, however, the filming occurs at random times, the couple cannot act in a way that they do not desire to be filmed by the web-camera. That is, they will be constantly avoiding actions that others should not see, that would evoke uncomfortable reactions by others, and perhaps seeking actions that let them appear favorably in the eyes of others. It then becomes dubitable whether the couple's desires are still the source of their behavior, or whether expectation of possibly being watched at any time leads them to act differently from what their desires would initially dictate. Perhaps they might even develop different desires as a consequence of the constant possibility of being watched. As Nissenbaum notes with reference to Bentham (1787, 2010) and Foucault (1995), "when we have internalized the gaze of the watchers and see our- selves through their eyes, we are acting according to their principles and not ones that are truly our own" (Nissenbaum 2010, 82).

Therefore it is not clear whether source control can originate from unspecific desires like the desire of being recorded at random moments without knowing when and where exactly, and it is not clear whether source control is at all possible in such panoptic situations. This means the source control account should be more explicitly defined to require at least some degree of specificity for the relevant desires, or at least more has to be said about which kinds of desires can be ascribed to individuals. If such unspecific desires are ruled out as proper sources for source control, then the source control can avoid this problem.

Turning to the more general discussion of privacy in relation to surveillance and government agencies, the following can be noted. Since surveillance is essentially like being recorded or watched by others at random times, individuals would have to hold the same kind of unspecified desires the couple in the web-camera example involving the computer bug would have to hold to maintain privacy. If unspecified desires and the desire to have these could establish source control, then practices like the state randomly surveilling individuals in a society on a large scale would not necessarily be diminishing privacy, if only all the citizens being surveilled desired this to be so and had the desire to have this desire.[12] However, since such desires are not a plausibly establishing source control, surveillance is clearly problematic. Moreover, the individuals have to be somehow causal for the flow of information occurring – for example their voting behavior might be seen as establishing the right causal connection.

[12] I am assuming here that the surveillance is done in the 'old school' fashion of individual human agents watching and tracking citizens, since the use of digital technology analyzing behavior and patterns without a human person reading every bit of information that is processed does not constitute actual information flow for some authors (Macnish 2018, 429 f; Menges 2020a, 45). However, as I will argue later, there is reason to see this kind of digital surveillance as involving access of personal information as well, and therefore the same argument applies to this style of surveillance.

However, voting seems to be only very distantly causal for the surveillance occurring in the way is does. Therefore, more would need to be said about the claim that voting is causal for government surveillance.

As we have seen, unspecified desires are quite problematic for a number of reasons, and it is not clear whether individuals can even meaningfully hold such desires. Moreover, in situations of constant surveillance (i.e. situations like a panoptic prison) it is no longer clear whether the desires grounding source control originate from the individual at all.

If those worries are justified, then an updated and plausible version of the source control account should clearly imply such large-scale surveillance diminishes privacy.

While for the limited access accounts this intuition is quite straightforward, the source control account does not provide a clear answer until we know what kinds of first and second order desires can plausibly form a foundation for exerting source control. If unspecified desires, like the desire to be filmed by a web-camera at unknown random times or the desire to be surveilled at random times by state authorities, are candidates for forming a basis for exerting source control, then this account has the implausible consequence of compelling its proponents to see privacy as intact in the two examples just given.

Of course, this worry can be met by the source control account by further spelling out what exactly source control is supposed to be. Excluding unspecified desires like the ones mentioned above might be a solution to this. Another way, which I am going to come back to in the next section, is to include facts about the context in which information is flowing into the definition of source control. Here, Nissenbaum's (2010) account of *contextual integrity* provides a lot of insight into which contextual facts matter in the assessment of an individual's state of privacy. On the one hand, source control might be grounded in different kinds of desires in different contexts. On the other hand, context will give important hints as to which desires we ascribe to individuals, and thereby help us avoid the problems arising from unspecified desires.

An important consequence that might arise if unspecified desires are excluded from forming a proper basis for source control is that current consent forms, which are mandatory for owners of websites in the EU, might not actually have the effect of avoiding a violation of privacy. That is, since individuals hardly ever spend enough time to read such policies, most would not be able to understand them, much less understand the consequences of what the collection and processing of their data implies (Tavani 1999), consenting to privacy policies virtually always rests on holding certain unspecified desires (and beliefs). If such desires are not a proper basis for source control, then consent forms do not ensure privacy as long as individuals do not read them, fully understand them, and mind the consequences their sharing of information might

have. While there is good reason to believe that this would be somewhat infeasible for individuals anyway (McDonald and Cranor 2008), there is at least some political effort to make such agreements more easily understandable for individual users.

In the next section, I am going to present a case in which initially it seems like the information subjects are holding an unspecified desire, but as we are going to see, referring to contextual facts makes all the difference to assessing the state of privacy.

### 2.3.2 Source Control and Shared Histories

As we have seen, there is reason to exclude what I have called unspecified desires as plausible foundations of source control. However, there might be cases in which similar unspecified desires are held and privacy might nonetheless be preserved. In the following I am going to present another variation of the web-camera example, which prompts us to draw the exact opposite conclusion. That is, in the following case, the intuition that privacy is maintained is rather strong, because information about the context helps us specify what initially looks like an unspecified desire.

From this I am going to conclude that context matters a lot for finding a plausible foundation for source control. The relationships between the two parties, their nature and roles in society, change a great deal about when privacy is maintained and when it is not.

The case I want to investigate here is the following. Imagine rather than a computer filming the couple, there is a person whom they sometimes tell to watch them. Assume the couple have had a long relationship with the individual watching them and know the individual respects their wishes. Moreover, the individual is always trying to be as responsive as possible to the couple's wishes, and the couple knows this. Now imagine one day the couple forgets to tell this person to watch them, but the person (perhaps remembering that in similar situations the couple has asked them to watch them before) goes on to watch the couple nonetheless, as is their desire, and they desire to have this desire. While watching the couple, the person is very attentive and tries to catch any hint of whether this time the situation might be different and there is reason to believe that the couple does not want to be watched this time.

In this case, just like in the web-camera example, the couple is not the direct source of the flow of information to the individual watching. However, an important distinction is the individual's attention and the history of the two parties, which might ensure the flow of information is grounded in the couple's desires nonetheless. It seems plausible that the person watching is not violating the couple's privacy given an adequate shared history. Still, it is not

clear how to explain this, since the individuals being watched are clearly not the source of the flow of information in the very moment it is happening.

One response to this is to insist that the person watching is violating the couples' privacy. However, my intuition is that privacy is preserved in such a situation, and that this is in line with how close relationships often work. A friend finding out personal information about me via another friend does not imply privacy being diminished, even though I never gave the first friend explicit permission to pass on the information.[13]

Another response of proponents of the source control account might be to accept that privacy is not diminished in cases involving unspecified desires, if certain conditions concerning the relationship of the two parties and their history or certain knowledge conditions are satisfied. One might be the right kind of source of an information flow if some kind of implicit consent has been given. A history or specific habit of sharing information might be enough for us to establish such a form of implicit consent. If the individual watching the couple has acquired the true belief that the couple wants to be watched by them from the couple in the right kind of way, then the couple can be seen as the right kind of source for the individual watching them without explicit consent.

Indeed, it is much more plausible for individuals to foresee what a good friend or person they have long known will do with their personal information than it is to foresee what state authorities, random individuals on the internet, or digital information processing systems will do with their personal information. Therefore, the important difference between the two kinds of cases involving (seemingly) unspecified desires lies in facts about the party receiving the information and the *relationship* between the two parties (sender and receiver).

On the one hand, this means the source control account is able to explain the differences between cases involving unspecified desires where the (behavior of) the receiving party and the relationship towards them is either unknown in some sense, and cases where the receiving party is well known and well understood by the individual whose personal information is flowing. That is so even if in neither of the two kinds of cases the sending party has initiated the specific flow of information.

On the other hand, as mentioned above, this line of argument is opening up a way of looking at privacy (violations) in a fashion akin to Helen Nissenbaum's conception of *contextual*

[13] I am not sure whether this intuition presupposes privacy to be more like intimacy (compare Inness, 1992), or whether it provides reason for us to see privacy as closely connected to intimacy. However, the intuition that source control is absent in this case is clearly less strong than in Menges' web-camera example. The friend is watching the couple for a reason, and this reason lies in the relationship between them and the couple.

*integrity*, which characterizes the right to privacy as "a right to *appropriate* flow of personal information" (Nissenbaum 2010, 127).[14] That is, when analyzing situations with respect to privacy it might be much more important to consider the relationship between (as well as properties of) the involved parties. In the discussion I have focused on in this section, little attention is usually paid to such facts. Who reads my diary, and for what reasons or in what role, can make a lot of difference to the state of privacy I am in.

While we might be able to formulate contextual facts in terms of necessary conditions for source control, the focus on individual desires makes it unnecessarily difficult to go beyond facts about the individual and her psychology when analyzing privacy. As we have seen, individuals do not always need to know when personal information is flowing and what exact information is flowing for privacy to be ensured. However, we need some other ways of conceptualizing source control in a way that enables it to deal with such cases. As long as this is not achieved, it seems more fruitful to focus on accounts like Nissenbaum's contextual integrity to make sense of the vast array of privacy threatening situations we find ourselves in. Still, defining source control in a more context sensitive way might be achievable.

As we have seen in this section, the account of privacy as source control faces some difficulties which cannot easily be resolved. Spelling out what being the right kind of source of an information flow means is crucial to presenting a plausible account of privacy. The above arguments have provided some good reasons to define source control in a context sensitive way, where being the right kind of source means something different depending on who is involved, what the relationship between the relevant parties is, and how the information is being transmitted. That is, source control can learn a lot from Nissenbaum's conception of contextual integrity.

Turning to a more normative discussion of privacy, economic theory can be helpful in analyzing the effects of holding private information in various contexts. As Wolff notes, models like his game theoretic model of "emergent privacy" can provide us with "an explanation for the difference between the norms of different social contexts" (2015, 156). In section 4 I am going to consult some economic approaches to analyzing the effects of (restricting) information flow, and compare the results these models find to the claim that privacy as control protects individuals' *bargaining power* (Menges 2020a, 42).

[14] Note that Nissenbaum is talking about the *right to* privacy here. However, according to Lundgren we might assume that there needs to be sufficient consistency between the right to privacy and the concept of privacy (Lundgren 2021, 671), i.e. since privacy is the object of the right to privacy, we can assume Nissenbaum defines the concept of privacy as *appropriate flow*.

## 2.4 Problems for Frankfurtian Source Control

When arguing in favor of defining privacy in terms of source control, Menges adopts an understanding of source control that is immediately based on Frankfurt's compatibilist theory of free willed action. While Menges emphasizes that this is only one possible account of source control, and much debate is needed for figuring out which understanding of source control is most promising, I nonetheless take the suggestion to adopt a Frankfurtian understanding of source control at face value for the moment. By reviewing some of the problems that come up for a Frankfurtian account of compatibilism in the free will debate, I am going to argue that such an account does not present a satisfactory definition of what 'being the right source' of a certain information flow means.

First, recall the argument involving the Frankfurt-style case introduced above. Instead of claiming that control requires alternative action, this line of argument provides good reason to see control over one's action to lie in being the source of the respective action. In the original debate, the argument then provides good reason to hold free will and determinism as compatible by showing that an individual does not need to have alternative paths of action available to her in order to be said to act freely and responsibly. Rather, as long as an individual is the source of a certain action, she can be held responsible for it. Analogously, being the source of a flow of information, rather than having alternative choices available, is what matters when figuring out whether an individual enjoys privacy. When information flows, then, as long as the individual is the source of this flow, she enjoys privacy.

Of course, the plausibility of this account depends on what sourcehood comes down to. Most importantly, we want to make sure that what an individual does is "grounded, *in the right way*, in them, and not in someone or something else" (Menges 2020b, 392). For Frankfurt this idea is expressed by referring to different kinds of desires an individual can have, and the interplay of those desires in light of the course of action that has been taken.

More explicitly, Frankfurt (1971) distinguishes between first-order desires, which are desires about performing certain actions, obtaining certain things or the like. These desires can either lead to action, or they remain hypothetical volitions of the individual. The former, the effective desires, are then said to form the individual's will (McKenna and Coates 2020, chap. 4.2).

The higher-order desires are such that they do not refer to external objects or actions, but merely to other desires. Second-order desires are desires about having certain first-order desires. Just like first-order desires need not be effective, second-order desires need not result in the

individual actually holding the relevant effective first-order desire. In this way, the individual has desires about the kind of desires comprising her will, i.e. her effective first-order desires.

This distinction allows Frankfurt to define free-willed action as actions that "issue from the will [the individual] wants" (ibid. 2020, chap. 4.2.1). That is to say, an individual acts freely if and only if the effective first-order desire they act on is the effective first-order desire they desire to have. Analogously an individual exerts source control if and only if, when information flows, it does so because of their will, and the desire to have this will.

The first problem this theory faces according to McKenna and Coates is the following. First, higher-order desires above second-order ones are perfectly compatible with Frankfurt's theory. That way, one can have third-order desires about which second order desires one wants to hold. Realizing that just like first-order desires second order desires can be contradictory, it becomes clear that sometimes third-order desires will be necessary to determine which second order desires reveal the agent's true volitions, just like second-order desires are intended to reveal which of the possibly conflicting first-order desires are truly revealing the agent's will. That is, just like sometimes "a person cannot be identified with her first-order desires because she can be alienated from them" (ibid. 2020, chap. 4.2.2), sometimes the second-order desires of a person will not be decisive in determining which desires they can be identified with, and the individual can be alienated from those as well by referring to third-order desires.

Now if we appreciate the fact that in theory an individual can hold conflicting second-order desires, and that they can form higher-order desires alienating them from one of the conflicting second-order desires, we can easily see how this becomes a potentially infinite chain of ever-higher-order desires never fully sufficient for determining whether the agent acted on their free will. Analogously, determining whether an agent was the right kind of source of an information flow cannot fully be determined by referring to first-, second- and even higher-order desires.

A second problem highlighted by McKenna and Coates concerns the origin of the first-, and second-order desires we employ for determining the agent's status of agency. We have already seen how determining the origin of one's desires can be problematic. Recall, the writings of Bentham (1787, 2010) and Foucault (1995) lead us to question whether individuals really act on their own desires in situations of surveillance, or perhaps in social situations in general. Much more plausibly, they are going to act the way they expect the watching authorities, or their social environment to desire them to act, and thereby leave their own true desires aside.

At first sight this would not be a problem for Frankfurt's conception, since he could simply claim that individuals in such situations are never acting freely.

However, internalizing an authority's desires also means making them one's own in a way

that is indistinguishable from the desires that are 'truly' one's own. If some of the desires an individual has originated from outside herself, then those desires are clearly not representative of herself. Frankfurt, however, offers us no tool to account for the origin of a desire, and moreover he seems to be committed to accepting that individuals act freely even if the relevant desires originated outside them. More strongly, even acting on desires implanted by hypnosis or similar techniques would result in free agency on this view. Analogously, individuals would enjoy privacy even if they acted on desires implanted by hypnosis.

It should be clear by now, that the Frankfurt style conception of source control faces some serious internal conceptional problems. On the one hand these arise because it is not clear at which level to stop when looking for desires we can truly identify an individual with, and therefore we might be referring to desires of an ever-higher order, possibly arriving at an infinite chain of desires we reference in search of the true self of an individual. On the other hand, the origins of the desires an individual has are not accounted for in this account. Therefore, desires one has as a consequence of hypnosis, drug use, or even societal influences could be referred to when looking for the individual's true self. While this is clearly problematic, Frankfurtian source control would have to accept such desires as a proper base for source control.

Now on the one hand, source control need not be defined in terms of first- and second-order desires. Menges explicitly mentions a number of possible accounts his conception of privacy could be based on (2020b, 393). On the other hand, these challenges might be overcome by further working on them. Still, this section has served to elicit the challenges that lie ahead in spelling out a plausible account of source control.

For now, I am going to leave the discussion at this, and turn to another point of interest in the recent discussion originating in Macnish (2018). Namely, I am going to examine whether information should be viewed as being accessed in cases of government surveillance like those practiced by NSA and GCHQ. I am going to argue against a view that both sides of the recent discussion share. Namely, I am going to argue against the claim that personal information should not be seen as accessed when it is processed by digital information processing systems like those run by the NSA and GCHQ. My claim is that information should be seen as accessed even when no individual human agent at those agencies receives it.

## 3. Access and Digital Data Analysis

Since the rise of digital computing technologies, human capabilities of collecting, storing and analyzing vast amounts of data have increased dramatically.[15] With the rise of sophisticated methods of Knowledge Discovery in Databases (KDD) – involving for example 'data mining' – the capability to sort, prepare, 'clean', as well as analyze data for particular patterns and even discover previously unknown relationships has increased immensely (Tavani 1999). While the kinds of informational advantages and disadvantages that arise from individuals having more or less access to such technologies are going to be of interest in section 4, where I am going to focus on an economic analysis of information-inequalities, in this section I want to focus on a different issue raised by practices of information collection and KDD.

According to Macnish (2018) the recent discussion between proponents of control accounts and limited-access accounts of privacy has important implications for settling the public debate evoked by the revelation of NSA documents by Edward Snowden in 2013 (see e.g. MacAskill et al. 2013). On the one hand, so Macnish claims, representatives of the NSA and GCHQ have taken a stance that is based on a limited-access account of privacy. As long as the data of most individuals is merely captured, stored and analyzed, their privacy is not diminished, they claim. Only when individual agents access information, which is supposed to happen based on warranted suspicion, will the privacy of the respective individual be diminished.

On the other hand, the side of the debate who see the practices of NSA and GCHQ as diminishing people's privacy supposedly understand privacy as control over personal information. Since collecting personal information already diminishes the amount of control individuals have over their personal information, the practices of the NSA and GCHQ diminish privacy.

However, the source control account of privacy is not drawing the same conclusion Macnish attributes to (other) control accounts. Since a lack of information-flow secures privacy on this account, proponents of source control accounts side with proponents of limited-access accounts in viewing NSA and GCHQ practices as maintaining the privacy of the affected individuals.

It is this conclusion I want to critically assess in the following passages. My focal point in

[15] Interestingly, reviewing various methods of how researchers attempted to measure the number of indexed sites on the WorldWideWeb based on the indices of search engines like Google and Bing, van den Bosch et al. (2016) conclude that measuring the amount of indexed web pages is a task that is not unproblematic. Most importantly, relying on commercial search engines like Google is not a reliable methodology, since the variations in their indices do not always link to the variations in the corpus they claim to index (Kilgarriff 2007). That is why I'm merely claiming capabilities have increased, rather than the actual amount of information or websites.

doing so will be the implicit claim that the technology involved in surveillance practices of NSA and GCHQ has a neutral role with regard to privacy. As I will argue, it is mistaken to believe that only human agents' directly reading information (or processing it in other ways 'in their brain') matters to determining whether this information has been accessed. More explicitly, I am going to argue that the collecting, storing and analyzing of personal information carried out by whatever technology the NSA and GCHQ employ should count as access in the relevant sense.

I claim that digital information processing systems are artifacts used by individual humans or collectives of humans. By using such systems to accumulate and process personal information, the whole systems involved in doing so (i.e. the collective made up from humans and the artifacts they use) can be seen as accessing information, even if information does not (immediately) flow to the human individuals or collectives using the systems.

### 3.1 Privacy, Artifacts and Digital Surveillance

Let me start by considering the interplay of humans and the artifacts they use in general, and the (moral) implications this interplay has. While in moral theory the focus mostly lies on individual human agency, the most helpful way for making sense of the complex interdependence of humans and machines in agencies like the NSA and GCHQ is to consider moral agency as distributed across humans and machines. Rather than only holding the individual and their mind to be of moral relevance in agency, this compels us to take into consideration their environment and how changes in the environment might be of moral relevance too.

The processes at play in agencies like the NSA and GCHQ can be understood as a cognitive process, where information is received, searched and analyzed for irregularities and salience, and on the basis of the results of this process decisions are made on which citizens to place under close scrutiny. Since the digital systems, helping the individual agents to achieve the task of finding criminals and potential threats, are deeply integrated in the whole process and relied upon by the agents, it seems implausible to consider only the human part of this complex system as the kind of recipient able to diminish privacy. Rather, as I am going to argue, the digital information processing units' receiving a certain bit of personal information might already constitute the kind of access necessary to diminish privacy.

In order to defend this claim, I am going to resort to the theory of *extended cognition*, which

is closely related to theories of embodied cognition (compare Shapiro and Spaulding 2021).[16] This way of viewing cognitive processes will help make sense of how to understand the use of tools or artifacts by individuals or collectives, like state agencies, with respect to privacy.

In their 1998 essay, Clark and Chalmers explore the idea of what they call the "extended mind". There, they argue that when humans use external objects or artifacts to aid their cognition, we should see these artifacts as part of an extended cognitive system, if certain criteria are fulfilled. These include the artifact being 'reliably coupled' (ibid. 1998, 11)), which is to say the artifact in question has to be reliable, trustworthy and easily accessible in repeated interactions. However, some argue many more dimensions have to be considered. These include, among others, some measure of intensity of information flow between artifact and human, and interpretation of the information being relatively easy (Heersmink 2017, 433).

To get a clearer grasp of why this is a quite plausible way of understanding cognitive processes, let us consider the following example of Otto found in Clark and Chalmers:

> Otto suffers from Alzheimer's disease, and like many Alzheimer's patients, he relies on information in the environment to help structure his life. Otto carries a notebook around with him everywhere he goes. When he learns new information, he writes it down. When he needs some old information, he looks it up. For Otto his notebook plays the role usually played by a biological memory (1998, 12).

In such a case, the authors conclude, the notes in the notebook play the same role memory plays for people with a functional memory, where whatever the notes say about e.g. the directions towards a certain concert hall will determine the route taken by Otto on his way to said concert hall. "When it comes to belief" they claim "there is nothing sacred about skull and skin" (ibid. 1998, 14). Rather, it is the role of a certain piece of information that characterizes it as a belief.

One important note is due at this point. We have to separate the discussion of whether artifacts can be part of an extended cognitive system from the question of whether these artifacts or the systems they are part of can thereby be attributed cognitive agency. As Heersmink notes, the fact that artifacts can be part of extended cognitive systems does not imply that they can be attributed cognitive agency. Rather, agency might still only be attributed to the system as a whole, but not the artifact alone (2017, 434 f).

[16] I am not going to defend theories of embodied or extended cognition rigorously here, since that would exceed the scope of this work. For a good starting point to some of the issues mentioned here see for example Adams and Aizawa (2008).

An extended cognitive system like that of Otto and his notebook might of course be involved in an action that is (potentially) threatening the privacy of some other individual. To illustrate this consider the following example:

> Imagine Otto lives in a future where his notebook is a digital device collecting not only what Otto enters into it, but also collects facts independently, based on how Otto has programmed it. The device has access to and frequently analyses personal information Linda has stored on her computer, and creates notes based on this information. These notes are available to Otto at all times.

Now the first intuition here should be shared by most parties in the privacy debate. That is, when Otto reads one of the notes his device has created, he is accessing personal information about another individual. However, how to treat the notes Otto does not read is what I want to focus on now.

Considering their claims about access in the case of NSA and GCHQ, Macnish and Menges would both claim that as long as Otto is not reading a certain note containing personal information, the information contained in the note has not been accessed, and privacy is not violated. This intuition becomes much less clear, though, when comparing the mechanism involved in such surveillance practices to the case of Otto and his smart device. As the idea of extended cognition is supposed to make clear, cognition of the individual does not end at the bodily borders, but rather is a question of function. The information contained in the unread note is functionally there, even if it is not in the conscious part of Otto's brain, and plays the same role it would if Otto were to memorize it. The extended cognitive system *has* the information, and that is what counts. Similarly, if some piece of personal information is being made into a note for some agent at the NSA, then it should not matter whether the agent actually reads that note. The information is there, it is being used, and the extended cognitive system that is the NSA has processed it even when it is not made into a note.

One important difference between the extended cognitive system of Otto and his device, and the NSA, however, is that the former system is centralized around the intentions of Otto, while the latter is decentralized, i.e. does not have a clear center around which it is organized (Heersmink 2017, 435).[17] Referring to the original example of Otto again, he notes that "Otto

[17] Heersmink is pointing to a distinction between extended cognition and distributed cognition here, which is discussed at length by Hutchins (2014).

and his notebook are an extended agent because when Otto uses his notebook, he uses information *he* made *himself* to realize *his* goals" (2017, 436). Conversely, decentralized systems like a laboratory in which many individuals work together are not extended cognitive agents if there is no center whose intentions are being realized.

For the NSA and GCHQ this might mean that, since they are complex organizations made up by many agents with different intentions, no center can be identified, and consequently no agency can be attributed to them as a whole. However, at least with respect to the laboratory example used by Heersmink there might be some relevant issues to be resolved before jumping to this conclusion.

First, the researchers in the laboratory might indeed have very different intentions that share no connection at all. However, if the lab team were to work on a project where everyone shares some goal (e.g. finding a specific medication to cure a specific disease) and they coordinate in order to achieve this goal, then it would become much less clear that there is no center around which this cognitive system is organized. The assumed amount of coordination in settings like a laboratory makes it plausible to assume some center of the whole endeavor, even if this center is not embodied by a human person.

Second, and closely related, it is not clear why we should restrict ourselves to holding only individual human beings as possible candidates for forming the center of an extended cognitive system. As Hutchins notes, this is a common assumption of extended mind theory (2014, 36). However, if all the center needs to provide is an intention that can be realized by the system, one could argue that collectives sometimes have intentions and therefore the center of an extended cognitive system need not be the material brain of a specific individual.[18]

For the case of agencies like the NSA and GCHQ this might mean that we can indeed see them as extended cognitive systems that possess agency. Therefore, the whole of these systems and not just the individual agents reading the information collected and processed by their tools are to be considered when investigating their endeavors with respect to privacy.

With respect to what Macnish calls the "New Controller" or simply NC in the diary case (2018, 428), one more remark is at hand. The NC and the NSA are taking control of the information in the diary and information about an individual's internet traffic respectively. However, unlike the NSA, the NC does not proceed to read the information immediately and so cannot be said to access it. Agencies like the NSA on the other hand use their tools to copy,

[18] While I am not going to be discussing conceptions of collective or shared intentionality here, a good starting points for this would be papers by Gilbert (1990), Gold and Sugden (2007), Searle (2002), or Tuomela and Miller (1985), all of which hold that collective intentions are part of our social world.

analyze and store the information they have found. Of these types of access, it is the analysis that should worry us the most here.[19]

By letting a machine analyze the information captured, the NSA does something much more invasive than what the NC of the diary does. They use a machine to take over a task that would have otherwise been done by an individual agent. If this agent were to analyze some personal information captured by the machine, there would be no doubt the agent would be accessing the information by doing so. Likewise, the software system must be seen as accessing personal information in a privacy diminishing way (compare Clark and Chalmers 1998, 9 and 13), if it fulfills the same function. If the individual agents in the NSA can be said to be part of an extended cognitive agent composed of these individual human agents, the machines, and the software they are using to analyze personal information, then we have to treat the analysis done by the software as equal to the NC opening my diary and analyzing it for whatever content they are interested in. The extended cognitive agent is accessing information even when its human part does not directly access the information. Privacy can therefore be diminished without human individuals directly accessing information.

At this point I want to recap the claim I have argued for. The argument I have presented gives us good reason to deny the plausibility of a claim both Macnish (2018, 29f) and Menges (2020a, 46) have made about access to personal information by government agencies like the NSA or GCHQ. Both claim that information collected by such agencies does not count as being accessed as long as no individual agent has read the relevant information. Therefore, privacy cannot be diminished as long as no individual agent in these agencies reads the information on limited-access accounts and the source control account of privacy. The case of Otto and his smart device provides a strong reason to refute this claim. It illustrates that sometimes humans and their artifacts are so closely intertwined that they form an extended cognitive agent able to diminish privacy even when the human part of the extended cognitive agent has not accessed the relevant personal information. That is to say, when artifacts collect, archive and analyze personal information as part of an extended cognitive agent, this should be seen as the information being accessed by the extended cognitive agent in a possibly privacy diminishing way.

[19] While I do believe that copying and storing might require a relevant form of access as well, the analysis of personal information seems to provide a much stronger example for my argument here. However, imagining the NC of my diary copying and storing the contents of my diary (e.g. by photographing it page by page) seems to be a case of my personal information being accessed by the NC. If the NSA does the same thing and we can hold them to be an extended agent like I am assuming the NC and her photo camera to be, then personal information is accessed by the NSA even if it is not analyzed or read by individual human agents.

To be perfectly clear, this argument applies not only to government agencies, but to any extended agent (individual or collective) that is collecting, processing, storing and perhaps selling personal information *without* having a human individual directly access this information (e.g. by reading it). In this way, operators of apps that collect information like location data or medical information from health trackers are accessing this information just as much as data-brokers collecting, buying, and selling information, or even offering some of it on search pages for the general public to access (see e.g. Sherman 2021a; 2021b). The fact that no human ever reads all of that information does not make such practices unproblematic with respect to privacy.

Over and above the argument about extended cognitive agents and access, I am going to argue that the lack of knowledge about how information is being collected in the digital realm is problematic for attributing source control to individuals when information flows.

Reviewing practices of data brokers and the like makes for another interesting point regarding the notion of *being the source* of a certain flow of information, which has been discussed in section 2. More specifically, data collection in the digital realm makes it much clearer that, in order to be the right kind of source of an information flow, individuals must have at least some knowledge of the situation they are in. Most centrally, they must know who is watching or listening to them when they act or talk.

When we realize that most of the activities of data-brokers – especially the collection of information by tracking users' phones, IP-adresses and so on – happen automatically without the average user ever witnessing what is going on, it is hard to see how the flow originated from the individual in the right kind of way. Of course, the individual is the one who opens a certain website and perhaps browses through an online store to find the perfect item to buy, but that cannot be sufficient for taking them to be the right kind of source of the flow of their personal information.

To make this clearer, imagine a person giving a very personal lecture (like the Dewey lecture of Peter Railton, which Menges refers to; 2020b, 390), not knowing who is in the room listening. In this case, source control is given since they desire to hold the lecture and want to desire to hold the lecture. So, the information flow to the crowd in the room is rooted in the person's desires in the right way for Menges. However, it seems completely off not to take the person's lack of knowledge about the crowd into account into account. Therefore, at least some knowledge about the situation or at least consideration of the situation has to be part of the desires and reasons that can plausibly make an individual be the right kind of source of an information flow.

If this is convincing, and if the relevant software's collecting information is sufficient for information being accessed, then the lack of knowledge about the practices of all kinds of parties collecting information on the internet, from our smartphones and other digital gadgets, makes it at least highly dubitable that we as users ever enjoy source control over our information in the digital realm. Rather, we are constantly surveilled by extended agents collecting, selling, analyzing and storing our personal information, without, for the most part, knowing who does all that, how they do it, and for what reasons. Government agencies are only one part of this problem.

Coming back to the claim that extended agents can diminish privacy, it is clear that this claim is surely in need of further defense, and one could argue that extended cognitive agents are not the kind of agent able to diminish privacy. However, it should also be clear that Menges and Macnish give no reasons why only individual human access of personal information is relevant to their accounts of privacy. Humans use tools all the time and this is seldom seen to diminish their agency, but rather to extend it beyond their prior capacities. I claim that the nonhuman parts of extended (cognitive) agents are able to diminish privacy even before the individual humans have received the processed information.

Moreover, exploring the ideas of collective intentionality prompts us to conclude that the center of an extended cognitive agent cannot only be found in material human brains, but also in collective or shared intentions held by groups of individuals. That is to say, even agencies like the NSA might be understood as an extended cognitive agent with respect to some of the processes and tasks they are undertaking. If such extended agents have a shared intention – like e.g. the intention to maintain national security – which all processes are focused on and coordinated around, then they might be seen as extended cognitive agents. If personal information is then accessed by any part of this extended cognitive agent, this access is of a relevant kind for limited access as well as source control accounts of privacy.

In the last section of this essay, I am now going to turn to some evaluative considerations. More specifically, I am going to investigate whether control accounts of privacy really protect individual bargaining power, as Menges (2020a, 42) claims. With the aid of economic models, I am going to argue this is not so.

## 4. Consequences of Keeping Information Private

So far making normative claims about the respective accounts of privacy has been avoided. In this section, however, I want to investigate the claim that privacy as source control (or control in general) protects our bargaining power and is therefore valuable (Menges 2020a, 42). While for Menges this claim is closely connected to securing oneself from social and political repression, I am going to focus on the economic aspect of the claim that privacy as source control secures the bargaining power of individuals.

At least initially this claim is surely reasonable. After all, controlling who gets to know personal information about us protects us from discrimination, including possible forms of economic discrimination. While I agree that control over who gets to know what about us, what they get to know, and at what exact time, will often be beneficial to the individual, I am specifically interested in cases where information does not flow, i.e. situations in which the individual chooses not to share information, it is not possible to transmit a certain piece of information, or if it is possible to transmit the information, the receiver cannot confirm it to be credible. In these cases, source control is maintained per definition. However, as we will see, bargaining power is arguably diminished in some of these cases. As some economic models have shown, the lack of information flow can lead to possible profits not being realized, or even no trade occurring at all. Based on this insight, I am going to argue that maintaining privacy as source control does not necessarily protect our bargaining power.

In order to do so, I am first going to clarify how bargaining power can be understood, and in which cases it can be seen as protected. While bargaining power is first and foremost a measure of the ability to capture profit from trade, my focus here will be to understand what we can say about bargaining power in situations where not all possible gains from trade are realized, or no trade occurs and therefore no profits are made at all.

Second, I am going to introduce the economic results my argument will be based on. The seminal "Market for Lemons" model (Akerlof 1970) is presented to introduce the effects of 'asymmetric information'. In this model, when one party possesses relevant information that cannot be (credibly) transmitted to the other party, profit for both parties will either be diminished or no trade will occur, and hence zero profit will be made. Closely akin to this analysis, mechanism design theory will provide a different perspective on the possible economic effects of individuals strategically controlling their personal information. Mechanism design theory focuses on problems of collective decision making, but importantly considers the additional assumption "that individuals' preferences are not publicly observable" (Mas-Colell

and Whinston 1995, 857). In a classical economic model situation, this leads to the result that no social choice mechanism can actually be Pareto-efficient and fulfill very plausible desiderata, which I am going present in this section.

Finally, on the basis of these economic analyses, I am going to argue that privacy as source control does not necessarily maintain bargaining power.

### 4.1 Bargaining Power

Bargaining power can be seen as an individual's ability to influence the outcome of a given economic transaction between buyer and seller, customer and firm, or employee and employer. An individual has bargaining power if she is in the position to influence such a transaction in a way that raises her profits from this transaction. This might mean the other party gets less in a zero-sum game, or otherwise might mean profits for both parties are raised.

One way of raising the profits made (or avoiding losses) from a certain exchange is to strike a careful balance between revealing and withholding relevant information. As an example, withholding information about a certain health condition I have might raise my ability to bargain for a higher wage with my employer (compare Menges 2020a, 42). If this is so, then my controlling the flow of personal information is coincident with raising my bargaining power.

Likewise, we can imagine cases in which revealing such information will be more beneficial. If my health insurer is lacking certain information about my health condition, then I will not be in the position to get the relevant benefits from them. For example, pressing my insurer to pay for a pair of new glasses will be fruitless unless the insurer knows I really need the glasses, which is usually proved by producing some prescription from a qualified physician.

One thing we can notice already is that there might be conflicting reasons for sharing and withholding personal information in bargaining situations. I might not want to share personal information, but still, if I want to get my glasses covered by insurance, I will have to share the information with them. In this case the decision to share involves a trade-off between concerns for privacy and a concern for my health. In the end, the desire to share might be stronger and information flows to the insurer. However, the mere choice to share should not be sufficient for claiming that source control is maintained. Rather, as we have seen, this choice must be based in the agent in the right kind of way.

Viewing bargaining situations and economic exchange in general might caution us to be careful in defining source control. People want many things for different reasons, but if individuals share personal information for purely economic reasons, then it is hard to see why

this sharing would not violate privacy as source control. As I have argued above, the idea here is that being the right kind of source of an information flow means the individual did not just share information for profit but has at least made some privacy-relevant considerations. Still individuals will sometimes share private information in part for economic reasons in a way that does not diminish their privacy.

This already shows how sometimes bargaining power and source control might come apart. Placing some minimal conditions on which reasons are necessary for maintaining source control will make source control and bargaining power come apart in some cases, if the relevant reasons are non-economic ones, such as showing a certain care for privacy-relevant considerations. In a certain bargain I might not enjoy source control if I did not consider the relevant reasons, however, my economic profit might still be raised through sharing personal information.

A second possible way in which source control and bargaining power might come apart will be at the center of the following section. In economic theorizing, a lack of information is usually considered a problem for achieving optimal outcomes. Not knowing the true value of the good one wants to buy, not knowing about the environmental circumstances of the area one is planning to buy a house in, etc. will usually make a person worse off from a trade than they could have been. Over and above this intuitively obvious importance of information, the models I am going to present in the next section show how not sharing (or not being able to share) relevant information might also make the person holding the information worse off than they could have been. Sometimes no trade will occur at all in such situations.

Here I want to briefly consider what that would mean for bargaining power. As mentioned before, bargaining power can be seen as the ability of parties to influence their payoff from a certain trade. So ultimately, being able to raise one's payoff (or avoiding loss) is at the heart of having bargaining power. Considering the possibility that lack of information on one side might lower the overall gains to be made from trade for all involved parties, we can say that bargaining power is diminished, since the maximum payoff the parties are able to make is lowered. This is even clearer in cases where no trade occurs. If no trade occurs, no profit can be made and there is no profit to bargain about. No trade occurring means the individuals have zero bargaining power. I therefore take bargaining power to be diminished in cases where less than optimal or no payoff is achievable.

Which exact situations evoke the consequences mentioned here will be explained thoroughly in the next section. The first are situations of asymmetric information, where one party has private information that is relevant to the trade. The second are situations where some mechanism determines an optimal outcome based on the parties' reported preferences, who

might or might not be telling the truth in order to influence the outcome.

### 4.2 Asymmetric Information

To get a better understanding of why the absence of information flow from one side of a bargain to the other sometimes has adverse effects I am going to consider a model analyzing the effects of asymmetric information in market situations presented by Akerlof (1970).[20] This model is going to provide good reason to believe that a lack of information flow from one side to the other can have negative economic effects for both sides of the market, and that therefore privacy as source control does not generally protect our bargaining power.[21] In the model situation, since no information flows, source control is maintained, but bargaining power is still diminished.

Now let me give a brief presentation of the model in question. Suppose there is a market for a certain good where the quality of individual items of the good lies in a continuum between "good" and "bad". The sellers of these goods have an estimate of the quality of the good that is "more accurate than the original estimate" (Akerlof 1970, 489), which is to say they know more about their particular item of the good than is common knowledge to all participants of the market.

For simplicity let us consider a market where there are just two types of goods, good ones and bad ones, and consumers are willing to pay a high price $p_H$ for good items, and a low price $p_L$ for bad items. The number of good items is denoted by $\lambda$, and the number of bad items is $(1-\lambda)$.

Now suppose consumers are not able to observe the quality of an individual good, but only know the average quality of goods on the market. They are willing to pay the value of the good they *expect* to get, which is the average value of all goods on the market, since consumers only know what the average quality of goods is in case all sellers enter the market. This implies the consumers are willing to pay $\lambda p_H + (1- \lambda)p_L$ for a unit of the good. However, sellers also have a

[20] Evidence that the effects we can observe in the model are also present in the real world can for example be found in Puelz and Snow (1994) for automobile insurance markets, Cutler and Zeckhauser (1998) for health insurance markets. Cohen and Siegelman (2010) provide a review of some of the arguments that have been made about the strength of empirical evidence for adverse selection in insurance markets, and arrive at the conclusion that this evidence does support the claim that adverse selection is of some concern in such markets.

[21] Note that Menges only claims privacy as control protects our bargaining power in a "less-than-fully-tolerant society" (2020a, 42). However, even in such societies it should be possible to find situations where not sharing information diminishes bargaining power. Unless Menges restricts his claim to the particular personal facts that are at the focus of discrimination, my argument still holds.

reservation prices $r_H$ or $r_L$ for good and bad items respectively. Below this price sellers will not be selling their goods. In the model situation, $r_H$ and $r_L$ are smaller than $p_H$ and $p_L$ respectively.

Now, if, given the expectation that both types of item are being traded, the price consumers are willing to pay is lower than the reservation price of sellers of good items, such sellers will refuse to sell their good items at $P= \lambda p_H + (1- \lambda)p_L$. Whenever $r_H > \lambda p_H+(1-\lambda)p_L$, therefore, the expectation that both types of items are being traded is not consistent with the behavior of sellers of good items. Given such a constellation of parameters, the situation in which all sellers enter trade is not 'stable.' We shall rather expect sellers of good item to withdraw from the market, and only expect sellers of bad items to stay. Although consumers would be willing to buy good items at a price equal to or even higher than what sellers of good items would accept ($p_H > r_H$), in the presence of information asymmetries the market mechanism may fail to guarantee that individuals will have an opportunity to implement such a trade.

This very brief summary of Akerlof's model is supposed to show a possible effect of withholding information in market situations. Although this lack of information flow clearly protects the seller's privacy under a source control account of privacy, the effects on their welfare are negative. Their bargaining power is not raised by the fact that relevant information is not or cannot be shared.

An important caveat of Akerlof's model is that it takes the relevant information to be nontransferable for the individual who has it (compare also Samuelson 1984, 996 f.). This restriction of the model, while clearly realistic in some situations, urges us to be cautious about the conclusions drawn with respect to source control. Most importantly, the option to transmit credible information would change the paths of action open to the parties and would therefore influence the results. However, the general problem remains if the parties involved in the exchange choose not to reveal relevant information. Because of this restriction of the model considered here, however, one can make out some ways for proponents of the source control account to respond to the argument presented above.

First, the source control account might be modified such that source control is only given when the individual acts on their deepest desires. In this way, whenever the individual does not share personal information in a situation where they desire to do so, source control might be diminished. In situations like those in the models presented here sharing information is not an option for the individual. Therefore, proponents of such a modified source control account would claim that, in some cases like those in the lemons model, where information cannot be shared, source control is diminished although no information flows.

However, this defense of source control presupposes that just like the flow of information

has to originate from the individual in the right kind of way, so does the lack of flow. If no information flows in cases where the individual desires to share this information, source control might not be maintained after all. Admittedly, in some situations this seems quite plausible if we take privacy to be like intimacy. In this case, not being able to establish intimacy would diminish privacy.

Still, reviewing this modification of the source control account with respect to the diary case discussed above, we can see how it also has counterintuitive implications. Assume I want a friend to read my diary and they fail to do so because it is written in a language they do not understand. It would be odd to say I am losing privacy in this case, while clearly something is going wrong with establishing intimacy (or whatever else the goal of my sharing might have been). Therefore, modifying the source control account to require source control over the lack of information flow does not seem to be a viable option to proponents of this account.

Moreover, it is not even clear such a modification would solve the problem. After all, whether source control were maintained in cases where no information flows would merely be a contingent fact dependent on how the flow of information is connected to the individual's desires, while in every situation that is sufficiently similar to the model situation, their economic profit will be less than possible or even zero. For this reason, and because of the counterintuitive implications of the modification, it is not a very promising way to answer the line of argument I am presenting here.

Setting aside the possible modification of the source control account, the claim I am therefore defending is that privacy as source control does not always protect our bargaining power and does not protect us from experiencing economic disadvantages. In cases where no information flows, we should hold privacy to be maintained. However, in cases like the ones we have seen in Akerlof's model, trade might not occur at all or at least only a suboptimal outcome might be achieved. The lack of information flow does not benefit the parties involved in the market.

Another possible defense against my argument that adverse selection shows how privacy as source control does not always protect bargaining power is the following. If we give up the claim that privacy is a neutral concept, and instead hold that privacy (as source control) is generally valuable[22], enjoying privacy would have some value that has to be accounted for. In particular, this would mean that the respective model would have to account for the utility derived from keeping information private. In a case where nothing is sold, the seller of a good,

[22] Compare the beginning of Section 2.

holding private information about the good, will derive some utility from the information staying private instead of deriving utility from selling the good. After all, we are only interested in the information subject, i.e. the seller, and the consequences of their information (not) flowing, so we should only be considering their level of utility from selling the good (which reveals the type to the buyer) or not selling the good.

Now one thing we can already say about the seller is that since they have been offering their good to potential buyers, the value they expect from selling at or above their reservation price will be greater than the value they derive from keeping the good. Therefore, not selling the good still leaves them worse off, even if we account for the fact that keeping the relevant information private is also valuable to the individual.

Moreover, it is not clear that the decision to share information for purely economic reasons maintains source control (as has been argued in section 2). If that were so, one might assume the reservation price of the sellers would change, since the additional value they would gain from maintaining source control would add to the amount of money they get in the respective utility function. However, as I have argued, sharing information for reasons that have nothing to do with privacy should not be considered to maintain privacy as source control. Therefore, considering the value of privacy as source control would not influence the reservation price of the seller, and even if it did, in some situations it will be high enough to produce the adverse selection effects we have seen above.

Considering source control as valuable to the individual would not change the model results, and the argument made above still stands.

While source control and bargaining power might sometimes coincide, the above argument has shown that in some situations this is not the case. When asymmetric information leads to market failure, bargaining power is diminished although source control will be maintained, since those individuals not selling do not initially reveal their type. These cases should be pressing proponents of the source control account to either work on accommodating them by further defining or modifying the source control account, or alternatively, to drop the claim that source control protects bargaining power.[23]

In cases where information flows, source control protects bargaining power to some extent. However, without further specifying the adequate foundation for source control, it might imply

[23] This conclusion would be different for the leeway version of the control account introduced and discussed in chapter 2. Since the seller has no way to credibly transmit information about the quality of her good in the model case, her options are limited and therefore leeway control is diminished. However, we have seen that this understanding of control is implausible for different reasons, which is why I focus on the source control account and its relation to bargaining power in this section.

that acting on solely economic (or at least privacy disregarding) reasons still maintains source control. This view is problematic for reasons discussed in section 2. Moreover, the fact that private information lowers bargaining power in some situations still stands. Recognizing this fact is sufficient for arguing that source control and bargaining power sometimes come apart.

In the next section I am going to consider a second kind of economic analysis, which reveals similar problems for the claim that source control protects bargaining power.

### 4.3 Mechanism Design

The theory of mechanism design is concerned with the problem of allocating resources among members of a given society according to some criteria (one of which is usually to achieve a Pareto efficient outcome), based on information that is initially only privately known by individuals in that society. The analysis of different allocation mechanisms allows us to compare them with respect to the chosen criteria, and select optimal mechanisms from them (compare Myerson 1989, 191).

A general result of mechanism design theory is that, under some constraints and for specific distributions of private information, no mechanisms will never guarantee Pareto efficiency (Myerson and Satterthwaite 1983). What exactly those constraints are, and how the result is relevant to the claim that source control protects bargaining power, will be at the center of this section. On a general note, in the model situations at least some information is initially only privately held by the individuals, and they are assumed to voluntarily reveal this information if it makes them better off, i.e. source control will be ensured at all times in the model setting.

First of all, let me introduce the general setting of Myerson and Satterthwaite. Their exemplary case focuses on a situation of *bilateral trade* in which one individual is selling an object and the other one wants to buy the object. Each of them has their own, privately known valuation of the object. The valuation of the other individual is only known as a random variable distributed according to some function.

Now the individuals are asked to state their valuations to some arbitrator. Based on those statements and some pre-defined mechanism, this arbitrator will allocate some outcome to each of them.[24] In our case, this outcome will involve the buyer paying a certain amount of money to the seller and receiving the object in return. In what is called a *direct bargaining mechanism,*

[24] Note that practically, the arbitrator invoked here is essentially identical to the mechanism. It does only what the predefined mechanism 'tells' it to do.

the individuals simply state their preferences, and without any other intermediate steps the outcome will be determined.

An important assumption is "that the valuations [of the respective individuals] have positive density over their respective intervals" (Myerson and Satterthwaite 1983, 273) and an intersection that is nonempty. That is to say, there is at least some probability that either both hold the same valuation, that the seller holds a higher valuation than the buyer, or that the buyer holds a higher valuation than the seller.

Now let me consider the constraints imposed on such a mechanism in the case considered here.

First, the mechanism should be *incentive compatible*, which is to say the mechanism should be such that an equilibrium strategy that maximizes the individuals' payoff is to tell the arbitrator their true valuation (Hurwicz 1972, 320). This constraint is important, since satisfying it will ensure that, once an individual has agreed to participating in the mechanism, they have no reason to withhold their (relevant) private information.

Second, interim *individual rationality* holds that on average individuals should expect to gain (or at least break even) from participating in the mechanism given the possible valuations of the other participant(s). Ex post individual rationality would be a stronger condition requiring that individuals gain or break even given all possible strategies of others in equilibrium, not just on average. This can be ensured by the mechanism possessing a *no-trade* option, which leaves individuals at their initial endowments (ibid. 1989, 144), so trade does not have to occur if it would make individuals worse off.

Finally, *Pareto efficiency* holds that there is no feasible allocation different from the one which was chosen, which would make all individuals at least as well or better off than they are in the chosen allocation and at least some individuals strictly better off than they are.

Now the theorem that follows from employing these constraints is that "in classical […] economic environments with a finite number of agents, there is no incentive compatible allocation mechanism, which possesses the no-trade option [i.e. is individually rational] and is Pareto-efficient" (ibid. 1989, 145).

Importantly for us at this point, this theorem states that the kinds of mechanisms just presented will never be realizing all possible gains that could be made from a certain transaction. However, we can also see that the individuals participating in the mechanism have private information about their valuation (i.e. their type) which they choose to reveal because said mechanism ensures it is advantageous to them, or else refrain from (honestly) participating in the mechanism. Either way, they are in control of whether or not they reveal information, so

privacy thus understood will be maintained.

Bargaining power on the other hand is at least somewhat diminished by the fact that not all possible gains from trade are realized. This result in turn follows at least partially from the fact that individuals can control the flow of (personal) information in the mechanism. If there were a way to know their valuations for sure without 'asking', the Pareto efficient allocation would clearly be much more straightforward to determine. Since the information is initially private, however, no mechanism satisfying the conditions stated above can reach Pareto efficiency.

Therefore, we have even more evidence that bargaining power and privacy as (source) control sometimes come apart, i.e. privacy as control does not always secure bargaining power.

Finally, in the next section I am going to briefly consider some important considerations for why we do not see some of the model effects in online environments where individuals 'pay' for the use of functional websites with their personal information.

### 4.4 Digital Information Collection

In the recently revived public discussion of privacy rights, the digital domain is of course one of the most important points of focus. With companies today collecting highly personal information, analyzing and even selling this information, privacy rights advocates are alerting the public again and again to stand up against such practices.

One important aspect of why the practices of digital companies like Google, Amazon, Meta (formerly known as Facebook) and the like are considered to be so problematic is of course the economic advantage they gain from collecting information about virtually all of society's individuals. However, in contrast to the Market for Lemons model considered above, the informational asymmetries present in the exchange of information between individuals using a website and the owners of the site, who receive the information, do not seem to lead to effects of adverse selection. In the following I am going to lay out possible reasons for why this is so.

First of all, let me consider how we could imagine the effects of adverse selection for the 'market' of websites, in which personal information can be considered the accepted currency. In such a market we could expect the amount of information individuals happily give away (representing their valuation) to vary according to the quality or functionality of a certain website. Owners, on the other hand, would have a certain reservation price, i.e. a minimum amount of personal information they want to get, below which they would not let an individual use the website. Now if we assume the users only get to know the quality of a website after having made their payment, it is easy to see how the dynamics of the adverse selection model

get going (compare section 3.2). When the proportion of bad websites is too high, individuals will only be willing to give away very little information, and consequently no good websites will be offered anymore.

Now from the everyday use of the internet it is obvious that this is not what is actually happening in online environments (compare Berendt, Günther, and Spiekermann 2005). Rather than observing the kind of market failure the Akerlof model would predict, we see that individuals give away private information more or less without considering what they get in return. One reason for this could be that individuals actually do not value their personal information, contrary to what many would state. Berendt et al. find some evidence for this, however the fact that individuals react with resentment when learning that information about them has been collected might show that there is something else going wrong in the actual behavior of individuals in online environments.

A second reason for the observed behavior of consumers in online environments might therefore be that they lack a substantial amount of knowledge about what happens with their data, how much it is worth to companies, and perhaps sometimes even that they are giving it away at all.[25]

Either way, it is highly unlikely individual customers consider the collection of information and their use of a website as a form of transaction at all. For providers of websites this means that they can basically 'harvest' information for very little costs and then sell or use this information to make profit in other ways.

One reason why we can still observe websites having useful functionality might be that providers need to offer at least some incentive for individuals to use their site rather than that of another provider. However, this means the functionality of websites should not be seen as payment for giving away information, but rather merely as an incentive. Individuals get paid for using a certain website instead of other sites, not for giving away information.

Now some privacy rights advocates urge users to demand payment for the information they provide as well, and see it as crucial that some awareness of the real worth of data is evoked in individuals. This is supposed to strengthen their position, and help to create a market for personal information, in which individuals can gain some profit rather than being totally

[25] Berendt et al. (2005) exclude the possibility that individuals are unaware that their information is being collected. At least in the EU this might be realistic, since providers of websites have to ask for consent when collecting information. However, only a small number of individuals might be assumed to be reading and understanding the consent forms, i.e. they do not ensure meaningful consent. Moreover, individuals might develop a habit of clicking either of the options in the consent form rather unconsciously in order to proceed to the website.

exploited. Awareness would lead to a market emerging where so far only exploitation exists.

However, given the current structure of internet use, or perhaps the internet in general, it is hard to see how this could be realized, and much less easy to see how the necessary change might be brought about.

The above argument can contribute to this task in a minimal way, by pointing out that online environments as we see them today demand a different kind of analysis from environments like the Market for Lemons model.

Moreover, in relation to state agencies or authorities collecting information, to talk of bargaining power might be misleading. For one thing, there is no immediate exchange between individuals and these authorities. Rather, the authorities justify their behavior on a promise of guaranteeing for example security for the subjects they are governing. Political justifications tend to be rather vague in this context, so knowing what is being offered in return for sharing information is hard to know. For another thing, evaluating whether or not individuals actually receive the kind of public good they have been promised is notoriously hard to assess, and perhaps receiving one good comes at the expense of other goods and freedoms the state should be providing. Assessing whether source control is actually valuable with regard to the goods social authorities should secure or promise to secure would require a much deeper analysis than Menges (2020a) provides, or can be provided here. For now, we can note that it is at least dubitable that this is so.

## Conclusion

As we have seen, privacy is not as clear of a concept as it would sometimes seem in public and political debates. Rather, in privacy debates many different issues converge. This makes it hard to isolate issues that concern privacy directly from those that concern important rights and freedoms, which might be closely connected to privacy.

For this reason, the first section of this essay was concerned with drawing some important conceptual distinctions that narrow down the context for the arguments in the remaining sections. Importantly, the focus here lays on merely defining the concept of privacy (rather than a right to privacy) with regard to personal information, that is not assumed to be inherently valuable. Limited-access accounts of privacy and control accounts of privacy have emerged to be the two main opponent views here.

In the recent discussion I have been following in the second section, two kinds of exemplary cases are used to test the intuitive plausibility of these two views. On the one hand, voluntary divulgence cases reveal that intuitions of proponents of either side of the debate are radically opposed. Proponents of limited-access accounts take privacy to be diminished in cases where individuals voluntarily share private information, while proponents of control accounts do not. As I have argued, these two stances can be at least somewhat reconciled by adopting a control view that allows for individuals to diminish their privacy even though they have full control. The important point is to adopt a control view that imposes conditions over and above having control. For example, individuals must have a certain degree of knowledge, consider at least some privacy-relevant reasons and so on.

On the other side of the debate, threatened-loss cases are supposed to show that control accounts of privacy are too narrow, since privacy is not diminished as long as information does not actually flow to others. However, the source control account of privacy can answer to this by taking control to be a determining factor only in case information flows at all.

The remainder of this section consisted in exploring the implications of new variations of the cases found in the literature. In doing so, I arrive at the conclusion that contextual variations of the cases matter a lot more than the recent debate might suggest. Therefore, a plausible account of privacy should find a way of incorporating contextual features. This means that authors in this recent philosophical debate should pay close attention to what other privacy scholars have found out about the context dependence of privacy.

In the third section I have turned my attention to a claim that unites both sides of the recent debate. Namely, I have argued against the claim that personal information is not being accessed by government agencies like the NSA or GCHQ, as long as no individual human agent accesses

the information. Reflecting upon the use of machines in our everyday life, it quickly becomes clear that attributing moral agency only to human individuals is highly implausible. We use tools and artifacts all the time to extend our capabilities and aid our agency. In case we closely depend on tools and artifacts to perform the tasks we want to perform, agency should be attributed to the whole system, rather than only the human part of it. Consequently, the machines and software used in government agencies is part of an extended agent, and the machines' accessing personal information *is* relevant to privacy. Even when no human agent reads the information, as soon as some part of the extended agent accesses it, it should be seen as accessed in a privacy-relevant sense.

If this argument is convincing, then many of the practices we can currently see in government agencies and private companies are a much bigger threat to privacy than the general public seems to be acknowledging. Companies accessing our data virtually anywhere we move on the internet are accessing personal information non-stop.

Finally, in the fourth section I shift the focus away from conceptual considerations, towards axiological considerations. Namely, I investigate the claim that privacy as source control protects bargaining power. When bargaining power is understood as the ability of individuals to raise their profits from bilateral trade, this claim will not be true in all situations. The Market-for-Lemons model gives a very clear illustration of why this is so. In situations where information that is relevant to the trade is asymmetrically distributed between buyers and sellers, the valuations of the less informed side will not be based on the true value of the good they get, but only on an average value they expect to get, given the probability they assign to the good being of a certain value. When the expected valuation of uninformed buyers is below the valuation of informed sellers of high-quality goods, the latter will not sell. As a consequence, at least some potential gains are not realized, or perhaps no trade occurs at all. Since bargaining power is the ability to increase one's share of technically feasible gains from trade, bargaining power is also diminished. However, since the relevant information did not flow, privacy is nonetheless maintained. Mechanism design theory provides another strong case for concluding that source control and bargaining power sometimes come apart. From the model considered above it becomes clear that, even though individuals have control over what information they share, no mechanism can guarantee their gains are maximized by doing so. Conversely, when they don't participate in the mechanism and keep their information private, their gains will also be less than optimal.

Moreover, mechanism design theory shows a similar result, where private information and control about when to reveal it truthfully leads to the problem that no mechanism determining

a social choice will be maximizing the individuals' welfare.

**References**

Adams, Frederick Ray, and Kenneth Aizawa. 2008. *The Bounds of Cognition*. Malden, Ma: Blackwell Publishing.

Akerlof, George. 1970. 'The Market for "Lemons": Quality Uncertainty and the Market Mechanism'. *The Quarterly Journal of Economics* 84 (3): 488–500.

Allen, Anita L. 1988. *Uneasy Access : Privacy for Women in a Free Society*. Totowa, N.J. : Rowman & Littlefield. http://archive.org/details/uneasyaccesspriv00alle.

Bentham, Jeremy. 2010. *The Panopticon Writings*. Edited by Božovič Miran. London: Verso Books.

Berendt, Bettina, Oliver Günther, and Sarah Spiekermann. 2005. 'Privacy in E-Commerce: Stated Preferences vs. Actual Behavior'. *Communications of the ACM* 48 (4): 101–6. https://doi.org/10.1145/1053291.1053295.

Bosch, Antal van den, Toine Bogers, and Maurice de Kunder. 2016. 'Estimating Search Engine Index Size Variability: A 9-Year Longitudinal Study'. *Scientometrics* 107 (2): 839–56. https://doi.org/10.1007/s11192-016-1863-z.

Cappelen, Herman. 2012. *Philosophy without Intuitions*. Oxford: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199644865.001.0001.

Chalmers, David J. 2014. 'Intuitions in Philosophy: A Minimal Defense'. *Philosophical Studies* 171 (3): 535–44. https://doi.org/10.1007/s11098-014-0288-x.

Clark, Andy, and David Chalmers. 1998. 'The Extended Mind'. *Analysis* 58 (1): 7–19.

Cohen, Alma, and Peter Siegelman. 2010. 'Testing for Adverse Selection in Insurance Markets'. *Journal of Risk and Insurance* 77 (1): 39–84. https://doi.org/10.1111/j.1539-6975.2009.01337.x.

Cutler, David M., and Richard J. Zeckhauser. 1998. 'Adverse Selection in Health Insurance'. *Forum for Health Economics & Policy* 1 (1). https://doi.org/10.2202/1558-9544.1056.

DeCew, Judith Wagner. 1986. 'The Scope of Privacy in Law and Ethics'. *Law and Philosophy* 5 (2): 145–73. https://doi.org/10.2307/3504687.

———. 1997. *In Pursuit of Privacy: Law, Ethics, and the Rise of Technology*. Ithaca: Cornell University Press.

Foucault, Michel. 1995. *Discipline & Punish: The Birth of the Prison*. Vintage Books. New York: Vintage Books.

Frankfurt, Harry G. 1969. 'Alternate Pssobilities and Moral Responsibility'. In *Free Will*, edited by Gary Watson, 2nd ed., 167–76. New York: Oxford University Press, 2003.

———. 1971. 'Freedom of the Will and the Concept of a Person'. *The Journal of Philosophy* 68 (1): 5. https://doi.org/10.2307/2024717.

Gavison, Ruth E. 1980. 'Privacy and the Limits of Law'. *The Jale Law Journal* 89 (3): 421–71.

Gilbert, Margaret. 1990. 'Walking Together: A Paradigmatic Social Phenomenon'. *Midwest Studies in Philosophy* 15 (1): 1–14. https://doi.org/10.1111/j.1475-4975.1990.tb00202.x.

Gold, Natalie, and Robert Sugden. 2007. 'Collective Intentions and Team Agency'. *The Journal of Philosophy* 104 (3): 109–37.

Goldman, Alvin I. 2015. Review of *Review of* A Naturalistic Epistemology: Selected Papers, by Hilary Kornblith, June. https://ndpr.nd.edu/reviews/a-naturalistic-epistemology-selected-papers/.

Heersmink, Richard. 2017. 'Distributed Cognition and Distributed Morality: Agency, Artifacts and Systems'. *Science and Engineering Ethics* 23 (2): 431–48. https://doi.org/10.1007/s11948-016-9802-1.

Hurwicz, Leonid. 1972. 'On Informationally Decentralized Systems'. In *Decision and Organization: A*

*Volume in Honor of Jacob Marschak*, edited by C. B. McGuire and Roy Radner, 297–336. Studies in Mathematical and Managerial Economics, V. 12. Amsterdam: North-Holland Pub. Co.

Hutchins, Edwin. 2014. 'The Cultural Ecosystem of Human Cognition'. *Philosophical Psychology* 27 (1): 34–49. https://doi.org/10.1080/09515089.2013.830548.

Inness, Julie C. 1992. *Privacy, Intimacy, and Isolation*. New York: Oxford University Press. http://search.ebscohost.com/login.aspx?direct=true&db=nlebk&AN=55481&site=ehost-live.

Kilgarriff, Adam. 2007. 'Googleology Is Bad Science'. *Computational Linguistics* 33 (1): 147–51. https://doi.org/10.1162/coli.2007.33.1.147.

Knobe, Joshua, and Shaun Nichols. 2017. 'Experimental Philosophy'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2017. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/win2017/entries/experimental-philosophy/.

Ledyard, John O. 1989. 'Incentive Compatibility'. In *Allocation, Information and Markets*, edited by John Eatwell, Murray Milgate, and Peter Newman, 141–51. London: Palgrave Macmillan UK. https://doi.org/10.1007/978-1-349-20215-7_15.

Lundgren, Björn. 2020. 'A Dilemma for Privacy as Control'. *The Journal of Ethics* 24 (2): 165–75. https://doi.org/10.1007/s10892-019-09316-z.

———. 2021. 'Confusion and the Role of Intuitions in the Debate on the Conception of the Right to Privacy'. *Res Publica* 27 (4): 669–74. https://doi.org/10.1007/s11158-020-09495-9.

MacAskill, Ewen, Gabriel Dance, Feilding Cage, Greg Chen, and Nadja Popovich. 2013. 'NSA Files Decoded: Edward Snowden's Surveillance Revelations Explained'. The Guardian. 1 November 2013. http://www.theguardian.com/world/interactive/2013/nov/01/snowden-nsa-files-surveillance-revelations-decoded.

Macnish, Kevin. 2018. 'Government Surveillance and Why Defining Privacy Matters in a Post-Snowden World'. *Journal of Applied Philosophy* 35 (2): 417–32. https://doi.org/10.1111/japp.12219.

Mainz, Jakob Thrane, and Rasmus Uhrenfeldt. 2021a. 'Too Much Info: Data Surveillance and Reasons to Favor the Control Account of the Right to Privacy'. *Res Publica* 27 (2): 287–302. https://doi.org/10.1007/s11158-020-09473-1.

———. 2021b. 'Privacy Rights, and Why Negative Control Is Not a Dead End: A Reply to Munch and Lundgren'. *Res Publica*, July. https://doi.org/10.1007/s11158-021-09524-1.

Marmor, Andrei. 2015. 'What Is the Right to Privacy?' *Philosophy & Public Affairs* 43 (1): 3–26. https://doi.org/10.1111/papa.12040.

Mas-Colell, Andreu, and Michael D. Whinston. 1995. *Microeconomic Theory*. Oxford Student Edition. New York, NY [u.a.]: Oxford Univ. Press.

McCloskey, H. J. 1980. 'Privacy and the Right to Privacy'. *Philosophy* 55 (211): 17–38. https://doi.org/10.1017/S0031819100063725.

McDonald, Aleecia M., and Lorrie Faith Cranor. 2008. 'The Cost of Reading Privacy Policies'. *I/S: A Journal of Law and Policy for the Information Society* 4 (3): 543–68.

McKenna, Michael, and D. Justin Coates. 2020. 'Compatibilism'. Edited by Edward N. Zalta. *The Stanford Encyclopedia of Philosophy (Spring 2020 Edition)*. https://plato.stanford.edu/archives/spr2020/entries/compatibilism/.

Menges, Leonhard. 2020a. 'Did the NSA and GCHQ Diminish Our Privacy? What the Control Account Should Say'. *Moral Philosophy and Politics* 7 (1): 29–48. https://doi.org/10.1515/mopp-2019-0063.

———. 2020b. 'A Defense of Privacy as Control'. *The Journal of Ethics*, October, 385–402. https://doi.org/10.1007/s10892-020-09351-1.

———. forthcoming. ‘Three Control Views on Privacy’. *Social Theory and Practice*. https://philpapers-org.uaccess.univie.ac.at/rec/MENTCV.

Myerson, Roger B. 1989. ‘Mechanism Design’. In *Allocation, Information and Markets*, edited by John Eatwell, Murray Milgate, and Peter Newman, 191–206. London: Palgrave Macmillan UK. https://doi.org/10.1007/978-1-349-20215-7_20.

Myerson, Roger B, and Mark A Satterthwaite. 1983. ‘Efficient Mechanisms for Bilateral Trading’. *Journal of Economic Theory* 29 (2): 265–81. https://doi.org/10.1016/0022-0531(83)90048-0.

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

Parent, W. A. 1983a. ‘Privacy, Morality, and the Law’. *Philosophy & Public Affairs* 12 (4): 269–88.

———. 1983b. ‘Recent Work on the Concept of Privacy’. *American Philosophical Quarterly* 20 (4): 341–55.

Puelz, Robert, and Arthur Snow. 1994. ‘Evidence on Adverse Selection: Equilibrium Signaling and Cross-Subsidization in the Insurance Market’. *Journal of Political Economy* 102 (2): 236–57. https://doi.org/10.1086/261930.

Pust, Joel. 2019. ‘Intuition’. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2019. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/sum2019/entries/intuition/.

Rössler, Beate. 2001. *Der Wert Des Privaten*. First. Frankfurt am Main: Suhrkamp Verlag.

Samuelson, William. 1984. ‘Bargaining under Asymmetric Information’. *Econometrica* 52 (4): 995–1005. https://doi.org/10.2307/1911195.

Scanlon, Thomas. 1975. ‘Thomson on Privacy’. *Philosophy & Public Affairs* 4 (4): 315–22.

Searle, John R., ed. 2002. ‘Collective Intentions and Actions’. In *Consciousness and Language*, 90–105. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511606366.007.

Shapiro, Lawrence, and Shannon Spaulding. 2021. ‘Embodied Cognition’. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2021. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/win2021/entries/embodied-cognition/.

Sherman, Justin. 2021a. ‘Data Brokers and Sensitive Data on U.S. Individuals’, 16.

———. 2021b. ‘Data Brokers Know Where You Are—and Want to Sell That Intel’. *Wired*, 2021. https://www.wired.com/story/opinion-data-brokers-know-where-you-are-and-want-to-sell-that-intel/.

Shoemaker, David. 2015. *Responsibility from the Margins*. First edition. Oxford, United Kingdom: Oxford University Press.

Solove, Daniel J. 2008. *Understanding Privacy*. Cambridge, Mass: Harvard University Press.

Tavani, Herman T. 1999. ‘Informational Privacy, Data Mining, and the Internet’. *Ethics and Information Technology* 1 (2): 137–45. https://doi.org/10.1023/A:1010063528863.

Thomson, Judith Jarvis. 1975. ‘The Right to Privacy’. *Philosophy & Public Affairs* 4 (4): 295–314.

Tuomela, Raimo, and Kaarlo Miller. 1985. ‘We-Intentions and Social Action’. *Analyse & Kritik* 7 (1): 26–43. https://doi.org/10.1515/auk-1985-0102.

Wolff, Ran. 2015. ‘Emergent Privacy’. *The Journal of Philosophy* 112 (3): 141–58.

**Abstract (English)**

In the recent past privacy has gained a lot of attention of the general public, mostly as a consequence of the so-called "Snowden Revelations", offering a strong picture of US and UK governing bodies surveilling most of today's internet traffic. The question whether this surveillance is invading individuals' privacy presents a starting point for the philosophical discussion forming the basis for this work. If privacy is limited access to personal information, then surveillance seems acceptable, as long as no human person reads the information that is being collected. If privacy is control over personal information, then surveillance diminished privacy.

This master thesis contributes to the philosophical debate around privacy in multiple ways. First, by revisiting two central kinds of cases form the literature, and sometimes presenting novel modifications of them, the plausibility of two kinds of control accounts is tested, and some remaining problems are pointed out. Most importantly, cases in which individuals voluntarily share information require some limiting conditions over and above having control over information.

Returning to the initial starting point, a common assumption of both sides of the recent debate is identified: Both sides assume that information is only accessed in the relevant sense, when human individuals read it, see it, or process it in some other way. By drawing on literature about extended agency, it becomes clear that this assumption is quite implausible. Humans use tools all the time and this is seldom seen to diminish their agency, but rather to extend it beyond their prior capacities. Therefore, in cases like digital surveillance, the nonhuman parts of extended agents are to be seen as accessing personal information before the individual humans have received the processed information. They can thus diminish privacy without human agents processing the relevant personal information within the boundaries of their body.

Finally, this thesis also evaluates a claim that has not only been implicit in much of the public debate around privacy, but can also be found in the philosophical literature. Namely, that privacy, understood as control about personal information, also protects the economic welfare of individuals because it protects their bargaining power. Drawing on economic analysis, however, it becomes clear that keeping information private is seldom seen as advantageous by the economist. Rather, asymmetries in relevant information often lead to market failure, i.e. feasible profits being diminished. Still, per definition, privacy is maintained in such situations.

**Abstract (Deutsch)**

In der jüngsten Vergangenheit hat das Thema Privatsphäre in der öffentlichen Debatte viel Aufmerksamkeit erfahren. Vor allem die der sogenannten „Snowden Enthüllungen", die ein starkes Bild davon lieferten wie staatliche Agenturen, wie NSA und GCHQ, heute den Datenverkehr im Internet überwachen. Die Frage, ob diese Praktiken die Privatsphäre der überwachten Individuen verletzen dient als Ausgangspunkt für diese Arbeit. Wird Privatsphäre als *limitierter Zugang* zu persönlichen Informationen verstanden, dann scheint digitale Überwachung akzeptabel, solange keine Personen die Information lesen. Wird Privatsphäre als *Kontrolle* über persönliche Informationen verstanden, dann verletzt die Überwachung und vor allem das Sammeln solcher Informationen die Privatsphäre der Individuen.

Diese Masterarbeit trägt auf mehrere Arten zur philosophischen Diskussion über Privatsphäre bei. Zunächst wird genauer beleuchtet ob Definitionen von Privatsphäre als Kontrolle wirklich unsere Intuitionen zu zwei zentralen Arten von Beispiel-Fällen aus der Literatur einfangen können. Aufbauend auf diesen Beispielen und Variationen dieser Beispiele werden Probleme für Kontroll-Definitionen aufgezeigt. Eine bedeutende Schlussfolgerung ist, dass das freiwillige Teilen von Information nur dann Privatsphäre schützt, wenn Individuen Voraussetzungen erfüllen, die über bloße Kontrolle hinausgehen.

Im nächsten Abschnitt kehren wir zurück zum Ausgangspunkt der Debatte, und betrachten eine gemeinsame Annahmen der beiden Seiten in der aktuellen philosophischen Debatte: Auf persönliche Information wird nur dann in einem relevanten Sinn *zugegriffen*, wenn menschliche Individuen diese lesen, sichten, oder in einer anderen Weise verarbeiten. Auf Literatur zur erweiterten Handlungsfähigkeit (extended agency) zurückgreifend wird klar dargestellt, dass diese Annahme nicht überzeugend ist. Werkzeuge sind ein ständiger Begleiter des Menschen. Die Verwendung dieser Werkzeuge wird jedoch nicht als beschränkender Faktor für menschliches Handeln gesehen, sondern als erweiternder. In Fällen von digitaler Überwachung sollten die verwendeten Maschinen demnach nicht als getrennt betrachtet werden, sondern als Teil eines erweiterten Akteurs, der aus menschlichen und technologischen Teilen besteht. Diese Akteure können in einem relevanten Sinn auf Informationen zugreifen, selbst wenn keine menschlichen Individuen die Information innerhalb der Grenzen ihrer Körper verarbeiten.

Schließlich soll auch eine weitere Annahme genau betrachtet werde, die in der öffentlichen Debatte implizit angenommen wird, aber auch in der philosophischen gefunden werden kann: Privatsphäre, als Kontrolle verstanden, schütze die Verhandlungsmacht von

Individuen, und somit ihr ökonomisches Wohlergehen. Durch einen Blick auf ökonomische Analysen, die sich mit Asymmetrien und Absenzen von relevanter Information befassen, wird klar, dass private Information in der Ökonomie selten als vorteilhaft gesehen wird. Vielmehr führen diese oft zu Marktversagen, also zur Verminderung der möglichen Profite. Per Definition ist die Privatsphäre der Individuen in solchen Fällen jedoch gewahrt.